

HEALTHINF 2008

International Conference on
Health Informatics

Proceedings

Volume I

Funchal, Madeira - Portugal · 28 - 31 January, 2008

TECHNICAL CO-SPONSORSHIP BY



CO-ORGANIZED BY



IN COOPERATION WITH



HEALTHINF 2008

Proceedings of the
First International Conference
on Health Informatics

Volume 1

Funchal, Madeira - Portugal

January 28 – 31, 2008

Co-organized by
**INSTICC – Institute for Systems and Technologies of Information,
Control and Communication**
and
UMA – University of Madeira

Technical Co-sponsorship by
IEEE EMB – Engineering in Medicine and Biology Society
and
WfMC – Workflow Management Coalition

In Cooperation with
AAAI – Association for the Advancement of Artificial Intelligence

Copyright © 2008 INSTICC – Institute for Systems and Technologies of
Information, Control and Communication
All rights reserved

Edited by Luís Azevedo and Ana Rita Londral

Printed in Portugal

ISBN: 978-989-8111-16-6

Depósito Legal: 268554/07

<http://www.healthinf.org>

secretariat@healthinf.org

HEALTHINF is part of BIOSTEC – International Joint Conference on Biomedical Engineering
Systems and Technologies

BRIEF CONTENTS

INVITED SPEAKERS.....	IV
SPECIAL SESSION CHAIRS	IV
ORGANIZING AND STEERING COMMITTEES	V
PROGRAM COMMITTEE	VI
AUXILIARY REVIEWERS	VI
SELECTED PAPERS BOOK	VII
OFFICIAL CARRIER.....	VII
FOREWORD.....	IX
CONTENTS.....	XI

INVITED SPEAKERS

Sérgio Cerutti

Polytechnic University of Milan

Italy

Kevin Warwick

University of Reading

U.K.

Fernando Henrique Lopes da Silva

University of Amsterdam

The Netherlands

Vipul Kashyap

Partners HealthCare System, Clinical Informatics R&D

U.S.A.

David Hall

Research Triangle Institute in North Carolina

U.S.A.

Albert Cook

University of Alberta, Faculty of Rehabilitation Medicine

Canada

SPECIAL SESSION CHAIRS

SPECIAL SESSION ON KNOWLEDGE DISCOVERY AND DECISION SUPPORT SYSTEMS

IN HEALTH INFORMATION SYSTEMS

Domenico Conforti, Università della Calabria, Italy

Manolis Tsiknakis, Foundation for Research and Technology - Hellas, Institute of Computer Science, Greece

Joerg Habetha, Philips Technologie GmbH, Research Group "Medical Signal Processing", Germany

Joël Bacquet, European Commission, Information Society & Media DG, Belgium

ORGANIZING AND STEERING COMMITTEES

CONFERENCE CO-CHAIRS

Ana Fred, IST- Technical University of Lisbon, Portugal
Joaquim Filipe, INSTICC / Polytechnic Institute of Setúbal, Portugal
Hugo Gamboa, Telecommunications Institute / Portugal
Jorge Cardoso, University of Madeira - UMA / Madeira, Portugal

PROGRAM CO-CHAIRS

Luís Azevedo, IST- Technical University of Lisbon / Portugal
Ana Rita Londral, ANDITEC/ Portugal

LOCAL CHAIR

Paulo Sampaio, University of Madeira - UMA, Portugal

PROCEEDINGS PRODUCTION

Paulo Brito, INSTICC, Portugal
Marina Carvalho, INSTICC, Portugal
Helder Coelhas, INSTICC, Portugal
Vera Coelho, INSTICC, Portugal
Andreia Costa, INSTICC, Portugal
Bruno Encarnação, INSTICC, Portugal
Bárbara Lima, INSTICC, Portugal
Vitor Pedrosa, INSTICC, Portugal
Vera Rosário, INSTICC, Portugal
Mónica Saramago, INSTICC, Portugal

CD-ROM PRODUCTION

Paulo Brito, INSTICC, Portugal

GRAPHICS PRODUCTION AND WEB DESIGNER

Marina Carvalho, INSTICC, Portugal

SECRETARIAT AND WEBMASTER

Marina Carvalho, INSTICC, Portugal

PROGRAM COMMITTEE

Osman Abul, TOBB University of Economics and Technology, Turkey

Arnold Baca, University of Vienna, Austria

Iddo Bante, University of Twente, the Netherlands

Jyh-Horng Chen, National Taiwan University, Taiwan, Province of China

Kay Connelly, Indiana University, U.S.A.

Amar Das, Stanford University, U.S.A.

Adrie Dumay, TNO, the Netherlands

Eduardo B. Fernandez, Florida Atlantic University, U.S.A.

Alexandru Floares, Oncological Institute Cluj-Napoca, Romania

Jose Fonseca, Universidade Nova de Lisboa, Portugal

Toshio Fukuda, Nagoya University, Japan

M. Chris Gibbons, Johns Hopkins University, U.S.A.

David Greenhalgh, University of Strathclyde, U.K.

Jin-Kao Hao, University of Angers, France

Tin Ho, Bell Labs, Alcatel-Lucent, U.S.A.

John H. Holmes, University of Pennsylvania School of Medicine, U.S.A.

Chun-Hsi Huang, University of Connecticut, U.K.

Benjamin Jung, Empirica - Gesellschaft für Kommunikatons - und Technologieforschung mbH, Germany

Stavros Karkanis, Technological Educational Institute of Lamia, Greece

Andreas Kerren, Växjö University, Sweden

Georgios Kontaxakis, Universidad Politécnica de Madrid, Spain

Nigel Lovell, University of New South Wales, Australia

Andrew Mason, University of Auckland, New Zealand

Boleslaw Mikolajczak, University of Massachusetts Dartmouth, U.S.A.

Ahmed Morsy, Cairo University, Egypt

Laura Roa, University of Seville, Spain

Jean Roberts, University of Central Lancashire, U.K.

Joel Rodrigues, Institue of Telecommunications / University of Beira Interior, Portugal

Marcos Rodrigues, Sheffield Hallam University, U.K.

George Sakellaropoulos, University of Patras, Greece

Meena Kishore Sakharkar, Nanyang Technological University, Singapore

Ovidio Salvetti, ISTI-CNR, Italy

Nickolas Sapidis, University of the Aegean, Greece

Sepe Sehati, Oxford Brookes University, U.K.

Boris Shishkov, IICREST / University of Twente, the Netherlands

Iryna Skrypnyk, University of Jyvaskyla, Finland

John Stankovic, University of Virginia, U.S.A.

Ron Summers, Loughborough University, U.K.

Adrian Tkacz, BioMedical-IT, Poland

Aristides Vagelatos, RACTI, Greece

Athanasios Vasilakos, University of Western Macedonia, Greece

Taieb Znati, University of Pittsburgh, U.S.A.

AUXILIARY REVIEWERS

Sara Colantonio, ISTI-CNR, Italy

Teduh Dirgahayu, University of Twente, the Netherlands

Ana Sofia Fernandes, Faculdade de Ciências e Tecnologia da Universidade Nova de Lisboa, Portugal

Hailiang Mei, University of Twente, the Netherlands

Davide Moroni, ISTI-CNR, Italy

John Sarivougioukas, General Hospital of Athens "G. Gennimatas", Greece

SELECTED PAPERS BOOK

A number of selected papers presented at HEALTHINF 2008 will be published by Springer, in a book entitled Biomedical Engineering Systems and Technologies. This selection will be done by the conference Chair and program co-chairs, among the papers actually presented at the conference, based on a rigorous review by the BIOSTEC 2008 program committee members.

OFFICIAL CARRIER



FOREWORD

This volume contains the proceedings of the *First International Conference on Health Informatics* (HEALTHINF 2008), organized by the Institute for Systems and Technologies of Information Control and Communication (INSTICC) and the University of Madeira, technically co-sponsored by the IEEE Engineering in Medicine and Biology Society (EMB) and the Workflow Management Coalition (WfMC), in cooperation with AAAI.

The purpose of the *International Conference on Health Informatics* is to bring together researchers and practitioners interested in the application of information and communication technologies (ICT) to healthcare and medicine in general and to the specialized support to persons with special needs in particular.

Databases, networking, graphical interfaces, intelligent decision support systems and specialized programming languages are just a few of the technologies currently used in medical informatics. Mobility and ubiquity in healthcare systems, standardization of technologies and procedures, certification, privacy are some of the issues that medical informatics professionals and the ICT industry in general need to address in order to further promote ICT in healthcare. In the case of medical rehabilitation and assistive technology the use of ICT has had important results in the enhancement of the quality of life, contributing to a full integration of all citizens in the societies they are also part of. HEALTHINF is a forum for debating all these aspects. Furthermore, this conference is also a meeting place for those interested in understanding the human and social implications of technology, not only in healthcare systems but in other aspects of human-machine interaction such as accessibility issues.

HEALTHINF is one of three integrated conferences that are co-located and constitute the International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC). The other two component conferences are BIOSIGNALS (International Conference on Bio-inspired Systems and Signal Processing) and BIODEVICES (International Conference on Biomedical Electronics and Devices).

The joint conference, BIOSTEC, has received 494 paper submissions from more than 40 countries in all continents. 65 papers were published and presented as full papers, i.e. completed work (8 pages/30' oral presentation), 189 papers reflecting work-in-progress or position papers were accepted for short presentation, and another 86 contributions were accepted for poster presentation. These numbers, leading to a "full-paper" acceptance ratio below 14% and a total oral paper presentations acceptance ratio below 52%, show the intention of preserving a high quality forum for the next editions of this conference.

The conference included a panel and six invited talks delivered by internationally distinguished speakers, namely: Sergio Cerutti, Kevin Warwick, F. H. Lopes da Silva, Vipul Kashyap, David Hall and Albert Cook. Their participation has positively contributed to reinforce the overall quality of the Conference and to provide a deeper understanding of the field of Biomedical Engineering Systems and Technologies.

FOREWORD (CONT.)

The proceedings of the conference will be indexed by several major indices including DBLP, INSPEC and ISI-Proceedings and it will also be submitted for indexing to EI. A book with the revised versions of a short list of selected papers from the conference will be published by Springer-Verlag in the new CS book series: Communications in Computer and Information Science (CCIS). Additionally, a special issue of the IEEE Transactions on Biomedical Circuits and Systems will be edited based on the very best papers of the conference.

The program for this conference required the dedicated effort of many people. Firstly, we must thank the authors, whose research and development efforts are recorded here. Secondly, we thank the members of the program committee and the additional reviewers for their diligence and expert reviewing. Thirdly, we thank the keynote speakers for their invaluable contribution and for taking the time to synthesise and prepare their talks. Fourthly, we thank the program chairs, Luis Azevedo and Ana Rita Londral, whose collaboration was much appreciated. Finally, special thanks to all the members of the INSTICC team, especially Marina Carvalho at the conference secretariat, and the local organising committee from the University of Madeira, especially Jorge Cardoso and Paulo Sampaio, whose collaboration was fundamental for the success of this conference.

This year, the organization will distribute two paper awards at the conference closing session: the best paper award and the best student paper award. The decision was mainly based on the paper classifications provided by the Program Committee.

We wish you all an exciting conference and an unforgettable stay in the lovely island of Madeira. We hope to meet you again next year for the 2nd HEALTHINF, details of which are available at <http://www.healthinf.org>.

Joaquim Filipe

INSTICC/Polytechnic Institute of Setúbal

CONTENTS

INVITED SPEAKERS

KEYNOTE LECTURES

- MULTIVARIATE, MULTIORGAN AND MULTISCALE INTEGRATION OF INFORMATION
IN BIOMEDICAL SIGNAL PROCESSING IS-5
Sergio Cerutti
- OUTTHINKING AND ENHANCING BIOLOGICAL BRAINS IS-9
Kevin Warwick
- ANALYSIS AND MODELS OF BRAIN EPILEPTIC ACTIVITIES IS-21
Fernando Henrique Lopes da Silva
- FROM THE BENCH TO THE BEDSIDE - The Role of Semantics in Enabling the Vision of
Translational Medicine IS-23
Vipul Kashyap
- THE CANCER INFORMATICS ECOSYSTEM - A Case Study in the Accretion of Federated Systems
based on Service Oriented Architectures, Semantic Integration and Computing Grids IS-25
David Hall
- ICT AND PERSONS WITH DISABILITIES - The Solution or the Problem? IS-27
Albert Cook

PAPERS

FULL PAPERS

- HOW TO EVALUATE HUMAN FACTORS AFFECTING WIRELESS BIOMEDICAL SENSORS -
Identifying Aspects of Patient Acceptance based on a Preliminary Clinical Trial 3
Rune Fensli and Egil Boisen
- A COMMUNITY-CENTERED ARCHITECTURE FOR THE DEPLOYMENT OF UBIQUITOUS
TELEMEDICINE SYSTEMS 9
Federico Cabitza, Marco P. Locatelli and Carla Simone
- PRE-DIABETES - An Informatics Research Agenda 17
Barbara M. Hayes and William Aspray
- PO@HEALTH - A Medical Training Telemedicine Case Study based on Ultrasound Images over an Hybrid
Power Line Network 28
Alécio Pedro Delazari Binotto, Flávio Ávila, Carlos Eduardo Pereira, Cirano Iochpe, Ilias Sachpazidis and Georgios Sakas
- ON EXTRACTION OF NUTRITIONAL PATTERNS (NPS) USING FUZZY ASSOCIATION RULE
MINING 34
M. Sulaiman Khan, Maybin Mueyba and Frans Coenen
- TRUSTED SMS - A Novel Framework for Non-repudiable SMS-based Processes 43
Antonio Grillo, Alessandro Lentini, Gianluigi Me and Giuliano Rulli

PROTECTING PRIVACY IN MEDICAL DATABASES - Efficient Local Generation of System-Wide Unique Health IDs <i>Peter Schartner and Martin Schaffer</i>	51
ON PSEUDONYMOUS HEALTH REGISTERS - While They Work As Intended, They Are Still Controversial in Norway <i>Herbjørn Andresen</i>	59
AUTHENTICITY AND INTEGRITY OF PORTABLE ELECTRONIC HEALTH RECORDS <i>Chung-Yueh Lien, Chia-Hung Hsiao, Lu-Chou Huang and Tsair Kao</i>	67
AUTHENTICATION OF PROFESSIONALS IN THE RTS E-HEALTH SYSTEM <i>André Zúquete, Helder Gomes and João Paulo Silva Cunha</i>	72
COMMUNICATION OF MEDICAL INFORMATION USING AGENTS <i>John McGrory, Frank Clarke, Jane Grimson and Peter Gaffney</i>	81
AN ONTOLOGICAL APPROACH TO REPRESENTING AND REASONING WITH TEMPORAL CONSTRAINTS IN CLINICAL TRIAL PROTOCOLS <i>Ravi D. Shankar, Susana B. Martins, Martin J. O'Connor and Amar K. Das</i>	87
MEDLINE ABSTRACTS CLASSIFICATION - Average-based Discrimination for Noun Phrases Selection and Weighting Applied to Categorization of MEDLINE Abstracts <i>Fernando Ruiz-Rico, José Luis Vicedo and María-Consuelo Rubio-Sánchez</i>	94
TOWARDS ON-DEMAND BIOMEDICAL KNOWLEDGE EXTRACTION <i>Vincenzo Lanza, M. Ignazia Cascio and Chun-Hsi Huang</i>	102
A BPM-BASED MOBILE U-HEALTH SERVICE FRAMEWORK <i>Dongsoo Han, In-Young Ko, Sungjoon Park, Minkyu Lee and Suntae Jung</i>	110
PROCESS MINING IN HEALTHCARE - A Case Study <i>R. S. Mans, M. H. Schonenberg, M. Song, W. M. P. van der Aalst and P. J. M. Bakker</i>	118
HEALTHCARE RISK MODELING FOR MEDICAID PATIENTS - The Impact of Sampling on the Prediction of High-Cost Patients <i>Sai T. Moturu, Huan Lin and William G. Johnson</i>	126
MEDICAL IMAGE UNDERSTANDING THROUGH THE INTEGRATION OF CROSS-MODAL OBJECT RECOGNITION WITH FORMAL DOMAIN KNOWLEDGE <i>Manuel Möller, Michael Sintek, Paul Buitelaar, Saikat Mukherjee, Xiang Sean Zhou and Jörg Freund</i>	134
SOFTWARE AGENTS REPRESENTING MEDICAL GUIDELINES <i>John McGrory, Frank Clarke, Jane Grimson and Peter Gaffney</i>	142
FORMAL ANALYSIS OF INTELLIGENT AGENTS FOR MODEL-BASED MEDICINE USAGE MANAGEMENT <i>Mark Hoogendoorn, Michel Klein, Zulfiqar Memon and Jan Treur</i>	148
 POSTERS	
INTELLIGENT SYSTEM FOR ASSISTING ELDERLY PEOPLE AT HOME <i>Juan Pascual, Miguel A. Sanz-Bobi and David Contreras</i>	159
ASSISTING WELLBEING - The Challenges of using Technology to Improve Wellbeing in Older Adults <i>Douglas Millward and Wendy Nicholls</i>	164

TELEHEMATOLOGY - ICT Solution of a Shared Digital Image Repository <i>Daniel Schwarz, Miroslav Penka, Ladislav Dusek and Petr Brabec</i>	168
A SURVEY OF INTEROPERABILITY IN E-HEALTH SYSTEMS - The European Approach <i>Marcos Da Silveira, Nicolas Guelfi, Jerry-David Baldacchino, Pierre Plumer, Marc Seil and Anke Wienecke</i>	172
INTEROPERABILITY IN SMART HOME MIDDLEWARE - The MPOWER Project <i>Sten Hanke and Thomas Fuxreiter</i>	176
WHO SHOULD ACCESS ELECTRONIC PATIENT RECORDS <i>A. Ferreira, L. Antunes, C. Pinho, C. Sá, E. Mendes, E. Santos, F. Silva, F. Sousa, F. Gomes, F. Abreu, F. Mota, F. Aguiar, F. Faria, F. Macedo, S. Martins and R. Cruz-Correia</i>	182
CLASSIFICATION OF PREMATURE VENTRICULAR BEAT USING BAYESIAN NETWORKS <i>Lorena S. C. de Oliveira, Rodrigo V. Andreão and Mário Sarcinelli-Filho</i>	186
USING A COMPLEX NUMERICAL AND MULTIMEDIA DATABASE IN INTERNAL MEDICINE - Pre-release Evaluations <i>Liana Stanescu, Dumitru Burdescu, Cosmin Stoica-Spabin, Anca Ion and Dorin Stanescu</i>	192
A LIFE SUPPORT NETWORK FOR ELDERLY PEOPLE LIVING IN A RURAL AREAS <i>Bayme Abaydulla, Jun Sasaki, Michiru Tanaka, Keizō Yamada and Yutaka Funyu</i>	196
PERFORMANCE ANALYSIS OF WIRELESS SYSTEMS IN TELEMEDICINE - Hybrid Network for Telemedicine with Satellite and Terrestrial Wireless Links <i>M. Luglio and F. Zampognaro</i>	200
DEALING WITH THE COMPLEXITIES WHEN IMPLEMENTING INFORMATION SECURITY PRACTICES IN HEALTHCARE ORGANIZATIONS <i>Heitor Gottberg and Ivan Torres Pisa</i>	205
HEALTH CARE PROCESS MODELLING AND IMPROVEMENT <i>Nadja Damij and Janez Grad</i>	209
A FIRST APPROACH FOR A REGIONAL WIDE VEPR <i>Pedro Vieira-Marques, Arthur Cunha, Luís Antunes, Ricardo Cruz-Correia and Altamiro Costa-Pereira</i>	215
FOR HOW LONG IS DATA FROM PREVIOUS ADMISSIONS ACCESSED BY HOSPITAL DOCTORS? <i>Ricardo João Cruz-Correia and Altamiro Costa-Pereira</i>	219
USING EXPLANATION FACILITIES IN HEALTHCARE EXPERT SYSTEMS <i>Keith Darlington</i>	223
MDFLUXO: OPHTHALMOLOGY EDUCATION WITH A PDA EFFICACY AND USABILITY EVALUATION <i>Vladimir Camelo Pinto, Thiago Martini da Costa, Marcelo Contardo Moscoso Naveira, Daniel Sigulem, Paulo Schor and Ivan Torres Pisa</i>	227
TELEMEDICINE TEST-BED - A Tool for Determination of Accuracy in Asynchronous Collaborative Method <i>Josceli Tenório, Cristina Lucia Feijó Ortolani, Paulo Roberto de Lima Lopes and Ivan Torres Pisa</i>	231
BRAZILIAN TELEMEDICINE NETWORKS <i>Ricardo Alfredo Quintano Neira, Leandro Galassi Zavitowski, Paulo Roberto de Lima Lopes, Henrique Manuel Lederman, Paulo Schor, Ivan Torres Pisa and Daniel Sigulem</i>	235
MEDICAL DIAGNOSIS ASSISTANT BASED ON CATEGORY RANKING <i>Fernando Ruiz-Rico, José-Luis Vicedo and María-Consuelo Rubio-Sánchez</i>	239

TOWARD MOBILE HEALTHCARE SERVICES BY USING EVERYDAY MOBILE PHONES <i>Akio Sashima, Yutaka Inoue, Takeshi Ikeda, Tomohisa Yamashita, Masayuki Ohta and Koichi Kurumatani</i>	242
ARE BETTER FEATURE SELECTION METHODS ACTUALLY BETTER? - Discussion, Reasoning and Examples <i>Petr Somol, Jana Novovičová and Pavel Pudil</i>	246
IMPLEMENTATION OF EPS_T2DM - Implementation of Early Prediction System for Type 2 Diabetes Mellitus <i>Dan Bi Kim, Ju-young Lee, Eun-jung Jang, Ji-young Lee and Taek Kim</i>	254
PET-CT IMAGING AND DIAGNOSIS SYSTEM FOLLOWING DOCTOR'S METHOD <i>Hiroshi Arisawa, Takako Sato and Sinya Hata</i>	258
SPECIAL SESSION ON KNOWLEDGE DISCOVERY AND DECISION SUPPORT SYSTEMS IN HEALTH INFORMATION SYSTEMS	
INITIAL RESULTS ON KNOWLEDGE DISCOVERY AND DECISION SUPPORT FOR INTRACRANIAL ANEURYSMS <i>Christoph M. Friedrich, Martin Hofmann-Apitius, Robert Dunlop, Ioannis Chronakis, Miriam C. J. M. Sturkenboom, Roelof Risselada, Baldo Oliva and Ferran Sanz</i>	265
MY-HEART PROJECT: ANALYSIS OF SLEEP AND STRESS PROFILES FROM BIOMEDICAL SIGNAL <i>Sergio Cerutti, Anna M. Bianchi and Martin O. Mendez</i>	273
A SEMANTIC GRID SERVICES ARCHITECTURE IN SUPPORT OF EFFICIENT KNOWLEDGE DISCOVERY FROM MULTILEVEL CLINICAL AND GENOMIC DATASETS <i>Manolis Tsiknakis, Stelios Sfakianakis, Stefan Rueping, Oswaldo Trelles, Thierry Siestang, Brecht Claerbout and Vasilis Virvilis</i>	279
DECISION SUPPORT AND IMAGE & SIGNAL ANALYSIS IN HEART FAILURE - A Comprehensive Use Case <i>Sara Colantonio, Massimo Martinelli, Davide Moroni, Ovidio Salvetti, Francesco Perticone, Angela Sciacqua, Franco Chiarugi, Domenico Conforti, Antonio Gualtieri and Vincenzo Lagani</i>	288
MYHEART - Fighting Cardio-vascular Diseases by Prevention and Early Diagnosis <i>Ralf Schmidt, Jörg Habetha and Matthew Harris</i>	296
ONTOLOGY BASED INTEGRATION OF DISTRIBUTED AND HETEROGENEOUS DATA SOURCES IN ACGT <i>Luis Martín, Alberto Anguita, Víctor Maojo, Erwin Bonsma, Anca Bucur, Jeroen Vrijnsen, Mathias Brochhausen, Christian Cocos, Holger Stenzhorn, Manolis Tsiknakis, Martin Doerr and Haridimos Kondylakis</i>	301
MEDICAL KNOWLEDGE REPRESENTATION WITHIN HEARTFAID PLATFORM <i>Dragan Gamberger, Marin Prcela, Alan Jović, Tomislav Šmuc, Gianfranco Parati, Mariaconsuelo Valentini, Kalina Kawecka-Jaszcz, Katarzyna Styczkiewicz, Andrzej Kononowicz, Antonio Candelieri, Domenico Conforti and Rosita Guido</i>	307
AUTHOR INDEX	315

**INVITED
SPEAKERS**

**KEYNOTE
LECTURES**

MULTIVARIATE, MULTIORGAN AND MULTISCALE INTEGRATION OF INFORMATION IN BIOMEDICAL SIGNAL PROCESSING

Sergio Cerutti

*Department of Bioengineering, Polytechnic University, Milano, Italy
sergio.cerutti@polimi.it*

Abstract: Biomedical signals carry important information about the behavior of the living systems under studying. A proper processing of these signals allows in many instances to obtain useful physiological and clinical information. Many advanced algorithms of signal and image processing have recently been introduced in such an advanced area of research and therefore important selective information is obtainable even in presence of strong sources of noise or low signal/noise ratio. Traditional stationary signal analysis together with innovative methods of investigation of dynamical properties of biological systems and signals in second-order or in higher-order approaches (i.e., in time-frequency, time-variant and time-scale analysis, as well as in non linear dynamics analysis) provide a wide variety of even complex processing tools for information enhancement procedures. Another important innovative aspect is also remarked: the integration between signal processing and modeling of the relevant biological systems is capable to directly attribute patho-physiological meaning to the parameters obtained from the processing and viceversa the modeling fitting could certainly be improved by taking into account the results from signal processing procedure. Such an integration process could comprehend parameters and observations detected at different scales, at different organs and with different modalities. This approach is reputed promising for obtaining an olistic view of the patient rather than an atomistic one which considers the whole as a simple sum of the single component parts.

BRIEF BIOGRAPHY

Sergio Cerutti is Professor in Biomedical Signal and Data Processing at the Department of Bioengineering of the Polytechnic University in Milano, Italy. In the period 2000-2006 he has been the Chairman of the same Department. His research interests are mainly in the following topics: biomedical signal processing (ECG, blood pressure signal and respiration, cardiovascular variability signals, EEG and evoked potentials), neurosciences and cardiovascular modelling. In his research activity he has put emphasis on the integration of information at different modalities, at different sources and at different scales in various physiological systems. Since 1983 he has taught a course at a graduate and a doc level on Biomedical Signal Processing and Modelling at Engineering Faculties (Milano and Roma) as well as at Specialisation Schools of Medical Faculties (Milano and Roma). He has been Elected Member of IEEE-EMBS AdCom (Region 8) in the period 1993-1996.

He is actually Fellow Member of IEEE and of EAMBES and Associate Editor of IEEE Trans BME. He is a member of the Steering Committee of the IEEE-EMBS Summer School on Biomedical Signal Processing: he was the local organiser of four Summer Schools held in Siena. He has been Visiting Professor at Harvard-MIT Division Health Science and Technology, Boston, USA for an overall period of 1 year. He is the Author of more than 400 international scientific contributions (more than 180 on indexed scientific journals).

1 INTRODUCTION

Biomedical signals and imaging carry important information about the behavior of the living systems under studying. A proper processing of these signals and images allow in many instances to obtain useful physiological and clinical information. Actually, many advanced algorithms of digital signal and image processing are at disposal and therefore

important selective information is now obtainable even in presence of strong sources of noise or low signal/noise ratio. In most of the cases it is not sure whether such sources might derive even by complex and unknown interactions with other biological systems whose implications could be important from the physiological or clinical standpoints. Traditional stationary signal analysis together with innovative methods of investigation of dynamical properties of biological systems and signals in second-order or in higher-order approaches (i.e., in time-frequency, time-variant and time-scale analysis, as well as in non linear dynamics analysis) provide a wide variety of even complex processing tools for information enhancement procedures in the challenging studying of a better explanation of many physiological and clinical phenomena.

2 INTEGRATION BETWEEN SIGNAL PROCESSING AND PHYSIOLOGICAL MODELING

Another important innovative aspect to improve the information content from biomedical data is constituted by the integration between signal processing and modeling of the relevant biological systems, thus directly attributing patho-physiological meaning to the model parameters obtained from the processing; and, viceversa, the modeling fitting could certainly be improved by taking into account the results from signal/image processing procedures.

3 MONOVARIATE AND MULTIVARIATE SIGNAL PROCESSING

Other kinds of integration may be fulfilled, taking into account more signals from the same system in a multivariate way (i.e. from a single-lead vs multichannel EEG or ECG analysis) and combining also the action of different systems such as autonomic nervous system, cardiovascular and respiratory systems, etc. Sleep is a formidable example of multiorgan involvement in both physiological (sleep staging and correlation with cardiorespiratory system) and pathological conditions (sleep apnea, sleep deprivation, restless leg syndrome and so on).

4 MULTISCALE APPROACH

Further, modern rehabilitation techniques (motor and/or cognitive) make use actually of objective indices obtained from the patient's biosignals and images to better "personalize" rehabilitation protocols (from EEG, EP's, ERP's, MRI, fMRI, NIRS, etc). In neurosciences such an integration process could comprehend parameters and observations detected also at different scales, from genome and proteome up to the single organ and to the entire body compartment. Examples will be described where an animal model (murine model) is developed by altering a gene putative to a determined pathology (i.e. epilepsy) and changes in EEG signals are studied (spike/wave occurrences and modifications in signal power bands). In clinical applications, it is worth mentioning the important data fusion which could be fulfilled by the integration of simultaneous EEG recordings and fMRI in some epileptic patients during inter-critical or critical events.

Finally, another important integration can be obtained along different observation scales. Traditionally, biological signal analysis is carried out at the level of organ or system to be investigated (i.e., ECG or EEG signal, arterial blood pressure, respiration and so on). It is very clear the advantage of correlating this information with that one obtained about the same system, but at different scale level, i.e. at cellular level or even at subcellular level (for example, analyzing possible genetic correlates or typical patterns of proteins or even DNA/RNA sequences). Biomedical engineering as a dedicated discipline may strongly contribute to this multiscale information processing

Along this approach line, even the long-QT syndrome, can be efficiently studied at different scale level: a mutation in a portion of gene SCN5A which presents a phenotype compatible to long-QT3 type, is known to produce an altered function of Na⁺ channels. Through a proper model which describes the functioning of ventricular cells is possible to evidence that this alteration may induce a prolongation of QT duration, as detected on ECG tracing. This event is further correlated with an increased risk of ventricular tachyarrhythmias. Hence, the path is completed: from the genetic expression up to the disease manifestation (Clancy and Rudy, 1999), (Priori et al., 2003). Many different signal processing and modeling are involved in this paradigmatic example: an integration along the various scales of observation may undoubtedly contribute to a better

understanding of the complex pathophysiological correlates.

A great effort is on course nowadays for creating very large databases and networking of models and technologies for integrating such information (Physiome project (Hunter et al., 2002), (Rudy, 2000) to be connected with Genome and Proteome projects and Virtual Physiological Human project – VPH – which is inserted into the activities of the 7th Framework Programme of EU).

Other examples are constituted by the studying of the profile of expressed proteins in 2D-gel supports, or after mass-spectrometry analysis, relative to a variety of pathologies (i.e. epilepsy, peripheral neuropathies or Amyotrophic Lateral Sclerosis (ALS), or in oncological studies) thus singling out the set of proteins which present a correlate with the pathology in respect to the control group.

This overall approach is reputed promising for obtaining an olistic view of the patient rather than an atomistic one which considers the whole as a simple sum of the single component parts.

REFERENCES

- G. Baselli, S. Cerutti, S. Civardi, A. Malliani and M. Pagani, Cardiovascular variability signals: towards the identification of a closed-loop model of the neural control mechanisms, *IEEE Trans Biomed Eng.*, vol.35(12):1033-46(1988).
- E. N. Bruce, *Biomedical Signal Processing and Signal Modelling*, J. Wiley (2001).
- C. E. Clancy and Y. Rudy, Linking a genetic defect to its cellular phenotype in a cardiac arrhythmia, *Nature*, vol. 400(6744):566-569 (1999).
- A. Cohen, *Biomedical Signal Analysis*, Vol. I, II, CRC Press (1986).
- P. Hunter, P. Robbins and D. Noble, The IUPS human physiome project, *Pflugers Arch-Eur. J. Physiol.*, vol. 445:1-9 (2002).
- R. I. Kitney, A nonlinear model for studying oscillations in the blood pressure control system, *J. Biomed. Eng.*, vol I,n. 2: 88-89 (1979).
- H. P. Koepchen, History of studies and concepts of blood pressure wave, in: *Mechanisms of blood pressure waves* (Miyakawa K., Koepchen, H.P., Polosa C. eds): 3-23, Springer-Verlag (1984).
- M. C. Mackey and J.C. Milton, Dynamical diseases, *Ann NY Acad.Sc.*, vol.504:16-32 (1987).
- C.-K. Peng, S. Havlin, J.M. Hausdorff, J.E. Mietus, H.E. Stanley and A.L. Goldberger, Fractal mechanisms and heart rate dynamics long-range correlation and their breakdown with diseases, *Journal of Electrocardiology*, vol.28: 59-65 (1995).
- S. G. Priori, P.J. Schwartz, C. Napolitano, R. Bloise, E. Ronchetti, M. Grillo, A. Vicentini, C. Spazzolini, J. Nastoli, G. Bottelli, R. Folli and D. Cappelletti, Risk stratification in the long-QT syndrome, *N Engl J Med.*, May 8, vol.348(19):1866-74 (2003).
- Y. Rudy, From Genome to Physiome: integrative models of cardiac excitation, *Ann. Biom. Engineer.*, vol 28: 945-950 (2000).
- Cerutti S, Esposti F, Ferrario M, Sassi R, Signorini MG, Long-term invariant parameters obtained from 24-h Holter recordings: a comparison between different analysis techniques, *Chaos*, vol.17(1):015108, (2007).
- Foffani G, Ardolino G, Rampini P, Tamma F, Caputo E, Egidì M, Cerutti S, Barbieri S, Priori A., Physiological recordings from electrodes implanted in the basal ganglia for deep brain stimulation in Parkinson's disease: the relevance of fast subthalamic rhythms, *Acta Neurochir Suppl*, vol. 93:97-9, (2005).

OUTTHINKING AND ENHANCING BIOLOGICAL BRAINS

Kevin Warwick
University of Reading, UK

Keywords: Brain-Computer Interface, Biological systems, Implant technology, Feedback control.

Abstract: In this paper an attempt has been made to take a look at how the use of implant and electrode technology can now be employed to create biological brains for robots, to enable human enhancement and to diminish the effects of certain neural illnesses. In all cases the end result is to increase the range of abilities of the recipients. An indication is given of a number of areas in which such technology has already had a profound effect, a key element being the need for a clear interface linking the human brain directly with a computer. An overview of some of the latest developments in the field of Brain to Computer Interfacing is also given in order to assess advantages and disadvantages. The emphasis is clearly placed on practical studies that have been and are being undertaken and reported on, as opposed to those speculated, simulated or proposed as future projects. Related areas are discussed briefly only in the context of their contribution to the studies being undertaken. The area of focus is notably the use of invasive implant technology, where a connection is made directly with the cerebral cortex and/or nervous system. Tests and experimentation which do not involve human subjects are invariably carried out *a priori* to indicate the eventual possibilities before human subjects are themselves involved. Some of the more pertinent animal studies from this area are discussed including our own involving neural growth. The paper goes on to describe human experimentation, in which neural implants have linked the human nervous system bi-directionally with technology and the internet. A view is taken as to the prospects for the future for this implantable computing in terms of both therapy and enhancement.

BRIEF BIOGRAPHY

Kevin Warwick is Professor of Cybernetics at the University of Reading, England, where he carries out research in artificial intelligence, control, robotics and cyborgs. He is also Director of the University KTP Centre, which links the University with Small to Medium Enterprises and raises £2.5 million each year in research income. As well as publishing 500 research papers, Kevin is perhaps best known for his experiments into implant technology. He has been awarded higher doctorates (DScs) both by Imperial College and the Czech Academy of Sciences, Prague. He was presented with The Future of Health Technology Award in MIT, was made an Honorary Member of the Academy of Sciences, St. Petersburg and in 2004 received The IEE Achievement Medal.

1 INTRODUCTION

Research is being carried out in which biological signals of some form are measured, are acted upon by some appropriate signal processing technique and are then employed either to control a device or as an input to some feedback mechanism (Penny et al., 2000), (Roitberg, 2005). In many cases neural signals are employed, for example Electroencephalogram (EEG) signals can be measured externally to the body, using externally adhered electrodes on the scalp (Wolpaw et al., 1990) and can then employed as a control input. Most likely this is because the procedure is relatively simple from a research point of view and is not particularly taxing on the researchers involved. However, reliable interpretation of EEG data is extremely complex – partly due to both the compound nature of the multi-neuronal signals being measured and the difficulties in recording such highly attenuated

In the last few years interest has also grown in the use of real-time functional Magnetic Resonance

Imaging (fMRI) for applications such as computer cursor control. This typically involves an individual activating their brain in different areas by reproducible thoughts (Warwick, 2007) or by recreating events (Pan et al., 2007). Alternatively fMRI and EEG technologies can be combined so that individuals can learn how to regulate Slow Cortical Potentials (SCPs) in order to activate external devices (Hinterberger et al., 2005). Once again the technology is external to the body. It is though relatively expensive and cumbersome.

It is worth noting that external monitoring of neural signals, by means of either EEG analysis or indeed fMRI, leaves much to be desired. Almost surely the measuring technique considerably restricts the user's mobility and, as is especially the case with fMRI, the situation far from presents a natural or comfortable setting. Such systems also tend to be relatively slow, partly because of the nature of recordings via the indirect connection, but also because it takes time for the individual themselves to actually initiate changes in the signal. As a result of this, distractions, both conscious and sub-conscious, can result in false indicators thus preventing the use of such techniques for safety critical, highly dynamic and, to be honest, most realistic practical applications. Despite this, the method can enable some individuals who otherwise have extremely limited communication abilities to operate some local technology in their environment, and, in any case, it can serve as a test bed for a more direct and useful connection.

The definition of what constitutes a Brain-Computer Interface (BCI) is extremely broad. A standard keyboard could be so regarded. It is clear however that various wearable computer techniques and virtual reality systems, e.g. glasses containing a miniature computer screen for a remote visual experience (Mann, 1997), are felt by some researchers to fit this category. Although it is acknowledged that certain body conditions, such as stress or alertness, can be monitored in this way, the focus of this paper is on bidirectional BCIs and is more concerned with a direct connection between a biological brain and technology, and ultimately a human and technology.

2 *IN VIVO* STUDIES

Non-human animal studies can be considered to be a pointer for what is potentially achievable with humans in the future. As an example, in one

particular animal study the extracted brain of a lamprey, retained in a solution, was used to control the movement of a small wheeled robot to which it was attached (Reger et al., 2000). The lamprey innately exhibits a response to light reflections on the surface of water by trying to align its body with respect to the light source. When connected into the robot body, this response was utilised by surrounding the robot with a ring of lights. As different lights were switched on and off, so the robot moved around its corral, trying to position itself appropriately.

Meanwhile in studies involving rats, a group of rats were taught to pull a lever in order to receive a suitable reward. Electrodes were then chronically implanted into the rats' brains such that the reward was proffered when each rat thought (one supposes) about pulling the lever, but before any actual physical movement occurred. Over a period of days, four of the six rats involved in the experiment learned that they did not in fact need to initiate any action in order to obtain a reward; merely thinking about it was sufficient (Chapin, 2004).

In another series of experiments, implants consisting of microelectrode arrays have been positioned into the frontal and parietal lobes of the brains of two female rhesus macaque monkeys. Each monkey learned firstly how to control a remote robot arm through arm movements coupled with visual feedback, and it is reported that ultimately one of the monkeys was able to control the arm using only brain derived neural signals with no associated physical movement. Notably, control signals for the reaching and grasping movements of the robotic arm were derived from the same set of implanted electrodes (Carmena et al., 2003), (Nicoletis et al., 2000).

Such promising results from animal studies have given the drive towards human applications a new impetus.

3 ROBOT WITH A BIOLOGICAL BRAIN

Human concepts of a robot may involve a little wheeled device, perhaps a metallic head that looks roughly human-like or possibly a biped walking robot. Whatever the physical appearance our idea tends to be that the robot might be operated remotely by a human, or is being controlled by a simple programme, or even may be able to learn with a

microprocessor/computer as its brain. We regard a robot as a machine.

In a present project neurons are being cultured in a laboratory in Reading University to grow on and interact with a flat multi-electrode array. The neural culture, a biological brain, can be electronically stimulated via the electrodes and its trained response can be witnessed.

The project now involves networking the biological brain to be part of a robot device. In the first instance this will be a small wheeled robot. The input (sensory) signals in this case will be only the signals obtained from the wheeled robot's ultrasonic sensors. The output from the biological brain will be used to drive the robot around. The goal of the project initially will be to train the brain to drive the robot forwards without bumping into any object. Secondly, a separate biological brain will be grown to be the thinking process within a robot head (called Morgui) which houses 5 separate sensory inputs.

What this means is that the brain of these robots will shortly be a biological brain, not a computer. All the brain will know is what it perceives from the robot body and all it will do will be to drive the robot body around or control the robot head respectively. The biological brain will, to all intents and purposes, be the brain of the robot. It will have no life, no existence outside its robotic embodiment.

Clearly this research alters our concept of what a robot is, particularly in terms of ethical and responsibility issues. If a role of animal research is to open up possibilities for future human trials, then in this case the research could well be opening a window on the ultimate possibility of human neurons being employed in a robot body. All the 'human' brain would know would be its life as a robot.

4 HUMAN APPLICATION

At the present time the general class of Brain-Computer Interfaces (BCIs) for humans, of one form or another, have been specifically developed for a range of applications including military weapon and drive systems, personnel monitoring and for games consoles. However, by far the largest driving force for BCI research to date has been the requirement for new therapeutic devices such as neural prostheses.

The most ubiquitous sensory neural prosthesis in humans is by far the cochlea implant (Fin and

LoPresti, 2003). Here the destruction of inner ear hair cells and the related degeneration of auditory nerve fibres results in sensorineural hearing loss. As such, the prosthesis is designed to elicit patterns of neural activity via an array of electrodes implanted into the patient's cochlea, the result being to mimic the workings of a normal ear over a range of frequencies. It is claimed that some current devices restore up to approximately 80% of normal hearing, although for most recipients it is sufficient that they can communicate to a respectable degree without the need for any form of lip reading. The typically modest success of cochlea implantation is related to the ratio of stimulation channels to active sensor channels in a fully functioning ear. Recent devices consist of up to 32 channels, whilst the human ear utilises upwards of 30,000 fibres on the auditory nerve. There are now reportedly well over 10,000 of these prostheses in regular operation.

Studies investigating the integration of technology with the human central nervous system have varied from merely diagnostic to the amelioration of symptoms (Warwick and Gasson, 2004). In the last few years some of the most widely reported research involving human subjects is that based on the development of an artificial retina (Rizzo, 2001). Here, small electrode arrays have been successfully implanted into a functioning optic nerve. With direct stimulation of the nerve it has been possible for the otherwise blind recipient to perceive simple shapes and letters. The difficulties with restoring sight are though several orders of magnitude greater than those of the cochlea implant simply because the retina contains millions of photodetectors that need to be artificially replicated. An alternative is to bypass the optic nerve altogether and use cortical surface or intracortical stimulation to generate phosphenes (Dobelle, 2000).

Most invasive BCIs monitor multi-neuronal intracortical action potentials, requiring an interface which includes sufficient processing in order to relate recorded neural signals with movement intent. Problems incurred are the need to position electrodes as close as possible to the source of signals, the need for long term reliability and stability of interface in both a mechanical and a chemical sense, and adaptivity in signal processing to deal with technological and neuronal time dependence. However, in recent years a number of different collective assemblies of microelectrodes have been successfully employed both for recording and stimulating neural activity. Although themselves of small scale, nevertheless high density

connectors/transmitters are required to shift the signals to/from significant signal processing and conditioning devices and also for onward/receptive signal transmission.

Some research has focussed on patients who have suffered a stroke resulting in paralysis. The most relevant to this paper is the use of a '3rd generation' brain implant which enables a physically incapable brainstem stroke victim to control the movement of a cursor on a computer screen (Kennedy, 2000), (Kennedy, 2004). Functional Magnetic Resonance Imaging (fMRI) of the subject's brain was initially carried out to localise where activity was most pronounced whilst the subject was thinking about various movements. A hollow glass electrode cone containing two gold wires and a neurotrophic compound (giving it the title 'Neurotrophic Electrode') was then implanted into the motor cortex, in the area of maximum activity. The neurotrophic compound encouraged nerve tissue to grow into the glass cone such that when the patient thought about moving his hand, the subsequent activity was detected by the electrode, then amplified and transmitted by a radio link to a computer where the signals were translated into control signals to bring about movement of the cursor. With two electrodes in place, the subject successfully learnt to move the cursor around by thinking about different movements. Eventually the patient reached a level of control where no abstraction was needed – to move the cursor he simply thought about moving the cursor. Notably, during the period that the implant was in place, no rejection of the implant was observed; indeed the neurons growing into the electrode allowed for stable long-term recordings.

Electronic neural stimulation has proved to be extremely successful in other areas, including applications such as the treatment of Parkinson's disease symptoms. With Parkinson's Disease diminished levels of the neurotransmitter dopamine cause over-activation in the ventral posterior nucleus and the subthalamic nucleus, resulting in slowness, stiffness, gait difficulties and hand tremors. By implanting electrodes into the subthalamic nucleus to provide a constant stimulation pulse, the over activity can be inhibited allowing the patient, to all external intents and purposes, to function normally (Pinter et al., 1999).

5 BRAIN WITHIN A BRAIN

Ongoing research, funded by the UK Medical Research Council, is investigating how the onset of tremors can be accurately predicted such that merely a stimulation current burst is required rather than a constant pulsing (Gasson et al., 2005: pp.16/1-16/4). This has implications for battery inter-recharge periods as well as limiting the extent of in-body intrusive signalling. The deep brain stimulator can be used to collect local field potential (LFP) signals generated by the neurons around the deep brain electrodes (Gasson et al., 2005: pp.16/1-16/4). Determining the onset of events can be investigated by using fourier transforms to transfer the time based signal to a frequency based spectrogram to determine the change in frequency at the critical time period. However, in addition to that, the frequency changes in the period of time immediately prior to the tremor occurrence can give important information.

Fig.1 shows the results of an initial attempt to train an artificial neural network to indicate not only that a Parkinsonian tremor is present but also that one is very likely to occur in the near future. The aim of this research is that, once a reliable predictor has been obtained, the stimulating pulsing will only be enacted when a tremor is predicted, in order to stop the actual physical tremor occurring before it even starts in the first place.

The bottom trace in Fig.1 shows emg (muscular) signals, measured externally, associated with movement due to the tremors. It can be seen that the tremors in this incident actually start at around the 45 to 50 second point. The trace just above this indicates the corresponding electrical data measured as deep brain Local Field Potentials in the Sub-Thalamic Nucleus of the patient involved. It can be witnessed how, in this case, the electrical data takes on a different form (in terms of variance at least) at around the 45 to 50 second point. The four top plots meanwhile indicate the outputs from 4 differently structured artificial neural networks, based on multi-layer perceptrons with different numbers of neurons in the hidden (middle) layer.

It can be seen how, for each network, the output of the network goes high (logic 1) at the 45 to 50 second point, to indicate the presence of a Parkinsonian tremor. This is all well and good, what is important however is that the output of the networks also briefly goes high around the 30




Figure 1: Time plot of the onset of a Parkinsonian tremor incident with corresponding artificial neural network indicators.

second point and this can be seen as an indication of the fact that a tremor will shortly occur. Ongoing research is involved with selection of the type and number of inputs to the network, presently these being based on the energy spectrum in different frequency ranges. The networks are also being tested on considerable amounts of resting data, that is long periods of brain activity where no tremors at all actually occur in patients. Clearly the aim is that a network will not give false predictions of tremors.

In fact false positive predictions are not so much of a critical problem. The end result with a false positive is that a stimulation may occur when it is not strictly necessary. In any event no actual tremor would occur, which is indeed a good outcome, however unnecessary energy would have been used – in fact if numerous false predictions occurred the intelligent stimulator would tend toward the present ‘blind’ stimulator. Effectively the occasional false positive prediction is perhaps not a problem, unless it became a regular occurrence. The good news is that results show that the network can be readily tuned to avoid false positives anyway.

6 GENERAL IMPLANT STUDIES

Some of the most impressive human research to date has been carried out using the microelectrode array, shown in Figure 2. The individual electrodes are

only 1.5mm long and taper to a tip diameter of less than 90 microns. Although a number of trials not using humans as a test subject have occurred (Branner and Normann, 2000), human tests are at present limited to two studies. In the second of these the array has been employed in a recording only role (Donoghue et al., 2002), (Donoghue et al., 2004), (Friebs et al., 2004), most notably recently as part of the ‘Braingate’ system. Essentially activity from a few neurons monitored by the array electrodes is decoded into a signal to direct cursor movement. This has enabled an individual to position a cursor on a computer screen, using neural signals for control combined with visual feedback. The first use of the microelectrode array (Figure 2) will be discussed in the following section as this has considerably broader implications which extend the capabilities of the human recipient.

A key selection point at the present time are what type of implant to employ, as several different possibilities exist, ranging from single electrode devices to multielectrode needles which contain electrode points at different depths to multielectrode arrays which either contain a number of electrodes which penetrate to the same depth (as in Figure 2) or are positioned in a banked/sloped arrangement. A further key area of consideration is the exact positioning of a BCI. In particular certain areas of the brain are, apparently, only really useful for monitoring purposes whilst others are more useful for stimulation.

Actually deriving a reliable command signal from a collection of captured neural signals is not necessarily a simple task, partly due to the complexity of signals recorded and partly due to time constraints in dealing with the data. In some cases however it can be relatively easy to look for and obtain a system response to certain anticipated neural signals – especially when an individual has trained extensively with the system. In fact neural signal shape, magnitude and waveform with respect to time are considerably different to the other signals that it is possible to measure in this situation.

If a greater understanding is required of neural signals recorded, before significant progress can be made, then this will almost surely present a major problem. This is especially true if a number of simultaneous channels are being employed, each requiring a rate of digitization of (most likely) greater than 20KHz in the presence of unwanted noise. For real time use this data will also need to be processed within a few milliseconds (100 milliseconds at most). Further, although many studies have looked into the extraction of command signals (indicating intent) from measured values, it is clear that the range of neural activity is considerable. Even in the motor area not only are motor signals present but so too are sensory, cognitive, perceptual along with other signals, the exact purpose of which is not clear – merely classifying them as noise is not really sufficient and indeed can be problematic when they are repeated and apparently linked in some way to activity.

It is worth stressing here that the human brain and spinal cord are linking structures, the functioning of which can be changed through electronic stimulation such as that provided via an electrode arrangement. This type of technology therefore offers a variety of therapeutic possibilities. In particular the use of implanted systems when applied to spinal cord injured patients, in whom nerve function is disordered, was described in (Warwick, 2004) as having the following potential benefits (among others):

1. Re-education of the brain and spinal cord through repeated stimulation patterns
2. Prevention of spinal deformity
3. Treatment of intractable neurogenic and other pain
4. Assisting bladder emptying
5. Improving bowel function
6. Treatment of spasticity
7. Improvement of respiratory function – assisting coughing and breathing

8. Reduction of cardiovascular maleffects
9. Prevention of pressure sores – possibly providing sensory feedback from denervated areas
10. Improvement and restoration of sexual function
11. Improved mobility
12. Improved capability in daily living, especially through improved hand, upper limb and truncal control

Sensate prosthetics is another growing application area of neural interface technology, whereby a measure of sensation is restored using signals from small tactile transducers distributed within an artificial limb (Fin and LoPresti, 2003). The transducer output can be employed to stimulate the sensory axons remaining in the residual limb which are naturally associated with a sensation. This more closely replicates stimuli in the original sensory modality, rather than forming a type of feedback using neural pathways not normally associated with the information being fed back. As a result it is supposed that the user can employ lower level reflexes that exist within the central nervous system, making control of the prosthesis more subconscious.

One final noteworthy therapeutic procedure is Functional Electrical Stimulation (FES), although it is debatable if it can be truly referred to as a BCI, however it aims to bring about muscular excitation, thereby enabling the controlled movement of limbs. FES has been shown to be successful for artificial hand grasping and release and for standing and walking in quadriplegic and paraplegic individuals as well as restoring some basic body functions such as bladder and bowel control (Grill and Kirsch, 2000). It must be noted though that controlling and coordinating concerted muscle movements for complex and generic tasks such as picking up an arbitrary object is proving to be a difficult, if not insurmountable, challenge.

In the cases described in which human subjects are involved, the aim on each occasion is to either restore functions since the individual has a physical problem of some kind or it is to give a new ability to an individual who has very limited motor abilities. In this latter case whilst the procedure can be regarded as having a therapeutic purpose, it is quite possible to provide an individual with an ability that they have in fact never experienced before. On the one hand it may be that whilst the individual in question has never previously experienced such an ability, some or most other humans have – in this

case it could be considered that the therapy is bringing the individual more in line with the “norm” of human abilities.

It is though also potentially possible to give extra capabilities to a human, to enable them to achieve a broader range of skills – to go beyond the “norm”. Apart from the, potentially insurmountable, problem of universally deciding on what constitutes the “norm”, extending the concept of therapy to include endowing an individual with abilities that allow them to do things that a perfectly able human cannot do raises enormous ethical issues. Indeed it could be considered that a cochlea implant with a wider frequency response range does just that for an individual or rather an individual who can control the cursor on a computer screen directly from neural signals falls into this category. But the possibilities of enhancement are enormous. In the next section we consider how far things could be taken, by referring to relevant experimental results.

7 HUMAN ENHANCEMENT

The interface through which a user interacts with technology provides a distinct layer of separation between what the user wants the machine to do, and what it actually does. This separation imposes a considerable cognitive load upon the user that is directly proportional to the level of difficulty experienced. The main issue it appears is interfacing the human motor and sensory channels with the technology. One solution is to avoid this sensorimotor bottleneck altogether by interfacing directly with the human nervous system. It is certainly worthwhile considering what may

potentially be gained from such an invasive undertaking.

Advantages of machine intelligence are for example rapid and highly accurate mathematical abilities in terms of ‘number crunching’, a high speed, almost infinite, internet knowledge base, and accurate long term memory. Additionally, it is widely acknowledged that humans have only five senses that we know of, whereas machines offer a view of the world which includes infra-red, ultraviolet and ultrasonic. Humans are also limited in that they can only visualise and understand the world around them in terms of a limited dimensional perception, whereas computers are quite capable of dealing with hundreds of dimensions. Also, the human means of communication, essentially transferring an electro-chemical signal from one brain to another via an intermediate, often mechanical medium, is extremely poor, particularly in terms of speed, power and precision. It is clear that connecting a human brain, by means of an implant, with a computer network could in the long term open up the distinct advantages of machine intelligence, communication and sensing abilities to the implanted individual.

As a step towards this more broader concept of human-machine symbiosis, in the first study of its kind, the microelectrode array (as shown in Figure 2) was implanted into the median nerve fibres of a healthy human individual (myself) in order to test *bidirectional* functionality in a series of experiments. A stimulation current direct onto the nervous system allowed information to be sent to the user, while control signals were decoded from neural activity in the region of the electrodes (Gasson et al., 2005:pp 365-375), (Warwick et al., 2003).




Figure 2: A 100 electrode, 4X4mm Microelectrode Array, shown on a UK 1 pence piece for scale.

In this way a number of experimental trials were successfully concluded (Warwick et al., 2004), (Warwick et al., 2005): In particular:

1. Extra sensory (ultrasonic) input was successfully implemented and made use of.
2. Extended control of a robotic hand across the internet was achieved, with feedback from the robotic fingertips being sent back as neural stimulation to give a sense of force being applied to an object (this was achieved between New York (USA) and Reading(UK))
3. A primitive form of telegraphic communication directly between the nervous systems of two humans was performed.
4. A wheelchair was successfully driven around by means of neural signals.
5. The colour of jewellery was changed as a result of neural signals – as indeed was the behaviour of a collection of small robots.

In each of the above cases it could be regarded that the trial proved useful for purely therapeutic reasons, e.g. the ultrasonic sense could be useful for an individual who is blind or the telegraphic communication could be very useful for those with certain forms of Motor Neurone Disease. However each trial can also be seen as a potential form of augmentation or enhancement for an individual. The question then arises as to how far should things be taken? Clearly enhancement by means of BCIs opens up all sorts of new technological and intellectual opportunities, however it also throws up a raft of different ethical considerations that need to be addressed directly.

8 ON STIMULATION

After extensive experimentation it was found that injecting currents below 80 μ A onto the median nerve fibers had little perceivable effect. Between 80 μ A and 100 μ A all the functional electrodes were able to produce a recognizable stimulation, with an applied voltage of 40 to 50 volts, dependant on the series electrode impedance. Increasing the current above 100 μ A had no apparent additional effect; the stimulation switching mechanisms in the median nerve fascicle exhibited a non-linear thresholding characteristic.

During this experimental phase, it was pseudo randomly decided whether a stimulation pulse was applied or not. The volunteer (myself), wearing a blindfold, was unaware of whether a pulse had been applied or not, other than by means of its effect in terms of neural stimulation. The user's accuracy in distinguishing between an actual pulse and no pulse at a range of amplitudes is shown in Figure 3.

In all subsequent successful trials, the current was applied as a bi-phasic signal with pulse duration of 200 μ sec and an inter-phase delay of 100 μ sec. A typical stimulation waveform of constant current being applied to one of the MEA's implanted electrodes is shown in Fig 4.

It was, in this way, possible to create alternative sensations via this new input route to the nervous system. Of the 5 enhancement features mentioned in the previous section, this one will be described, as an example, in further detail. Background information on the other enhancements can be found in a number of references, e.g. (Gasson et al., 2005:pp 365-375), (Warwick et al., 2003), (Warwick et al., 2004), (Warwick and Gasson, 2004).

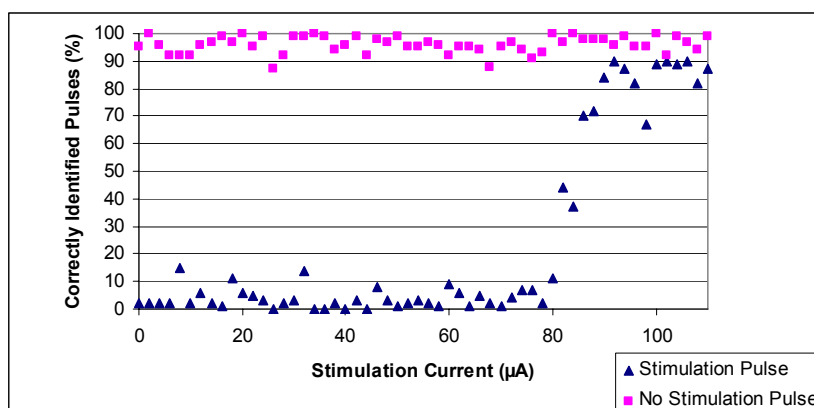


Figure 3: Effect of stimulation amplitude on the number of correctly identified pulses and absence of pulses (over 100 trials).




Figure 4: Voltage profile during one bi-phasic stimulation pulse cycle with a constant current of $80\mu\text{A}$.

It must be reported that it took 6 weeks for my brain to repetitively recognize the stimulating signals accurately. This time period can be due to a number of contributing factors:

- (a) The team had to learn which signals (what amplitude, frequency etc.) would be best in order to bring about a recognizable stimulation.
- (b) The recipient's brain had to learn to recognize the new signals it was receiving.
- (c) The bond between the recipient's nervous system and the implant was physically changing (becoming stronger).

9 EXTRA SENSORY EXPERIMENT

An experiment was set up to determine if the human brain is able to understand and successfully operate with sensory information to which it had not previously been exposed. Whilst it is quite possible to feed in such sensory information via a normal human sensory route, e.g. electromagnetic radar or infra-red signals are converted to visual, what we were interested in was feeding such signals directly onto the human nervous system, thereby bi-passing the normal human sensory input.

Ultrasonic sensors were fitted to the rim of a baseball cap (see Figure 5) and the output from these sensors, in the form of a proportional count, was employed to bring about a direct stimulation of the nervous system. Hence when no objects were in the vicinity of the sensors, no stimulation occurred, and as an object moved close by so the rate of stimulation pulses being applied increased in a linear fashion up to a pre-selected maximum rate. No increase in stimulation occurred when an object moved closer than 10cm to the sensors.

The ultrasonic sensors were open type piezoelectric ceramic transducers with conical metal resonators and operated at 40 KHz. These were used in a pair, one for transmit and one for receive, to give maximum sensitivity for small and distant objects. The most useful range for the experimentation was found to be 2 – 3m, this being also dependent on the size of object. A simple microcontroller was programmed to perform the echo ranging on the pair of transducers, and provide the range to the first detectable object only. This was translated into a stimulation pulse train, which operated on a single pin of the electrode array. Pins on the array had been tested for their suitability for stimulation by the earlier experimentation in which the recipient identified the presence or absence of stimulation pulse trains at various amplitudes and repetition frequencies.



Figure 5: Experimentation and testing of the ultrasonic baseball cap.

It was found that very little learning was required for the new ultrasonic sense to be used effectively and successfully – merely a matter of 5/6 minutes. This said it must be remembered that it had already taken several weeks for the recipient's brain to successfully, accurately recognize the current signals being injected.

As a result, in a witnessed experiment, the recipient, whilst wearing a blindfold, was able to move around successfully within a cluttered laboratory environment, albeit at a slower than normal walking pace. The sensory input was “felt” as a new form of sensory input (not as touch or movement) in the sense that the brain made a direct link between the signals being witnessed and the fact that these corresponded in a linear fashion to a nearby object.

10 CONCLUSIONS

External input-output interfaces with human and animal brains have been studied for many years. These are sometimes referred to as Brain-Computer Interfaces (BCIs) even though the interface may be external to the (human) body and its sensorimotor mechanism. In this paper an attempt has been made to put such systems in perspective. Emphasis has been placed on such interfaces that can be obtained

by means of implanted devices through invasive surgery and actual direct neural connections. In particular a number of trials in this area have clearly shown the possibilities of monitoring, stimulating and enhancing brain functioning.

Although there is no distinct dividing line it is quite possible as far as humans are concerned to investigate BCIs in terms of those employed for direct therapeutic means and those which can have an enhanced role to play. It is clear that the interaction of electronic signals with the human brain can cause the brain to operate in a distinctly different manner. Such is the situation with the stimulator implants that are successfully used to counteract, purely electronically, the tremor effects associated with Parkinson's disease. Such technology can though potentially be employed to modify the normal functioning of the human brain and nervous system in a number of different ways.

The same stimulator, with slightly different positioning, has been shown to elicit feelings of sadness or happiness in the recipient. Given the nature of the intelligent stimulator described here it would appear to be possible to monitor, in real time, a human brain with a computer brain, and for the computer brain to predict when the human is going to feel sad – quite some time before they actually feel sad. In theory a signal could then be injected at

that time to make them feel happy, or at least to stop them actually ever feeling sad in the first place. Maybe this could be regarded as an electronic anti-depressant. There are of course questions about recreational use here – but this would need a deep brain implant which might well prove to be rather too onerous for most people.

Perhaps understandably, invasive BCIs are presently far less well investigated in University experiments than their external BCI counterparts. A number of animal trials have though been carried out and the more pertinent have been indicated here along with the relevant human trials and practice. In particular the focus of attention has been given to the embodiment of grown neural tissue within a technological body. Whilst only 1,000 or so neurons are involved this presents an interesting research area in a number of ways. But once the number of such neurons used increases 1,000 or 1,000,000-fold, it also raises enormous philosophical and ethical issues. For example is the robot ‘thinking’ and what rights should it have?

The potential for BCI applications for individuals who are paralysed is enormous, where cerebral functioning despite generate command signals is functional despite the motor neural pathways being in some way impaired – such as in Lou Gehrig’s disease. The major role is then either one of relaying a signal of intention to the appropriate actuator muscles or to reinterpret the neural signals to operate technology thereby acting as an enabler. In these situations no other medical ‘cure’ is available, something which presents a huge driver for an invasive implant solution for the millions of individuals who are so affected. Clearly though, bidirectional signalling is important, not only to monitor and enact an individual’s intent but also to provide feedback on that individual’s resultant interaction with the real world. For grasping, walking and even as a defensive safety stimulant, feedback is vital. This paper has therefore focussed on such studies.

Where invasive interfaces are employed in human trails, a purely therapeutic scenario often exists. In a small number of instances, such as use of the microelectrode array as an interface, an individual has been given different abilities, something which opens up the possibilities of human enhancement. These latter cases however raise more topical ethical questions with regard to the need and use of a BCI. What might be seen as a new means of communication for an individual with an extreme form of paralysis or a new sensory input for

someone who is blind, opening up a new world for them, can also be seen as an unnecessary extra for another individual, even though it may provide novel commercial opportunities. What is therapy for one person may be regarded as an enhancement or upgrading for another.

Whilst there are still many technical problems to be overcome in the development of BCIs, significant recent experimental results have indicated that a sufficient technological infrastructure now exists for further major advances to be made. Although a more detailed understanding of the underlying neural processes will be needed in the years ahead, it is not felt that this will present a major hold up over the next few years, rather it will provide an avenue of research in which many new results will shortly appear through trials and experimentation, possibly initially through animal studies although it must be recognised that it is only through human studies that a full analysis can be made and all encompassing conclusions can be drawn. Nevertheless the topic opens up various ethical questions that need to be addressed and as such, research in this area should, I believe, only proceed in light of a pervasive ethical consensus.

ACKNOWLEDGEMENTS

The Author would like to acknowledge the considerable assistance and input of the Consultant Neurosurgeons Mr. Peter Teddy, Mr. Amjad Shad, Mr. Ali Jamous and Mr. Tipu Aziz and researchers Iain Goodhew, Mark Gasson, Ben Whalley and Ben Hutt. Ethical approval for the author’s research was obtained from the Ethics and Research Committee at the University of Reading, UK and with regard to the neurosurgery aspect, the Oxfordshire National Health Trust Board overseeing the Radcliffe Infirmary, Oxford, UK.

REFERENCES

- Branner, A. and Normann, R. (2000) A multielectrode array for intrafascicular recording and stimulation in the sciatic nerve of a cat, *Brain Research Bulletin*, Vol.51, pp.293-306.
- Chapin, J.K. (2004) Using multi-neuron population recordings for neural prosthetics. *Nature Neuroscience*, 7, 452-454.
- Carmena, J., Lebedev, M., Crist, R., O’Doherty, J., Santucci, D., Dimitrov, D. Patil, P., Henriquez, C. and Nicolelis, M. (2003) Learning to control a brain-

- machine interface for reaching and grasping by primates, *Plos Biology*, Vol.1, Issue.2, article number e2.
- Dobelle, W. (2000) Artificial vision for the blind by connecting a television camera to the visual cortex, *ASAIO J*, Vol.46, pp.3-9.
- Donoghue, J. (2002) Connecting cortex to machines: recent advances in brain interfaces, *Nature Neuroscience Supplement*, Vol.5, pp.1085-1088.
- Donoghue, J., Nurmikko, A., Friehs, G. And Black, M. (2004) Development of a neuromotor prosthesis for humans, Chapter 63 in *Advances in Clinical Neurophysiology, Supplements to Clinical Neurophysiology*, Vol.57, pp.588-602.
- Finn, W. and LoPresti, P. (eds.) (2003) *Handbook of Neuroprosthetic methods*, CRC Press.
- Friehs, G., Zerris, V., Ojakangas, C., Fellows, M. and Donoghue, J. (2004) Brain-machine and brain-computer interfaces, *Stroke*, Vol.35, Issue.11, pp.2702-2705.
- Gasson, M., Hutt, B., Goodhew, I., Kyberd, P. and Warwick, K. (2005) Invasive neural prosthesis for neural signal detection and nerve stimulation, *Proc. International Journal of Adaptive Control and Signal Processing*, Vol.19, No.5, pp.365-375.
- Gasson, M., Wang, S., Aziz, T., Stein, J. and Warwick, K. (2005) Towards a demand driven deep brain stimulator for the treatment of movement disorders, *Proc. 3rd IEE International Seminar on Medical Applications of Signal Processing*, pp.16/1-16/4.
- Grill, W. and Kirsch, R. (2000) Neuroprosthetic applications of electrical stimulation, *Assistive Technology*, Vol.12, Issue.1, pp.6-16.
- Hinterberger, T., Veit, R., Wilhelm, B., Weisopf, N., Vatine, J. and Birbaumer, N. (2005) Neuronal mechanisms underlying control of a brain-computer interface, *European Journal of Neuroscience*, Vol.21, Issue.11, pp.3169-3181.
- Kennedy, P., Bakay, R., Moore, M., Adams, K. and Goldwaith, J. (2000) Direct control of a computer from the human central nervous system, *IEEE Transactions on Rehabilitation Engineering*, Vol.8, pp.198-202.
- Kennedy, P., Andreasen, D., Ehirim, P., King, B., Kirby, T., Mao, H. and Moore, M. (2004) Using human extracortical local field potentials to control a switch, *Journal of Neural Engineering*, Vol.1, Issue.2, pp.72-77.
- Mann, S. (1997) Wearable Computing: A first step towards personal imaging, *Computer*, Vol. 30, Issue.2, pp. 25-32.
- Nicolelis, M., Dimitrov, D., Carmena, J., Crist, R., Lehew, G., Kralik, J. and Wise, S. (2003) Chronic, multisite, multielectrode recordings in macaque monkeys, *Proc. National Academy of the USA*, Vol.100, Issue.19, pp.11041-11046.
- Penny, W., Roberts, S., Curran, E., and Stokes, M. (2000) EEG-based communication: A pattern recognition approach, *IEEE Transactions on Rehabilitation Engineering*, Vol. 8, Issue.2, pp. 214-215.
- Pinter, M., Murg, M., Alesch, F., Freundl, B., Helscher, R. And Binder, H. (1999) Does deep brain stimulation of the nucleus ventralis intermedius affect postural control and locomotion in Parkinson's disease?, *Movement Disorders*, Vol.14, Issue.6, pp.958-963.
- Reger, B., Fleming, K., Sanguinetti, V., Simon Alford, S., Mussa-Ivaldi, F. (2000) Connecting Brains to Robots: an artificial body for studying computational properties of neural tissues, *Artificial life*, Vol.6, Issue.4, pp.307-324.
- Rizzo, J., Wyatt, J., Humayun, M., DeJuan, E., Liu, W., Chow, A., Eckmiller, R., Zrenner, E., Yagi, T. and Abrams, G. (2001) Retinal Prosthesis: An encouraging first decade with major challenges ahead, *Ophthalmology*, Vol.108, No.1.
- Roitberg, B. (2005) Noninvasive brain-computer interface, *Surgical Neurology*, Vol.63, Issue.3, p.195.
- Warwick, K. (2004) *I Cyborg*, University of Illinois Press.
- Warwick, K., Gasson, M., Hutt, B., Goodhew, I., Kyberd, P., Andrews, B., Teddy, P., Shad, A. (2003) The application of implant technology for cybernetic systems. *Archives of Neurology*, 60 (10), pp.1369-1373.
- Warwick, K., Gasson, M., Hutt, B., Goodhew, I., Kyberd, P., Schulzrinne, H. and Wu, X. (2004) Thought Communication and Control: A First Step Using Radiotelemetry, *IEE Proceedings on Communications*, Vol.151, No. 3, pp 185-189.
- Warwick, K., Gasson, M., Hutt, B. and Goodhew, I. (2005) An Attempt to Extend Human Sensory Capabilities by means of Implant Technology. *Proc. IEEE Int. Conference on Systems, Man and Cybernetics, Hawaii*.
- Wolpaw, J., McFarland, D., Neat, G. and Forheris, C. (1990) An EEG based brain-computer interface for cursor control, *Electroencephalogr. Clin. Neurophysiol.*, Vol.78, pp.252-259.
- Pan, S., Warwick, K., Gasson, M., Burgess, J., Wang, S., Aziz, T. and Stein, J. (2007) Prediction of parkinson's disease tremor onset with artificial neural networks", *Proc. IASTED Conference BioMed 2007, Innsbruck, Austria*, pp.341-345.
- Warwick, K. (2007) The promise and threat of modern cybernetics, *Southern Medical Journal*, Vol.100, Issue.1, pp.112-115.
- Warwick, K. and Gasson, M.N. (2004) Practical Interface Experiments with Implant Technology, in "Computer Vision in Human-Computer Interaction", Sebe, N., Lew, M.S and Huang, T.S. (eds.), *Lecture Notes in Computer Science*, Vol. 3058, pp.7-16. Xie, S., Yang, Z. and Yang, Y. (2004) Brain-computer interface based on event-related potentials during imitated natural reading, *International Journal of Psychology*, Vol.39, Issue.5-6, Supplement S., p.138.
- Yoo, S., Fairney, T., Chen, N., Choo, S., Panych, L., Park, H., Lee, S. and Jolesz, F. (2004) Brain-computer interface using fMRI: spatial navigation by thoughts, *Neuroreport*, Vol.15, Issue.10, pp.1591-1595.
- Yu, N., Chen, J. and Ju, M. (2001) Closed-Loop Control of Quadriceps/Hamstring activation for FES-Induced Standing-Up Movement of Paraplegics, *Journal of Musculoskeletal Research*, Vol. 5, No.3, pp.173-184.

ANALYSIS AND MODELS OF BRAIN EPILEPTIC ACTIVITIES

Fernando Henrique Lopes da Silva
University of Amsterdam, The Netherlands

Abstract: The essence of epilepsy is the sudden occurrence of a qualitative change in the behaviour of neuronal networks of some specific areas of the brain. In general we may assume that neuronal networks possess multistable dynamics. We may simplify this concept considering the case that a neuronal network may display, at least, two dynamical states: an interictal state characterised by a normal on-going neural activity, as revealed in the Electroencephalogram or Magnetoencephalogram (EEG, MEG), that may be apparently random, and another one – the ictal state - that is characterised by the sudden occurrence of synchronous oscillations, most commonly with large amplitude. The latter becomes manifest as a paroxysmal change of behaviour and /of the state of consciousness of a patient, i.e. an epileptic seizure. In the terminology of the mathematics of non-linear systems, we may state that a neuronal network behaves as a bistable system with two attractors, to which the system converges depending on initial conditions and on the system's parameters. We propose schematically that the transition between the normal on-going to the seizure activity can take place according to three basic models: Model I – a transition may occur due to random fluctuations of some system's parameters. These transitions are thus unpredictable. Models II and III – a transition may result from a gradual change of some unstable parameters, either due to endogenous (model II) or exogenous (model III). In these cases the change of parameter values causes a deformation of the attractor resulting in a transition from the basin of the attractor corresponding to the normal state, to the attractor corresponding to the seizure dynamical state. Some experimental findings obtained in different cases of epilepsy, both in human and in animals, are compatible with each of these 3 models. Some examples of these cases are illustrated.

BRIEF BIOGRAPHY

Fernando Henrique Lopes da Silva received his Medical Degree from the University of Lisbon in 1959, got his Ph.D. from the University of Utrecht in 1970, and in 1980 was appointed Full Professor of General Physiology at the Faculty of Science at the University of Amsterdam (since 2002 part of the Swammerdam Institute for Life Sciences). From 1993 to 2000 he was Director of the newly created Institute of Neurobiology of the University of Amsterdam, and member of the Scientific Directorate of the Graduate School Neurosciences Amsterdam. In 2000, when he reached the retirement age of 65, he became Emeritus Professor of the same University, and has at present a freelance contract with the Swammerdam Institute for Life Sciences.

Since 1970, he supervised a large number of student trainees from different Universities and Faculties: Medical, Biology, Sciences, (Bio-medical) Engineering. Supervised 65 Ph.D. students (up to December 2006).

His research interests are centred on the biophysical aspects of electrical activity of the brain and the functional organization of neuronal networks, namely of the cerebral cortex and the limbic system, with a special interest in the generation and functional significance of brain rhythmic activities. He published more than 220 papers in peer-reviewed journals and contributed Chapters to 10 multi-authored books (of 6 he is co-editor), among which the Handbook "Electroencephalography: Basic principles, clinical applications and related fields", Niedermeyer, E. and Lopes da Silva, F.H. (Eds), published by Lippincott, Williams and Wilkins, Baltimore; 5 Editions: 1982, 1987, 1993, 1998, 2004. In addition he contributed chapters to the Encyclopedia of Neuroscience (George Adelman, Barry H. Smith. Eds), Elsevier Science, 2003 (3rd edition), to the Encyclopedia of the Human Brain (Ed. V. S. Ramachandran), Academic Press, 2002, and to The Handbook of Brain Theory and Neural Networks (Ed. Michael A. Arbib), The MIT Press, 2003 (2nd edition).

Selection of Scientific Awards

- 1975 He received the Winkler Medal from the Netherlands Association for Neurology for scientific contributions in the field of neurosciences.
- 1985 Elected member of the Royal Netherlands Academy of Arts and Sciences.
- 1990 "Lord Adrian" Lecturer at the 12th World Congress of Electroencephalography and Clinical Neurophysiology in Rio de Janeiro, Brazil.
- 1992 Honorary President of the VIIth European Congress of Clinical Neurophysiology, Budapest, Hungary.
- 1995 Honorary Life Member of The British Society for Clinical Neurophysiology (Formerly The EEG Society), London, United Kingdom.
- 1997 Doctor Honoris Causa of the University of Lisbon (Portugal).
- 1997 Special "Berger" Lecturer at the 14th International Congress of EEG and Clinical Neurophysiology in Florence, Italy.
- 1999 Recipient of the Herbert H. Jasper Award, selected by the American Clinical Neurophysiology Society for his "lifetime of outstanding contributions to the field of clinical neurophysiology."
- 2000 Recipient of the 'Storm van Leeuwen/Magnus Prize' of the Dutch Society of Clinical Neurophysiology.
- 2000 Honorary member of the Portuguese Society of Electroencephalography and Clinical Neurophysiology.
- 2002 Recipient of the Ragnar Granit Prize for his work on the field of Bioelectromagnetism.
- 2002 Doctor Honoris Causa of the University of Porto (Portugal).
- 2004 Recipient of the first Prize "Universidade de Coimbra" for a (sic) "person of Portuguese nationality who has made a particular relevant and innovative contribution in the fields of culture or science."

General Honors

- 2000 High Officer of the Order of Santiago da Espada, for outstanding achievements in the field of Science/Art/Literature, awarded by the President of the Republic of Portugal.
- 2001 Knight of the Order of the 'Nederlandse Leeuw' awarded by the Queen of the Netherlands in appreciation for his achievements in science.

FROM THE BENCH TO THE BEDSIDE

The Role of Semantics in Enabling the Vision of Translational Medicine

Vipul Kashyap

Partners HealthCare System, Clinical Informatics R&D, USA

Abstract: Biomedical research and healthcare clinical transactions are generating huge volumes of data and information. At the same time, the results of biomedical research in the form of new molecular diagnostic tests and therapies are being increasingly used in the context of clinical practice. There is a critical need to speed "translation" of genomic research insights into clinical research and practice. In this talk, we will discuss challenges faced by a healthcare enterprise in realizing the vision of Translational Medicine, such as:

- The need to create structured and semantic representations of genotypic and phenotypic data such as clinical observations and molecular diagnostic tests.
- The need for cost-effective and incremental data integration for combining genotypic and phenotypic information at the point of care.
- The need for actionable decision support for suggesting molecular diagnostic tests and therapies in the context of clinical care.
- The need for knowledge update, propagation and consistency to keep abreast of the rapid pace of knowledge discovery being witnessed in the life sciences, a crucial pre-requisite to reduce the cost of knowledge acquisition and maintenance.

Semantics-based approaches to address the above-mentioned challenges, including the applicability of semantic web standard (RDF, OWL, Rules); and issues related to the value proposition of these technologies will be presented.

BRIEF BIOGRAPHY

Vipul Kashyap, PhD is a Senior Medical Informatician in the Clinical Informatics Research & Development group at Partners HealthCare System and is currently the chief architect of a Knowledge Management Platform that enables browsing, retrieval, aggregation, analysis and management of clinical knowledge across the Partners Healthcare System. Vipul received his PhD from the Department of Computer Science at Rutgers University in New Brunswick in the area of metadata and semantics-based knowledge and information management. He is also interested in characterization of the value proposition of semantic technologies in the enterprise context. Before coming to Partners, Vipul has held positions at MCC, Telcordia (Bellcore) and was a fellow at the National Library of Medicine. Vipul has published 2 books on the topic of Semantics, 40-50 articles in prestigious conferences and journals; and has participated in panels and presented tutorials on the topic of semantic technologies. Vipul sits on the technical advisory board of an early stage company developing semantics-based products, and represents

Partners on the W3C advisory committee and the HealthCare Information Technology Standards Panel (HITSP).

THE CANCER INFORMATICS ECOSYSTEM

A Case Study in the Accretion of Federated Systems based on Service Oriented Architectures, Semantic Integration and Computing Grids

David Hall

Research Triangle Institute in North Carolina, USA

Abstract: Information technology is playing an increasingly critical role in health and life sciences research due to the profound expansion in the scope of research projects in the post-genomic age. Robust data management and analysis systems are becoming essential enablers of these studies. Driven by funding agency requirements, funding opportunities, and grass roots organizing, efforts are underway to develop standards and technologies to promote large-scale integration of publicly-funded systems and databases including infrastructure developed for individual studies. Predicted benefits include an enhanced ability to conduct meta-analyses, an increase in the usable lifespan of data, a funding agency-wide reduction in the total cost of IT infrastructure, and an increased opportunity for the development of third party software tools. This presentation will critically examine efforts towards developing publicly-accessible interoperable and distributed production systems in the health and life sciences via ontologies, formal metadata, service oriented architectures, and grid computing models with a focus on several projects under the direction of the author in the area of cancer informatics.

BRIEF BIOGRAPHY

David Hall is a Senior Software Project Leader at RTI International based in North Carolina, USA. He leads teams of up to 30 developers implementing computer systems that support large biomedical and biotechnological research enterprises in cancer research, drug discovery, genetic epidemiology, and plant biotechnology. His area of interest is the practical application of bioinformatics and medical informatics methods, technologies, and standards in the development of production software. Particular topics of interest include data visualization, semantic integration, systems integration, and high performance computing. Recent clients include the US National Institutes of Health, GlaxoSmithKline, Syngenta, and Duke University. Data systems developed by David's group manage clinical and research data for nearly one million patients. Applications include data warehouses, metadata registries, workflow systems, high resolution image databases, analytical applications, and web services. David is currently Principal Investigator of the Informatics Support Center for the National Cancer Institute's Breast and Colon Cancer Family Registries. He holds a Ph.D. in Genetics from the University of Georgia and a B.S. in Computer Science from Wake Forest University.

ICT AND PERSONS WITH DISABILITIES

The Solution or the Problem?

Albert M. Cook

*Faculty of Rehabilitation Medicine, University of Alberta, Edmonton Alberta, Canada
al.cook@ualberta.ca*

Keywords: Assistive technologies, information and computer technologies, persons with disabilities.

Abstract: In order to lead full and productive lives, persons with disabilities need to have the same access to information and communication systems as the rest of the population. Advances in information and communication technologies (ICT) are occurring quickly, and the capability of technologies to meet the needs of persons with disabilities is growing daily. Future developments in assistive technologies (AT) and the successful application of these technologies to meet the needs of people who have disabilities are dependent on exploitation of these ICT advances. AT also involves the development of specialized interfaces such as the brain computer interface (BCI), adaptive interfaces that accommodate for changes in the user's physical skills, cognitive interfaces that allow understanding of the human technology interface by individuals with intellectual disabilities and systems that accommodate for user needs based on environmental sensing (e.g., GPS interfaces) and downloading of profiles to meet specific user needs. Universal Design (or design for all) calls for the design of products and environments to be usable by all people, to the greatest extent possible, without the need for adaptation or specialized design. In the physical world this often means ramps, curb cuts and other adaptations to the built environment to accommodate individuals who have disabilities. In the ICT world the barriers to access are technological, and the goal for ICT universal design is to have an environment with enough embedded intelligence to be easily adaptable to the varying cognitive, physical and sensory skills of a wide range of individual's in order to meet their productivity, leisure and self care needs. If ICT advances are not adaptable enough to be accessible to persons with disabilities it will further increase the disparity between those individuals and the rest of the population leading to further isolation and economic disadvantage. On the other hand, availability of these technologies in a transparent way will contribute to full inclusion of individuals who have disabilities in the mainstream of society.

BRIEF BIOGRAPHY

Dr. Albert Cook is Professor of Speech Pathology and Audiology and currently Dean of the Faculty of Rehabilitation Medicine and Chair of the Health Sciences Council at the University of Alberta. Dr. Cook has worked with interdisciplinary teams to develop assistive devices and to assess the effectiveness of technology being used by persons with disabilities. Dr. Cook is also associated with the I CAN Centre the Glenrose Rehabilitation Hospital. He was formerly Professor of Biomedical Engineering at California State University, Sacramento where he established the graduate program in biomedical engineering and directed it for 12 years. He also served as Co-Director of the Assistive Device Center in Sacramento, California,

helping over 500 persons with disabilities to identify and obtain assistive technologies.

He received his Bachelor of Science in Electrical Engineering at the University of Colorado, a Masters in Bioengineering and his doctorate from the University of Wyoming He is a member of Tau Beta Pi, Phi Kappa Phi and Gold Key honorary societies.

Dr. Cook co-authored with Janice Polgar, OTR, Cook and Hussey's *Assistive Technologies: Principles and Practice* 3rd edition, published in October 2007 by Elsevier. He has co-edited three other textbooks with John Webster and others and has written numerous chapters in rehabilitation and biomedical engineering texts and monographs.

Dr. Cook's research interests include augmentative and alternative communication, biomedical instrumentation and assistive technology design, development and evaluation. His most recent

research has focussed on the use of robotics with young children who have severe disabilities to develop and assess cognitive and linguistic skills. He has US and foreign patents and numerous publications and conference presentations in these areas. He has been principal investigator on research and training grants in augmentative communication, assistive technologies and biomedical engineering. Dr. Cook is Past-President and Fellow of RESNA, a major professional society for assistive technology practitioners in North America. He has also served in national United States positions in the Institute of Electrical and Electronic Engineers Engineering in Medicine and Biology Society, the American Society for Engineering Education, the Biomedical Engineering Society, the International Society for Augmentative and Alternative Communication and the Association for the Advancement of Medical Instrumentation. Dr. Cook is a registered professional engineer (electrical) in California.

1 ICT AND PERSONS WITH DISABILITIES TECHNOLOGY AND PROGRESS

Societal Progress requires change much of which is accomplished through advances in technology. In his book, *A Short History of Progress*, Ronald Wright (2004) points out that this characteristic has been true for millions of years as societies have advanced through greater utilization of technology.

Wright goes on to describe the problems that technology typically creates such as over consumption, environmental ruin, and separation of classes. These problems are amplified for people who have disabilities, and they lead to a gap in the access to work, self care and community participation for persons with disabilities compared to the general population. Since people with disabilities often depend on technologies for societal participation, the lack of availability of accessible technology or the obsolescence of accessible technologies isolates them further. This is an extension of the concept of the “digital divide” that separates people along socioeconomic lines based on their access to ICT. I refer to it as the “disability gap”.

2 ADVANCES IN INFORMATION AND COMMUNICATION TECHNOLOGIES (ICT)

The 21st Century is characterized by a continuous move from a machine-based to a knowledge based economy (Ungson & Trudel, 1999). In this shift, the basis of competence is changing to knowledge skills from machine skills. Information currently amounts to 75% of value added to products. This will continually increase, and connectivity will be the key to business success. There is also a move from a regional or national scope of business influence to a global scope, in which virtual networks dictate organizational structures.

Key players in business development are becoming communication suppliers with the move from host-based to network based systems. Telephone, cable TV and internet service providers control commercial growth. Along with these changes networks will become more graphically-based moving increasingly from text-based systems. In order to lead full and productive lives, persons with disabilities need to have the same access to this new information and communication system as the rest of the population.

2.1 What Can we Expect from Technology in the Next 20 Years?

The cost of information technology is continually dropping for comparable or increased computing power and speed. There is also a greater understanding of the biological/physical interface for the control of computers. The human computer interface (HCI) is being developed to be more human-like, more user oriented and more intelligent-providing additional capabilities for searching, processing and evaluating information.

There are a number of changes that are likely to occur over the next few years (Applewhite, 2004). There will be an increase in automated transactions between individuals and organizations enabling people to complete all transactions without face-to-face interactions. It is expected that we will achieve equalized access to the web and information between the developed and developing world. Embedded systems will dramatically increase with application such as “intelligence in the doorknob” that recognizes the owner and doesn’t require key manipulation. We are likely to see much greater

understanding of the biological to physical interface for the control of computers.

2.2 Changes in Mainstream Tech with AT Implications

There are many examples of emerging mainstream technologies with potential for assisting people with disabilities to access ICT systems. A few of these are described in this section.

Display-based assistive technologies present an array of choices for a user to select from (Cook & Polgar, 2007). This is often referred to as scanning since the choices are highlighted sequentially and then chosen using some sort of gross movement. One of the problems associated with this approach is that there must be a physical display for making selections. This often requires the overall system to be larger and more bulky or places a display between a user and a communication partner. A new development is a direct retinal display that creates an image that overlays the view of a real object (Lewis, 2004). The retinal display is low powered because it is shined on the retina directly. Scanning light into the eye allows the image to overlay an object such as a communication partner's face, enabling eye contact and small size. The scanning array could be the retinal image, since the display scans across the retina power levels can be kept low for safety.

Another development is 3-D displays that create a more intuitive view of objects, events and activities (Lewis, 2004). This type of display may be helpful to individuals who have cognitive disabilities. It might also create new challenges for individuals with visual limitations.

Embedded automatic speech recognition is being developed for PDAs because of the need for keyboards with more and more functions and the limitations of very small keyboards (Kumagai, 2004). This feature could be very useful to reduce individuals who have limited hand function or for those who cannot see the keyboard to make entries.

3 MEETING THE ICT NEEDS OF PERSONS WITH DISABILITIES

Over the centuries, our ability to make tools is what distinguishes us as human, but our tools ultimately control us by making us dependent on them (Wright, 2004). This dependence is less optional for people who have disabilities

3.1 Impact of Technology Advances on People who have Disabilities

Technology advances increase the gap between people who have disabilities and those who don't (Wright, 2004). All societies become hierarchical with an upward concentration of wealth (including aggregations of technology tools) that ensures that "there can never be enough to go around", and this disparity contributes to the "digital divide" and the "disability gap". As advances occur more quickly, the gap widens faster and the people who are poor and/or disabled lose out even more completely and faster. This is a characteristic of cultural and societal "progress" over centuries—technology drives change, and creates both positive and negative outcomes in the process.

The prognosis is not good for people with disabilities unless there is considerable effort to keep them connected to ICT and thereby to commerce, employment and personal achievement. There are two fundamental approaches to this problem: (1) make mainstream technologies accessible to people who have disabilities, or (2) design special purpose technologies specifically for people with disabilities. The former approach is referred to as *universal design* or *design for all*. The second approach utilizes *assistive technologies*.

3.2 Implications for Assistive Technologies

Access to ICT for people with disabilities is a significant global problem, and it has major implications for assistive technologies. There is a constant challenge to keep ICT systems accessible to persons who have disabilities as mainstream advances occur and adaptations become potentially incompatible with the new systems. Communication technologies change rapidly, and each change may result in the need to re-design accessible interfaces. We are closer to the goal of having assistive technology adaptations available when the mainstream consumer product ships, but there are still many problems with "workarounds" necessary to make mainstream operating systems, productivity software and internet access accessible to people with disabilities.

Development and maintenance of access to ICT must be driven by the needs of people with disabilities. Developments which broaden the scope, applicability and usability of the human technology

interface will be driven, at least in part by the needs of people who have disabilities.

The Internet (e-mail and chat rooms) have the advantage of anonymity, and this can be a major benefit to individuals who have disabilities. Because the person's disability is not immediately visible, people who have disabilities report that they enjoy establishing relationships with people who experience them first as a person and then learn of their disability. For example, Blackstone, (1996) describes some of the advantages of e-mail for individuals who have disabilities. Since the receiver of the message reads it at a later time composition can be at a slower speed. The person with a disability can communicate with another person without someone else being present, establishing a greater sense of privacy than situations in which an attendant is required. It is also possible to work from any location-avoiding some transportation problems

3.3 Universal Design

Increasingly, commercial products are being designed to be usable by all people, to the greatest extent possible, without the need for adaptation or specialized design (NC State University, The Center for Universal Design, 1997).

3.3.1 General Principles of Universal Design

Features are built into products to make them more useful to persons who have disabilities (e.g., larger knobs; a variety of display options--visual, tactile, auditory; alternatives to reading text--icons, pictures) are built into the product. This is much less expensive than modifying a product after production to meet the needs of a person with a disability. The North Carolina State University Center for Universal Design, in conjunction with advocates of universal design, have compiled a set of principles of universal design, shown in Box 1. This center also maintains a Web site on universal design (www.design.ncsu.edu/cud).

3.3.2 Universal Design for ICT

In universal design for ICT the barriers are technological rather than political and economic barriers that characterize architectural and commercial product design (Emiliani, 2006). The goal of universal design for ICT is to have an environment with enough embedded intelligence to be easily adaptable. The features of future information services are that there will be no clearly

predefined service and little distinction between interpersonal communication and access to information. Services will need to be highly interactive, inherently multimedia, sensory multimodal (i.e., access via auditory or visual means is equally possible). To achieve this cooperation between users or representatives of users is critical in a variety of contexts of use. The overall goal is to have access to information involving communities of users with a wide range of motor, sensory and cognitive skills.

ONE: EQUITABLE USE

The design is useful and marketable to people with diverse abilities.

TWO: FLEXIBILITY IN USE

The design accommodates a wide range of individual preferences and abilities.

THREE: SIMPLE AND INTUITIVE USE

Use of the design is easy to understand, regardless of the user's experience, knowledge, language skills, or current concentration level.

FOUR: PERCEPTIBLE INFORMATION

The design communicates necessary information effectively to the user, regardless of ambient conditions or the user's sensory abilities.

FIVE: TOLERANCE FOR ERROR

The design minimizes hazards and the adverse consequences of accidental or unintended actions.

SIX: LOW PHYSICAL EFFORT

The design can be used efficiently and comfortably and with a minimum of fatigue.

SEVEN: SIZE AND SPACE FOR APPROACH AND USE

Appropriate size and space is provided for approach, reach, manipulation, and use regardless of user's body size, posture, or mobility.

Box 1: Principles of Universal Design From North Carolina State University, The Center for Universal Design, 1997.

In addition to Universal Design for ICT, access to capabilities of mainstream technologies includes individualized assistive technologies that are easily – customized. This in return requires an increased understanding of the biological/physical interface for the control of assistive technologies and expanded availability of embedded systems networks.

3.4 A Working Definition of Assistive Technologies

The *International Classification of Functioning, Disability and Health* (ICF) is a system developed by the World Health Organization (WHO) that is designed to describe and classify health and health related states. These two domains are described by

body factors (body structures and functions) and individual and societal elements (activities and participation) (WHO, 2001). The ICF recognizes two contextual factors that modify health and health related states: the environment and personal factors (WHO, 2001). Environmental elements include assistive technologies in relation to activities of daily living, mobility, communication, religion and spirituality as well as in specific contexts such as education, employment and culture, recreation and sport (WHO, 2001). Other environmental elements such as access to public and private buildings, and the natural and built outdoor environments, also have implications for assistive technologies.

A commonly used definition of assistive technology is from the Technical Assistance to the States Act in the United States (Public Law (PL) 100-407): *Any item, piece of equipment or product system whether acquired commercially off the shelf, modified, or customized that is used to increase, maintain or improve functional capabilities of individuals with disabilities.*

3.4.1 Hard and Soft Technologies

Odor (1984) has distinguished between hard technologies and soft technologies. Hard technologies are readily available components that can be purchased and assembled into assistive technology systems. The main distinguishing feature of hard technologies is that they are tangible. On the other hand, soft technologies are the human areas of decision making, strategies, training, concept formation, and service delivery as described earlier in this chapter. Soft technologies are generally captured in one of three forms: (1) people, (2) written, and (3) computer (Bailey, 1997). These aspects of technology, without which the hard technology cannot be successful, are much harder to obtain. Soft technologies are difficult to acquire because they are highly dependent on human knowledge rather than tangible objects. This knowledge is obtained slowly through formal training, experience, and textbooks such as this one. The development of effective strategies of use also has a major effect on assistive technology system success. Initially the formulation of these strategies may rely heavily on the knowledge, experience, and ingenuity of the assistive technology practitioner. With growing experience, the assistive technology user originates strategies that facilitate successful device use. There is a false belief that progress is solely driven by “hard” technological change. The gap between the general public and persons with

disabilities can only be closed by gains in both soft and hard technologies

3.4.2 Mainstream Technologies to Specially Designed Technologies: A Range of Options

As illustrated in Figure 1, the needs of people with disabilities can be met in a number of ways. Off the shelf “standard” (i.e., mainstream technologies) commercially available devices (especially those designed using the principles of universal design) can often be used by people with a variety of disabilities. For example, standard personal computers designed for the general population are often used by persons with disabilities. Sometimes these need to be modified however, to make them useable. Another type of commercially available device is one that is mass-produced but specifically designed for individuals with disabilities (*special commercially available devices*). These devices often need to be modified to meet the needs of a specific individual. Our goal is to reduce the amount of modification necessary and to make mainstream technologies as accessible as possible. However, there will always be a portion of the disabled population that will require specifically designed assistive technologies.




Figure 1: This diagram shows the progression from commercially available devices for the general population and commercially available devices for special populations to modified devices and custom devices. From Cook and Polgar, (2007).

3.5 The Human Technology Interface for ICT

3.5.1 General Concepts

It is estimated that as many as 40 million persons in the United States alone have physical, cognitive, or sensory disabilities (Lazzaro, 1999). The world-wide impact is significantly larger. If these people are to compete on an equal basis with non-disabled individuals, then it is extremely important that the internet be accessible to all. As the internet becomes more and more dependent on multimedia representations involving complex graphics, animation, and audible sources of information, the challenges for people who have disabilities increase. In order for access to the Internet to be useful to people with disabilities, the accessibility approach must be independent of individual devices. This means that users must be able to interact with a *user agent* (and the document it renders) using the input and output devices of their choice based on their specific needs. A **user agent** is defined as software to access Web content (www.w3.org/wai). This includes desktop graphical browsers, text and voice browsers, mobile phones, multimedia players, and software assistive technologies (e.g., screen readers, magnifiers) that are used with browsers. The person with a disability interacts with technology through the Human Technology Interface (HTI) (Cook and Polgar, 2007).

The graphical user interface (GUI) has both positive and negative implications for persons with disabilities. The positive features are those that apply to non-disabled users (e.g., use of icons, recognition rather than recall memory, screen icons for the same task look the same, operations such as opening and closing files are always the same). The GUI is the standard user interface because of its ease of operation for novices and its consistency of operation for experts. The latter ensures that every application behaves in basically the same way. People with motor disabilities may not have the necessary physical (eye-hand coordination) and visual skills to navigate the GUI. Modification of the GUI to allow specialized access (see Figure 1) can also be more challenging for GUI-based operating systems.

As networks are expanded and more devices (e.g., cell phones, PDAs) have open architectures, it will be possible to download profiles, adaptations and special instructions that enable adaptable systems to be developed to meet the needs of people

who have disabilities. Some examples are (1) trainable hearing aids that adjust automatically to the environments in which they are used; (2) a “Smart House” that assesses occupants current state and the state of various home utilities to aid with common activities of daily living, provides feedback should residents become disoriented or confused and report medical emergencies automatically; an orientation and direction finding device that senses the current location (via GPS) and gives directions to a desired location for individuals who cannot read maps because of visual or cognitive disabilities.

3.5.2 Access for Motor Impairment

There are a significant number of people who cannot effectively use standard keyboards, mouse controls or switches. It is likely that we will see a much greater understanding of the biological/physical interface for the control of computers in the future (Applewhite, 2004).

One approach that may offer promise is the brain computer interface (BCI). BCI systems may be grouped into a set of functional components including the input device, amplification, feature extraction, feature translation and user feedback (Mason and Birch, 2003). Signals are mathematically analyzed to extract features useful for control (Fabiani, Mcfarland, Walla, and Pfurtscheller 2004). Features or signals that have been used include slow cortical potentials, P300 evoked potential, sensorimotor rhythms recorded from the cortex and neuronal action potentials recorded within the cortex). A typical task for a user is to visualize different movements or sensations or images.

Another approach to cursor control is the use of a digital camera and image recognition software to track a particular body feature to control an on-screen mouse cursor (Betke, Gips and Fleming, 2002). The most easily tracked feature is the tip of the nose, but the eye (gross eye position not point of gaze), lip, chin and thumb have also been used. Non-disabled subjects used this approach and found that the camera mouse was accurate but slower than a typical hand-controlled mouse. Using an on-screen keyboard the camera mouse was half as fast as a regular mouse in a typing task, but the accuracy obtained was equivalent on each system. More and more computers have built-in cameras, so the camera mouse requires only software to capture the body feature image and interpret its movement as mouse commands. This may lead to wider application of this technique.

There are many other approaches that are used to provide access to and control over technologies for people with severe motor disabilities (Cook and Polgar, 2007) \. These range from keyboards of various type, to automatic speech recognition to mouse and mouse emulators systems to single and multiple switches.

3.5.3 Access for Cognitive Impairment

Cognitive disabilities include a wide range of skills and deficiencies. Learning disabilities typically involve significant difficulties in understanding or in using either spoken or written language, and these difficulties may be evident in problems with reading, writing, mathematical manipulation, listening, spelling or speaking (Edyburn, 2005). These limitations make it increasingly difficult to access complicated Web sites that may include flashing pictures, complicated charts, and large amounts of audio and video data. While there are assistive technologies that are specifically designed to address these areas (discussed later in this chapter), many of the technological tools are useful for all students, and are part of instructional technology (Ashton, 2005). Even the so-called assistive technologies have features (e.g., multimedia, synthetic speech output, voice recognition input) that are useful to all learners.

For individuals with acquired cognitive disabilities due to injury (e.g., traumatic brain injury) or disease (e.g., stroke (CVA) or dementia) changing features such as font size, background/foreground color combinations, contrast, spacing between words, letters and paragraphs and using graphics can all improve access to screen-based information. Another technological concept for these individuals is a cognitive prosthesis, which is a custom-designed computer-based compensatory strategy that directly assists in performing daily activities¹. It may also include additional technologies such as a cell phone, pager, digital camera or low tech approaches

Persons with intellectual disabilities have difficulties with memory, language use and communication, abstract conceptualization, generalization and problem identification/problem solving. Characteristics of the HTI that are important for these individuals include simplicity of operation, capacity of the technology to support repetition, consistency in presentation, and inclusion of

multiple modalities (e.g., speech, sounds and graphical representations) (Wehmeyer, Smith and Davies, 2005).

An example of technology designed for cognitive needs is the Planning and Execution Assistant and Trainer (PEAT). It is a PDA-based personal planning assistant designed to assist individuals with cognitive disorders due to brain injury, stroke, Alzheimer's disease, and similar conditions (Levinson, 1997). PEAT employs artificial intelligence to automatically generate plans and also to revise those plans when unexpected events occur. PEAT uses a combination of manually entered schedules and a library of stored scripts describing activities of daily living (e.g., morning routine or shopping). Scripts can be used for both planning and for execution. Planning involves a simulation of the activity with key decision points presented and prompts (auditory and visual) supplied necessary to aid the individual through the planning process. The plan to be executed can be either the stored script or a modified script based on the simulation. The PEAT artificial intelligence software generates the best strategy to execute the required steps in the plan (LoPresti, Mihailidis, and Kirsch, 2004). PEAT also automatically monitors performance, and corrects schedule problems when necessary.

3.5.4 Access for Auditory Impairment

Since web pages are a mixture of text, graphics, and sound, people who are deaf may be prevented from accessing some information unless alternative methods are available. The primary approach for these individual is the use of the Microsoft Synchronized Accessible Media Interchange (SAMI), which allows authors of Web pages and multimedia software to add closed captioning for users who are deaf or hard of hearing. This approach is similar to the use of closed captioning for television viewers. The W3C WAI SMIL (www.w3.org/WAI) is designed to facilitate multimedia presentations in which an author can describe the behavior of a multimedia presentation, associate hyperlinks with media objects, and describe the layout of the presentation on a screen

Trainable hearing aids adjust automatically to the environments in which they are used through access to embedded information networks. This allows automatic adaptation to changing noise levels and environments.

¹ Institute for Cognitive Prosthetics, <http://www.brain-rehab.com/definecp.htm>

3.5.5 Access for Visual Impairment

The W3C WAI user agent guidelines are based on several principles that are intended to improve the design of both types of user agents. The first is to ensure that the user interface is accessible. This means that the consumer using an adapted input system must have access to the functionality offered by the user agent through its user interface. Second, the user must have access to document content through the provision of control of the style (e.g., colors, fonts, speech rate, and speech volume) and format of a document. A third principle is that the user agent help orient the user to where he is in the document or series of documents. In addition to providing alternative representations of location in a document (e.g., how many links the document contains or the number of the current link), a well-designed navigation system that uses numerical position information allows the user to jump to a specific link. Finally, the guidelines call for the user agent to be designed following system standards and conventions. These are changing rapidly as development tools are improved.

Communication through standard interfaces is particularly important for graphical desktop user agents, which must make information available to assistive technologies. Technologies such as those produced by the W3C include built-in accessibility features that facilitate interoperability. The standards being developed by the W3C WAI provide guidance for the design of user agents that are consistent with these principles. The guidelines are available on the W3C WAI Web page (www.w3.org/wai).

3.5.6 Other ICT Access

Cellular telephones are becoming more powerful with capabilities approaching that of personal computers. This expanded capability will provide significant advantages for people with disabilities, especially those with low vision or blindness. Three changes will be particularly valuable to people who have disabilities: (1) standard cell phones will have sufficient processing power for almost all the requirements of persons with visual impairments, (2) software will be able to be downloaded into these phones easily, (3) wireless connection to a worldwide network will provide a wide range of information and services in a highly mobile way (Fruchterman, 2003). Because many of these features will be built into standard cell phones the cost will be low and reachable by persons with disabilities. A major advance will occur if the cell

phone industry moves away from proprietary software to an open source format providing the basis for a greater diversity of software for tasks such as text-to-speech output, voice recognition and optical character recognition in a variety of languages. Many applications for people with disabilities will be able to be downloaded from the internet. With expanded availability of embedded systems, it will be possible for a user to store their customized programs on the network and download them as needed from any remote location.

Downloading a talking book program into a cell phone can provide access to digital libraries for persons who are blind. Outputs in speech or enlarged visual displays can be added as needed by the user. With a built-in camera and network access a blind person could obtain a verbal description of a scene by linking to on-line volunteers who provide descriptions of images. These applications will depend on the increasing application of universal design in information technology products (Tobias, 2003). These applications include ATMs, cell phones, vending machines and other systems that are encountered on a daily basis (Tobias, 2003).

4 INFRASTRUCTURE FOR FUTURE ACCESSIBILITY

The infrastructure for future accessibility consists of: (1) an expanded, smarter and more available "real" and "virtual" internet, (2) Home automation systems that are smarter and have greater interconnectivity, (3) universal design principles that are applied more widely, (4) alternative approaches for accessing information technologies, and (5) special-purpose assistive technologies.

The Infrastructure for future accessibility will depend on several factors. These include: Web-based virtual systems, home automation, universal design for ICT, alternatives for accessing information technologies and special-purpose assistive technologies. In addition there is a continuing need for the development of soft technology tools.

If ICT advances are not adaptable enough to be accessible to persons with disabilities it will further increase the disparity between those individuals and the rest of the population leading to further isolation and economic disadvantage. On the other hand, availability of these technologies in a transparent way will contribute to full inclusion of individuals who have disabilities in the mainstream of society.

5 CONCLUSIONS

The move to the information age offers great promise for persons with disabilities. It also holds great threats for persons with disabilities. Constant vigilance is required to insure that information technologies remain accessible and responsive to the needs of persons with disabilities. The future for persons with disabilities will not be driven by advances in technology, but rather by how well we can take advantage of those advances for the accomplishment of the many tasks of living that require technological assistance

6 SUMMARY

Anticipated changes in technologies coupled with the focus on the social aspects of disability, provide a significant opportunity for major advances in the degree to which individuals with disabilities can participate in all aspects of life, including work, school, leisure and self care.

Technological advances will be particularly important as the percentage of the population that is elderly rises. Concepts from universal design will be important in ensuring that this segment of the population remains active and is able to participate in society. This new group of elderly individuals will also be more experienced with computers and other technologies than their predecessors and they may well demand greater performance and adaptability from both assistive technologies and mainstream ICT (e.g., telephones, internet communication).

The percentage of individuals with long-term disabilities who join the over 65 age group will also increase. These individuals will have been long-term users of assistive technologies, and their experience will have major implications for developments to meet future needs.

While much of what I have described is conjecture, it is based on modest extrapolation from the current state of the art. There are some things that we know with a high degree of certainty. We know that computer systems will be faster, have more memory be smaller and be less expensive for the same or greater functionality. We also know that the communication channel bandwidth will continue to increase allowing much more information and much more sophisticated information processing. Finally, it is clear that people with disabilities will continue to assert their right to fully participate in society.

Technological advances also raise questions for people who have disabilities. The most important of these is whether accessibility will keep pace with technological developments. For example, will assistive technologies for input and output be compatible with the user agents and operating systems of tomorrow. A second major question is whether the needs of persons with disabilities will be a driving force in future technological developments. Will people who have disabilities have to adapt to the existing technologies based on characteristics for non-disabled people or will universal design become a greater reality? In the latter case, adaptations will become less important and accessibility will become the rule rather than the exception.

For people who have disabilities, there are significant implications of emerging information processing technologies. If not closely monitored, these could result in less rather than more access to the new information economy for persons with disabilities. Despite the wider use of universal design principles, there will still be a need for effective assistive technology design and application if individuals with disabilities are to realize the full potential of the new information age.

REFERENCES

- Applewhite A. (2004). 40 years: the luminaries. *IEEE Spectrum*, 41(11), 37-58.
- Ashton T.M. (2005). Students with learning disabilities using assistive technology in the inclusive classroom, pp. 229-238, in Edyburn D., Higgins K. & Boone R. (eds.): *Handbook of Special Education Technology Research and Practice*, Whitefish Bay, Wisconsin: Knowledge by Design, Inc.
- Bailey, R.W: (1996). *Human performance engineering*, ed 2, Upper Saddle River, NJ, , Prentice Hall.
- Betke M., Gips J. & Fleming P: (2002). The camera mouse: Visual tracking of body features to provide computer access for people with severe disabilities. *IEEE Trans. Neural Systems and Rehabilitation Engineering*, 10(1): 1-10.
- Blackstone S. (1996). The Internet: what's the big deal *Augment Commun News* 9(4):1-5.
- Cook A.M. & Polgar J.M. (2007). *Cook and Husey's Assistive Technologies: Principles and Practice*, 3rd ed St. Louis: Elsevier.
- Edyburn D.L. (2005). Assistive technology and students with mild disabilities: from consideration to outcome measurement, pp 239-270, in Edyburn D, Higgins K. & Boone R. (eds.): *Handbook of Special Education Technology Research and Practice*, Whitefish Bay, Wisconsin: Knowledge by Design, Inc.,

- Emiliani, P.L. (2006). Assistive technology (AT) versus Mainstream Technology (MST): The research perspective. *Tech Disab*, 18, 19-29.
- Fabiani G.E., Mcfarland D.J., Wolpaw J.R., & Pfurtscheller G (2004). Conversion of EEG activity into cursor movement by a brain-computer interface (BCI). *IEEE Trans. Neural Systems and Rehab Engr*, 12: 331-338.
- Fruchterman J.R: (2003) In the palm of your hand: A vision of the future of technology for people with visual impairments. *J Vis Impair Blindness*, 97(10), 585-591.
- Kumagai, J. (2004).Talk to the machine. *IEEE Spectrum*, 39(9): 60- 64.
- Lazzaro J.L. (1999). Helping the web help the disabled, *IEEE Spectrum* 36(3): 54-59
- Levinson R.L. (1997): The planning and execution assistant and trainer. *J. head trauma rehabilitation*, 12(2): 769-775.
- Lewis, J.R. (2004). .In the eye of the beholder, *IEEE Spectrum*, 41(5)24- 28.
- LoPresti E.F., Mihailidis A. & Kirsch N. (2004): Assistive technology for cognitive rehabilitation: State of the art. *Neuropsychological Rehabilitation*, 14(1): 5-39.
- Mason S.G. & Birch G.E: (2003).A general framework for brain-computer interface design. *IEEE Trans. Neural Systems and Rehab Engr*, 11, 70-85.
- Odor P. (1984): Hard and soft technology for education and communication for disabled people, *Proc Int Comp Conf*, Perth, Western Australia,
- Tobias J: (2003). Information technology and universal design: An agenda for accessible technology. *J Vis Impair Blindness*, 97(10), 592-601.
- Ungson G.R. and Trudel J.D. (1999). The emerging knowledge-based economy. *IEEE Spectrum*, 36(5):60-65.
- Wehmeyer M.L., Smith S.J., Palmer SB, Davies D.K. & Stock S.E: (2004). Technology use and people with mental retardation. *International Review of Research in Mental Retardation*, 29: 291-337.
- World Health Organization (2001): *International classification of functioning disability and health-ICF*, Geneva, World Health Organization
- Wright, RA. (2004). *A Short History of Progress*, Toronto: Anansi Publications.

FULL PAPERS

HOW TO EVALUATE HUMAN FACTORS AFFECTING WIRELESS BIOMEDICAL SENSORS

Identifying Aspects of Patient Acceptance based on a Preliminary Clinical Trial

Rune Fensli¹ and Egil Boisen²

¹*Faculty of Engineering and Science, University of Agder, Grimstad, Norway*

²*Department of Health Science and Technology, Aalborg University, Aalborg, Denmark*
rune.fensli@uia.no, eb@hst.aau.dk

Keywords: Quality of life, Patient Acceptance of Health Care, Ambulatory Monitoring, Telehomecare, Usability.

Abstract: In this paper, we highlight some important aspects of how to evaluate patient acceptance with wearable sensors recording vital signs information, used in a telehomecare environment. Questions of human factors and patient satisfaction need to be addressed, where the patient is able to carry out daily life activities in his own environment. We compare results from a preliminary clinical trial with patients using a wireless ECG sensor for three days out-of-hospital service, to available published results from telehomecare projects, and propose important aspects and plans for future investigations. It is important not only to observe the patient, but also the possible changes in the family situation when a sick patient is transferred to his own home for active treatment. At the same time, emotional barriers and stigmatisation are challenging factors where time is needed to let the patient adopt this new situation. Therefore, measures should be on an on-going basis with long-term use of the technological equipment in order for the patient to integrate this into his body scheme as well as daily activities. Of special importance will be the evaluation of the communication between the patient and the health professionals, as quick feedback from the doctor to the patient on his or her own measurements is of utmost importance.

1 INTRODUCTION

Several research projects have been focusing on wearable biomedical sensors and their benefits for ambient assisted living, where the patients are remotely monitored by different sensors placed on the body for vital signs recording (Jones et al., 2006). However, little efforts have been done investigating how the patients experience and manage this new technology.

Chronically ill patients experience a greater degree of freedom and are more involved in the treatment with daily monitoring of vital information during hospitalization in their own home, than with the traditional treatment procedures at hospital. Introducing advanced medical technology in the patient's own home will influence the patient's situation as it makes empowerment and self-management possible (Barlow et al., 2002).

At the same time, coordinated follow-up and new workflow procedures for the health-care services need to be implemented in order to give the patient

satisfactory support by virtual visits in his home (Wootton and Kvedar, 2006). However, this support also must be integrated in the self-monitoring of vital signs information performed by the patients, with understandable interpretations of the results.

The primary goal of developing wireless ECG sensors is to find ways of monitoring the everyday life of the patient as closely as possible. In other words, the rationale is to increase the quality of data being collected by decreasing the impact on the patient's everyday life when he or she is being monitored. When implementing such equipment it is therefore necessary to allow the patient to integrate it into his or her daily activities. Also, we argue that this integration concerns a person's body scheme, being a dynamic, and typically unaware, representation of the positions of one's body parts (Haggard and Wolpert, 2005). To monitor this process, measures should therefore be based on long-term use of the wireless equipment in an everyday life situation.

In this paper we will focus on how the technology acceptance can be monitored in order to highlight obstacles and possible improvements both in

the technology solutions and in coordination with the health care services. Based on experiences from a preliminary clinical trial, we try to identify some important aspects of patient acceptance and present plans for future projects involving observations of user acceptance.

1.1 Usability and Telehomecare

Home telehealth is a growing field of patient treatment and follow-up, but most research studies have focused on the technology, and so questions of human factors and patient satisfaction need to be addressed. A systematic review of studies of patient satisfaction with telemedicine was done by Mair and Whitten, arguing that *“available research fails to provide satisfactory explanations of the underlying reasons for patient satisfaction or dissatisfaction with telemedicine”*. They found that the studies concerning patient satisfaction mainly used simple survey instruments and that many of the studies had only a few participants (Mair and Whitten, 2000).

According to Friedman & Wyatt, usability studies can be useful in the evaluation of new biomedical equipment, with the aim of observing speed of use, user comments and completion of simple tasks (Friedman and Wyatt, 2006). They describe field function studies as useful in the validation of prototypes or released versions of new biomedical equipments, but the trials should be conducted in situations with real use of the equipment, and with the aim of observing speed and quality of data collected and accuracy of advice given from the devices.

The principles of user-centered design can be useful when designing new telehomecare devices. Adlam and colleagues describe how user evaluation can be implemented in the design process (Adlam et al., 2006). They start with discovering the “real problem” and the users’ requirements, which can be accomplished with a simple prototype demonstrating the actual functionality of a device. However, interacting with real users in their own environments will be a challenging task, and prototypes with limited functionality can give restrictions on the use of the new solution to be developed, resulting in preliminary tests, which do not reflect a more complicated use situation, i.e. the daily activities of the patients.

Kaufmann and colleagues developed a system design for both usability testing and usability inspection in the patient’s home for a diabetes telemedicine system. First, they used a cognitive walkthrough to identify goals and sequences of actions to anticipate potential user problems. Second, a field usability testing was performed as a series of tasks to be ac-

complished by the subjects in their home and closely followed by semi-structured interviews to reveal problems and barriers to efficient and safe use of the system (Kaufmann et al., 2003). A similar approach has been suggested by Kushniruk and Patel, who have developed a low-cost portable usability testing solution intended to be used by patients in their homes (Kushniruk and Patel, 2004).

In a study of the patient’s perspectives on high-tech homecare technology, Lehoux found that the user-acceptance was shaped by different types of anxiety, which mainly was related to the actual equipment and the procedures. At the same time, the patients complained of reductions in daily activities as well as feelings of stigmatization, as they tended to withdraw from social activities in order to hide the medical equipment from the eyes of visitors. For a patient wearing a permanent catheter, this will alter the patient’s body image (Lehoux, 2004). Lehoux found that good usability is dependant on competence, where lay people are supposed to use high-tech medical devices, as well as on technical and human dimensions, where the technology is integrated into the patients’ private and social lives.

Hopp and colleagues measured the outcome for patients receiving telehealth home-care and used a questionnaire at baseline and after six months, where they used a modified version of the SF-36 to measure Health Related Quality Of Life (HRQOL) (Hopp et al., 2006). In addition they used separate questions to ask about satisfaction with the telehealth equipment for the intervention group, with questions from the National Ambulatory Care Survey and modified for use in evaluating telehealth services. They found a high degree of satisfaction with the telemedicine equipment, but few patients reported that their family members had been taught how to use the equipment installed in their homes.

2 OBJECTIVES

In this paper, we present some ideas of how to evaluate human aspects of telehomecare solutions, which are based on preliminary results from a clinical study where patients have used wearable sensors for a three day period of out-of-hospital service. We focus on the impact on everyday life during the patients’ use of the wireless ECG-sensor; the experiences with daily behaviour as well as general patient satisfaction. The important question is which factors to include in future evaluations, and how to observe the degree of the patients’ influence on new telehomecare services, together with their perception

and possible adoption of the new technology.

3 METHODS

Patients referred to long-term ambulatory “Holter” arrhythmia procedures at the outpatient clinic at Sørlandet Sykehus HF in Arendal, Norway, were asked to participate in a study wearing our newly developed wireless ECG-sensor (Fensli et al., 2005). By signing the informed consent form, they participated in the study during their ordinary arrhythmia investigation. The inclusion criteria was patients with suspected arrhythmias, and the exclusion criteria was patients with dementia who were anticipated not being able to handle the equipment and contribute in filling out the required questionnaires.

After signing the informed consent form, the patients were given information of the research project and they received several questionnaires to be filled in during the time they used the wireless ECG-system as a usability study. Because influences from participating in a three day trial of the new recording solutions can be a bias in the evaluation of patient acceptance, a reference group with patients undergoing a “normal” routine investigation at the hospital using conventional “Holter” monitoring equipment (Huntleigh Healthcare) was established.

During the period from November 2006 to May 2007, 11 patients were enrolled in the study, and 25 patients were included in the reference group. The questionnaires focused on several topics, among them questions defining the Dimensions: Hygienic Aspects (3 items), Physical Activity (6 items), Skin Reactions (2 items), Anxiety (3 items), and Equipment (5 items) (Fensli et al., 2008). After a three day period of arrhythmia investigation with the use of the wireless sensor, 4 patients in the intervention group were followed-up with qualitative interviews, in order to discover more general experience with regard to the use of a wearable sensor.

4 EXPLORATORY RESULTS

In the questionnaire, the patients are asked about their experiences with the use of the equipment, where we have used an 11-point semantic differential scale. The patients should evaluate their experience of using the wireless equipment in terms of their agreement to the statements describing actual situations of the equipment usage. For some of the questions the scale was ranging from “0 – Extremely

problematic” to “10 – Without any problems”. Some items described a statement with scale values ranging from “0 – I completely disagree” to “10 – I completely agree”. In addition, they filled in some general characteristics, such as gender and age. The results are given in Table 1.

Table 1: Patients’ general characteristics and results from the calculated dimensions based on the questionnaires, Mean (SD). Calculations of Cronbach’s alpha for the dimension Hygienic Aspects was 0.83, for Physical Activity 0.88, for Skin Reactions 0.65, for Anxiety 0.80 and for Equipment 0.86.

General characteristics and Dimensions	Wireless sensor (n=11)	Reference group (n=25)
Gender: Man/woman	6 / 5	7 / 18
Age	40.2 (19.4)	56.4 (13.2)
SAI	8.2 (1.0)	7.6 (1.9)
Hygienic Aspects	8.6 (1.6)	6.6 (2.9) *
Physical Activity	9.2 (0.8)	8.0 (2.8)
Skin Reactions	6.3 (2.5)	8.6 (2.4) *
Anxiety	9.0 (1.3)	7.3 (3.5)
Equipment	8.0 (1.7)	7.5 (2.8)

* significant difference $p < .05$

The internal consistency, as calculated by Cronbach’s alpha, showed acceptable levels. The construct validity was evaluated by confirmatory factor analysis, giving reasonable factor loading according to our expectations. Calculations of a Sensor Acceptance Index (SAI) as an aggregated score showed a tendency towards a higher score for the wireless group compared to the reference group; however this difference was not at a significant level. For the dimension Hygienic Aspects regression analyses showed a significant difference between the two groups ($F(1,34)=4.51$; $p < 0.05$), with a higher score for the wireless group. The dimension Skin Reactions showed a significantly higher score for the reference group ($F(1,31)=5.95$; $p < 0.05$). With regard to the dimensions Anxiety and Physical Activity, the wireless sensor showed higher scores; however, this was not a significant difference.

In the interviews with four of the patients in the intervention group, we tried to discover some general experiences from the patients’ use of the wireless recording equipment. All patients reported some anxiety because of what they found to be a degree of uncertainty, as they did not receive any feedback from the recording system of the measurements made. They expected a quick feedback from the hospital, and two of the patients expressed the need for patient influence, while one of the respondents

said that he was not concerned with influence.

The hygienic factor focused on actual tasks related to the patient's ability to perform body wash and use of the equipment during the night while asleep, in order to obtain information about any problems relating to the wearable recording sensor. The survey showed a significant difference for the hygienic aspects, and the wireless solution obviously was preferable, since it is easy to wear and does not involve any hindrances like cables. This was confirmed both by responses to the open questions in the questionnaire and by statements in the interviews, as the patients generally expressed high satisfaction with the wireless solution compared to the existing "Holter" system. They felt free to carry out daily activities without any hindrance.

With respect to the equipment used, one patient complained about the "Holter" recorder, and said that he had "a feeling of being a living medical instrument", because of all the cables he had to wear. With regard to the wireless sensor, he said: "The wireless sensor was comfortable to wear, and most of the time I just forgot I was wearing this system". He said that the sensor after a while became "a part of me".

Another patient said the wireless system made it possible for her to participate in physical exercise. It was much easier to wear, especially during the night, and it did not prevent her from being able to take a shower. The Holter equipment was troublesome with all the cables, and made the hygienic activities more problematic. In her view, the differences in use between those two systems were huge, and they can not be compared at all.

One of the patients expressed some dissatisfaction with wearing this equipment, and she wanted to hide the equipment from other people. Similar expressions of embarrassment were also found in responses to the open questions in the questionnaire, and even if the calculated difference was not at a significant level, it showed a trend toward more anxiety regarding the use of the Holter equipment.

5 DISCUSSION

The results presented above are of interest in the evaluation of patient experience with wearable sensors attached to the chest for three days of continuous monitoring. Even though this clinical trial was limited both in terms of time and the number of patients, some interesting aspects have been discovered in relation to integration into the patients' everyday life. By comparing results and experiences

from the clinical trial with available published results from other telehomecare projects, it has been possible to discover some general aspects and point out ideas for future investigations.

First, our findings point to the issue of stigmatization, as some patients wish to hide the wireless recorder from the eyes of other people. This was similar to the findings by Myers et al., who studied the impact of home-based monitoring on patients with congestive heart failure, and followed up patients for a 2 month period. They found it necessary to train patients in telemonitoring procedures on an ongoing basis (Myers et al., 2006). During their study, 13.5% of patients withdrew due to anxiety or because they did not "like" the telemonitoring procedures or equipment. Their experiences of patient perception and ability to learn how to use the equipment indicate that the emotional barriers and stigmatization were a challenging factor and time was needed to allow the patients to adapt to this new situation. Our findings, however, also showed that when patients feel like the sensor is becoming "a part of me", as expressed by one of the patients, the stigmatization factor does not seem to represent any problems for the patient during daily activities, and can be adopted within his or her body scheme.

Second, our findings point to a need of constant feedback from the system or the health professionals to the patients. In this study we did not implement a daily reporting schema between the patient and the hospital, which was probably the reason why the patients expressed some worries about what the technology was measuring in terms of irregularities in their heart beats. Even though they trusted the equipment, they would like to see the results and the doctor's evaluation of the results when they felt irregular heart beats. These findings were in line with our experiences from an earlier study where patients underwent a 24-hour Holter recording procedure (Fensli et al., 2004), and quick feedback from the doctor was evaluated to be of utmost importance. However, although the patients appreciated good information during the research project, they expressed some uncertainty with respect to from whom they would receive an answer concerning arrhythmia findings. Their misperception of the health care sector as a "clear and strictly coordinated service", capable of giving them the desired follow-up, shows that organizational issues will be of utmost importance when introducing new telemedical solutions. If the co-ordination within the health care sector is not clearly defined, questions from the patients will not be correctly addressed and there will easily be situations where patients will suffer from not having

received the required feedback to the situation at hand.

Third, during the interviews the patients expressed overall good confidence with using the wearable sensor, mostly because of the ease of use during daily activities, which was confirmed by the significant difference calculated for the Hygienic Aspects. As a general measure of satisfaction (SAI), we found a relatively high score at 8.2 for the intervention group and 7.6 for the reference group. According to the intentions of home-based health care as expressed by Barlow and colleagues, those expressions of satisfaction were somewhat as expected (Barlow et al., 2002). This has also been confirmed by Whitten and Mickus in their study of patients with congestive heart failure in addition to chronic pulmonary disease, finding that the patients were satisfied with the technology (Whitten and Mickus, 2007). Wootton and Kvedar have also pointed to the importance of changes in the health care services (Wootton and Kvedar, 2006), and their findings are also in line with what was reported by Scalvini and colleagues in their study of chronic heart patients and the effects of home-based telecardiology (Scalvini et al., 2005). In our study, the scores for Anxiety were relatively high indicating a low degree of anxiety, with 9.0 in the intervention group and 7.3 in the reference group. The age of patients in the reference group was slightly higher and consisted of more female patients, which may represent a bias. However, being confident with using the wearable equipment combined with a feeling of safety is important to the patients.

Patient acceptance of home hospitalization equipment on a long-term basis does not seem to have been given the necessary attention in previous studies of telehomecare. Following a systematic study of observed effects in home telemonitoring of patients with diabetes, Jaana and colleagues found that studies should be extended in time and involve larger samples of patients in order to generalize the findings and obtain sufficient validity (Jaana and Pare, 2007). Long-term evaluation may probably also discover some underlying reasons for the feelings of anxiety as reported by Lehoux.

We therefore propose future studies to follow the patients' use of wearable sensors and telehomecare equipment for a relatively long time in order for the patient to adopt the technology into his/her daily activities and body scheme. Attention should be paid to the patients' ability to carry out hygienic activities such as body wash or even taking a shower, and it should be possible to participate in physical sports activities while using wearable sensor recorders.

Also, not only the patient but also his/her spouse or partner should be taken into account and given the necessary information about how the treatment should be performed and how to interact with the health care services in a coordinated manner using tele-medical equipment. As proposed by Kushniruk and Patel, multi-method evaluations can be important, and even the use of video-recording can be performed in the patients' home (Kushniruk and Patel, 2004). Hence, a multi-method approach can be employed, where questionnaires can provide some background information about the patients, their behaviour, and acceptance of the technology. In addition, interviews in the patient's home can be combined with video recordings and, later on, analysis of the process data to obtain a more thorough understanding of obstacles and barriers to the use of such solutions. Evaluating the communication between the patient and health care services will also be of special importance, and new e-health-based forms of communication should be investigated. As quick responses are required by the patients, quality factors in the communication between the patient and the health care service should be observed.

6 CONCLUSIONS

Implementation of wireless sensors for vital signs recording for the use in home hospitalization can be a great benefit for the patient, as this gives a feeling of freedom compared to ordinary hospital stay. During the use of wireless ECG-sensors, the patients in this study were satisfied with the ease of use in a daily life situation, and gave a higher score in the factor Hygienic Aspects than patients in the reference group. With respect to the factor Anxiety no significant differences were discovered. However, the use of wearable medical equipment can also affect the patients' everyday life situation in a negative manner, where they tended to hide the equipment from the eyes of other people, and they expressed anxiousness for not using the telemedical equipment in a correct manner. But at the same time, the patients also expressed confidence with the system and tended to adopt its use into their daily life.

During the interviews, the patients, however, expressed worries of not having immediate feedback and responses to irregular heart beats, and their expectations of instant follow-up by the health care sector were revealed. This can be a great challenge where necessary coordinated routines and workflow within the health care sector should be defined and established before implementing new telemedical

services. The feedback channels to the patients' questions and recordings should be timely defined and validated. In order to investigate the use of telemedical technology in everyday life situations, research studies should be performed in a relatively long-term manner. Also, by employing a multi-method approach, such studies should focus on observing to which degree the patient adopts this new technology into his or her everyday life as well as body scheme or body image. Additionally, such studies should take into account the interaction between the patient and the system, as well as the interaction between the patient and the health care professionals. Finally, it seems necessary to also investigate how the patient's partner or spouse experiences the technology in daily use.

ACKNOWLEDGEMENTS

The study is supported by the Research Council of Norway. The clinical trials are done in close cooperation with Sørlandet Sykehus HF, Arendal and the company WPR Medical AS. The authors thank Torstein Gundersen, Ellen Ytrehus and Åse Løsneløkken at Sørlandet Sykehus, and Eirik Aanonsen at WPR Medical for good assistance with the patients during clinical trials. We also thank Kjersti Nag, Hjørdis Løvdaal Gulseth and Hedda Løvland for their assistance with interviewing the patients.

REFERENCES

- Adlam, T., Orpwood, R. & Dunn, T., 2006. User evaluation in pervasive healthcare. IN BARDRAM, J. E., MIHAILIDIS, A. & WAN, D. (Eds.) *Pervasive Computing in Healthcare*. CRC Press, Inc. Boca Raton, FL, USA.
- Barlow, J., Wright, C., Sheasby, J., Turner, A. & Hainsworth, J., 2002. Self-management approaches for people with chronic conditions: a review. *Patient Educ Couns*, 48, 177-87.
- Fensli, R., Gundersen, T. & Gunnarson, E., 2004. Design Requirements for Long-Time ECG recordings in a Tele-Home-Care Situation, A Survey Study. *Scandinavian Conference in Health Informatics*. Arendal, Norway.
- Fensli, R., Gunnarson, E. & Gundersen, T., 2005. A Wearable ECG-recording System for Continuous Arrhythmia Monitoring in a Wireless Tele-Home-Care Situation. *The 18th IEEE International Symposium on Computer-Based Medical Systems*. Proceedings ed. Dublin, Ireland, IEEE CNF.
- Fensli, R., Pedersen, P. E., Gundersen, T. & Hejlesen, O., 2008. Sensor Acceptance Model - Measuring Patient Acceptance of Wearable Sensors. *Methods of Information in Medicine*, Accepted for publication.
- Friedman, C. P. & Wyatt, J. C., 2006. *Evaluation Methods in Biomedical Informatics*, Springer-Verlag.
- Haggard, P. & Wolpert, D. M., 2005. Disorders of body schema. IN FREUND, H.-J. (Ed.) *Higher-order Motor Disorders: From Neuroanatomy and Neurobiology to Clinical Neurology*. NY, Oxford University Press.
- Hopp, F., Woodbridge, P., Subramanian, U., Copeland, L., Smith, D. & Lowery, J., 2006. Outcomes associated with a home care telehealth intervention. *Telemedicine Journal and E-Health*, 12, 297-307.
- Huntleigh Healthcare, Medilog AR4 Digital Holter Recorder.
- Jones, V., vanHaltem, A., Dokovsky, N., Koprnikov, G., Peuscher, J., Bults, R., Konstantas, D., Widya, I. & Herzog, R., 2006. Mobihealth: Mobile Health Services based on Body Area Networks. IN ISTEPANIAN, R. S. H., LAXMINARAYAN, S. & PATTICHIS, C. S. (Eds.) *M-Health: Emerging Mobile Health Systems*. Springer.
- Jaana, M. & Pare, G., 2007. Home telemonitoring of patients with diabetes: a systematic assessment of observed effects. *Journal of Evaluation in Clinical Practice*, 13, 242-253.
- Kaufmann, D. R., Patel, V. L., Hilliman, C., Morin, P. C., Pevzner, J., Weinstock, R. S., Goland, R., Shea, S. & Starren, J., 2003. Usability in the real world: assessing medical information technologies in patients' homes. *Journal of Biomedical Informatics*, 36, 45-60.
- Kushniruk, A. W. & Patel, V. L., 2004. Cognitive and usability engineering methods for the evaluation of clinical information systems. *Journal of Biomedical Informatics*, 37, 56-76.
- Lehoux, P., 2004. Patients' perspectives on high-tech home care: a qualitative inquiry into the user-friendliness of four technologies. *BMC Health Services Research*, 4, 28.
- Mair, F. & Whitten, P., 2000. Systematic review of studies of patient satisfaction with telemedicine. *British Medical Journal*, 320, 1517-1520.
- Myers, S., Grant, R. W., Lugn, N. E., Holbert, B. & Kvedar, J. C., 2006. Impact of Home-Based Monitoring on the Care of Patients with Congestive Heart Failure. *Home Health Care Management & Practice*, 18, 444.
- Scalvini, S., Capomolla, S., Zanelli, E., Benigno, M., Domenighini, D., Paletta, L., Glisenti, F. & Giordano, A., 2005. Effect of home-based telecardiology on chronic heart failure: costs and outcomes. *Journal of Telemedicine & Telecare*, 11, 16-18.
- Whitten, P. & Mickus, M., 2007. Home telecare for COPD/CHF patients: outcomes and perceptions. *J Telemed Telecare*, 13, 69-73.
- Wootton, R. & Kvedar, J. C., 2006. *Home Telehealth: Connecting Care Within the Community*, RSM Press.

A COMMUNITY-CENTERED ARCHITECTURE FOR THE DEPLOYMENT OF UBIQUITOUS TELEMEDICINE SYSTEMS

Federico Cabitza, Marco P. Locatelli and Carla Simone

Università degli studi di Milano-Bicocca, Via Bicocca degli Arcimboldi 8, 20126 Milano, Italy
{cabitza,locatelli,simone}@disco.unimib.it

Keywords: Ubiquitous Computing, Home Care, Telemedicine, Hypertension, Awareness.

Abstract: In this paper, we present an ubiquitous and pervasive computing architecture, CASMAS, aimed at supporting cooperation among the members of a community and their devices. We also show how CASMAS can be augmented by the WOAD framework, which was developed independently to model and express coordination mechanisms in document-mediated communities. We take the distributed hypertension monitoring case as an exemplifying and sufficiently complex scenario to show the feasibility and advantages of the our *semantically informed modular* approach. The scenario is then declined in terms of architectural components and cooperation-oriented mechanisms that are shared between the devices and entities of the designed community of care.

1 INTRODUCTION

The technological trend toward Ubiquitous/Pervasive Computing allows for an easy and cheap connectivity among the various actors involved in the care of patients with chronic diseases: the patient, her family doctor and relatives, the caring facilities of the city where they live, and so on. Being connected however is not enough. What it is needed is to make those connected people a social network that could provide the patient with a communication environment where she can feel *safe*. This is obviously true for elder people who can otherwise experience loneliness and fragility with respect to their disease but it also holds true for people who are engaged in full time activities involving work, family care, leisure and so on. What is different is the kind of *social network* that is needed in the various situations to deal with the normal monitoring of the disease development and to face unexpected conditions when ad hoc treatments which can involve several actors at the same time are needed. On the other hand, even if the physical conditions are the same, patients and their relatives are not alike when confronting those events: individual and family history can generate different psychological states that require a strongly personalized communication to avoid that these factors have a negative impact on the caring process. The communication dimension is a necessary factor for this to happen but it is not suffi-

cient: for a patient to feel safe, for a family to feel confident and for a doctor to feel able to achieve the best possible results, they all need to be part of a caring process where they cooperate to achieve a common goal. The technology has to be challenged against this richer view, in that technological solutions should not only be conceived from an engineering point of view to make connectivity robust, fast, protected since it involves sensible data, and so on, but equally from the point of view of the disciplines that take cooperation *through* and *in interaction with* the technology as their main concerns. In this respect, Computer Supported Cooperative Work (CSCW), Computer Human Interaction (CHI) and what has been recently called Computer Supported Cooperative Care (CSCC) (Consolvo et al.,) offer a rich set of indications that are mostly derived from observations and studies in the field and from experimental practices oriented to the evaluation of the impacts of technology introduction both in work and every day life situations. Too often indeed technology adoption is considered a “secondary” problem to be considered “after” that the technology has been already put there. We like to take the opposite attitude: since technology can be “easily” put there, we have to anticipate its impacts on the basis of a stratification of experiences (sometimes of failures) that the above mentioned studies offer to the research and development of new solutions. This paper is a step in this direction.

1.1 Our Main Contribution

Telemedicine (as any other tele-*something* involving cooperation) is not a technological field that aims to reproduce the current social relations among the involved actors nor to change them in a predefined way by “optimizing” communication patterns. Both goals have been proved impossible: technology becomes an effective support if its developers and users understand that technology is unable, on the one hand, to reproduce the reality without losing some of its peculiar and tacit aspects; and, on the other hand, to constrain human behavior to prescriptive rules that can hold just in a virtual world. In telemedicine (as in any other application domain), technology has to harmonize with real life in both directions: i.e., toward both the offering of new possibilities and capabilities, and the modification of the current ones without disrupting or disregarding existing individual and social behaviors. For these reasons, we like to speak of *communities of care*, which are constituted by the patient, close relatives, selected health assistants and nurses, the family doctor and possibly involved specialists: practitioners are members of a number of communities, one for each of their patients. Accordingly, in order to conceive of a supportive technology, we claim that design has to be based on an application layer (i.e., a layer on top of the unavoidable one that guarantees physical connectivity) that reflects the nature and needs of those communities. To this aim, we propose a framework that incorporates the notion of community as a first class object and it allows to express the community’s cooperation mechanisms abstractly, i.e., at the semantic level of articulation work. From the technological point of view, the framework recognizes the need of allowing for a high degree of flexibility and adaptability to the actual domain and current situation since it separates the cooperation and communications concerns of all the entities involved in a telemedicine systems by taking a *semantically informed modular* approach. These characteristics make our approach quite different from that recently adopted to make computation “ambient aware” (Baker et al., 2007). Indeed, our proposal is not only aimed at making the members of a community of care aware of what happens in the enlarged environment. Rather, it is aimed at making these distributed and heterogeneous actors aware of the local conventions by which they can make sense of environmental conditions, consider them as meaningful parts of a patient-centered logic environment of care and trust, and correlate those conditions with timely and effective behaviors whenever they occur. Consequently, for our implementation, we do not adopt third-party middlewares that

are mainly devoted to context awareness, such as the Context Toolkit (Salber et al., 1999) or the Java Context Awareness Framework (JCAF) (Bardram, 2005); these can provide “only” the *context abstraction and reasoning* part of our model. Conversely, we need, and achieve, support for both coordination and propagation of awareness information, in addition to the ability to enhance the behavior of any application by providing it with a larger access to its context of use (Dey et al., 2001).

Our approach will be illustrated by means of a reference scenario that is described in the next section. The following sections present the approach and illustrate it through a subset of functionalities we focus on for the sake of simplicity. We chose the domain of hypertension monitoring as the reference domain where to place our technological proposal since high blood pressure is recognized as one of the most dangerous silent killer (Hoel and Howard, 1997) and the main concern of physicians involved in its treatment is that many hypertensive persons are unaware of the importance of following monitoring and treatment with commitment in a lifelong perspective (WHO/ISH, 1999).

2 AN INTEGRATED CARE SCENARIO

Let us imagine that a patient, Ms Dorothy, has been inserted in the Hypertension Monitoring Programme (HMP) by her family doctor, Dr. Robert. HMP patients suffer for a kind of hypertension that requires both lifestyle change and medication. For this reason, Ms Dorothy is provided with a series of devices aimed at making continuous monitoring easier and seamless with respect to her daily life at home: these devices are, namely, (1) an electronic agenda; (2) an automatic drug dispensing machine; (3) an automatic blood pressure machine; (4) an electronic paper form; and (5) a simple mobile phone (her own was fine). All these devices are able to transmit and receive information from the communication network, although in different ways (e.g., via GPRS, cabled ADSL, wi-fi and GSM) and have been selected to require as much as natural interaction with the patient, which in many cases can lack the necessary skills and practice to use personal computers, palmtops and complex smart-phones.

In order to take her blood pressure, Ms Dorothy uses the Automatic Blood Pressure Machine (ABP). She can use it on her own since the process is completely automatic and does not require a specific technological skill. Today, Ms Dorothy has just taken

her blood pressure and her values are 140/90 mmHg. Then, she switches on the electronic paper form (EPF) to send those values to her family doctor, Dr Robert. The EPF is an A5-sized touch screen using E Ink's technology¹ that is thinner than a postcard, has the clarity of a traditional paper sheet, can be bent with no distortion and can be inscribed with a regular stylus. Once switched on, the EPF presents Dorothy with a regular form, where her latest pressure values, the current date and her identity data (e.g., name, social security number) are copied in the corresponding fields of the form automatically. As any other HMP patient, also Ms Dorothy is supposed to fill in additional fields by hand. The most important is the 'further remarks' field: this is where she is asked to annotate the pressure data with some notable past event or condition that could help the doctor make sense of the pressure measurement (e.g., "I had a discussion with a neighbor the morning before I took the pressure"). While the patient is left free to either jot down her remarks or not when all is well, the system invites her to add a further justification to collected data whenever it has detected unexpected pressure values, with respect to either previous data, current trends or medication regimens.

Since today Ms Dorothy exhibits a blood pressure that is fairly high for a person under pharmacological treatment and with her risk factors (i.e., overweight, smoker), the system also automatically imports from Dorothy's personal agenda those items that she had labeled as 'work overload' (i.e., a series of meetings in the previous two weeks), according to some conventions agreed upon with Dr Robert: specifically, they agreed upon the need to note into an electronic agenda Ms Dorothy's daily engagements and, whenever reasonable, to characterize her schedules in terms of simple categories of events, like 'passive sport' (e.g., watching a football match on tv), 'active sport' (e.g., working out at the gym), but also 'business meeting', 'office assignment delivery' and any event that could be associated with stressful states, discomfort and anxiety. These conventional data allow the doctor to find specific correlations between high pressure peaks and risky behaviors and to isolate the actual risk factors of a specific person in order to identify more discriminating dietary and less generic practice restrictions and hence a better and patient-focused treatment. Obviously, these schedule data are strictly confidential and Ms Dorothy can remove the entries reported in the EPF. If she leaves them, Dr Robert can read them only after that Ms Dorothy also gives explicit consent by signing the form. Since the pressure values are high and Dorothy was used to be an inten-

sive smoker, the system also asks Ms Dorothy to fill in the form and to report how many cigarettes she has smoked in the last week (if any), as well as any event that could justify these values. Also this rule has been agreed upon by Dr Robert and Ms Dorothy. She reports she did not take any cigarette in that time lapse and jots down in the remark field that she has often had evening headaches, fatigue and anxiety. Moreover, since Dr Robert has prescribed her a low-calory diet, the form displays a field where to fill in the current weight also. Indeed, the EPF form can change its structure according to the doctor's requests and additional fields can be presented to patients to be filled in; in this case, after that Ms Dorothy has annotated her weight on the EPF, Dr. Robert can also assess if the dietary regimen is yielding its fruits and give her some feedback on that. To put it briefly, the EPF form is a regular document that can hold extra-data beyond what regularly fed in by digital devices. It is used to consolidate those health data, in that to sign it implies giving an explicit consent for their management. Besides the reasons of legal accountability, the form is also used to have patients get an active approach in monitoring their own blood pressure, since trend awareness and active inclusion in the monitoring process can give patients the necessary motivation to change her lifestyle if this is the case. The EPF form is also used to enable asynchronous communication between patients and their doctor via either typed or handwritten messages. Asynchronous messaging is used in order to reduce the number of phone calls that could interrupt doctors during their work. This kind of messaging is particularly appreciated by Ms Dorothy since the EPF represents a written source of information to rely on for those doubts that do not require vis-a-vis or phone talks. Handwriting with a regular stylus is allowed to enable patients that do not have - or are not confident with - personal computers and keyboards to write messages and send them online. As last thing before signing the form, Ms Dorothy writes down in the question field whether she can have some herb tea before going to bed. Then she puts her signature at the bottom of the form and, in doing so, the form content is sent to the Dr Robert officially. Besides being an input device, the EPF is also a flexible output device. The form can also serve to reproduce the official headed notepaper where the doctors jots down drug prescriptions and puts her signature. In this way, the form can be used to buy drugs at the pharmacist. Likewise, prescriptions can be updated by the doctor even remotely without any effort by the patient. The mobile phone is the other main output device at patients' home. The mobile phone is used to convey small messages to remind the pa-

¹<http://www.eink.com/products/index.html>

tient, for instance, that she has to take the pills, that the family doctor has changed the therapy remotely (and it is available at the EPF) or that he has just sent her a message.

The doctor, Dr Robert, is informed of the incoming message from her patient while he is going upstairs, headed to his office. Indeed, also family doctors are endowed by a couple of smart devices that are supplied for the HMP to convey alerts when they are not at their desktop or are engaged in some important administrative task. These alerts regard the occurrence of extremely high pressure values of a patient or, less critically, the notification that a patient has sent a message via the electronic form. Then Dr Robert's mobile phone vibrates and a short message tells him that Ms Dorothy has written him. He cannot recall exactly who she is and puts away the mobile as he opens the door of his office. When he sits at his desk, he switches on a small, flat monitor and sees a couple of pulsating circles on it. This monitor is much alike a digital picture frame like the Widget-Station². This smart frame (SF) runs various "widgets", i.e., small, user-tuned applications that perform a variety of tasks, like displaying current agenda, calendar, website and news feeding. One of these widgets is devoted to represent the patients involved in the HMP programme. Each patient is a small dot and the color indicates whether she is under drug therapy or not and whether her last pressure values are on average or not. If the dot's border winks, the corresponding patient has sent the doctor a message that has been also forwarded to his email client and to his mobile if he is out of office. If pressure values of a patient become critical, the corresponding dot enlarges to become a small quivering circle whose color indicates the seriousness of the case (e.g., red, orange, yellow). By touching the circle, the doctor can have the electronic patient record of the related patient displayed, as well as a timeline overview of the pressure data collected so far and the related medications (in a manner quite similarly to that used in the Danish TMBP project (Bardram et al., 2005)). Once seated, Dr Robert touches the only circle whose border is also winking and the station displays the name 'Ms Dorothy'. He then tips again on the circle and maximizes in a two-side view both the last page of Dorothy's personal record and her current pressure trend. Since the trend over a two month period does not show any significant pressure reduction, Dr Robert decides that the previous therapy is not enough and an additional drug is necessary. He then changes the current drug regimen in the Ms Dorothy's file within his electronic patient record.

²<http://www.emtrace.com/widgetstation/>

Finally, he opens his mail client and finds the message from Ms Dorothy quite immediately amidst the daily spam since the system has assigned the message a high priority as it comes from a patient of the HM programme. Dr Robert answers Ms Dorothy that herb teas are fine as long as they do not contain theine. This message will be displayed on the EPF of Ms Dorothy in a thread-like fashion inside the question field the next time that she will switch the form on. In a while, a number of devices at Ms Dorothy's side react to the new drug prescription. Ms Dorothy's mobile rings once and a short message notifies her that a message and a new prescription have been dispatched from Dr Robert. The prescription can be displayed at full screen on the EPF form so that it can be presented to the pharmacist at due time.

Also the Automatic Drug Dispensing Machine (ADD) receives the prescription data. According to these data, the ADD provides the patient with the right drugs by dispensing just the right dosage and by allowing the patient to open the drug-till only at the right interval of time. When the patient withdraws the drug pill, the ADD records the drug administration as accomplished. In our vignette, the ADD detects that the new drug is not loaded in its tills. Consequently, both Ms Dorothy and her closest relatives are notified of it by mobile phone. Specific relatives or friends that were pointed out as the patient's closest helpers are also notified whenever the patient has skipped two drug administrations in a row or if she presents very high values of hypertension. This is done to commit the relatives in paying attention to their dears and in reminding them due commissions at due time.

3 THE CASMAS ARCHITECTURE

In the architecture that we defined in (Cabitza et al., 2006), ubiquitous-computing systems are viewed as constellations of dynamically defined *communities* of human and technological *entities* that interact through *cooperation* mechanisms and coordinate themselves on the basis of *awareness* information related to the community context. Communities are identified by points of aggregation, called *community fulcra*: entities gather around them to access the information that characterizes the community and contributes to its definition. The structure of an entity and how it relates to the other elements of the model is described in Figure 1. Each entity — composed of inner elements as depicted in Figure 1.a — is depicted as a rounded rectangle, and can be linked to multiple fulcra to enable intra- and inter-community coordination

(see Figure 1.b) and to awareness *topological spaces* (represented as graphs, space site in Figure 1.a) to promote awareness information sharing. More specifically, a *Coordination agent* (C-agent, C in Figure 1.a) is instantiated for each connection to a fulcrum, i.e., for each of the communities in which the entity participates; an *Awareness agent* (A-agent, A in Figure 1.a) is instantiated for each connection to an awareness space.

An example of technological entity is the ABP device; this is typically connected to a fulcrum of an entity that represents a person, Ms Dorothy in this case, so as to be part of the community of devices that are associated to that person. The notation used to depict this situation is showed in Figure 1.c; in this case, a C-agent is created when an entity connects to another one and it docks to its own *entity fulcrum*. Figure 2 represents the set of entities corresponding to the patient community of care described in Section 2: in this case the community encompasses both human and technological entities.

Once created, all C-agents are provided with generic inference capabilities and with a set of entity-specific facts (i.e., declarative data structures) and mechanisms (i.e., rules). By being connected to a *community fulcrum* (Figure 1.b), C-agents associated to entities can share information (facts) and acquire community-specific behaviors (mechanisms) that are either defined at design-time or injected into the fulcrum by other entities at run-time. To this aim, the model provides mechanisms, for instance to either insert facts into fulcra (i.e., *assertion*) or to move facts from a source's to a destination's fulcrum (i.e., *translation*), by which to define the behavior of C-agents (Cabitza et al., 2006).

The behaviors of C-agents might be influenced by the *degree of participation* of entities in communities, according to additional context information that is supplied by the part of the model that manages awareness promotion (on the basis of the above mentioned awareness spaces and A-agents populating them). This part of the model has been illustrated and used in a previous work (Locatelli and Simone, 2006) and will not be further described since the degree of participation is not used in the considered scenario. Instead, we give more details on the part of the architecture that computes the awareness information to be promoted, in a simplified way, to entities that occupy isolated nodes of the awareness spaces (i.e., without using the space topology). From now on, we will use entities without any further reference to their internal structure since they support the semantically based modular approach mentioned in the introduction.




Figure 1: Notation used to represent the entities.




Figure 2: The involved entities in the scenario.

3.1 The Woad Framework

Within the CSMAS architecture, the WOAD framework is what we used to model community conventions and mechanisms of awareness provision. WOAD is a conceptual and design-oriented framework that we proposed (Cabitza and Simone, 2007b; Cabitza and Simone, 2007a) to provide a set of high-level concepts – like those of *documental artifact*, *fact*, *fact space*, and *fact interpreter* – that guide the design of a rule-based reference architecture for context-aware and coordination-oriented electronic document systems. In the CSMAS architecture, the community fulcrum can be seen as a fact space and one of the agents can play the role of fact interpreter. This agent would be supposed to import into its private fulcrum (i.e., its working memory) the WOAD mechanisms of the related community of care and then apply them to the content of the community fulcrum so as to produce consistent awareness information. In WOAD, two main abstract kinds of components are distinguishable: a *fact space*, i.e., a common and shared repository where contextual information – i.e., what we call *facts* – is stored; and a *fact interpreter*, i.e., an inference agent that is able to react to the content of the fact space and produce new contextual information, usually meta-data by which *situational awareness* (Endsley, 1995) information is associated to data. Situational awareness is any information about either an event or new contextual condition which the system can convey in any way (e.g., visual, aural, textual) to support the user become aware of it, know what is going on and figure out what to do consequently. The WOAD language is a high-level programming interface by which to express mechanisms of awareness provision and conventional patterns of data production and routing. These mechanisms are intended to reflect the local conventions that practitioners and patients can agree upon about

how to manage data and interventions within a given community. WOAD mechanisms can be shared into a fact space and be used by any fact-interpreter sharing this space to (a) provide suitable awareness information to support human actors in articulating their activities of integrated care; and (b) to process the content of a document according to locally defined and agreed patterns of coordination. WOAD encompasses a set of both static and dynamic constructs by which the designer can express either contextual, organizational or procedural knowledge about a work arrangement in a declarative manner, that is by focussing on the expression of what a system should do, rather than worrying about how it really accomplishes it. These static data structures and dynamic behaviors are expressed by two specific constructs: *facts* and *mechanisms*, respectively. In the WOAD language, whatever is given the suffix *-fact* (e.g., *activity-fact*, *relation-fact* and *awareness-fact*) is a *key-value* data structures, by which the programmer can characterize the relevant entities of a documental domain by simply assigning a value to specific attributes. The WOAD language provides designers with templates (i.e., *entity-facts*) for the most generic categories of articulation work (cf. (Simone and Divitini, 1997)), like those of actor, activity and artifact; yet, by means of its *extends* primitive, it also allows for the definition of domain-specific entities (such as patient, doctor and care activity) that inherit from and specialize those general categories. Mechanisms can be seen as simple conditional statements, like *if-then* rules: they produce some output according to the actions expressed in their *consequent* (the *then* part) whenever specific contextual conditions, which are expressed in their *antecedent* (the *if* part), are evaluated true by matching their patterns with the internal state of a fact interpreter (i.e., usually but not necessarily, the working memory of the rule-based system). Mechanisms are hence intended to make explicit the relationship between some contextual conditions regarding either the existence or the content of some facts within the *fact space* and some behavior (i.e., functionality) that the system should exhibit in that particular situation, whenever it occurs. In our application domain, the production of suitable awareness information is conveyed as attached to some contextual or documental data. In the next section we will see two examples of information flows derived from the scenario depicted in Section 2.

4 APPLICATION TO THE HM SCENARIO

The scenario described in Section 2, identifies a community, called Patient Care Community (PCC), that contains the following logical entities interacting through the *PCC fulcrum* (see Figure 2): *Patient* (we denote entities with this *typographical notation*), *Doctor*, and *Relative* i.e., entities managing the information pertaining to each person involved; *ADD*, *Agenda*, *ABP*, *E-paper Form*, and *Mobile phone* are linked to *Patient*; *Smart Frame*, *Mobile phone*, *EPR*, and *Email gateway* are linked to *Doctor*; finally, *Mobile phone* is linked to *Relative*. Moreover, there is a framework-specific entity, *FI-WOAD*, that is the WOAD *Fact Interpreter*; this entity perceives the facts published by the entities that take part to the patient care community and infers on them. The patient care community provides the involved entities with proper mechanisms to coordinate and exchange information as illustrated in the scenario. In the following, we give some examples of mechanisms involving *Patient*, *FI-WOAD*, *Doctor* and *Relative*. *Patient* is in charge of making available in the PCC's fulcrum the fact representing that the form has been signed (generated by *E-paper form*), and the fact that a specific drug is unavailable (generated by *ADD*): this transfer of information is realized by means of the translation mechanism, which accomplishes a mere copy of information from the source's to the destination's fulcrum.

By using the same mechanism, *Doctor* (associated to Dr Robert) can put into the PCC fulcrum mechanisms expressing criteria to evaluate pressure values for the patient at hand and identify "critical values" accordingly. In this way, *FI-WOAD* can acquire them to assess when pressure values are critical for the patient at hand and publish both alerting and reminding facts, in particular those on critical blood-pressure values and, in a similar way, on new prescriptions.

Let us now consider a more articulated situation. According to the conventions established between Dr Robert and Ms Dorothy, her *Agenda* publishes in the *Patient's* fulcrum the events labeled as "work overload" (see Figure 3.a, step 1a). When Ms Dorothy takes her blood pressure using the *ABP*, *ABP* publishes this information into the *Patient's* fulcrum (step 1b); *E-paper form* reacts to this information (step 2) and shows the form to Ms Dorothy so that she can fill in the pieces of information required by the EPF's structure that is dynamically set

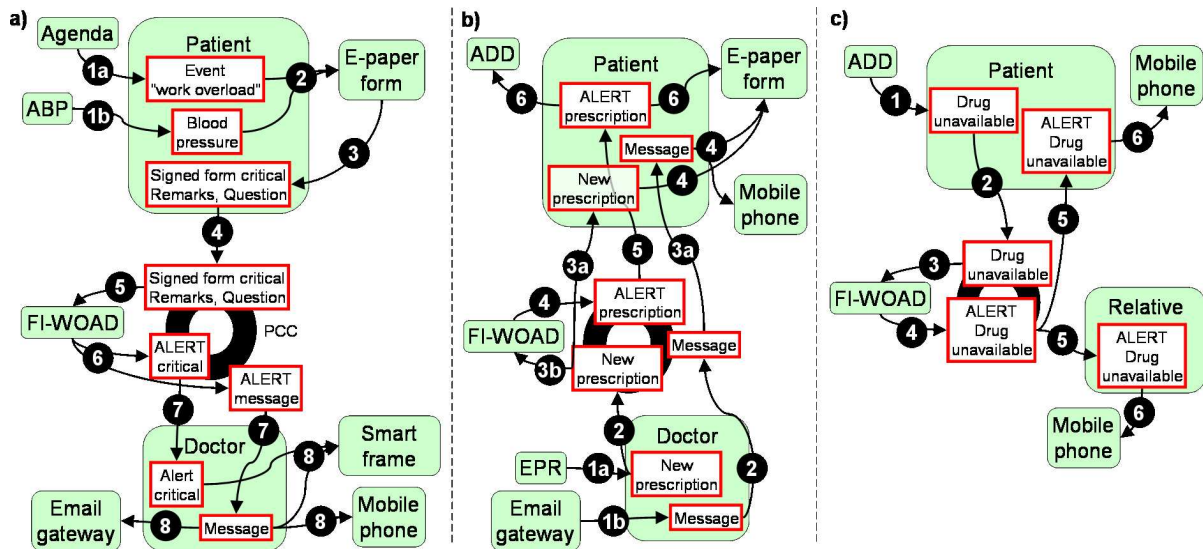


Figure 3: The step sequences for the three main cases depicted in the scenario.

by Dr Robert for her, such as her weight, and write her remarks and questions. Ms Dorothy signs the form to submit the information and *E-paper form* publishes the form to the *Patient's* fulcrum (step 3); then *Patient* publishes the signed form into the PCC fulcrum (step 4). Once the signed form is published into the community fulcrum, all the entities in charge of managing the form's content can perceive it; in our case, *FI-WOAD* perceives the form (step 5) and, due to its mechanisms, it asserts an alerting fact both for the critical pressure values and for the message that contains the Ms Dorothy's question (step 6). Since Dr Robert must be warned of alerts, *Doctor* reacts on these facts and translates the information into its fulcrum (step 7) to make it available to the entities of its private community. The alert about critical values is rendered only by the SF that executes the HM awareness widget; the message activates other three entities (step 8): *Email gateway* translates the message into an email for Dr Robert, *Smart Frame* provides awareness about the received message, and *Mobile phone* notifies Dr Robert of the incoming message since he is not in his office.

Dr Robert evaluates the received information and writes the new prescription in the EPR; once Dr Robert is done, *EPR* asserts the new prescription into the *Doctor's* fulcrum (see Figure 3.b, step 1a); moreover, Dr Robert replies the question of Ms Dorothy using his email client. This event is caught by *Email gateway* that extracts the information and asserts the corresponding message in *Doctor's* fulcrum (step 1b). *Doctor* translates both the message and the new prescription into the PCC ful-

crum (step 2) to make them available to the community. Then *Patient* translates this information into its fulcrum (step 3a) and, concurrently, *FI-WOAD* reacts on the new prescription (step 3b) and asserts an alert to the community (step 4). Now *Patient* translates also the alert into its fulcrum (step 5). *E-paper form* reacts to all this information because it has to (a) convey Dr Robert's reply, (b) show Ms Dorothy the new prescription, and obviously (c) alert her by providing awareness of the changes occurred (step 4 and 6). Due to the importance of these changes, Ms Dorothy is also notified by her mobile phone that a new prescription has been made available. Also the *ADD* is programmed to react on changes of prescription (step 6): it checks then whether the drug is available in its tills or not. Consequently, *ADD* detects that the new drug is not available and asserts this information into the *Patient's* fulcrum (see Figure 3.c, step 1). *Patient* translates this information in the community fulcrum (step 2), *FI-WOAD* perceives it (step 3) and asserts an alert (step 4). In this situation, the relative selected by Ms Dorothy is involved too: *Relative* perceives the alert fact (step 5) and translates the corresponding information into its fulcrum so that the relative's devices (such as her mobile phone) can warn her. The same happens for Ms Dorothy when she is informed by her mobile phone (step 5 and 6).

5 CONCLUSIONS

The paper presented the CASMAS architecture. This architecture is aimed at the development of applica-

tions that take advantage from ubiquitous/pervasive computing to support cooperation among the members of a community. The paper also showed how CASMAS could be used in the construction of mechanisms that are useful to enhance cooperation in a specific domain, namely within communities of care that take patients as their “fulcrum”. Two are the main features of the proposed approach. On the one hand, the identification of the components is guided by the idea that they all have access to information (i.e., local conventions and mutual awareness) that constitutes the glue of the community but they are also left free of using the shared information autonomously in order to contribute to the community’s good functioning and behavior. On the other hand, the mechanisms supporting cooperation are defined in terms of a modular and declarative approach that defines them in an abstract manner as reactive behaviors with respect to the changes of the shared information mentioned above. These two features implement the *semantically informed modular* approach and support the application flexibility and adaptability to the context that we mentioned in Section 1: indeed, each behavior (from a whole component up to the atomic constituents of a single mechanism) is aimed at the goal of supporting cooperation and awareness provision between the members of the community (of care). Hence, these community-specific behaviors can be naturally externalized and appropriated by the members, also while they are interacting with the community designers, in order to adapt the application to the single context of usage. This context is always locally characterized by the kinds of patients and their diseases, their family relationships, the local caring structure and conventional practices and so on. In addition, the above mentioned separation of concerns applies to the devices too (seen as a special kind of community “members”) so as to define a clear interface toward any underlying layer that guarantees the general purposes of the communication infrastructure. The future work is aimed at making the integration between the CASMAS architecture and WOAD framework more transparent and at improving its usability from the designer’s point of view. In doing so, we would provide an integrated framework that responds to the basic requirements of cooperation within communities (of care) and that can be fully validated in the field.

REFERENCES

- Baker, C. R. et al. (2007). Wireless sensor networks for home health care. In *AINAW’07 proceedings*, pages 832–837, Washington, DC, USA. IEEE Computer Society.
- Bardram, J. E. et al. (2005). The java context awareness framework (jcaf) - a service infrastructure and programming framework for context-aware applications. In *Lecture Notes in Computer Science*, volume 3468, pages 98–115. Springer.
- Bardram, J. E. et al. (2005). Designing for transformations in collaboration: a study of the deployment of home-care technology. In *GROUP ’05 proceedings*, pages 294–303, New York, NY, USA. ACM Press.
- Cabitza, F., Locatelli, M. P., and Simone, C. (2006). Designing computational places for communities within organizations. In *CollaborateCom’06 proceedings*, Atlanta, Georgia, USA, pages 1–10.
- Cabitza, F. and Simone, C. (2007a). A language for the executable specification of coordinative functionalities for electronic document systems. In *CHIItaly’07 proceedings*, Padova, Italy, June, 28–30.
- Cabitza, F. and Simone, C. (2007b). “...and do it the usual way”: fostering awareness of work conventions in document-mediated collaboration. In *ECSCW’07 proceedings*, Limerick, Ireland, 24–28 September.
- Consolvo, S. et al. (2004). Technology for care networks of elders.
- Dey, A. K. et al. (2001). A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction*, 16(2–4):97–166.
- Endsley, M. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1):32–64.
- Hoel, D. and Howard, R. (1997). Hypertension—stalking the silent killer. *Postgrad Med Journal*, 101(116–121).
- Locatelli, M. P. and Simone, C. (2006). Supporting care networks through an ubiquitous collaborative environment. In Nugent, C. and Augusto, J., editors, *Smart Homes and Beyond*, volume 19, pages 204–211. Assistive Technology Research, IOS Press.
- Salber, D., et al. (1999). The context toolkit: Aiding the development of context-enabled applications. In *CHI’99 proceedings*, pages 434–441.
- Simone, C. and Divitini, M. (1997). Ariadne: Supporting coordination through a flexible use of the knowledge on work processes. *Journal of Universal Computer Science*, 3(8):865–898.
- WHO/ISH (1999). *Guidelines for the management of hyper-tension*, volume 1. Blood Press.

PRE-DIABETES

An Informatics Research Agenda

Barbara M. Hayes

School of Informatics, Indiana University, 535 W. Michigan Street, Indianapolis, IN, USA
bmhayes@iupui.edu

William Aspray

School of Informatics, Indiana University, 901 E. 10th Street, Bloomington, IN, USA
waspray@indiana.edu

Keywords: Pre-diabetes, informatics research, public health and community informatics, knowledge dissemination, management alerts, information alerts, decision support, clinical guidelines, health literacy, technological literacy, patient feedback systems, pharmacy feedback systems, laboratory feedback systems, interface design, reminder systems, information quality, consumer informatics, security, privacy.

Abstract: The paper sets out a research agenda for practitioners of the relatively new, interdisciplinary field of informatics who wish to improve the health experience for people who have susceptibility to diabetes – a condition known as pre-diabetes. Using information technology tools and methods, but with sensitivity to the social and organizational complexities of the health care system, the article suggests addressing a set of problems that will improve the lives of patients and their friends and families, as well as making the provision of pre-diabetes care more effective and cost-efficient. Topics include public health and community informatics, knowledge dissemination, information alerts, decision support, clinical guidelines, health literacy, patient, pharmacy, and laboratory feedback systems, interface design, reminder systems, consumer informatics, and privacy and security issues.

1 INTRODUCTION

The healthcare industry is one of the last to heed the call of the information revolution. Its relationship with information technology is deeply conflicted. Healthcare leaders recognize that transformational IT will decrease the risk of many common errors, streamline workflow and, in some instances, save money. Advanced computational techniques may uncover genetic predispositions to disease and identify new and more targeted pharmacologic agents. Despite these promises, the healthcare industry faces tremendous challenges as it integrates information technology into healthcare delivery. The cost of the transformation is high. Patient privacy is in jeopardy. And, by its very nature, IT threatens to disrupt the treasured, traditional relationship between physician and patient.

Healthcare is shaped strongly by the interaction of human need, economics, social demographics, and the complex organization of the healthcare community. These interacting variables make

healthcare a subject well suited to study by the relatively new academic discipline of social and organizational informatics. This field of study, which is pursued in approximately twenty universities in the United States and a few other universities (including the University of Edinburgh, City University in London, National University of Singapore, Singapore Management University, and the IT University of Copenhagen) is not to be confused with the similarly named programs in Europe that are focused primarily on computer hardware and software. The practitioners of this new informatics discipline are faculty members with an interdisciplinary mindset, a social science methodology, and a familiarity with IT and its applications.

It is difficult for a healthcare practitioner to find out about research in social and organizational informatics not only because the literature is scattered, but also because it often appears in places where the healthcare community might not typically look. Nevertheless, healthcare information system

designers, healthcare policy makers, Medical and Nursing Informatics researchers, and teachers need to know about the discipline because it can improve technological solutions in healthcare and quality of life for patients. Since IT development and deployment are becoming important adjuncts in the treatment of chronic diseases such as diabetes, awareness of social and organizational informatics research is especially important to those who develop, deploy, and assess those technologies. Examples of these technologies include e-mail and Internet-based support, consumer-centered personal electronic health records, home monitoring systems, telemedicine, decision support aids, and online interventions.

Healthcare insurers and providers have made substantial investments in IT in order to make their care more effective and cost-efficient. Less attention has been given to using IT to improve the lives of patients and their families. This paper uses one stage of a widespread and expensive illness, Type 2 diabetes, as a means of examining ways in which IT can be used to improve the lives of patients when social and organizational factors are considered in the design and delivery of care.

Pre-diabetes lends itself to social and organizational informatics study because tools for managing that syndrome may be found in several areas of IT: public health and community informatics, knowledge dissemination and management, decision support, health literacy and technological literacy, feedback systems, interface design, information quality, consumer informatics, and security. There is a particular need to disseminate research on IT design and management that takes into full consideration the way IT affects individuals and organizations. IT designers often lack an understanding of the environments in which their work will be deployed, particularly in the multifaceted world of healthcare.

In an ongoing research program, the authors are looking at the social and organizational informatics issues related to every stage of diabetes, from the public health issues, to the diagnosis of the disease, to the self-care issues that face most diabetes patients as they live with the disease, to complications of the disease such as loss of eyesight or heart or kidney problems, to end-of life issues for the diabetes patient.

We have chosen pre-diabetes, a syndrome associated with Type 2 diabetes, as the focus of this paper. Individuals with pre-diabetes have blood glucose levels that are higher than normal, but not high enough to qualify for a diagnosis of diabetes.

The American Diabetes Association (ADA) now estimates that there are 54 million people in the United States who have pre-diabetes (American Diabetes Association, 2007c). We will examine several information and IT challenges associated with identifying pre-diabetes to allow informatics researchers who are unfamiliar with health care to “witness” the social and organizational factors in the ebb and flow of information around this syndrome.

To that end, the first part of this paper is organized in sections, each of which is preceded by a question. The body of the section then provides information about the topic, some examples of how the question has been addressed, and some examples of current challenges to date. The authors hope that this format will stimulate informatics researchers to create innovative research agendas that can provide ever-improving answers to these critical questions relating information to pre-diabetes. In the last portion of the paper, we offer our own suggestions for new research.

2 PUBLIC HEALTH AND COMMUNITY INFORMATICS

How do researchers educate and persuade the public to act on important new information about a syndrome -- in this case, pre-diabetes, in which higher than normal glucose levels and insulin resistance are present but do not qualify for a diagnosis of diabetes?

New information about diabetes is frequently incorporated into the medical literature. The general public may read about new medical studies in the newspaper, hear about them on the evening news or encounter them on web sites, but the vast majority of those studies are of interest only to the provider community and then only as background information. Occasionally, however, a major shift in thinking occurs. In 2002, the Department of Health and Human Services and the ADA issued position statements to the press on two conditions linked to an increased risk for developing diabetes. The term pre-diabetes was used to describe these conditions.

Patients with pre-diabetes have either impaired fasting glucose (IFG) or impaired glucose tolerance (IGT) (American Diabetes Association, 2007b). Research shows that some long term cardiac and circulatory damage may already be occurring during pre-diabetes (American Diabetes Association, 2007c). While healthcare providers had known about IFG and IGT for some time, the results of three major randomized controlled studies in different countries with different populations had concluded

at nearly the same time that individuals could prevent or delay diabetes with changes in diet and exercise (Narayan, Imperatore, Benjamin, & Engelgau, 2002). On that basis, the ADA recommended screening overweight people 45 years of age or older to detect those with impaired glucose tolerance or impaired fasting glucose (American Diabetes Association & National Institute of Diabetes and Digestive and Kidney Diseases, 2003). Those with pre-diabetes became candidates for diabetes prevention interventions.

A risk test for pre-diabetes is available on the Association's web site at <http://www.diabetes.org/risk-test.jsp> (American Diabetes Association, 2007a). People with pre-diabetes are slowly becoming insulin resistant. Medications exist to reduce insulin resistance, but more emphasis is put on weight loss, healthy diet, and exercise, which also reduce insulin resistance (Diabetes Prevention Research Group, 2002).

The identification of a new condition, syndrome or infectious agent triggers an effort to educate the public about (1) the existence of the condition (2) its symptoms (3) screening tools (4) treatments, if they exist and (5) prevention, if prevention is possible. The identification and "naming" of pre-diabetes signalled a shift from identifying people with diabetes to identifying people with pre-diabetes. The public policy ramifications of such a shift are substantial. Every time a new syndrome is identified and a recommendation is made for screening, new costs are added to the nation's healthcare bill.

The ADA and the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK) discussed five conditions that should be met before attempting to prevent a disease. They are (1) that the disease to be prevented is an important public health problem and affects a significant population (2) that the early history of the disease is understood well enough to measure its progression (3) that safe, predictable, acceptable tests exist to identify the pre-disease state (4) that safe and reliable methods exist to help prevent or delay the disease and (5) that it is cost effective to find individuals at high risk and treat them. The ADA and NIDDK argued that all five conditions had been met by cumulative research (Diabetes Prevention Research Group, 2002).

National and state public health agencies, along with diabetes advocacy groups, began to communicate the shift in emphasis and information to the public. They employed traditional media (television, radio and print) and websites to engage the public in learning about this pre-cursor to diabetes. The shift was important because it

conveyed hope. Lifestyle changes really do reduce risk. For example, the Diabetes Prevention Program study concluded that individuals with pre-diabetes who lost 5% to 10% of their total body weight and exercised could lower their risk substantially (58%) (Diabetes Prevention Research Group, 2002)

3 KNOWLEDGE DISSEMINATION AND MANAGEMENT/INFORMATION ALERTS

How do researchers undertake informing an entire community of healthcare providers who are already in practice about a new diagnostic entity and the appropriate screening and treatment of that entity?

When the discussion about screening for pre-diabetes intensified, multiple avenues were already in place to educate providers who routinely came into contact with diabetic patients. Providers were asked to screen individuals (1) who are overweight and at least 45 years of age and (2) who are under 45 and have one of several other risk factors, such as membership in high risk ethnic groups, high blood pressure, a close relative with diabetes, and others factors. Lifestyle recommendations were fairly simple: increase physical activity and achieve weight loss.

In the United States, primary care physicians (usually internal medicine specialists or family physicians) see the majority of patients before a diagnosis of diabetes is established. All physicians have a sizeable information burden, but primary care physicians bear the additional burden of having to stay aware of medical research in nearly every area of medicine since they are the first line diagnosticians. They need tools with which to screen, organize, absorb, and implement the substantial amounts of new medical information created each year.

Physicians receive information from a variety of resources, including colleagues, conferences, medical journals available in print or on the Internet, online texts and repositories such as UpToDate, handheld decision support tools such as Epocrates, web-based decision support tools such as Isabel, online databases such as PubMed, and alerts from Federal and state public health agencies. Many physicians receive information from pharmaceutical representatives who come to discuss products, although there is concern about this practice in the

United States on the presumption that such information is biased toward increasing sales.

State licensure boards and most professional societies require physicians to complete a certain number of mandatory continuing medical education hours each year. In addition, most medical specialties require re-certification. For example, family practitioners recertifying through the American Board of Family Medicine (ABFM) interact with online clinical simulations that may include content about diabetes.

There is often considerable lag time between the introduction of new information or recommendations and the subsequent formation of clinical guidelines, the necessary adaptation of paper or electronic record systems, and the adoption of new practice patterns in the office and clinic. While some physicians see patients in large group practices or hospital settings with information infrastructures that support the rapid transmission of information, others are solo or small group practitioners with severe time constraints and paper record keeping systems that are not easily updated.

4 DECISION SUPPORT AND CLINICAL GUIDELINES

How can new information be reviewed and inserted into existing information systems to help providers identify those patients who are at risk for developing pre-diabetes?

With the identification of a new diagnostic entity or clinical recommendation, existing information systems must be revised to prompt providers to screen and treat those patients at risk. Procedures must exist to vet new clinical information and determine whether and how it will be integrated into routine data collection, treatment activity, and quality measures. As in the case of screening people who may pre-diabetic, a general consensus about *what* should occur emerged after the results of several large clinical trials became public. In the United States, *how* to implement those recommendations is usually left to individual health systems and individual physicians.

While all physicians make claim to expertise in medical care, there will be “experts among experts” (in this case, those with recognized expertise in diabetes) who step forward to provide authoritative guidance in implementing new information. Many healthcare information systems, both paper and electronic, include some component of expert

guidance. Clinical advisory committees assemble to work with technical staff to approve changes for the paper or electronic health record. These groups may adopt existing clinical guidelines or develop their own guidelines. Such guidelines always have cost implications, so administrators and financial staff may also be found on advisory committees. The National Guideline Clearinghouse of the Agency for Healthcare Research and Quality (AHRQ) provides a searchable database of evidence-based guidelines at <http://www.guideline.gov/>.

After new guidelines and procedures are agreed upon, the designated experts must educate their colleagues about new forms or procedures and persuade them that new behavior is in order. That often presents a challenge, since practicing physicians are very busy and their practice patterns are difficult to change. Large groups can provide incentives to change by capturing compliance data and feeding it back to providers. Solo practitioners may find it easier to change their own behavior, but face a greater burden of staying up to date on their own.

5 HEALTH LITERACY AND TECHNOLOGICAL LITERACY

Patients have different degrees of health literacy and computer and technological literacy. If digital tools become important ways to communicate with patient and prompt them to action, what factors must be considered in the design and delivery of those health messages?

Health literacy has been defined in various ways. Practically speaking, it is the ability of a patient to take in health information, comprehend it, and take appropriate action to protect and preserve his or her health. In 2007, the Joint Commission published a white paper on improving health literacy in order to protect patient safety. It recommended raising awareness across organizations of the impact of health literacy and English proficiency on patient safety and emphasizing patient-centered communication across the entire span of care (Joint Commission, 2007). The report noted that many individuals with chronic medical conditions also have low health literacy. It recommended specific techniques, such as “teach back” and “show back” techniques to assess patient understanding; limiting information provided to two or three important points at one time; using drawings, models and devices; and giving patients information about all of

their diagnoses, medications, test result, and plans for follow-up care. Studies have shown that diabetic patients with low health literacy were less likely to achieve good control over their blood sugars and more likely to have vision impairment (Schillinger et al., 2002).

Computer literacy or technological literacy is a topic more often discussed in traditional academic circles. There are active, ongoing debates in education about what students need to learn about computing, technology, and strategies for acquiring information in the digital world in order to become informed and effective adults. Providers, insurers, and hospitals are interested in using digital tools to communicate with patients because they present opportunities to target information, reinforce it, and reduce costs. However, patients have very different degrees of computer or technological literacy as well as health literacy. There will be no “one-size-fits-all” disease information. It is far more likely that patients’ comfort with technology will need to be assessed in the same way that their health literacy will need to be assessed, before they can simply be referred, for example, to web-based education materials or cell phone reminder systems.

6 PATIENT, PHARMACY, AND LABORATORY FEEDBACK SYSTEMS

How can IT systems develop patient, pharmacy and laboratory feedback systems that might assist in the treatment of patients with pre-diabetes?

Healthcare providers may prescribe smoking cessation, diet, exercise, medications, and laboratory tests for pre-diabetic patients. However, there is a significant body of literature on non-adherence to medical recommendations. Relatively few patients are successful at making substantive change, particularly when those changes involve ingrained habits. In the case of weight loss, adherence to recommendations can be measured by a decrease in the patient’s weight, but it is still difficult to determine if the patient is following a healthy diet, exercising regularly, and using medication appropriately.

There is great interest in extending communication and a sense of connection between the diabetes provider and the patient with pre-diabetes or diabetes outside the walls of the exam room. Web-based tools such as interactive risk assessments, exercise diaries, and diet planners are

widely available for patients. Virtual health coaches are being developed to help patients adhere to medication and exercise recommendations. A few providers and systems have established two-way communication through the Internet on password-protected sites so that risk scores and patients’ exercise and food plans can be integrated into reporting mechanisms that give feedback to providers on patients’ actions. This type of communication has not been universally implemented for two reasons: (1) concerns about protecting patient privacy and complying with Health Insurance Portability and Accountability Act (HIPAA) regulations and (2) concerns about overwhelming the provider with information that he or she will not have time to read. If a physician has information and doesn’t act upon it, his or her legal liability may increase.

There is also interest in mobile communication devices for extending the relationship. Patients and providers, who move from room to room, are both mobile. The portability of cell phones and Personal Digital Assistants (PDAs) offer many advantages. For example, if the patient is recording diet choices, he or she may record and transmit them via cell phone application while at the dinner table rather than waiting to access a desktop or laptop computer. The assumption is that such information will be more accurate since it is reported so close to the event.

Compliance with medication is a particular concern. In a recent study of heart attack patients, researchers were surprised to find that one out of eight patients quit taking important medications – aspirin, beta blockers, and statins – within one month of discharge from the hospital. Those same patients were three times more likely to die during the next year than those who adhered to recommended medications (Ho et al., 2006). Pharmacy benefits manager Medco Health Solutions predicts that spending on diabetes therapies will increase up to 20% each year between 2007 and 2009 (Medco Health Solutions, 2007). One recent study found that physicians discussed cost, coverage, or purchase logistics of medications in just one-third of interviews when prescribing new medications (Tarn et al., 2006). Many patients are embarrassed to admit to their physician that they cannot afford to pay for medication.

Physicians need to know whether patients are actually filling prescriptions as directed or showing up for laboratory tests. If the patient is filling prescriptions at an in-house pharmacy (one owned by the clinic, hospital, or health plan) that information may be captured and fed back into the

record fairly easily. It will not be added to the record easily if the patient fills the prescription at an “outside” pharmacy. The physician’s only recourse is asking the patient directly. Self-report is not the best way to gather that information, since many patients have difficulty facing a physician if they have not complied entirely with his or her recommendations. At present, the capacity to give providers feedback on whether a patient actually filled or refilled a prescription exists, but it exists with some risk to the patient’s privacy.

It is somewhat easier to know if the patient has completed a laboratory test. The convention is that the results are sent directly to the physician. The patient must ask the physician to interpret the test results. The fact that a patient actually reported for a laboratory test is most easily added to the record if the patient is sent to a laboratory owned by the clinic, hospital, or health system.

7 INTERFACE DESIGN ISSUES AND REMINDER SYSTEMS

Can easy-to-use, inexpensive, reliable reminder systems be devised that will assist patients with a variety of lifestyles to remember office visits and medication?

Patients who are pre-diabetic need regular visits to assess their progress. If the initial treatment interventions do not produce the desired weight loss and increase in exercise, other recommendations may be made. In most cases, results from routine laboratory tests, which are private medical information, will stay in the medical record until the patient’s next visit. If the patient wants the results, the patient needs to return to the office.

Physician offices employ a number of strategies for reminding patients about the time of their next visit. Most patients leave the office with a printed reminder of the date of their next visit. Some offices provide a telephone prompt a few days before the visit, although such calls add to office overhead. Some offices charge a fee for missed appointments to motivate patients to keep appointments.

The Veterans Administration and a few private insurers are beginning to “push” information out to patients on their appointment times and lab tests (Ferris, 2007). Some private medical providers and facilities also see the Internet as a way for patients to view their information on line. Authentication technologies will make it easier to assure that only

the patient can see his or her own personal medical information.

Information may be captured about whether a patient filled a prescription, but it will still be difficult for physicians to assess if patients are actually taking medication or using it as prescribed. One company has produced a pill bottle that uses Short Message Service (SMS) to track how often pills are taken and send a reminder to the patient’s phone if a dose is missed (www.simpill.com).

8 INFORMATION QUALITY AND CONSUMER INFORMATICS

What tools can be developed to assure consumers that they are accessing the highest quality health information as they interact with the Internet and other information resources? How can search engine results lead patients to reputable information? How can patients avoid “health mythology” propagated by participatory tools such as chat rooms and blogs that may transmit information with little basis in fact?

The Internet affords patients with access to a computer an incredible number of tools with which to research their risk factors and conditions. The amount of information can be overwhelming. On a single day in April 2007, a Google search using the word “diabetes” returned 92,500,000 “hits.” Yahoo returned 78,900,000 results. Microsoft’s search engine, MSN, returned 18,462,447 results.

An increasing number of U.S. citizens research their medical conditions online. The major search companies are well aware of that fact. Steve Case, the founder of AOL, has launched Revolution Health, a health web site that will coach subscribers on their health, store their health information, match them with doctors, and help them with insurance claims (Freudenheim, 2007). Google introduced a health information subscription service in 2006 (Modern Healthcare, 2006). Microsoft bought a health information search engine in 2007 (Lohr, 2007). Most of the major search engines are actively engaged in a race to produce more relevant, focused results. Google, Microsoft, and Yahoo all have test sites that collect and display large amounts of information in intuitive ways. Some return definitions first and then categorize results. For the search term “diabetes,” one search engine grouped clusters of information into these categories: care, research, management, control, risk, centers, types, drugs, and supplies. Video search engine

Blinkx.com retrieved over 14,000 videos with diabetes content.

Unfortunately, there is no consensus on how to evaluate the quality of the information cited (Eysenbach, Powell, Kuss, & Sa, 2002). Online health information varies in quality; patients are vulnerable to misinformation and fraud if they are unable to evaluate the quality of the material accessed. Information gained through participatory tools such as chat rooms or blogs may be inaccurate information, giving patients false hope or diverting them from evidence based treatment. Websites may expose them to worthless or even harmful diet pills and exercise equipment for which unreasonable claims are made.

Operational definitions of quality are still needed, although rating tools are beginning to emerge. The Health on the Net Foundation offers the HONcode designation for health web sites that follow its standards of quality. The Foundation has a policing system that is designed to help developers monitor their own compliance to the code, as well as remain responsive to user concerns. The policing procedures can be initiated by individual site users or by the Foundation itself.

Reputable organizations try to provide quality information on the Internet. Some of these include the American Diabetes Association, The National Diabetes Educational Program, The National Diabetes Information Clearinghouse, and the National Institute of Diabetes and Digestive and Kidney Diseases.

There is a growing body of research on information seeking behavior. Researchers are beginning to build a set of tools and techniques with which to examine patient interaction with healthcare materials available on the Internet. Several studies have shown that online health information has a positive influence on patients' ability to cope with serious illness (Mills & Davidson, 2002; Ziebland et al., 2004). Access to disease information online has also been linked to reduced anxiety and increased perceptions of self-efficacy (Ybarra & Suman, 2006). There is still much to be learned about search strategies, information retrieval, demographic differences, and subsequent actions over the course of a chronic illness.

While information available on the Internet has helped to equalize the power in physician-patient relationships, it has two consequences that are less positive. One is the annoyance some physicians feel when patients question their judgment and recommendations. The other is the time that must be spent evaluating and responding to the patient's

attempt to gather information and participate in his or her healthcare.

9 SECURITY AND PRIVACY

How can systems be designed to secure patients' confidential information? How can stigmatizing information be kept private so that patients feel they can confide in their physicians? How can patients be persuaded to allow their treatment information to be collected and be assured it will not be used to penalize them at a later date, for example, by cutting off care for those with pre-existing conditions?

In the past two years, there have been literally millions of accidental and intentional breaches of patient privacy through lost laptops, inadequate storage procedures, and outright fraud. Some recent news stories provide examples: hackers accessed personal data for 14,000 Pentagon employees through health insurance records (Pulliam, 2006); the loss of 130,000 Aetna records by Aetna when backup tapes were stolen in a burglary (Zeller Jr, 2006).

Few incentives exist to encourage insurers, hospitals, and providers to tighten their security. Although mandatory notification of data loss and financial penalties are being considered by several states, there are few real penalties for compromising or losing medical information. Further, the Health Information Portability and Accountability Act (HIPAA), which was the impetus for spending millions of dollars to redesign systems to achieve compliance with governmental privacy directives, has had little impact. *"In the three years since Americans gained federal protection for their private medical information (through HIPAA), the Bush administration has received thousands of complaints alleging violations yet not imposed a single civil fine and has prosecuted just two criminal case"*(Stein, 2006).

The Healthcare Information Management and Systems Society and Phoenix Health Systems published a survey on HIPAA compliance in the summer of 2006. The survey noted that only 56% of providers had implemented the security standards and that a substantial portion of providers (22%) and payers (13%) remained non-compliant with the privacy regulations. The report further suggested that even those who were compliant had significant implementation gaps and that there may be a core group of entities covered by the law that cannot or will not implement the privacy standards at all

(Healthcare Information and Management Systems and Society 2006).

Many people disagree about whether business or government presents a bigger threat to privacy. While it might be assumed that government-sponsored healthcare programs have stricter privacy standards, when data from millions of U.S. veterans were contained on a stolen disk in May of 2006, the VA waited two weeks before reporting the loss. Several weeks went by before all the details of the situation came to light (Stout & Zeller Jr., 2006).

Medical privacy is also important because some diseases are stigmatizing. At different times in history, different diseases have had greater or lesser amounts of stigma attached to them. At one time, cancer was a stigmatizing illness. Some people thought cancer was contagious and that people who had it were to be avoided. Others thought death was inevitable and stopped visiting those afflicted because it was "too depressing." Even today, individuals with lung cancer may be blamed for their disease on the assumption that they must have been smokers.

Diabetes has been relatively free of stigma, although that is starting to change as the association between obesity and diabetes becomes clearer. Many people in the United States see the condition of being overweight or obese as a sign of lack of discipline or laziness. Corporate wellness programs are beginning to reward people for losing weight and maintaining the reduced weight.

Patients with pre-diabetes or diabetes have to provide information to their physician in order to receive treatment. Courts in the United States have guarded doctor-patient privilege as essential to the greater societal good. Many people find it difficult to confide in a physician. In order to encourage open and honest exchanges, patients have been assured that information about their medical records would be kept confidential. The assumption is that society benefits when patients are treated, because the potential spread of the disease and cost of its treatment have been kept in check.

There are times when all the information and power in the relationship resides with the physician. An example is the treatment of an unconscious patient in the emergency room. Society accords the physician the responsibility of gathering information and making decisions on the patient's behalf. There are also situations of shared information and power. An example is an acute illness during which the patient consults the physician but is unlikely to argue or negotiate about the prescribed treatment, such as a dosage of an antibiotic over a certain

number of days. In pre-diabetes, the power resides with the patient and the physician is in a consulting role. The physician, diabetes nurse, or dietician recommends lifestyle management techniques, but the patient has to implement them on a daily basis.

Unfortunately, healthcare providers are no longer able to guarantee privacy when they act for the unconscious patient, advise the acutely ill patient, or consult with the chronically ill patient. The worst case scenario is that the patient's own medical data cause his or her insurance company to drop insurance coverage. Patients who have chronic illness fear losing their healthcare coverage. Some stay at jobs they dislike because they would lose coverage for pre-existing conditions if they changed jobs and had to obtain new insurance. If self-employed, they run the risk that their insurance company declares them "uninsurable" and terminates their policy. If insured by government program, they may lose access to the latest treatments or experimental treatments.

In a very real sense, the information patients give to their providers for treatment and their insurance company for reimbursement may well be used against them. In the United States, the courts do not compel a defendant to testify against himself, but healthcare information systems are used to do exactly that. If patients choose to lie about their conditions, they also run the risk of losing coverage. At present, there is no way for healthcare providers to guarantee that patients will not be penalized for their honesty in providing information to the medical record.

Of course, a patient is not required to use his or her insurance benefits, but most cannot cover the cost of treating a chronic illness without doing so. If a patient wants to use insurance benefits, he or she allow information about the diagnosis to be shared with the insurance company. The insurer then decides whether or not to reimburse for care based on the terms of the policy. Insurance companies battling rising healthcare costs may use data analysis to limit access to care. They employ information systems to control costs and increase profitability. They review data to forestall unnecessary spending.

Some companies put burdensome procedures in place, deny care, or insist on lesser care in the time period before sustained evidence of efficacy can be added to the medical research. For example, insulin pumps became available in 1979. They are useful for some patients, but expensive. Some insurance companies would not cover them at first; others required providers to provide written justification of the need for an insulin pump.

10 CONCLUSIONS

Throughout the United States, medical and nursing practitioners are rapidly becoming interested in solving some of the information challenges described here, as they exist specifically for pre-diabetes and generally for every other disease and syndrome. Healthcare professionals are integrating technology into the everyday delivery of care. As their familiarity and level of comfort increase, they will seek IT support for their patients as well.

Clinicians, often led by early adopters of technology, are commissioning IT applications from commercial providers and academic research teams to solve day to day medical problems. The application developers are often unfamiliar with the realities of healthcare. These applications tend to be stand-alone; they rarely generalize well to wider use.

Individual public health, medical and nursing researchers are seeking major grant funding for large scale development of IT solutions to healthcare conundrums. In “Toward an Informatics Research Agenda: Key People and Organizational Issues,” Kaplan et al. present a research agenda model that

addresses individual, institutional, trans-organizational, and transnational concerns, aligning them with the social science disciplines that may be brought to bear on their exploration. Those disciplines include cognitive psychology, social psychology, sociology, and cultural anthropology (Kaplan, Brennan, Dowling, Friedman, & Peel, 2001). In this context, and with these observations of pre-diabetes, in Table 1 we suggest some areas of productive research.

At present, many healthcare IT solutions fall short of their intentions because patients and providers do not respond to those solutions as anticipated. Social and organizational factors are often at the core of those unanticipated, unsatisfactory responses. Money is being spent that does not result in real human benefit. The emphasis needs to shift from the construction of specific technologies to human and organizational behavior in interaction with those technologies. Informatics researchers should lead the way toward incorporating a respect for and expectation of social science research in IT development.

Table 1: Pre-Diabetes Informatics Research.

<p>Public health and community informatics</p> <ul style="list-style-type: none"> • Identify effective tools, such as community dashboards, to educate individuals and communities about pre-diabetes and the importance of life change styles to reduce the risk of diabetes; develop those tools in ways that help communities set priorities for spending on such activities as screening. • Design and fund public health and community information systems that allow data about pre-diabetes to flow between and among Federal, state, and local agencies, as well as advocacy groups, clinicians, and individuals. • Explore and evaluate information gathering patterns in individual communities; design reliable, predictable information pathways for publicizing new information in the domain of public health. • Find ways for important public health messages to rise above the “noise” of a media-saturated environment; develop and place screening tools in the media that healthcare consumers already use, rather than trying to train them to use new technologies; develop the ability to target individuals with personalized messages about their specific risk for pre-diabetes.
<p>Knowledge dissemination and management/information alerts</p> <ul style="list-style-type: none"> • Identify tools for primary care physicians with which they can effectively screen, organize, absorb, and implement the substantial amounts of new medical information created each year; consider ways to effectively manage the information burden placed on busy clinicians. • Study search strategies and patterns of information seeking in domain experts such as physicians and nurses; develop knowledge dissemination patterns that fit into the existing work flow rather than disrupting it. • Identify ways to reduce the lag time between the introduction of new information or recommendations about pre-diabetes and diabetes, and the subsequent formation of clinical guidelines, the necessary adaptation of paper or electronic record systems, and the adoption of new practice patterns in the office and clinic. • When presenting clinicians with new information, assist them with clear information on how to implement screening and treatment recommendations at the point of care.

<p>Decision support and clinical guidelines</p>
<ul style="list-style-type: none"> • Study existing processes for the institutional adoption of new standards of care; develop models for identifying the social and organizational barriers that slow the implementation of new expert knowledge; develop systems that will alert organizations and institutions to new evidence-based medical research and assist them in implementing it, rather than relying solely on practitioners to come into contact with and absorb that new knowledge on a clinician by clinician basis. • Identify effective ways to marshal current expert knowledge about diabetes and pre-diabetes, update existing medical information systems to reflect that knowledge, and provide communication technologies, strategies, and incentives for busy primary care physicians to pay attention to these recommended changes in practice.
<p>Health literacy and technological literacy</p>
<ul style="list-style-type: none"> • Identify effective means to test the health literacy and technology literacy of pre-diabetes and diabetes patients and match them with web sites and other communication means that are suited to their particular health and technology literacy. • Develop rubrics for measuring whether best practices in promoting health literacy are being incorporated into healthcare technologies. • Understand the ways in which healthcare consumers gather information about their health and the health of their families, the ways they use technology to support that process, and the thresholds or decision points that prompt them to take action, such as scheduling an appointment or attempting to change a habit or behavior.
<p>Patient, pharmacy and laboratory feedback systems</p>
<ul style="list-style-type: none"> • Identify technologies and procedures that will extend communication and a sense of connection between the diabetes provider and the patient with pre-diabetes or diabetes outside the walls of the exam room. These might include, for example, web-based tools such as interactive risk assessments, exercise diaries, and diet planners, virtual health coaches, or two-way communication systems using PDAs between patients and providers that allow that patient to record and report data and ask questions. • Design systems to provide feedback to providers about how well patients are adhering to the prescribed treatment (medications, diet, exercise, further tests or medical consultations); those systems should be sensitive to patient privacy issues and avoid overwhelming providers with data that does not contribute to decision-making or increases their legal liability. • Create and evaluate systems that optimize the capture of patient adherence data even if that data exists across multiple organizations.
<p>Interface design issues and reminder systems</p>
<ul style="list-style-type: none"> • Develop and evaluate treatment support systems that “follow” the patient into his or her work and home environments; patients should be able to choose from among a number of support systems based on their individual profiles and preferences. • Design inexpensive, reliable reminder systems that are patient-specific; avoid generic applications that require the patient to wade through information or reminders that are not specific to his or her situation.
<p>Information quality and consumer informatics</p>
<ul style="list-style-type: none"> • Develop tools to assure consumers that they are accessing the highest quality health information as they interact with the Internet and other information resources. • Develop operational standards of quality of information and tools that allow both website developers and users to rate the quality of their information on the site. • Redesign search engines to lead patients to reputable information and away from “health mythology” propagated by participatory tools such as chat rooms and blogs that may transmit information with little basis in fact.
<p>Security and privacy</p>
<ul style="list-style-type: none"> • Develop legal, economic, technological, and social means to increase the privacy of patients’ confidential information; develop technologies and protocols that protect the trust and tradition of the doctor-patient relationship. • Design systems that ensure that stigmatizing information is kept private so that patients feel they can confide in their physicians; allow treatment information can be collected without fear it will be used to penalize patients at a later date, for example, by cutting off care for those with pre-existing conditions.

REFERENCES

- American Diabetes Association. (2007a). Diabetes Risk Test [Electronic Version] from <http://www.diabetes.org/risk-test.jsp>.
- American Diabetes Association. (2007b). How to Tell If You Have Pre-Diabetes [Electronic Version] from <http://www.diabetes.org/pre-diabetes/pre-diabetes-symptoms.jsp>.
- American Diabetes Association. (2007c). Pre-Diabetes [Electronic Version]. Retrieved June 14, 2007 from <http://www.diabetes.org/pre-diabetes.jsp>.
- American Diabetes Association, & National Institute of Diabetes and Digestive and Kidney Diseases. (2003). The Prevention or Delay of Type 2 Diabetes. *Diabetes Care*, 26 Supplement(1), S62-S69.
- Diabetes Prevention Research Group. (2002). Reduction in the evidence of type 2 diabetes with life-style intervention or metformin. *N Engl J Med*, 346(6), 393-403.
- Eysenbach, G., Powell, J., Kuss, O., & Sa, E. (2002). Empirical studies assessing the quality of health information for consumers on the World Wide Web: a systematic review. *JAMA: Journal of the American Medical Association*, 287(20), 2691-2700, 2723-2694.
- Ferris, N. (2007, January 24). Vets can see EMR data. *Government Health IT News*.
- Freudenheim, M. (2007, April 17, 2007). AOL Founder Hopes to Build New Giant Among a Bevy of Health Care Web Sites. *New York Times*.
- Healthcare Information and Management Systems and Society, P. H. S. (2006). *US Healthcare Industry HIPAA Compliance Survey Results: Summer 2006*.
- Ho, P. M., Spertus, J. A., Masoudi, F. A., Reid, K. J., Peterson, E. D., Magid, D. J., et al. (2006). Impact of Medication Therapy Discontinuation on Mortality After Myocardial Infarction. *Arch Intern Med*, 166(17), 1842-1847.
- Joint Commission. (2007). "What Did the Doctor Say?": *Improving Health Literacy to Protect Patient Safety* (White Paper).
- Kaplan, B., Brennan, P. F., Dowling, A. F., Friedman, C. P., & Peel, V. (2001). Toward an Informatics Research Agenda: Key People and Organizational Issues. *J Am Med Inform Assoc*, 8(3), 235-241.
- Lohr, S. (2007, February 27, 2007). Microsoft to Buy Health Information Search Engine. *New York Times*.
- Medco Health Solutions. (2007, May 17). Unmanaged, diabetes drug costs could surge nearly 70 percent within three years. *Drug Trends Extra*.
- Mills, M. E., & Davidson, R. (2002). Cancer Patients' Sources of Information: Use and Quality Issues *Psycho-Oncology*, 11, 371-378.
- Modern Healthcare. (2006). Google introduces health-info subscription service [Electronic Version]. Retrieved May 14.
- Narayan, K. M. V. a. p., Imperatore, G. m. e., Benjamin, S. M. e., & Engelgau, M. M. c. e. a. s. b. (2002). Targeting people with pre-diabetes: Lifestyle interventions should also be aimed at people with pre-diabetes. *BMJ*, 325(7361), 403-404.
- Pulliam, D. (2006, May 1). Hackers access personal information on TRICARE servers. *Government Executive*.
- Schillinger, D., Grumbach, K., Piette, J., Wang, F., Osmond, D., Daher, C., et al. (2002). Association of health literacy with diabetes outcomes. *JAMA: Journal of the American Medical Association*, 288(4), 475-482.
- Stein, R. (2006, June 5). Medical Privacy Law Nets No Fines. *Washington Post*.
- Stout, D., & Zeller Jr., T. (2006, March 24). Agency Delayed Reporting Theft of Veterans' Data. *New York Times*.
- Tarn, D. M., Paterniti, D. A., Heritage, J., Hays, R. D., Kravitz, R. L., & Wenger, N. S. (2006). Physician communication about the cost and acquisition of newly prescribed medications. *American Journal of Managed Care*, 12(11), 657-664.
- Ybarra, M. L., & Suman, M. (2006). Help seeking behavior and the Internet: A national survey. *International Journal of Medical Informatics*, 75(1), 29-41.
- Zeller Jr, T. (2006, December 18). An Ominous Milestone: 100 Million Data Leaks. *New York Times*.
- Ziebland, S., Chapple, A., Dumelow, C., Evans, J., Prinjha, S., & Rozmovits, L. (2004). How the internet affects patients' experience of cancer: a qualitative study. *BMJ*, 328(7439), 564.

PO@HEALTH

A Medical Training Telemedicine Case Study based on Ultrasound Images over an Hybrid Power Line Network

Alécio Pedro Delazari Binotto, Flávio Ávila, Carlos Eduardo Pereira

CETA - Center of Excellence in Advanced Technologies SENAI, Av. Assis Brasil 8450, Porto Alegre, Brazil
UFRGS - Federal University of Rio Grande do Sul (Informatics and Electrical Engineering Departments), Av. Bento Gonçalves 9500, Porto Alegre, Brazil
abinotto@inf.ufrgs.br, favila@ceta.senairs.org.br, cpereira@ece.ufrgs.br

Cirano Iochpe

PROCEMPA - Information and Communication Technology Company of Porto Alegre, Av. Ipiranga 1200
Porto Alegre, Brazil
UFRGS - Federal University of Rio Grande do Sul, Av. Bento Gonçalves 9500, Porto Alegre, Brazil
ciranoi@procempa.com.br

Ilias Sachpazidis, Georgios Sakas

IGD - Fraunhofer Institute for Computer Graphics, Fraunhoferstrasse 5, Darmstadt, Germany
isachpaz@igd.fhg.de, gsakas@igd.fhg.de

Keywords: Telemedicine, Power Line Communication, Medical Training, Medical Images, Collaborative Interactivity.

Abstract: The growth of fast internet, including the recent advance on using PLC (Power Line Communication) for reaching rural and remote areas in Brazil, and the state-of-the-art of image compression methods allowed rapid teleconsultations and medical training based on medical images. At the present time, one of the challenging problems in telemedicine is the real-time teleconsultation in case of emergency and for the medical training at remote regions where the internet access is precarious. In this paper, we present the kick-off of the applied project PO@Health, which merges the European-Latin American T@lemed Project (telemedicine based on ultrasound images) and the Brazilian PLC Restinga Project (communication via Power Line in a remote district). The platform used for medical teleconsultations and residents training is able to perform both on-line (in real-time) and off-line image-based teleconsultations over the Internet connection. For the ultrasound cases, the platform is being adapted to work with the DICOM medical images synchronized with the physician hand position images of the examination in order to increase the diagnosis precision. In addition, we describe the hybrid network over PLC that is being used by the telemedical platform.

1 INTRODUCTION

Actually, telemedicine is currently growing as a research topic for the improvement of health services, increment of the population life quality, and also for remote medical training, including the residents. Based on this advance, the qualified health care of excellence medical centres is extended to remote and undeserved areas with difficult access and precarious communication.

Most telemedicine applications are massive image-based (examinations obtained from

ultrasound, CT, ray-X, resonance, etc), being used not just for teleconsultation, but also for a simple second opinion, repository research or remote training through the Internet. However, this strategy depends completely on the Internet bandwidth and its reliability and security.

The communication through the electric net grows very quickly all over the world (Motorola, 2007). Designated PLC (Power Line Communication) in Europe and BPL (Broadband Power Line) in the United States, the data communication through electric power nets is already an alternative that competes and/or

complements the wireless communication systems, satellite and wired applications, like cable TVs (Opera, 2007).

Based on the telemedicine and using PLC as part of the communication channel, this paper will present the ongoing work and its preliminary results performed at the PO@Health Project, located in the capital city of Porto Alegre (Rio Grande do Sul State, Brazil) and it is organized as follows. It starts with a brief motivation and description of the medical scenario. We follow with previous work on projects in the telemedicine and PLC areas and a detailed explanation of the proposed platform (emphasizing ultrasound examinations) and the network communication. Then, we present the medical, technical, and social results achieved during the first pilot. Finally, we discuss the estimated benefits and directions of the on-going and for future work.

2 MOTIVATION

At Porto Alegre, capital city of Rio Grande Do Sul State in Brazil, the Maternal-Infantile Hospital Presidente Vargas (HPV) is a medical referral center involving pregnancy. The public hospital assists a vast part of the population in the city who lacks of specialized maternal/medical infrastructure. Most of these patients come from remote districts just for the accomplishment of routinely ultrasound examinations and for accompaniment of pregnancy evaluation.

Restinga is the poorest and the most remote district of Porto Alegre, having more than 80.000 inhabitants, with a population density of 23 inhabitants/ha, and occupying more than 20.000 homes. The growth tax between 1991 and 2004 was from 5,6% per year and the medium monthly income of the answerable for the domiciles is 3,03 minimum wages.

The district counts just with a small health center and lacking of specialist physicians and basic medical devices, including ultrasound equipment. The health center sums up an average of 300 patient transfers to HPV per month for basic ultrasound examinations, being more than the half in the field of obstetric/gynaecologic. This represents a risk for the patient since this kind of transfer involves transportation risks in bad conditions streets, traumas and other financial costs for simple exams. In addition, it overflows the HPV capacity with patients that, in the majority of the cases, could be assisted in their own district by available General Practitioners, the residents, for example, helped by

an expert doctor, using a basic structure with ultrasound and internet connection. In fact, most of the cities have a basic structure of general doctors who, very often, cannot give a final and correct diagnosis without a second medical opinion or assistance/discussion, even remotely over the internet.

Nevertheless, there was no fast internet connection structure at Restinga since the year of 2006, when a power line communication pilot project started to be implemented, offering fast internet communication in the plug-ins of four public buildings.

The next section will go in detail over two previous projects, one based in telemedicine and the other based on PLC, which are the basis for the just-started applied project PO@Health (telemedicine over PLC).

3 PREVIOUS AND PARALLEL EXPERIENCES

3.1 T@lemed

Following the telemedicine concepts, the T@lemed Project (T@lemed, 2007) was developed by the authors and was based on a teleconsultation platform, named TeleConsult, which allows a medical store-and-forward image-based telediagnosis in real-time on-line mode or either off-line.

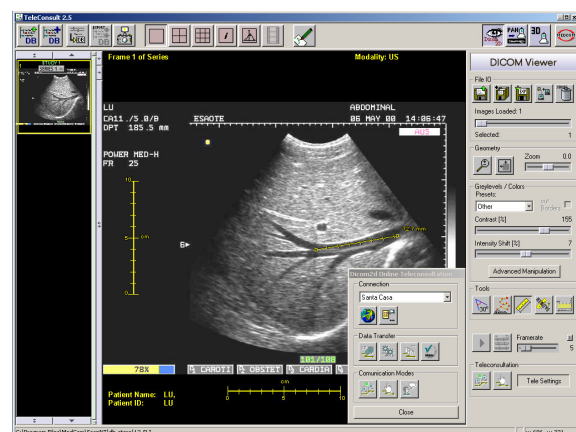


Figure 1: TeleConsult main interface.

The TeleConsult software platform is based on TeleInViVo (Kontaxakis et al., 2000), which is a telemedicine workstation used in isolated areas such as islands, rural areas and crisis situation areas. The TeleInVivo system integrates in one custom-made

device a portable PC with telecommunication capabilities and a light and portable 3D ultrasound station, combining low price, low weight, mobility and a wide range of non-radiating examinations. The integrated workstation used advanced techniques able to collect 3-dimensional ultrasound data of patients, which were presented on (Sakas and Hartig, 1992), (Sakas, 1993) and (Sakas et al., 2000).

For T@lemed and for this case study project, the reason to work with ultrasound data is based on its support to a very large range of applications (Ferrer-Roca et al., 2001), varying from gynaecology and abdominal scans to cardiological examinations and it is currently the only economically and practically affordable imaging modality. However, the platform can deal with any other DICOM images acquired either through the DICOM network. Figure 1 shows an overview of the TeleConsult interface.

As an example of its functionalities, digital annotations can be made by the generalist in the medical images and sent to the specialist physician, aiming to delineate some region of interest to be argued. The data sending can be carried through an off-line connection, where messages (images + annotations + first opinion + other crucial data to the diagnosis) are sent in determined moment (at night, for example) and later on, in another moment, the medical specialist performs the diagnosis or opinion; or through an on-line connection. In this last way, depending on the bandwidth, the data are transmitted in few seconds and collaborative discussion (annotations + chat + voice + measurements + interaction), is carried out in real-time. Figure 2 depicts the annotation interaction.




Figure 2: TeleConsult collaborative annotations.

In the scope of T@lemed, the doctors from four remote cities of the Rio Grande do Sul State were connected over wired internet (512Kbps) with the specialist center (Santa Casa Hospital) in the capital

city of Porto Alegre. Figure 3 explore the simple network configuration.

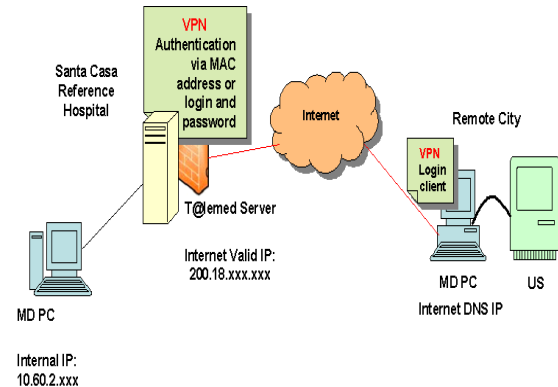


Figure 3: T@lemed network configuration.

3.2 PLC Restinga

The PLC Restinga Pilot Project, PLC network in the Restinga district, arises to supply an economical gap promoted by wire telecommunication companies, to attend deprived communities. Concerning the work of (Borges, 2005), the digital inclusion goal in Brazil is to look for the population (or at least its great majority), independently of age, sex, income, race, ethnic origin, exceptionality level or geographical location, to be able to receive access to tools, services, and necessary technological abilities in the new economy. The PLC technology implementation cost and installation, using the medium tension net for data transmission, could be cheaper than the costs of available technologies, being the sustainability and cost effectively one point of study in this on-going work performed by part of the authors from this article.




Figure 4: The PLC network at a whole.

Basically, we can classify the segments of communication networks via PLC in 3 areas (see Figure 4):

- **Medium tension:** interval between the electric power company substation and the transformer of low tension that serves the final consumers;
- **Last mile:** interval of electric net in between the transformer of low tension and the consumer's residence;
- **Last inch:** interval of electric low tension net located in the consumer's dependences.

Taking advantage of the fiber optic network from the Information and Communication Technology Company of Porto Alegre, which is interlinked to the optic ring from the State Company of Electric Energy, a PLC network was formed, beginning from the CEEE substation, located in one of the extremities of the Restinga neighborhood. Illustrated on Figure 5, four different points are connected, chosen from its geographical position and lack of digital services: (1) public primary school; (2) district administrative center; (3) professional primary school; (4) health center.

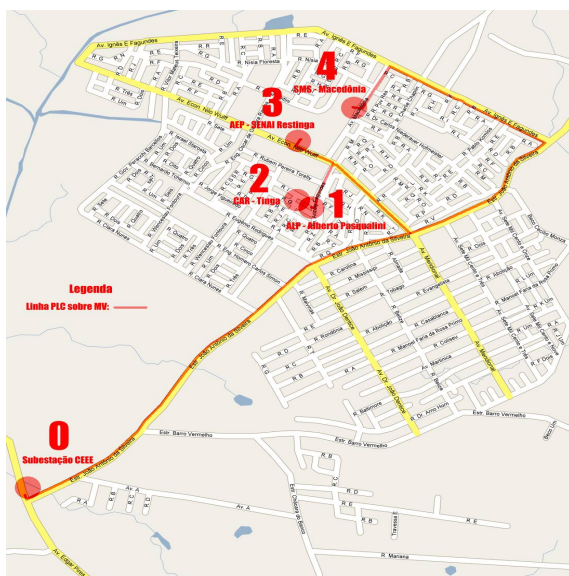


Figure 5: PLC medium tension map at Restinga.

The pilot reaches a linear extension of approximately 3.5 kilometers, transmitting data in high-speed (45Mbps, since we use the first generation of equipments) on the electric network of energized medium tension of 13.8 kV. The project foresees the implantation of several services, taking maximum advantage of the communication speed made available by the system. Thus, its complementarily with the telemedicine application.

The implementation of this net was only possible due to an optic fiber channel located at the CEEE substation (point 0 in the Figure 5). Starting

from this point, the sign from the optic fiber is injected in the medium tension net throughout capacitive couplers (Figure 6a). The PLC sign goes direct trough the electrical line, with acceptable losses, in distances of up to 1200 meters, where regenerators modems (Figure 6b) are installed aiming the system sign losses reconstitution. Repeating modems (Figure 6b) are also installed in these points to overlap the maneuver keys and the derivations in the medium tension net. In the extremities special modems (Figure 6c) are used, called HE (Head End), that receives the PLC signal from the medium tension line and re-inject this signal in the electric net of low tension (127V/220V). The signal that arrives in the assisted points through the low tension is extracted from the power plug using a modem for low tension (Figure 6d), designated as CPE. Finally, the communication system among the modems HE and CPE are point-to-point, i.e., for each modem HE just exists one CPE modem connected.




Figure 6: PLC equipments: (a) capacitive coupler; (b) regenerator and repeating modem; (c) HE modem; (d) low tension modem.

4 PO@HEALTH PROPOSED PLATFORM

The “Health at Porto Alegre” (PO@Health) Pilot Project is on its user requirements phase and has the goal to merge the experiences described on Section 3 (telemedicine over hybrid PLC network) with the foreseen significant differences bellow:

- **Ultrasound training for residents:** the focus on this pilot is on the resident physicians distance training. This means that the platform should offer an interface with fast connectivity where the expert physician is able to guide interactively the resident physician on image


acquisition and thereafter on the diagnosis. The study case increases on its importance since the Brazilian law authorizes just medical physicians to operate ultrasound;

- **Interactive tele-acquisition:** as the focus is on residents, the expert doctor has to visualize the *ongoing-acquirement* of images in real-time or at least with acceptable interactivity and not allow just a store-and-forward collaborative opinion or telediagnosis;
- **Transducer location image:** ultrasound is the only medical device which produces medical images that are completely human-dependent on its acquirement. Due to that situation, the ultrasound transducer position on the patients body is a crucial information for the exam interpretation. Hence, the platform must also afford the medical images synchronized with the transducer position;
- **PLC and telemedicine economical model improvement:** the pilot will deploy an feasibility study for PLC based on telemedicine direct costs (for example, costs of exams, communication, transportation, etc) and non-direct costs (for example, psychological traumas, savings due prevention, etc);
- **Prevention control:** on the medical point of view, the pilot should also contribute on development of prevention control methods for remote regions, since they do not need to wait for the symptoms because telemedicine is a way to improve the capillarity of expert knowledge.


5 PRELIMINARY RESULTS

To illustrate PO@Health, only for demonstrating the project conceptual meaning, the Figure 7 depict the first prototype using a low-cost portable ultrasound and a *screen sharing* application running on the background of the developed ultrasound viewer. Different from T@lemed, here we used common sharing software to perform the initial tests regarding the real-time acquisition and a web-camera window for the transducer position. As expected, the screen sharing consumed a high bandwidth of dedicated 4 Mbps for screen transmitting and VOIP. However, the images update was not performed in real-time allowing basic training, but not capable for exams aiming the pulsation detection, like, for example, the fetus hearts pulse, which is very important to recognize if

the fetus is sleeping or dead in the case of motionless.



(a)



(b)

Figure 7: PO@Health using a portable ultrasound (a) and a *screen sharing* application for the ultrasound images and transducer position (b).

This test will be useful for a future image-quality and bandwidth consume benchmark, comparing the final stream-based solution with other possibilities.

6 CONCLUSIONS AND FUTURE WORKS

This short paper presented the applied research for the general conception of a telemedicine system which is being developed/adapted based on ultrasound images and on PLC data communication for resident physicians practical distance training. The study case scenario is on the Restinga remote district of Porto Alegre city, where there is a PLC-based network implementation going on. We believe that the use of PLC as the channel for data transferring will be feasible and bring great capillarity for telemedicine services, which deal with large data.

As current work, we are performing the benchmark tests using a stream-based application sharing (RRD Streaming, 2007) which is a general solution to stream desktop contents to remote

locations in real-time, focusing on live video, presentations, simulations, 3D visualizations, and dynamically changing scenes applications. Depending on the results, RRD Streaming will be the basis for the TeleConsult adaptation to attend the main user requirements (synchronized interactive tele-acquisition module and transducer position in real-time).

Performance measurements on the PLC network at Restinga are also a current work, verifying in loco that the first generation of equipments will allow us to reach up to 45Mbps in an optimistic situation. We also intend to update some PLC equipments to the second generation, allowing theoretically up to 200Mbps and look over the performance of a mixed network composed by first and second generation equipments.

Further studies on PLC and telemedicine feasibility will also be performed, trying to expose it's complementarity and thus the PLC capillarity.

RRD Streaming. In <http://www.igd.fhg.de/igd-a2/rrds/index.php>. Access in July 2007.

Sakas, G, Hartig, J., 1992. Interactive Visualization of Large Scalar Voxel Fields. In Proceedings IEEE Visualization'92 Conference. IEEE Press, p.29-36.

Sakas, G., 1993. Interactive volume rendering of large fields. In Visual Computing, n.9, v.8, p.425-438.

Sakas, G., Walter, S., Grimm, M., Richtscheid, M., 2000. Free hand acquisition, reconstruction and visualization of 3D and 4D ultrasound. In Radiologe, v40, p.295-303.

T@lemed, 2007. In <http://www.alis-telemed.net>. Access in June 2007.

ACKNOWLEDGEMENTS

We would like to thanks the following institutional partners for partial financial support and fruitful discussions: Porto Alegre City Hall, Science and Technology Ministry of Rio Grande do Sul State, Getúlio Vargas Hospital physicians, and CEEE (State Company of Electric Energy).

REFERENCES

Borges, A. A., 2005. Tecnologia PLC. In <http://www.teleco.com.br/emdebate/aderbal01.asp>. Access in November 2006.

Ferrer-Roca, O., Vilarchao-Cavia, J., Troyano-Luque, J. M., Clavijo, M., 2001. Virtual Sonography through the Internet: volume compression issues. In Journal of Medical Internet Research. <http://www.jmir.org/2001/2/e21/>.

Kontaxakis, G., Walter, S., Sakas, G., 2000. EU-TeleInViVo: an integrated portable telemedicine workstation featuring acquisition, processing and transmission over low-bandwidth lines of 3D ultrasound volume images. In Information Technology Applications in Biomedicine, Proceedings of IEEE EMBS International Conference, p. 158 – 163.

Motorola, 2007. BPL: Otimismo impulsionando o crescimento. In <http://www.connectwithmotowi4.com/index.cfm?canopy=solutions.flstory&aid=59&lang=Portuguese>. Access in May 2007.

Opera, 2007. Open PLC European Research Alliance. In <http://www.ist-opera.org>. Access in May 2007.

ON EXTRACTION OF NUTRITIONAL PATTERNS (NPS) USING FUZZY ASSOCIATION RULE MINING

¹ M. Sulaiman Khan^{1,2}, Maybin Muyeba¹
School of Computing, Liverpool Hope University, Liverpool, UK
khanm@hope.ac.uk, muyebam@hope.ac.uk

² Frans Coenen²
Department of Computer Science, University of Liverpool, UK
frans@csc.liv.ac.uk

Keywords: Association rules, nutritional patterns, fuzzy association rules, nutrients.

Abstract: This paper proposes a framework for mining market basket data to generate Nutritional Patterns (NPs) and a method for analysing generated nutritional patterns using Fuzzy Association Rule Mining. Edible attributes are filtered from transactional input data by projections and are then converted to Recommended Dietary Allowance (RDA) numeric values. The RDA database is then converted to a fuzzy database that contains expended normalized fuzzy attributes comprising of different fuzzy sets. Analysis of nutritional information is performed either from normal generated association rules or from a converted fuzzy transactional database. Our approach uses prototype support tool that extract Nutritional Patterns (NPs) and signifies the level of nutritional content in an association rule per item. The paper presents various performance tests and interestingness measures to demonstrate the effectiveness of the approach and concludes with experimental results and discussion on evaluating the proposed framework.

1 INTRODUCTION

Association Rule Mining (ARM) (Agarwal, 1993) is a popular data mining technique that has been used to determine customer buying patterns from market basket data. General association rules are of the form $X \rightarrow Y$, which means customers who buy X also buy Y, with given support and confidence measures. A support measure is used to determine the number of transactions that include all items in the antecedent (X value) and the consequent (Y value) parts of the rule, while a confidence measure is the ratio of support to the number of transactions that include all items in the antecedent. The discovered rules indicate patterns of associating items. Such rules can be helpful in shelf arrangements, advertisement, sales promotion etc. However, nowadays health concerns are becoming increasingly important to a large community of people including health practitioners, sporting organizations, governments and recently supermarkets. In data mining, association rules have been used to determine buying patterns (to the shop owner's benefit) but not nutritional pattern in general (to the customers health benefit).

People have recently become "healthy eating" conscious, but largely they are unaware of qualities, limitations and above all, constituents of food. For example, how often do people who buy baked beans bother with nutritional information other than looking at expiry dates, price and brand name? Unless the customer is diet conscious, there is no explicit way to determine nutritional requirements and consumption patterns. There are many dietary schemes and programmes that individuals follow that helps them determine how healthy they are but do not critically analyse nutritional elements that may affect their health. It is known that certain nutritional chemical elements when taken in large quantities do alter genetic material of a person, but also other elements are known to be more important for health than others.

Nowadays, nutritional information is usually labeled on supermarket products but is not used to determine actual nutritional patterns of every given customer transaction. This information would be useful for individual customers own health evaluation, supermarkets own reports on likely healthy buying patterns and many health related

organizations including government health ministries. As modern society is concerned with health issues, association rules can be used to determine nutritional patterns by analysing product nutritional information, using market basket data. The approach signifies the level of nutritional content in an association rule per item.

Most algorithms in the literature have concentrated on improving performance through efficient implementations of the modified *Apriori* algorithm (Bodon, 2003), (Lee, 2003), (Coenen, 2004), (Wang, 2004). Although improving performance and efficiency of various ARM algorithms is important, determining Nutritional Patterns (NPs) from customer transactions and association rules is also important. Extracting health related information using association rules from market basket data has mostly been overlooked.

In this paper we propose a fuzzy based approach for extracting nutritional patterns using fuzzy association rule mining, where a transactional database is converted into a database that contains the average RDA of nutrient values per item. This database is then converted into a fuzzy database with fuzzy attributes, according to the nutrients intake. The fuzzy database contains the actual contribution of nutrients per transaction or fuzzy membership degrees in fuzzy sets for each particular item (e.g. values 0.0, 0.3, 0.5, 0.2, 0.0 for fuzzy attributes very low, low, ideal, high and very high respectively) as shown in the figure 1.

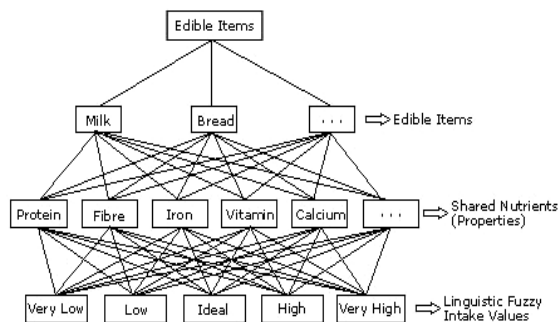


Figure 1: Edible items, Nutrients & Fuzzy Intervals.

We show the effectiveness of this new method by applying it on different datasets. Our contributions are that edible attributes in market basket data are used with an RDA table, a fuzzy normalization process and correlation analysis produce effective rules and records good performance.

The paper is organised as follows: section 2 presents background and related work; section 3 gives a problem definition; section 4 discusses the

proposed methodology; section 5 reviews experimental results, and section 6 concludes the paper with directions for future work.

2 BACKGROUND AND RELATED WORK

Many applications of association rule mining have been proposed in medical domain (Xie, 2005), (Yuanchen, 2006), (Lavrac, 1996), (Delgado, 2000) but most of the researches in the literature have concentrated on improving performance through efficient implementation than producing effective rules (Bodon, 2003), (Lee, 2003), (Coenen, 2004), (Wang, 2004). Again, in almost all ARM algorithms, thresholds (both confidence and support) are crisp values. This support specification may not suffice for queries and rule representations that require generating rules that have linguistic terms such as “low/Ideal/High for protein intake” etc. Fuzzy approaches (Chen, 2002), (Wai, 1999), (Xie, 2005), (Guenesei, 2001) deal with quantitative attributes (Srikant, 1996) by mapping numeric values to Boolean values. Detailed overviews for fuzzy association rules are given in (Chen, 2002), (Wai, 1999), (Dubois, 2006).

Effective and efficient fuzzy algorithms supporting the mining process, i.e. the extraction of interesting associations from a database, have received less attention in the fuzzy community. This might be explained to some extent by the fact that, for fuzzy extensions of association analysis, standard algorithms can often be used or at least adapted in a relatively straightforward way. Still, some contributions have been made in this field. For instance, (Muyeba, 2006) describe the process of fuzzy association rules to obtain healthy buying patterns using binary Apriori algorithm.

Mining nutrient associations among itemsets is a new type of ARM technique which attempts to investigate Nutritional Patterns (NPs) by analysing nutrition consumption patterns. In (Xie, 2005), fuzzy associations are presented, where a reduced table is used to effectively minimise the complexity of mining such rules. The authors also present mining for nutrients in the antecedent part of the rule, but it is not clear how the fuzzy nutrient values are aggregated and largely, how membership functions are used. Our algorithm’s ultimate goal is to determine customers’ buying patterns for healthy foods, which can easily be evaluated using RDA standard tables. Other related work deals with

building a classifier using fuzzy ARs in biomedical applications (Yuanchen, 2006).

Fuzzy association rules have been used for medical data mining (Xie, 2005), (Lavrac, 1996), (Delgado, 2000), but we propose a novel approach to determine whether customers are buying healthy food, which can easily be evaluated using required daily allowance (RDA) standard tables.

3 PROBLEM DEFINITION

In this section a sequence of formal definitions is presented to: (i) describe the concept of fuzzy association rule mining and (ii) the fuzzy approach adopted by the authors. The normalization process for Fuzzy Transactions (T') and rules interestingness measures will also be discussed later in this section.

3.1 Fuzzy Association Rules

Mining fuzzy association rules is the discovery of association rules using fuzzy sets such that quantitative attributes can be handled (Dubois, 2006). A fuzzy quantitative rule represents each item as (item, value) pair. Fuzzy association rules are thus expressed in the following form:

If X is A satisfies Y is B

For example if (age is young) \rightarrow (salary is low) where age and salary represents X and Y , and young and low are the discretised/linguistic values for attributes age and salary respectively representing A and B .

In the above rule, $X = \{x_1, x_2, \dots, x_n\}$ and $Y = \{y_1, y_2, \dots, y_n\}$ are itemsets, where $X \subset I, Y \subset I$, and $X \cap Y = \emptyset$. Sets $A = \{f_{x1}, f_{x2}, \dots, f_{xn}\}$ and $B = \{f_{y1}, f_{y2}, \dots, f_{yn}\}$

contain the fuzzy sets associated with the corresponding attributes in X and Y , for example (protein, low), (protein, ideal), (protein, high). The semantics of the rule is that when 'X is A' is satisfied, we can imply that 'Y is B' is also satisfied, which means there are sufficient records that contribute their votes to the attribute fuzzy set pairs and the sum of these votes is greater than the user specified threshold which could be crisp or fuzzy.

For a given database D with transactions $T = \{t_1, t_2, t_3, \dots, t_n\}$ with items $I = \{i_1, i_2, i_3, \dots, i_{|I|}\}$ and converted fuzzy transactions

$T' = \{t'_1, t'_2, t'_3, \dots, t'_n\}$ with attributes $P = \{p_1, p_2, p_3, \dots, p_{|P|}\}$ and the fuzzy sets $F = \{fp_1, fp_2, \dots, fp_m\}$ associated with each attribute in P .

Table 1: Set of ordinary transactions.

D	i_1	i_2	i_3
t_1	1	0	1
t_2	0	1	0
t_3	1	1	1
...

A fuzzy transaction is a special case of transformed ordinary transaction (table 1) and nonempty fuzzy subset of P where $T' \subseteq P$. In table 2 an item p_j and transaction t'_k contains a value v (membership degree) in $[0, 1]$. The membership degree of p_j in t'_k is $t'_k(p_j(v))$. Without loss of generality, we also define edible set of items $E \subseteq I$ where any $i_j \in E$ consists of quantitative nutritional information $\bigcup_{p=1}^{|P|} i_j^p$, where each i_j^p is given as standard RDA numerical ranges and consists of $|P|$ nutrients.

Table 2: Set of edible fuzzy transactions.

E	fp	fp	fp	F	F	fp	fp	fp	fp	fp
	i	i	i	p_1	p_1	2	2	2	2	2
	$(v$	$(v$	$(v$	$(v$	$(v$	$(v$	$(v$	$(v$	$(v$	$(v$
	$1)$	$2)$	$3)$	$4)$	$5)$	$1)$	$2)$	$3)$	$4)$	$5)$
t'_1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
	0	7	3	0	0	0	0	8	2	0
t'_2	1.	0.	0.	0.	0.	1.	0.	0.	0.	0.
	0	0	0	0	0	0	0	0	0	0
t'_3	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
	0	0	9	1	0	0	0	8	2	0
...

Each quantitative item p_j is divided into various fuzzy sets $f(p_j)$ and $m(l, v)$ denotes the membership degree of v in the fuzzy set l , $0 \leq m(l, v) \leq 1$ as shown in table 2.

3.1.1 Fuzzy Transactions Normalization Process

As mentioned above each quantitative item p_j in t'_k is divided into various fuzzy sets $f(p_j)$ and $m(l, v)$ denotes the membership degree of v in the fuzzy set l , $0 \leq m(l, v) \leq 1$. For each fuzzy transaction $t' \in E$ (edible items), a normalization process to find significance of an items contribution to the degree of support of a transaction in order to guarantee a partition of unity is given by the equation (1):

$$m' = \frac{m(l, t'_k(p_j))}{\sum_{i=1}^{f(p_j)} m(l, t'_k(p_j))} \quad (1)$$

Without normalisation, support of an individual fuzzy item could increase in a transaction. The normalisation process ensures fuzzy membership values for each nutrient are consistent and are not affected by boundary values.

3.1.2 Fuzzy Support and Confidence

The problem of mining fuzzy association rules is given following a similar formulation in (Kuok, 1998). To generate Fuzzy Support (FS) value of an item set X with fuzzy set A , we use the equation (2):

$$FS(X, A) = \frac{\sum_{t_i \in T} \prod_{x_j \in X} m(t'_i[x_j])}{|E|} \quad (2)$$

A quantitative rule represents each item as <item, value> pair. In the above equation we have used arithmetic mean averaging operator for fuzzy nutrients aggregation of candidate itemsets in a transactional database and used multiplication “mul” operator for fuzzy union of candidate items in a transaction. *min* or *max* operators can also be used but *mul* provides us the simplest and reasonable results as shown in table 3. In case when the fuzzy transactions are not normalized *mul* is more suitable because it takes the degrees of all items in a transaction into account.

Table 3: Effect of fuzzy mul operator.

i_1	i_2	i_3	i_4		Max	Min	Mul
0.2	0.6	0.7	0.9	→	0.9	0.2	0.075
0.9	0.8	0.5	0.6	→	0.9	0.5	0.216
0.7	0.0	0.75	0.8	→	0.8	0.0	0.0
0.3	0.9	0.7	0.2	→	0.9	0.2	0.037

For a rule $\langle X, A \rangle \rightarrow \langle Y, B \rangle$, the fuzzy confidence value (FC) where $X \cup Y = Z, A \cup B = C$ is given by equation (3):

$$FC(\langle X, A \rangle \rightarrow \langle Y, B \rangle) = \frac{\sum_{t'_i \in T'} \prod_{z_j \in X} m(t'_i[z_j])}{\sum_{t'_i \in T'} \prod_{x_j \in X} m(t'_i[x_j])} \quad (3)$$

where each $z \in \{X \cup Y\}$. For our approach, $X, Y \subset E$, where E is a projection of edible items from D . Depending on the query, each item i_j specified in the query and belonging to a particular transaction, is split or converted into $|P|$ nutrient parts $\bigcup_{p=1}^{|P|} i_j^p, 1 \leq j \leq |I|$. For each

transaction t , the bought items contribute to an overall nutrient p by averaging the total values of contributing items i.e. if items i_3, i_4 and i_7 are in a transaction t_1 and all contain nutrient $p=5$ in any proportions, their contribution to nutrient 5 is $\sum \frac{|i_j^5|}{3}, j \in \{3,4,7\}$. These values are then

aggregated into an RDA table with a schema of nutrients (see table 5, section 4) and corresponding transactions. We use the same notation for an item i_j with nutrient p, i_j^p as item or nutrient p_j in the RDA table. Given that items p_j are quantitative (fuzzy) and we need to find fuzzy support and fuzzy confidence as defined, we introduce membership functions for each nutrient or item since for a normal diet intake, ideal intakes for each nutrient vary. However, five (5) fuzzy sets for each item are defined as {very low, low, ideal, high, very high} based on expert analysis on nutrition.

Based on this analysis, examples of fuzzy membership functions for nutrient Protein is shown in figure 1. There are many different types of membership function and the type of representation

of the membership function depends on the nature of the fuzzy set. In figure 2 the functions assume a trapezoidal shape since nutrient values in excess or in deficiency mean less than ideal intake according to expert knowledge. Ideal nutrients can assume value 1 naturally, but this value could be evaluated computationally to 0.8, 0.9 in practical terms.




Figure 2: Fuzzy membership functions.

$$\mu(x, \alpha, \beta, \gamma, \delta) = \begin{cases} 1 & , x \in [\beta, \gamma] \\ \frac{(x - \alpha)}{(\beta - \alpha)} & , x \in [\alpha, \beta] \\ \frac{(\delta - x)}{(\delta - \gamma)} & , x \in [\gamma, \delta] \\ 0 & , otherwise \end{cases} \quad (4)$$

Equation 4 (Paetz, 2002) represents all nutrient membership degrees of a nutrient value “x”. The input database value x has “ideal” values of a nutrient between β and γ , with lowest value α and highest value δ . The task is to determine a membership value of x in equation 4.

Note that equation 4 gives values equal to $m(l, v)$ in equations 1, 2 and 3. We can then handle any query after a series of data transformations and fuzzy function evaluations of associations between nutritional values. For missing nutrient values or so called “trace” elements, the fuzzy function evaluated zero degree membership.

3.2 Interestingness Measures

Measures of interestingness other than standard support and confidence are required in order to evaluate the quality of fuzzy association rules. The quality measure for a rule to be interesting is called certainty factor (). A rule can be considered interesting if the fuzzy set union of antecedent and the consequent has enough significance and the rule has adequate certainty. A measure of significance for a rule is similar to equation (3) and we have adopted it as the confidence of a rule. The certainty

factor is determined by computing the fuzzy correlation of antecedent and the consequent of the rule. We have used Pearson’s product-moment correlation coefficient between attributes which is different from the general statistical usage of correlation because in association rule mining $X \Rightarrow Y \neq Y \Rightarrow X$.

The correlation $Corr(X, Y)$ between two variables X and Y with expected values $E(X)$ and $E(Y)$ and standard deviations σ_x and σ_y is defined as:

$$Corr(X, Y) = \frac{Cov(X, Y)}{\sigma_x \sigma_y}$$

$$Corr(X, Y) = \frac{E(X \cdot Y) - E(X) \cdot E(Y)}{\sqrt{E(X^2) - E(X)^2} \sqrt{E(Y^2) - E(Y)^2}}$$

where E is the expected value of the variables and cov is covariance. We can transform the above correlation equation to find the certainty factor between two or more fuzzy attributes and can calculate fuzzy correlation as:

$$Corr_{Fuzzy}(\langle X, A \rangle, \langle Y, B \rangle) = \frac{Cov(\langle X, A \rangle, \langle Y, B \rangle)}{\sqrt{Var(X, A)} \sqrt{Var(Y, B)}} \quad (5)$$

The value of correlation ranges from -1 to +1. Value -1 means no correlation and +1 means maximum correlation. In our problem, only positive values can be considered as the degree of relation. As the certainty value increases from 0 to 1, the more related the attributes are and consequently the more interesting they are. Therefore if the rule “IF Protein is low THEN Vitamin A is high” holds, then the certainty value should be at least greater than zero. This could mean customers prefer to buy more vitamin related items to protein ones and the HBP value is simply the certainty value obtained (see section 6.1)

4 PROPOSED METHODOLOGY

The proposed methodology consists of various phases, each of which is evaluated using fuzzy sets for quantitative attributes (Nutrients) as mentioned earlier. We have developed an algorithm called Fuzzy Healthy Association Rule Mining algorithm (FHARM). FHARM can deal with other kinds of transactional and relational databases to generate

fuzzy association rules using quantitative attributes. We have discovered two techniques to obtain Nutritional Patterns as described in the next sections.

4.1 Nutritional Fuzzy ARM Mining

To mine from the transactional file (table 4), input data is projected into edible database on-the-fly thereby reducing the number of items in the transactions and possibly transactions too. The latter occurs because some transactions may contain non-edible items which are not needed for nutrition evaluation. This new input data is converted into an RDA transaction file (table 5) using RDA table (definition 6) with each edible item expressed as a quantitative attribute and then aggregating all such items per transaction (see definition 2, equation 1).

Table 4: Market Basket Data.

TID	Items
1	X, Z
2	Y
3	X, Y, Z
4	..

Table 5: Converted RDA transactions.

TID	Pr	Fe	Ca	Cu
1	45	150	86	28
2	9	0	47	1.5
3	54	150	133	29.5
4

At this point, two solutions may exist for the next mining step. One is to discretised nutrients into intervals and converts RDA Transactions into discretised transactions with boolean values (Table 6) for each nutrient and corresponding value for each interval per nutrient. Each transaction then (table 6), will have repeated fuzzy values {very low, low, ideal, high, very high} for each nutrient present in every item of that transaction. Table 6 actually shows only two nutrients.

Table 6: Discretised (Boolean) transaction file.

TID	Protein (Pr)					Iron (Fe)				
	VL	L	Ideal	H	VH	VL	L	Ideal	H	VH
1	0	1	0	0	0	0	0	1	0	0
2	1	0	0	0	0	1	0	0	0	0
3	0	0	1	0	0	0	0	1	0	0
4

Thus nutrients can have only values [0, 1] and only 1 intake value out of five in table 6, which represents its complete membership in that interval. The discretised boolean data can be mined by any binary type association rule algorithm to find frequent item sets and hence association rules. This approach only gives us, for instance, the total support of various fuzzy sets per nutrient and not the degree of support as expressed in equations 2 and 4.

The other approach (which we have adopted) is to convert RDA transactions (table 5) to linguistic values for each nutrient and corresponding degrees of membership for the fuzzy sets they represent above or equal to a fuzzy support threshold. Each transaction then (table 7), will have repeated fuzzy values {very low, low, ideal, high, very high} for each nutrient present in every item of that transaction.

Table 7: Linguistic (Fuzzy) transaction file.

TID	Protein (Pr)					Iron (Fe)				
	VL	L	Ideal	H	VH	VL	L	Ideal	H	VH
1	0.0	0.7	0.3	0.0	0.0	0.0	0.0	0.8	0.2	0.0
2	1.0	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0
3	0.0	0.0	0.9	0.1	0.0	0.0	0.0	0.8	0.2	0.0
4

Table 7 actually shows only two nutrients. A data structure is then used to store these values (linguistic value and degree of membership) and large itemsets are found based on the fuzzy support threshold.

To obtain the degree of fuzzy support, we use equations 2 and 4 on each fuzzy set for each nutrient and then obtain ARs (Nutritional Patterns) in the normal way e.g.

IF Protein intake is High AND Vitamin A intake is Low THEN Fat intake is High.

4.2 Rule Query on Nutrient Associations

To mine a specific rule, $X \rightarrow Y$, for nutritional content, the rule base (table 8) is scanned first for this rule and if found, converted into an RDA table (table 9) otherwise, the transactional database is mined for this specific rule. The latter involves projecting the database with attributes in the query, thus reducing the number of attributes in the transactions, and mining as described in 4.1.

In the former case, NPS are generated and the rule is stored in the new rule base with appropriate support, for example [proteins, ideal] \rightarrow [carbohydrates, low], 35%. A rule of the form “Diet Coke \rightarrow Horlicks, 24%” could be evaluated to many rules including for example, [Proteins, ideal] \rightarrow [Carbohydrates, low], 45%; where, according to rule representations shown in section 3, X is “Proteins”, A is “ideal” and Y is “Carbohydrates”, B is “low” etc. The same transformation to an RDA table occurs and the average value per nutrient is calculated before conversion to membership degrees or linguistic values. Using equations 2, 3, 4 and 5, we evaluate final rules \expressed as linguistic

values. The following example shows a typical query as described in 4.1 where TID is transaction ID, X,Y, Z are items and P (protein), Fe (Iron), Ca (calcium), Cu (Copper) are nutritional elements and support (Supp) and confidence (Conf) is given:

Rules	Support
X→Y	24%
Y→Z	47%
X,Y→Z	33%
..	..

	Pr	Fe	Ca	Cu	..
X→Y	20	10	30	60	..
..

Nutritional Pattern
 X→Y [Proteins, Very Low] → [Carbohydrates, Low],
 Supp=45%, Conf=20%;

5 EXPERIMENTAL RESULTS

In order to show the defined frameworks effectiveness, we performed experiments using the prototype tool (figure 3) using T10I4D100K dataset containing simulated market basket data [generated by the IBM Almaden Quest research group] (Agrawal, IBM). The data contains 100K transactions and 1000 items. We considered 600 edible items out of the 1000 and used a real nutritional standard RDA table to derive fuzzy sets. Nutritional Patterns are then generated from T10I4D100K dataset using methodology as described in section 4.

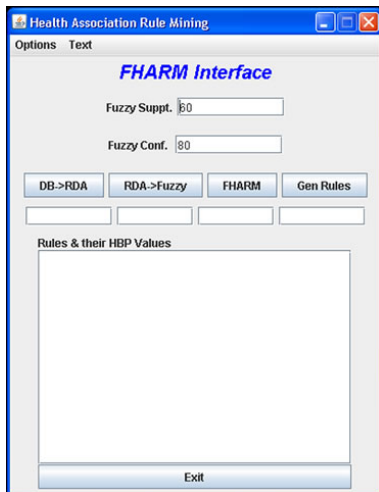


Figure 3: Prototype Tool.

5.1 Experiment One

Experiment for method mentioned in 4.1 shows how our approach produces NPs in terms of interesting rules. We use all the 27 nutrients with T10I4D100K dataset.

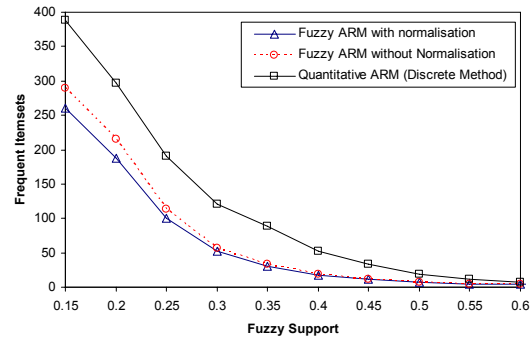


Figure 4: Number of Interesting Rules using fuzzy correlation.

Figure 4 shows difference between number of frequent items produced by fuzzy method (with and without normalisation) and discrete method (discretised boolean data). From the results, it is clear that the approach with boolean data produces more and irrelevant rules than the fuzzy approach. This is because using discrete method we cannot get the actual membership degree of nutrients intake for different intervals in each transaction thus considering full membership in any of the interval.




Figure 5: Number of Interesting Rules using fuzzy confidence.

Figure 5 and Figure 6 shows the number of interesting rules using user specified fuzzy confidence and fuzzy correlation values respectively. Correlation has not been applied to quantitative ARM algorithm due to the boolean data and so only Fuzzy approaches (with and without normalisation) have been shown in figure 6.




Figure 6: Number of Interesting Rules using fuzzy correlation.

Figure 6 presents less and more interesting rules than figure 5 because it uses the correlation value for evaluation of interestingness between the antecedent and the consequent of a rule. The experiments show that normalization before applying correlation yields significantly less rules. In addition, the novelty of the approach is in being able to analyse nutritional content of itemsets or rules. Some interesting rules produced by our approach are as follows:

IF *Protein* intake is *Ideal* THEN *Carbohydrate* intake is *low*.

IF *Protein* intake is *Low* THEN *Vitamin A* intake is *High*.

IF *Protein* intake is *High* AND *Vitamin A* intake is *Low* THEN *Fat* intake is *High*.

Depending on expert analysis, these rules are useful in analysing customer buying behaviour concerning their nutrition. Again, using fuzzy aggregation operators for these linguistic rules could show real aggregated values for each rule and hence their strength if measured against fuzzy thresholds.

5.2 Experiment Two

We also implemented the algorithms for analysing rule queries and calculating fuzzy support and fuzzy confidence for methodology described in section 4.2. For missing nutrient values or “trace” elements, the fuzzy function evaluated zero degree membership. We run classical ARM AprioriTFP algorithm on the data to produce a rule base. Some of the rule queries are as follows:

Rule 1: Milk \rightarrow *Honey*, *Support=29%*

The rule is evaluated accordingly (see 4.2) as

44% - Very Low in [Calcium Cholesterol Fats Iodine Magnesium Manganese Phosphorus Sodium VitaminA VitaminC VitaminD VitaminK]

3% - Low in [VitaminB12]

14% - Ideal in [Fiber Protein VitaminB6 Zinc]

7% - High in [Niacin VitaminE]

29% - Very High in [Biotin Carbohydrate Copper Folacin Iron Riboflavin Selenium Thiamin]

Rule 2: Cheese, Eggs \rightarrow *Honey*, *Support=19%*

37% - Very Low in [Calcium Fats Iodine Magnesium Phosphorus VitaminA VitaminB12 VitaminC VitaminD VitaminK]

3% - Low in [Carbohydrate]

22% - Ideal in [Manganese Protein Sodium VitaminB6 VitaminE Zinc]

3% - High in [Cholesterol]

33% - Very High in [Biotin Copper Fiber Folacin Iron Niacin Riboflavin Selenium Thiamin]

Rule 3: Jam \rightarrow *Milk*, *Support=31%*

48% - Very Low in [Calcium Cholesterol Fats Iodine Iron Magnesium Phosphorus] etc.

It is surprising to see that for most rules (at least these shown here), calcium purchases from calcium rich products like milk and cheese are very low. Contrary, Biotin (Vitamin H, rules 1 and 2) deficiency that causes cholesterol, loss of appetite, hair loss etc is very high possibly because it is found in egg yolks and milk (dry skimmed). These inferences could be useful in real data applications.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we presented a novel framework for extracting nutritional patterns (NPs) from customer transactions. Projections are made on input data into edible attributes to find fuzzy association rules using nutrients actual membership degrees. Standard health information for each nutrient is provided as fuzzy RDA data. Fuzzy support and confidence as well as correlation are used as interestingness measures. A user can extract interesting HBP (Healthy Buying Patterns) rules from the transactions or from a given rule in the rule base.

In future, we intend to evaluate our approach on real and larger customer data. The determination of our presented approach for nutrient analysis is also important and viable future work, which should involve expert knowledge in evaluating the real value of nutritional patterns in health terms. Defuzzification of rules will bring new insights to the real healthy trends described in the paper. Similarly, non-coded food items (those without nutritional information) like fruits, vegetable in general etc need to be fuzzified in some other way and evaluated as shown.

Overall, the approach presented here is effective, efficient and could be very useful for both the customer and health organizations.

REFERENCES

- Agrawal, R., Srikant, R., IBM. Quest Synthetic Data Generator. IBM Almaden Research Center, http://www.almaden.ibm.com/cs/projects/iis/hdb/Projects/data_mining/datasets/data/assoc.gen.tar.Z
- Agrawal, R., Imielinski, T., Swami, A. N., 1993 Mining Association Rules Between Sets of Items in Large Databases. *In Proceedings of the 1993 ACM SIGMOD Conference on Management of Data, (1993), 207-216.*
- Bodon, F., 2003. A Fast Apriori Implementation. *IEEE ICDM Workshop on Frequent Itemset Mining Implementations (FIMI'03), Melbourne, Florida, USA, Vol. 90.*
- Lee, C.-H., Chen, M.-S., Lin, C.-R., 2003. Progressive Partition Miner: An Efficient Algorithm for Mining General Temporal Association Rules. *IEEE Transactions on Knowledge and Data Engineering, Vol. 15, No. 4, 1004 - 1017.*
- Chen, G., Wei, Q., 2002. Fuzzy Association Rules and the Extended Mining Algorithms. *Information Sciences-Informatics and Computer Science: An International Journal archive, Vol. 147, No. (1-4), 201 - 228.*
- Wai, A.-H., Chan, K., 1999. Farm: A Data Mining System for Discovering Fuzzy Association Rules. *In Proceedings of the 18th IEEE Conference on Fuzzy Systems, 1217-1222.*
- Srikant, R., Agrawal, R., 1996. Mining Quantitative Association Rules in Large Relational Tables. *In Proceedings of ACM SIGMOD Conference on Management of Data. ACM Press, 1-12.*
- Dubois, D., Hüllermeier, E., Prade, H., 2006. A Systematic Approach to the Assessment of Fuzzy Association Rules. *Data Mining and Knowledge Discovery Journal, Vol. 13, No. 3, 167-192.*
- Xie, D. -W., 2005. Fuzzy Association Rules discovered on Effective Reduced Database Algorithm. *In Proceedings of IEEE Conference on Fuzzy Systems.*
- Yuanchen, H., Tang, Y., Zhang, Y-Q., Synderman, R., 2006 Adaptive Fuzzy Association Rule Mining for Effective Decision Support in Biomedical Applications. *International Journal Data Mining and Bioinformatics, Vol. 1, No. 1, 3-18.*
- Gyenesei, A., 2001. A Fuzzy Approach for Mining Quantitative Association Rules, *Acta Cybernetical, Vol. 15, No. 2, 305-320.*
- Paetz, J., 2002. A Note on Core Regions of Membership Functions, Proc. of the 3rd Int. Symp. of Medical Data Analysis (ISMDA 2002), LNCS Vol. 2526, 42-52.
- Coenen, F., Leng, P., Goulbourne, G., 2004. Tree Structures for Mining Association Rules. *Journal of Data Mining and Knowledge Discovery, Vol. 15, No. 7, 391-398.*
- Wang, C., Tjortjis, C., 2004. PRICES: An Efficient Algorithm for Mining Association Rules. *In Proceedings of the 5th Conference on Intelligent Data Engineering Automated Learning, Lecture Notes in Computer Science Series, Vol. 3177, Springer-Verlag, 352-358.*
- Lavrac, N., Keravnou ET, Zupan B., 1996. Intelligent data analysis in medicine and pharmacology. *First International Workshop on Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-96). Dordrecht: Kluwer Academic Publishers.*
- Delgado, M., Sánchez, D., Vila, MA., 2000. Acquisition of fuzzy association rules from medical data. *In: Barro, MarõAn R, editors. Fuzzy logic in medicine. Wurzburg (Wien): Physical-Verlag, in press.*
- Muyeba, M., M. Sulaiman Khan, M. Malik, Z., Tjortjis, C., 2006. Towards Healthy Association Rule Mining (HARM), A Fuzzy Quantitative Approach, *In Proc. 7th Conf. Intelligent Data Engineering Automated Learning, LNCS Series, Vol. 4224, Springer-Verlag, 1014-1022.*
- Kuok, C., Fu, A., Wong, H., 1998. Mining Fuzzy Association Rules in Databases. *ACM SIGMOD Record, 27, 41-46.*

TRUSTED SMS

A Novel Framework for Non-repudiable SMS-based Processes

Antonio Grillo, Alessandro Lentini, Gianlugi Me

Department of Computer Science and Engineering, University of Rome Tor Vergata, Rome, Italy
antonio.grillo@hotmail.it, lentini.alessandro@gmail.com, me@disp.uniroma2.it

Giuliano Rulli

Spike Reply s.r.l., Reply S.p.A., Rome, Italy
g.rulli@reply.it

Keywords: Confidentiality, data security, sms, ECDSA, elliptic curves.

Abstract: The exponential growth of the Short Message Service (SMS) use has led this service to an indispensable tool for social, marketing and advertising messaging. Moreover, mobile devices such as smartphones, handsets and PDAs represent an enabling factor for distributing digital content. Mobile devices are quickly becoming Personal Trust Device (PTD); mobile devices embed personal data, which allow sending/receiving confidential information from/to the PTD. This paper aims to introduce Trusted-SMS, a novel framework to exchange secure SMS. This system is composed by three main entities: the Service Supplier, which publishes and delivers services; the End User, which chooses and eventually pays for a specific service, that belongs to the service-set offered by a Service Supplier; the Certification Authority (CA) which represents the trusted entity shared by the Service Supplier and the End User. The CA plays the role of the Certification Authority. The main requirements of the overall system are strictly non-repudiability, user friendliness and platform portability. The security requirement includes customer transaction authentication, confidentiality, integrity and non repudiation, in an environment composed of heterogeneous networks and devices, with different security weaknesses. Trusted-SMS allows exchanging SMS digitally signed with Elliptic Curve Digital Signature Algorithm. SMS digitally signed are useful in many scenarios, such as commercial transaction, production of delegation from a remote site and provisioning of e-healthcare services. The signature is fully contained in a single SMS; the size of a digital signature amount to fifty bytes leaving more than one hundred bytes (110 bytes) for the SMS payload. Moreover the application of Elliptic Curve Integrated Encryption Schema cryptographic algorithm, which is based on the same credentials needed by the digital signature algorithm, allows protecting the payload from intrusions.

1 INTRODUCTION

Short message services (SMS), thanks to its relative simplicity and ease of use, similar to the e-mail, keeps on growing in penetrating the market throughout the world. In fact, during the 2000, just 17 billion SMS messages were sent; in 2001, the number was up to 250 billion, and 500 billion SMS messages were sent in 2004 (Kivimaki and Fomin, 2001). At an average cost of USD 0.10 per message, this generates revenues in excess of \$50 billion for mobile telephone operators and represents close to 100 text messages for every person in the world.

Text messaging has become so popular that advertising agencies are now considering it as an irrinuciable

vehicle to keep the customer up to dated about the new offers and products (Dickinger et al., 2004). The services providing bulk text messages are also becoming a popular way for clubs, associations, and advertisers to quickly reach a group of opt-in subscribers. This advertising has proven to be extremely effective, but some insiders worry that advertisers may abuse the power of mobile marketing and it will be considered spam.

Moreover, the proliferation of SMS attracts malware writers that adapt phishing and other activities based on social engineering techniques (e.g. spoofing become SMS-spoofing) in order to manipulate people into performing actions or divulging confidential information (van der Merwe et al., 2005).

This paper aims to describe a novel framework to exchange secure SMS: Trusted-SMS. Due to the well known, or yet undiscovered, security weaknesses of SMS, it's important, for some kind of applications, to enforce security layers, such as confidentiality, integrity and non repudiation. The goals of this paper are concerned with:

- a Expressing the requirements of the Trusted-SMS framework;
- b Analyzing and designing a framework meeting those needs;
- c Characterizing that framework, with a sketch of a case study.

In the remaining of the paper, the Section 2 sketches the current level of development of modern methods and products that involve a secure SMS exchange. The Section 3 describes Trusted-SMS framework; this section deals with framework needs and goals in required features, architectural design and some technical details. The Section 4 presents a prototype system compliant to Trusted-SMS framework as case study. Section 5 presents some conclusions, and points to the future trends.

2 STATE OF THE ART

The huge success of some modern network protocols, not designed with security in mind, poses many concerns about fraudulent intrusions and undiscerning violation of the privacy. As for the TCP/IP case, the SMS suffer the lack of many features that are desirable or needed on an insecure network, as the GSM can be considered. For this reason, the companies offer some security-enhancing products, which allow to protect private information exchanged via SMS. The available solutions can be divided into two different models: "Peer-to-Peer" (P2P) and "Client-Server" (CS); the former focuses on the SMS exchange which involves only mobile devices, while the latter focuses on the SMS exchange which occurs between a mobile device, as the Client, and a remote computer, as the Server.

Many solutions belong to the P2P model, such as Message in a Bottle (Chirico, 2007), Spider-SMS (Barbi, 2007), Kryptext (Kryptex, 2007), MultiTasker (MultiTasker, 2007) and CryptoSMS (CryptoSMS, 2007). Products as Fortress-SMS (FortressSMS, 2007) follow the CS model, this can be view as a generalization of products based on P2P model. The level of security of the previous solutions is strictly related to the key-distribution schema and the security algorithm applied. Miabo and Spider SMS cannot be used

for sending and receiving information as part of business transactions and personal communications. They don't fit security needs because the former relies on a PGP schema as public key infrastructure and the latter relies on a random key generator.

Fortress-SMS relies on a Rijndael engine with cipher block chaining (Advanced Encryption Standard compliant) as encryption algorithm; CryptoSMS uses three overlapping strong encryption schemes, employing both block and stream ciphers. Fortress-SMS and CryptoSMS are not fully documented about their execution times; to the best of our understanding encryption schemas adopted by these products require long processing times on mobile devices (Waadt et al., 2005).

Many products allow both the SMS digital signature and the SMS encryption; Trusted-SMS allows SMS digital signature or alternatively SMS signature on SMS encrypted.

Furthermore, the GSM specifications (3GPP, 2007) don't define any mechanism for ensuring integrity of SMS content and authentication of SMS sender. SMS digitally signed can be used to avoid SMS-tampering, ensuring integrity and authentication of the sender (Center, 2007). SMS encrypted are suitable to avoid unauthorized access to SMS content.

3 TRUSTED-SMS FRAMEWORK

Trusted-SMS is a complex framework for allowing secure SMS exchange between involved entities. This framework enables a user to send or receive SMS messages digitally signed with standard Elliptic Curve Digital Signature Algorithm (ECDSA) and optionally encrypted with standard Elliptic Curve Integrated Encryption Schema (ECIES). According to the digitally signed SMS, this framework can be used to avoid SMS from tampering ensuring integrity and non repudiation: the recipient can detect the tampering during the SMS verification process. With encrypted SMS, this framework can be used to avoid SMS from unauthorized access. If someone intercepts an encrypted SMS he cannot eavesdrop the SMS content.

Trusted-SMS is composed by three main entities:

Service Supplier: this entity is responsible for providing services;

End User: this entity is the consumer of the supplied services;

Certification Authority: this entity is responsible to ensure trusting to the involved entities.

3.1 Required Features

The expected use cases and the resulting requirements suggest the following list of features, trying to depict the Trusted-SMS framework. For the sake of simplicity, we divide the framework required features in three subsets, which are related to the previously identified entities:

Service Supplier entity features:

1. each Service Supplier owns a Large Account number; clients of a Service Supplier use its Large Account number as its identifier;
2. the Service Supplier has to be able to generate, send and receive different kind of SMS;
3. each SMS, which Service Suppliers have to send to Customers, must be ECDSA-Signed using the CA private key (KPrivCA);
4. each SMS, which Service Suppliers receive from a Customer, must be ECDSA-Verified using the Customer public key (KPubMSISDN);
5. the Service Supplier has to manage a large number of information as end user personal data which has to be stored in a persistent memory storage;
6. the Service Supplier wants to preserve SMS content also from CA access.

End User entity features:

1. the subsystem placed on the End User platform, a smartphone, has to be realized as a MIDP 2.0 java mobile application;
2. End User platform has to be able to manage different kinds of SMS;
3. each End User is characterized by a key-pair: a public key(KPubMSISDN) and a private key (KPrivMSISDN);
4. each SMS, which Customers have to send to Service Suppliers, must be ECDSA-Signed by client application using its own private key (KPrivMSISDN);
5. each SMS, which Customers receive from Service Suppliers, must be ECDSA-Verified by client application using the CA public key (KPubCA);
6. the subsystem placed on the End User platform must be defined on a parameters set which defines the elliptic curve; that elliptic curve is used for generating/verifying ECDSA signatures (this curve is shared with server application);
7. the subsystem placed on the End User platform must be defined on a parameters set which defined the hash function; that hash function is

used for creating a digest which must be signed with the identified signature algorithm;

Certification Authority entity features:

1. a single server realizes the Certification Authority (CA); each Service Supplier knows this server.
2. the CA must be able to sign SMS using ECDSA with its own private key (KPrivCA);
3. the CA must be able to verify SMS using ECDSA with an appropriate public key (KPubMSISDN);
4. the CA has to manage a large number of confidential information (e.g. end users' keys), which has to be stored in a persistent memory storage;
5. each keypair is characterized by a temporal validity;
6. the CA subsystem offers a graphical user interface in order to allow administration of some services;
7. inserting/deleting end users in/from CA repository has to be an administration task;
8. updating end user status (e.g. able/disable) has to be an administration task; a disabled user is temporally excluded from whole system, this feature is activated when end user's devices are lost or stolen, when contract's temporal validity is expired and is in line for renewed;
9. updating end user key status (e.g. able/disable) has to be an administration task, an end user with a disabled key cannot sign SMS while he can continue to verify information sent to its device.
10. inserting/deleting service suppliers in/from CA repository has to be an administration task;
11. updating service supplier status (e.g. able/disable) has to be an administration task.

3.2 Architectural Design

The architectural design evolves naturally from identified required features. Figure 1 proposes the overall architecture of Trusted-SMS framework. Despite of previous definition of the framework as result of three main entities, figure 5 shows six components:

Published Services: it represents the common space where Service Suppliers interact with End Users advertising services and products.

Service Supplier: it represents the subsystem that realizes the Service Supplier entity features previously identified.

Certification Authority: it represents the subsystem that realizes the Certification Authority entity features previously identified.

End User: it represents the subsystem that realizes the End User entity features previously identified.

SMS Gateway: it represents a generic Short Message Service Center (SMSC).

CA Administrator: it represents the front end to CA administration services.

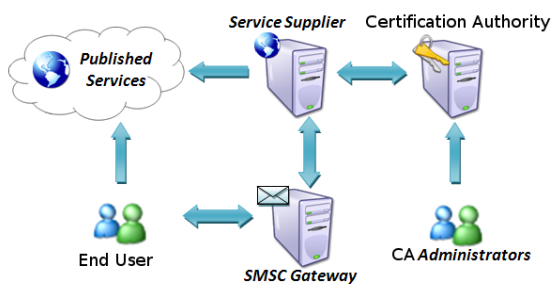


Figure 1: Overall architecture of Trusted-SMS framework.

This paper goes on to describe accurately identified subsystem.

Service Supplier subsystem: The Service Supplier subsystem can be divided in 4 main components (Fig. 2):

- a component that acts as general purpose end user interface, publishes provided services;
- a component that acts as interface with the sms gateway, sends and receives SMS to/from end users;
- a component that acts as interface with the Certification Authority in order to request a signature for a text or to verify an SMS signed by an user;
- a component that can be view as Service Supplier core, generates SMS text, interacts with others components in order to process and store end user requests and responses.

An end user request for a specific service could generally done by interacting with the service supplier directly or by a dedicated interface. The service supplier subsystem collects end user requests, automatically converts the request into a SMS text, requests to CA a signature for this text and finally forwards the SMS to the end user through the SMS gateway. In order to preserve confidentiality of the SMS content the signature is requested for a hash of that content. Reaction to an end user response involves the

same service supplier subsystem components. The response is received through the SMS gateway interface and it is stored into the database. The service supplier interacts with the CA to verify the appended signature. The result of the verifying test is recorded in the database. If the requested service counts a kind of confirmation, the end user is informed by a new SMS communication or a message on the service supplier interface.

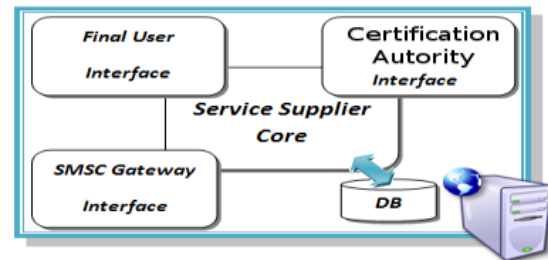


Figure 2: Structure of Service Supplier subsystem.

End User subsystem: The End User interacts with the remaining system by two different ways.

In the first case the End User chooses a specific service from the set of the published services; the Service Supplier of that service asks for confirmation. The End User has to confirm its choice to make use of requested service.

In the second case, the End User receives a communication about a service directly from the Service Supplier; only if necessary the End User answers to that communication. Service Supplier in this case can send service-communication without any explicit End User requests. In the first case described above, the End User realizes its interactions using two different channels; the first channel is used to select a service, the second one is used to accept the delivery of the selected service. In the second case these channels are coincident.

Separating the capability of requesting one service from the capability of accepting or denying the delivery of that service ensures an high level of security. If those channels rely on different technologies a potential attack becomes a very hard challenge. The main software component of the End User entity is represented by a midlet: a Java 2 Micro Edition (J2ME) application that can be executed on a personal mobile device characterized by low computational capability, such as a smartphone or a pda. The midlet has two main responsibilities; it has to verify the signatures that are carried by the incoming SMS and it has to generate the signatures that are carried by the outgoing SMS according to the CA and End User credential. Each midlet deployed on a mobile device has

to be provided with the CA public key and the End User private key; the former is used to verify CA signatures, while the latter is used to generate End User signatures.

End User credential (i.e. its cryptographic keys) are included in the midlet resources and protected by using a symmetric key injected into the midlet code. Received and sent SMS are stored in a secure separate repository, different from the usual mobile phone's SMS repository, secured and encrypted by user's PIN. Thus, if someone snaffles your mobile phone full of your private SMS he cannot read them because they are protected inside the repository.

Figure 3 shows the End User subsystem pointing out that three units made up this subsystem; SMS I/O Interface, which manages the SMS sending and receiving activities, and Graphical User Interface, which leads the interaction with the User, are coordinated by the End User Core. In order to store application data End User Core interacts with Record Management System(RMS). Cryptographic operations (i.e. signature generation or verification) are performed by the Core unit. Service Supplier Interface represents the second unit; listing all services offered by supplier and requesting for a service delivery represent the main issues that the second unit has to resolve.

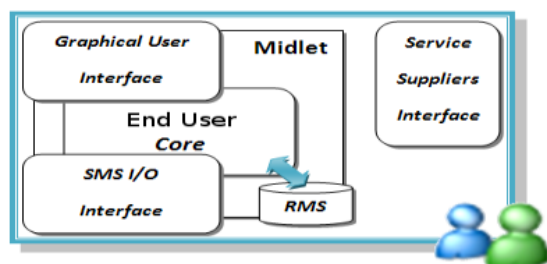


Figure 3: Structure of End User subsystem.

Certification Authority subsystem: The Certification Authority represents the trusted entity; End User and Service Supplier share the same confidence towards the CA entity. The CA subsystem has two main communication partners: the CA Administrator and the Service Supplier.

Interactions with CA Administrator are intended for managing users registered to Trusted-SMS framework. Service Supplier interacts with CA subsystem in order to obtain the generation/verification of signature carried by SMS.

The CA subsystem is supported in all its activities by a database, which allows the permanent storage of needed information. Database keeps information about End Users, Service Suppliers and CA itself. End Users information concern with personal

data (e.g. name, surname) and personal credential(e.g. public key, public-key-validity flag, user-enabled flag). In order to insert new users, delete or enable/disable existent user, enable/disable users keys those information are accessed in a "write-mode" by CA Administrator with related services. In order to verify End User generated signatures, credential information are accessed in a "read-mode" by services related to the Service Supplier. CA information concern with CA credential and CA Administration privileges; the former are accessed by Service Supplier related services in order to generate a signature, the latter are used in authentication process while acquiring CA management control.

Finally, Service Supplier information are related with supplier data, which are accessed by Service Supplier related services in order to verify if that Service Supplier is licensed to use these services.

Figure 4 proposes the structure of the Certification Authority subsystem. Certification Authority Core provides cryptographic and database services to the other modules of this subsystem. These modules manage the communication with CA partners.

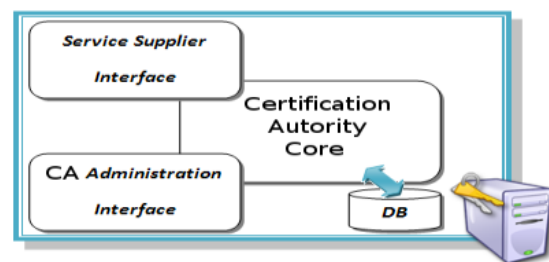


Figure 4: Structure of Certification Authority subsystem.

3.3 Technical Features

Elliptic curve cryptography (ECC) (Martinez et al., 2005) is a family of public-key algorithms that can provide shorter key lengths and, depending upon the environment and application in which it is used, improved performance over systems based on integer factorization and discrete logarithms (Higgins and Clement, 2001). ECC has its security based on a difficult mathematical problem. An elliptic curve can be thought of as a mathematical structure in which certain operations can be defined. These operations provide a one-way function that can be used to produce efficient cryptographic systems. The one-way function used in ECC is called the elliptic curve discrete logarithm problem (ECDLP). The ECDLP is similar to the one-way function on which DSA (of Standard et Technology, 2007) and Diffie-Hellman are based, and hence, elliptic curve analogs of each of

these algorithms have been defined. The most popular signature scheme which uses elliptic curves is called the Elliptic Curve Digital Signature Algorithm (ECDSA) (Jonson et al., 2001). The Secure Hash Algorithm (SHA) (F.I.P.S., 2002) hash functions refer to five FIPS-approved (F.I.P.S., 2007) algorithms for computing a condensed digital representation (known as a message digest) that is, to a high degree of probability, unique for a given input data sequence (the message).

Trusted-SMS relies on ECDSA to sign the exchanged data, thus avoiding unnoticed modification. The SHA1 hash function makes a digest of the data. Reducing a message in a digest lowers computational capabilities needed by the signing task, as we need, since the mobile devices involved in the system are very resource bounded. Trusted-SMS uses ECDSA with a 192 bit Elliptic Curves (EC) with a consequent 48 byte signature value. The Keys size of 192 bit agrees with the size of an SMS that is 140 bytes for binary and 160 bytes for textual SMS. In comparison to a 1024 bit RSA key, ECC (Elliptic Curve Cryptography) provides shorter keys, shorter encrypted messages and faster private key operations. Comparing the two cryptographic systems, a 1024 bit RSA key is considered to have the same security of a 160 bit ECC key. The cellular networks can be broken into two chief components - the radio, or "air interface" and the wired backbone (out of scope of this paper). The air interface of a cellular network can be divided into control and traffic channels. Control channels (CCHs) are used for call setup and SMS delivery. Traffic channels (TCHs) are used for the duration of voice calls. It helps to think of control channels as a very small portion of radio frequency that allow cellular towers to send information pertaining to call setup, SMS delivery and network conditions (such as the availability of traffic channels) to mobile phones. The Short Message Service - Point to Point (SMS-PP) is defined in GSM recommendation 03.40 (3GPP, 2007). GSM 03.41 (3GPP, 2007) defines the Short Message Service - Cell Broadcast (SMS-CB) which allows messages (advertising, public information, etc.) to be broadcast to all mobile users in a specified geographical area. Messages are sent to a Short Message Service Centre (SMSC) which provides a store-and-forward mechanism. It attempts to send messages to their recipients. If a recipient is not reachable, the SMSC queues the message for later retry. Transmission of the short messages between SMSC and phone can be done through different protocols such as Signaling System 7 (SS7) within the standard GSM MAP framework or TCPIP within the same standard. Messages are sent with the additional MAP operation for-

ward_short_message, whose payload length is limited by the constraints of the signalling protocol to precisely 140 bytes (140 bytes = 140 * 8 bits = 1120 bits). In practice, this translates to either 160 7-bit characters, 140 8-bit characters, or 70 16-bit characters. General Packet Radio Service (GPRS) is a packet-switched technology. GPRS that is an extension of GSM can be used as the bearer of SMS. Spoofing from a mobile is impossible unless you can forge GPRS radio traffic (Levy and Arce, 2004).

4 A CASE STUDY

Realizing a prototype developed with open source tool we illustrate the validity of our framework. We utilized Trusted-SMS framework in a specific health-care environment. E-Health involving information and communication technologies that are at the service of a wide range of actors in the health sector, from doctors, hospital managers, nurses, data processing specialists, social security administrators to - of course - patients and citizens, is a good test-bed for our framework.

The Certification Authority performs its tasks by using Web Services. The W3C Web service definition refers to clients and servers that communicate using XML messages that follow the SOAP-standard (Lai et al., 2005). The CA publishes two different lists of web services, each one for its communication partners: the CA Administrator and the Service Supplier. The permanent storage of information is achieved using a mysql database and it is protected by symmetric key encryption (Zoratti, 2006).

Service Supplier publishes its services on an interface, from which end users can request the delivery of a specific service. Once the end user requests a service the related request is stored into a mysql database. For each request received the supplier use the signature generation service in order to produce the related SMS. Each received SMS is drawn from the database and it is processed in order to verify its signature validity. The Service Supplier subsystem utilizes the signature verification/generation services exported by CA as web services.

A SMS gateway is realized by an external hardware component that behaves as a gateway between GSM mobile phone network and web applications (AreaSX, 2007). This component provides the SMS sending/receiving capability using simple GET or POST HTTP transactions over Ethernet.

The End User subsystem is mainly represented by a Java 2 Micro Edition application compliant to MIDP 2.0 and CLDC 1.1 and installed on different models of

Nokia mobile phones: series 40 equipped with Symbian OS 8.0 (i.e. 6630, 6680), series 60 2nd edition equipped with Symbian OS 8.1 (i.e. N70), and series 60 3rd edition equipped with Symbian OS 9.1 (i.e. N73, E70). Bouncycastle Crypto APIs offer support for realizing cryptographic operations. This APIs are an open source solution, which works with everything between J2ME and JDK 1.6 (Bouncycastle, 2007).

Next, we present the description of two use cases, both derived from different ways of the end user interaction with the framework described above.

In first use case, the End User (a patient) has a preliminary contact with a medical professional for receiving a complete check-up. As soon as the physicians have the results for physical examinations, he sends a SMS signed to the patient. The SMS carries personal information about patient health status protecting them with some cryptographic techniques. Signing the SMS content ensures patient on sender authenticity; moreover cryptographic protection guarantees patient privacy. In this way physicians can inform immediately the patient about the examination, if necessary they can ask for a secondary care.

The second use case, which refers to a more complex situation, will be described using a schematic representation (Fig. 5). In particular this case-study depicts the whole scenario including all interaction from the preliminary registration phase. In order to reach clearness in exposition, description is organized as a finite sequence of steps, which refers to Figure 5.




Figure 5: Structure of Certification Authority subsystem.

1. The CA-Administrator add End Users and Service Suppliers to the framework, this represents the zero step of the happy scenario of the case study. This step is executed for every End User/Service Supplier, who decides to subscribe to Trusted-SMS services.
2. The Service Supplier that delivers a set of services publishes them on its interface (such as a web site).
3. The End User (a patient) has an initial contact with a medical professional for receiving health examination.

4. The End User chooses a specific service (such as request for its own case history, specific examination results, information about personal cure, reservation for a secondary care, buying a medical examination, etc.) from that interface and
5. its request reaches the Service Supplier.
6. The Service Supplier produces an SMS related to the request received.
7. Utilizing the signature generation service offered by Certification Authority, the Service Supplier signs the SMS content
8. The CA checks both the Service Supplier credentials and the End User credentials; only if the credentials are valid, CA begins the signature generation.
9. The CA sends the signed response to the requesting Service Supplier.
10. The Service Supplier forwards the digitally signed SMS to SMS Gateway, (10) which delivers it to the End User.
11. The End User verifies the signature carried by the SMS, if necessary accepts or refuses the service delivery and generate the signature for response SMS.
12. The response SMS is sent to the SMS Gateway,
13. which forwards it to the Service Supplier.
14. The Service Supplier requests to utilize the digital signature verification process.
15. The CA processes the Service Supplier request by checking the credentials of entity involved and
16. answer to Service Supplier with the status of verification process.
17. The Service Supplier realizes the End User choice (cfr. step 10 accept or refuse) if the End User digital signature was successfully checked.
18. The Service Supplier publishes on the web site the transaction status
19. The End User can control the effective status of the service delivery. Steps 18 and 19 are conditional on the specific requested service.

5 CONCLUSIONS

The growth of mobile technology has opened various opportunities, both in marketing and in M-Commerce applications; technologies such as cell phone networks are becoming integrated with other systems such as the Internet. The SMS have been very popular among mobile phone users because it is silent, fast, and cheap. Unfortunately, this may lead to new security risks, for example when conducting phishing attacks using mobile technology and especially the SMS.

In this paper we proposed a framework to exchange secure SMS with respect to: (a) *security*, communication between principals are based on well cryptographic techniques, (b) *traceability*, transaction are not kept anonymous, (c) *usability and convenience*, cost of deployment and management for all subsystems involved (consumer, service suppliers, trusted service provider) are acceptable.

Furthermore, in order to set Trusted-SMS into a specific e-healthcare scenario, we made a prototype developed with open source tool. Since current trends in mobile phone technology move towards a direction of miniaturization and higher computational and graphical performance, allowing a complete transaction procedure in less than a minute, we believe that Trusted-SMS prototype can show the validity of our framework in the field.

Finally, the case study covers only one of the possible uses of Trusted-SMS framework. The systems managing private information, or systems that schedule a booking/reservation procedure (e.g. system that can be used by theatre, stadium cinemas, airline company, university), or involving money transfers, can benefit of security features provided by Trusted-SMS framework.

REFERENCES

- 3GPP (2007). 3rd generation partnership project.
- AreaSX (2007). Sms machine/http.
- Barbi, L. (2007). Spidersms - sending and reception of encrypted sms.
- Bouncycastle (2007). The legion of the bouncycastle.
- Center, C. C. R. (2007). Sms spoofing - q-a with csrc staff.
- Chirico, U. (2007). Miabo - messages in a bottle.
- CryptoSMS (2007). Cryptosms - protecting your confidential sms messages.
- Dickinger, A., Haghirian, P., Murphy, J., and Scharl, A. (2004). An investigation and conceptual model of sms marketing. In *System Sciences, Proceedings of the 37th Annual Hawaii International Conference on*.
- F.I.P.S. (2002). Secure hash standard, fips publication.
- F.I.P.S. (2007). Federal information processing standards.
- FortressSMS (2007). Fortresssms - phone based application to send and read encrypted sms text messages.
- Higgins, J. Z. L. and Clement, M. (2001). Performance of finite field arithmetic in an elliptic curve cryptosystem. *Security Technology, CCST '05. 39th Annual 2005 International Carnahan Conference on*, pages 249 – 256.
- Jonson, D., Menezes, A., and Vanstone, S. (2001). The elliptic curve digital signature algorithm.
- Kivimaki, A. and Fomin, V. (2001). What makes a killer application for the cellular telephony services? *Standardization and Innovation in Information Technology, 2nd IEEE Conference*, pages 25 – 37.
- Kryptex (2007). Kryptex - send and receive encrypted text sms.
- Lai, K. Y., Phan, T. K. A., and Tari, Z. (2005). Efficient soap binding for mobile web services. *Local Computer Networks 30th Anniversary, The IEEE Conference on*, pages 218 – 225.
- Levy, E. and Arce, I. (2004). Interface illusions. In *IEEE Security and Privacy*, pages 66 – 99.
- Martinez, V. G., Avila, C. S., Garcia, J. E., and Encinas, L. H. (2005). Elliptic curve cryptography: Java implementation issues. *Local Computer Networks 30th Anniversary, The IEEE Conference on*, pages 238 – 241.
- MultiTasker (2007). Multitasker - messaging made easy.
- of Standard et Technology, N. I. (2007). Digital signature standard.
- van der Merwe, A., Seker, R., and Gerber, A. (2005). Phishing in the system of systems settings: mobile technology. In *Systems, Man and Cybernetics, IEEE International Conference on*, pages 228 – 232.
- Waadt, A., Bruck, G., Jung, P., Kowalzik, M., Trapp, T., and Begall, C. (2005). Qos monitoring for professional short-message-services in mobile networks. *Wireless Communication Systems, 2nd International Symposium on*, pages 228 – 232.
- Zoratti, I. (2006). Mysql security best practices. *Crime and Security, The Institution of Engineering and Technology Conference on*, pages 183 – 198.

PROTECTING PRIVACY IN MEDICAL DATABASES

Efficient Local Generation of System-Wide Unique Health IDs

Peter Schartner and Martin Schaffer
Institute of Applied Informatics, System Security Group
Klagenfurt University, 9020 Klagenfurt, Austria
p.schartner@syssec.at, m.schaffer@syssec.at

Keywords: Privacy protection, anonymity, linkability, pseudonyms, nyms, medical registers, medical databases.

Abstract: In this paper we will introduce a replacement for linkable unique health identifiers: locally generated system-wide unique digital pseudonyms. The presented techniques are based on a novel technique called *collision free number generation* which is discussed in the introductory part of the article. After this brief introduction, we will pay attention onto two specific variants of collision-free number generation: one based on the RSA-Problem and one based on the Elliptic Curve Discrete Logarithm Problem. The main part of the article focuses on two applications of unique digital pseudonyms: centralized medical records and anonymous medical databases.

1 INTRODUCTION

In this paper we will present an application of digital pseudonyms in the scope of unique health identifiers (health IDs) and digital medical records. So far, unique health IDs are issued by some sort of centralized agency, in order to guarantee the system-wide uniqueness of the employed identifiers. This enables linking between the Health ID (and indirectly the patient's name) and the medical record. Hence, there are severe privacy concerns which are documented by the presence of "Health IDs & Privacy" in the media (Institute for Health Freedom, 2000; CNN, 2000) or even special web-sites which are protesting against such identifiers (Medical Privacy Coalition, 2007).

It is obvious that in some scenarios (e.g. electronic prescription) there must be a link between the health ID and the individual's name, but there are several scenarios, where such a link is not necessary or is even not wanted. Consider centralized medical records or anonymized databases for certain diseases (e.g. cancer registers). In these scenarios, the only requirement is that all data concerning a specific person is added to the right medical record. So, the real identity (i.e. the unique Health Identifier) can be replaced by a digital pseudonym, which has again to be unique (Pfitzmann and Köhntopp, 2001).

A critical problem is, how one can efficiently generate provably unique pseudonyms. The straightforward solution is to issue these pseudonyms by a centralized instance, which we have to trust completely. This, however, is an unrealistic assumption since there is no single instance whom one can trust. The ideal solution would be, that each individual generates the pseudonym on his own. Without cross-checking, however, there is a risk that some pseudonyms are chosen to be identical. This drawback can be overcome by using the technique of collision-free number generation for the establishment of pseudonyms. Pseudonyms are then

1. *generated locally*, i.e. derived from a system-wide unique identifier, and they are
2. *system-wide unique* without being linkable to the original identifier, from which they are derived.

The remainder of this paper is organized as follows. After going through some preliminaries we will briefly describe a scheme, that allows the individual to locally generate almost random numbers, which are globally (in particular system-wide) unique (Schartner and Schaffer, 2005; Schaffer et al., 2007). These numbers can be used as pseudonyms and hence they can be used to replace the original health ID. The individual's name and the corresponding ID (thus the

medical record) are then computationally unlinkable. After describing the original scheme with two practical implementations, we will present two application scenarios. In the first, we will discuss centralized databases which store a medical record for each individual in order to improve the quality and security of ongoing and later treatments. In the second, we will discuss databases for medical records of specific diseases which are used to statistically investigate them.

2 PRELIMINARIES

Henceforth, l_x denotes the bit-length of x .

AES Encryption (NIST, 2001). Let k be a random secret AES-key. The AES-encryption of a binary message m is given through

$$\text{AES}_k(m) = c$$

where c is the computed ciphertext. For simplicity, the decryption process is denoted similarly:

$$\text{AES}_k^{-1}(c) = m$$

It is assumed that if the message is large, the encryption is carried out by using an appropriate mode of operation with a standardized padding scheme.

RSA Encryption (Rivest et al., 1978). Let n be the product of two safe (or strong) primes p and q . The public key e is chosen at random (or system-wide static), such that $\text{gcd}(e, \phi(n)) = 1$. The corresponding private key d is computed such that $ed \equiv 1 \pmod{\phi(n)}$ holds. Given (e, n) , a message $m \in \mathbb{Z}_n$ is encrypted through the function $\text{RSA}_{(e,n)}$, defined as follows:

$$\text{RSA}_{(e,n)}(m) = m^e \text{ MOD } n$$

A ciphertext $c \in \mathbb{Z}_n$ can be decrypted through the function $\text{RSA}_{(d,n)}^{-1}$ defined as follows:

$$\text{RSA}_{(d,n)}^{-1}(c) = c^d \text{ MOD } n$$

The security of the scheme relies on the problem of factoring n and computing e -th roots modulo n . The problem of finding m , given c and (e, n) is sometimes called the RSA-Problem and defined as follows:

Definition 2.1 Let $n = pq$, where p and q are primes, $e \in \mathbb{Z}_{\phi(n)}^*$, $m \in \mathbb{Z}_n$ and $c = m^e \text{ MOD } n$. The *RSA-Problem* is the following: given c and (e, n) , find m .

Like all state-of-the-art implementations of RSA we employ optimal asymmetric encryption padding OAEP (Jonsson and Kaliski, 2002), which is a special form of padding, securing RSA against chosen ciphertext attacks.

Elliptic Curve Cryptography (ECC). To speed up computations, discrete logarithm-based cryptosystems are often run over an elliptic curve group, or in particular on a subgroup of prime order. A good introduction to ECC can be found in (D. Hankerson, 2004). We use scalar multiplication in such groups as a one-way function. The corresponding intractable problem is defined in the following:

Definition 2.2 Let $E(\mathbb{Z}_p)$ be an elliptic curve group, where p is an odd prime. Let $P \in E(\mathbb{Z}_p)$ be a point of prime order q , where $q \mid \#E(\mathbb{Z}_p)$. The *Elliptic Curve Discrete Logarithm Problem (ECDLP)* is the following: given a (random) point $Q \in \langle P \rangle$ and P , find $k \in \mathbb{Z}_q$ such that $Q = kP$.

To avoid confusion with ordinary multiplication we henceforth write $\text{SM}(n, P)$ to express the scalar multiplication nP in $E(\mathbb{Z}_p)$. It is currently believed that the ECDLP using $l_p \approx 192$ and $l_q \approx 180$ is secure against powerful attacks like Pollard's rho algorithm (D. Hankerson, 2004).

ECC Point Compression (IEEE, 2000). A point on an elliptic curve consists of two coordinates and so requires $2l_p$ bits of space. It is a fact that for every x -value at most two possible y -values exist. Since they only differ in the algebraic sign, it suffices to store only one bit instead of the whole y -value. We therefore define the point compression function $\text{CP} : E(\mathbb{Z}_p) \rightarrow \mathbb{Z}_p \times \{0, 1\}$ as follows:

$$\text{CP}((x, y)) = (x, y \text{ MOD } 2)$$

Storing (x, \tilde{y}) instead of (x, y) requires only $l_p + 1$ bits of space. To uniquely recover (x, y) from $(x, \tilde{y}) \in \mathbb{Z}_p \times \{0, 1\}$ one needs a decompression function $\text{DP} : \mathbb{Z}_p \times \{0, 1\} \rightarrow E(\mathbb{Z}_p)$, such that

$$\text{DP}((x, \tilde{y})) = (x, y)$$

How y is uniquely recovered depends on the kind of elliptic curve that is used and on how p is chosen. For instance, consider an elliptic curve defined through the equation $y^2 \equiv x^3 + ax + b \pmod{p}$, where a, b are some appropriate public domain parameters. If $p \equiv 3 \pmod{4}$, one can recover y from (x, \tilde{y}) through one modular exponentiation:

$$y_{0,1} \equiv \pm(x^3 + ax + b)^{(p+1)/4} \pmod{p}$$

where $y = y_{\tilde{y}}$. This is a very efficient way to compute square roots modulo a prime and is also used in the Rabin cryptosystem.

ElGamal Encryption (ElGamal, 1985). In this paper, we use a particular variant of the ElGamal encryption scheme. Let $E(\mathbb{Z}_p)$ be the elliptic curve group as described above and P a point of order

q . Then $d \in_R \mathbb{Z}_q$ denotes the private key whereas $Q = SM(d, P)$ denotes the corresponding public key. If done straightforwardly according to (ElGamal, 1985), a message $M \in E(\mathbb{Z}_p)$ can be encrypted by first choosing a random $r \in \mathbb{Z}_q$ and then computing the ciphertext-pair $(A, B) = (SM(r, P), M + SM(r, Q))$. Given (A, B) and d , M can be obtained through $B + SM(-d, A)$. However, it is not trivial to map a binary message m to a point M on an elliptic curve. An efficient method to overcome this drawback is to choose a cryptographic hash-function $\mathcal{H} : E(\mathbb{Z}_p) \rightarrow \{0, 1\}^{l_m}$, where l_m denotes the bit-length of the binary message m , and define the ElGamal encryption function ElG_Q :

$$\text{ElG}_Q(m) = (SM(r, P), m \oplus \mathcal{H}(SM(r, Q))), \quad r \in_R \mathbb{Z}_q$$

Given the ciphertext-pair $(A, B) \in E(\mathbb{Z}_p) \times \{0, 1\}^{l_m}$ and d , m can be obtained through ElG_d^{-1} :

$$\text{ElG}_d^{-1}((A, B)) = \mathcal{H}(SM(-d, A)) \oplus B$$

The correctness of this ElGamal variant is obvious.

3 COLLISION-FREE NUMBER GENERATION

In (Schaffer et al., 2007) we proposed a general method for generating system-wide unique numbers in a local environment, whilst preserving the privacy of the generating individual. The described generator, called collision-free number generator (CFNG), fulfils the following requirements:

R1 (Uniqueness). A locally generated number is system-wide unique for a certain time-interval.

R2 (Efficiency). The generation process is efficient regarding communication, time and space.

R3 (Privacy). Here we distinguish two cases:

1. *Hiding*: Given a generated number, a poly-bounded algorithm is not able to efficiently identify the corresponding generator.
2. *Unlinkability*: Given a set of generated numbers, a poly-bounded algorithm is not able to efficiently decide which of them have been generated by the same generator.

In the following subsection we summarize two general principles of collision-free number generation, formerly introduced in (Schaffer et al., 2007).

3.1 General Construction

For efficiency reasons, an identifier-based approach is used. Every generator is (once) initialized with a

system-wide unique identifier, which we denote by UI . The idea is, to derive several unique numbers from UI , such that none of them is linkable to UI .

Uniqueness Generation. As a first step a routine is needed, which derives a unique number u from UI in every run of the generation process. We call such a routine *uniqueness generator* and denote it by UG . UG can be designed as follows:

$$u = \text{UG}(), \quad u = UI || cnt || pad$$

where pad is a suitable padding for later use of u and cnt is an l_{cnt} -bit counter initialized by $cnt \in_R \{0, 1\}^{l_{cnt}}$ and incremented modulo $2^{l_{cnt}}$ in every round. So, the output of UG is unique for at least $2^{l_{cnt}} - 1$ rounds. So far, the privacy is not preserved, since UI is accessible through u .




Figure 1: Collision-Free Number Generator CFNG1.

Uniqueness Randomization. A first step to provide privacy is to transform u such that the resulting block looks random. Hereby, we use an injective function f_r , where r is chosen from a set R . The idea is that u is randomized by r and hence we call f_r the *uniqueness randomization function*. We suggest to either use an injective mixing-transformation for f_r (e.g. symmetric encryption) or an injective one-way function based on an intractable problem (e.g. the Discrete Logarithm Problem). In both cases, r is chosen at random. The output of f_r is obviously not guaranteed to be unique: let $u, u', u \neq u'$ and r, r' random, where $r \neq r'$. Then $f_r(u) = f_{r'}(u')$ may hold since two different injective functions f_r and $f_{r'}$ map on the same output space. On the other hand, this problem cannot happen if $r = r'$, since f_r is injective. To generate a unique block o , sufficient information about the chosen function f_r has to be attached to its output. Hence, o can be defined as $o = f_r(u) || r$. The concatenation of two blocks by writing them in a row is an unnecessary restriction. The bits of the two blocks can be concatenated in any way. This leads to the following construction (cf. Figure 1):

$$o = \pi(f_r(u), r), \quad u = \text{UG}(), \quad r = \text{Pre}()$$

where π is a (static) bit-permutation function (or bit-permuted expansion function) over the block $f_r(u) || r$

and the generation of r is done by a routine called *pre-processor*, denoted by Pre . The main-task of Pre is the correct selection of r . This can include a key generation process if f_r is an encryption function.

Theorem 3.1. *Let u be a system-wide unique number, f_r be an injective function and $r \in R$. Furthermore, let π be a static bit-permutation function or bit-permuted expansion function. Then $o = \pi(f_r(u), r)$ is system-wide unique, for all $r \in R$.*

Proof. Let $o' = \pi(f_{r'}(u'), r')$ and $u \neq u'$. The case where $r \neq r'$ obviously guarantees that the pairs $(f_r(u), r)$ and $(f_{r'}(u'), r')$ are distinct. Now consider the case where $r = r'$. Since $u \neq u'$ holds per assumption $f_r(u) \neq f_{r'}(u')$ holds due to the injectivity of f_r and the fact that $r = r'$. Thus, we have $(f_r(u), r) \neq (f_{r'}(u'), r')$ for all $r, r' \in R$. It remains to show that $o \neq o'$. This is obviously the case, because π is static and injective. \square

Privacy Protection. Given $o = \pi(f_r(u), r)$, computing $\pi^{-1}(o) = (a, b)$ is easy since π is public. Concerning f , two cases might occur concerning the obtainable information about u and hence about UI :

1. Computing $f_b^{-1}(a) = u$ is computationally hard. Then we are done, since obtaining u is hard.
2. Computing $f_b^{-1}(a) = u$ is easy. Then an extension of the basic construction is necessary, which is sketched in the following.

To keep finding u hard, we suggest using an injective one-way function g to hide $\pi(f_r(u), r)$. We call g the *privacy protection function*. The new output o is defined as follows (cf. Figure 2):

$$o = g(\pi(f_r(u), r)), \quad u = \text{UG}(), \quad r = \text{Pre}()$$

The extended construction still outputs unique numbers, which is shown by the following corollary.




Figure 2: Collision-Free Number Generator CFNG2.

Corollary 3.1. *Let u be a system-wide unique number, f_r be an injective function and $r \in R$. Furthermore, let π be a static bit-permutation function or bit-permuted expansion function and g an injective one-way function. Then $o = g(\pi(f_r(u), r))$ is system-wide unique, for all $r \in R$.*

Proof. Let $o' = g(\pi(f_{r'}(u'), r'))$ with $u \neq u'$. By Theorem 3.1 $\pi(f_r(u), r) \neq \pi(f_{r'}(u'), r')$. Since g is injective, $o \neq o'$ for all $r, r' \in R$. \square

3.2 Practical Approaches

In the following two examples are given, how CFNG1 and CFNG2 can be implemented for the generation of system-wide unique pseudonyms. The first one is based on the RSA-Problem and the second is based on AES and the ECDLP. For both variants, we will only discuss the principle operation, but no details (e.g. bit lengths) are given.

In order to guarantee the proper generation of the pseudonyms we assume that each individual has been provided with a system-wide unique identifier UI (which could be the original health identifier or the ICCSN () of a smartcard). Henceforth, we assume that all computations are done in a smartcard.

3.2.1 RSA-based CFNG1

In order to generate a system-wide unique pseudonym N , the smartcard generates an RSA key pair consisting of the public exponent e , the random private exponent d and the random modulus n (consisting of two appropriately chosen large prime factors). Now, the smartcard calculates the pseudonym N as follows:

$$N = \text{RSA}_{(e,n)}(u) || (e, n)$$

Setting $f = \text{RSA}$, $r = (e, n)$ and $\pi = ||$, Theorem 3.1 can be applied and thus N is system-wide unique. Moreover, the encryption process is hard to invert, given only N . Hence, obtaining u from N is infeasible for a poly-bounded algorithm which leads to a sufficient privacy-protection.




Figure 3: RSA-based Pseudonym Generation.

To keep the bit-length of N as short as possible, every user might use the same public key e . In this case, there is no need to embed e in N :

$$N = \text{RSA}_{(e,n)}(u) || n$$

Figure 3 shows the principle design of the RSA-based CFNG. Here Pre denotes a key generator, which generates the key (e, n) and padding material for the encryption process. The system parameters includes the bit-length of n and of the padding.

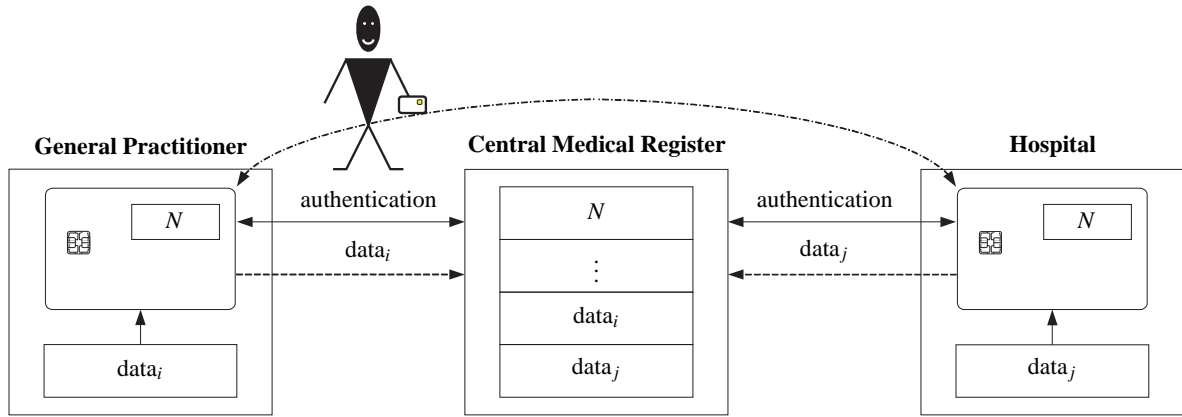


Figure 4: Centralized Medical Registers.

3.2.2 ECC-based CFNG2

In this variant, we employ a symmetric encryption algorithm to conceal the identity of the patient. As above, the randomly chosen key used for this encryption process (this time k) has to be concatenated to the output the ciphertext. However, this construction is not sufficient for privacy protection, since one can easily invert the encryption process (it is symmetric). According to the design of CFNG2, one has to choose a privacy protection function g . For efficiency, we use scalar multiplication in an elliptic curve group as a one-way function. The pseudonym N is defined as:

$$N = SM(AES_k(u)||k, P)$$

Figure 5 shows the principle design of the ECC-based generation process. Here, the system parameters include the system-wide constant point P , and the bit lengths of the primes p and q and the key k .




Figure 5: ECC-based Pseudonym Generation.

Setting $f = AES$, $r = k$ and $g = SM$, Corollary 3.1 can be applied and system-wide uniqueness is achieved. Notice that $AES_k(u)||k$ is highly random and hence inverting the ECC point multiplication process is hard due to the ECDLP. Thus, the requirement for privacy is fulfilled sufficiently for practical use.

4 UNIQUE PSEUDONYMS IN MEDICAL REGISTERS

4.1 Centralized Medical Records

In the first scenario, we employ a centralized medical database which stores a medical record for each individual. From the individual's point of view it is essential, that

- there is no link between the medical record and the real identity,
- no unauthorized person gets read or write access to the medical record, and
- additional data which is stored in the course of time is accumulated in the correct medical record.

4.1.1 Initialization

In this application we will employ smartcards which are equipped with an ECC-based CFNG2 for efficiency reasons. Hence, the digital pseudonym is of the form $N = CP(SM(AES_k(u)||k, P))$, where $u = UI||cnt||pad$ (cf. Section 3.1) and UI the patient's identifier. Notice that we apply the point compression function CP to make N as short as possible. Whenever additional medical data has to be stored in the medical record, this pseudonym will unambiguously identify the record belonging to the holder of the pseudonym.

4.1.2 Data Storage

Prior adding some data to his medical record, the patient has to prove that he owns the necessary rights. To protect the privacy of the patient, the authentication process (cf. Figure 4) must not include secret information on server side (this excludes symmetric encryption based authentication). An idea is to consider

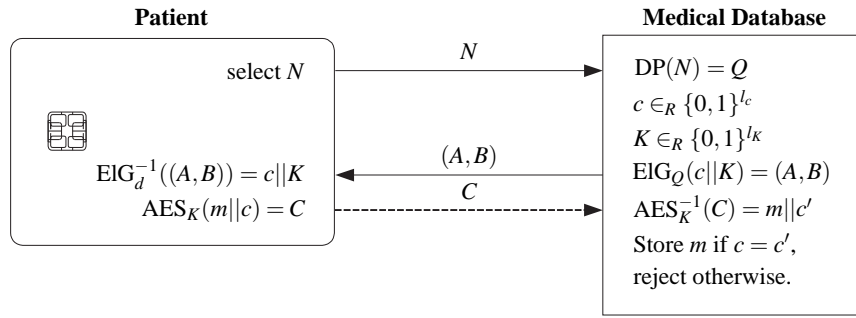


Figure 6: Authentication and Data Storage.

$Q = (x, y)$, where $N = CP(Q)$, as the public key of the ElGamal encryption scheme described in Section 2. The corresponding private key is $d = AES_k(u) || k$, since $N = CP(Q)$ and $Q = SM(AES_k(u) || k, P)$. This enables an indirect asymmetric authentication and key-exchange protocol, that provides anonymity of the patient, confidentiality of the sent medical data and freshness of the sent messages (cf. Figure 6):

1. To start the upload of medical data, the patient's smartcard contacts the server that stores the medical records, and sends his pseudonym N .
2. The server generates a random challenge c of length l_c and a session key K of length l_k , such that $l_c + l_k = l_m$ and l_k with respect to the involved symmetric encryption algorithm (here AES). Then he encrypts $c || K$ with Q using EIG_Q as described in Section 2. The resulting ciphertext-pair (A, B) is returned to the smartcard.
3. The smartcard decrypts (A, B) by using EIG_d^{-1} with the private key d and retrieves c and K . The medical data m (extended by the server's challenge) is encrypted to $C = AES_K(m || c)$ and sent to the server.
4. The server decrypts the encrypted message, compares the received challenge to the sent one and updates the medical record if they are identical. Otherwise he rejects the update request.

The authentication process only succeeds, if the patient knows d , such that $Q = SM(d, P)$. The probability of computing a correct C (that contains the same challenge as generated by the server) without knowing d is negligible.

4.1.3 Data Retrieval

In principle, the data retrieval process is similar to the data storage process. At the beginning the patient needs to authentication himself to the server. Thereby, the server generates a session key. The data retrieval

message is encrypted with the session key and provides the server with some data to identify the requested section of the medical record. The server selects this section, encrypts it with the session key and returns it to the patient's smartcard. Here, the message is decrypted and the medical data can be processed by the specific application.

4.1.4 An Extension

So far, the pseudonym is of the form $N = CP(SM(AES_k(u) || k, P))$, and provides full access to the medical record. The medical record could be split into several (unlinkable) parts, if we append SID_i (the identifier of section i) to UI and generate

$$N_i = CP(SM(AES_{k_i}(u_i) || k_i, P))$$

where u_i contains $UI || SID_i$. A pseudonym N_i will now identify and grant access to a specific section of the medical record, which enables access control and unlinkability on a finer level.

4.1.5 Security & Efficiency

The pseudonym is generated by a collision-free number generator that protects the privacy of the generating instance. In the current context this means that computing UI from a pseudonym is computationally hard, i.e. requires solving the ECDLP. The generation process is very efficient since only one symmetric encryption and one scalar multiplication is necessary. The authentication process reveals no information about UI . The user can only respond correctly (apart from some negligible cheating probability) to the challenge, if he knows the pre-image of the pseudonym. The authentication process only requires the exchanging of three messages between server and smartcard. The smartcard only performs one scalar multiplication, one run of a hash algorithm and some symmetric encryption. The server performs one run of the point decomposition function DP which equals

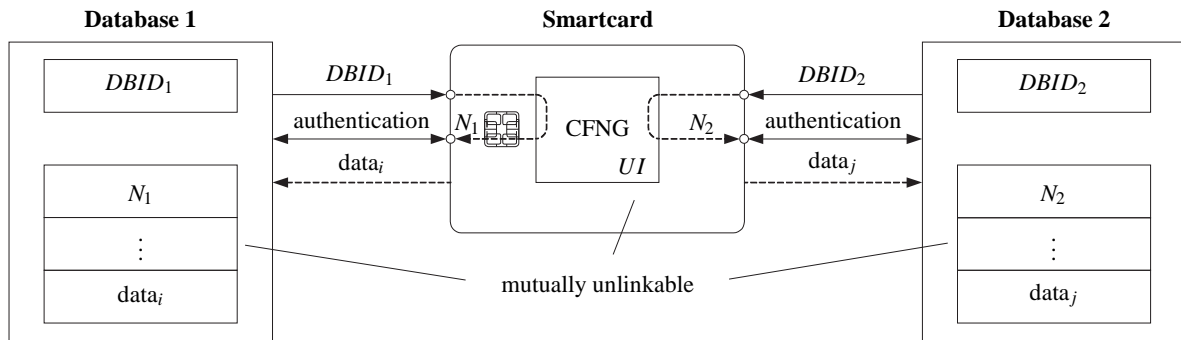


Figure 7: Anonymous Medical Databases.

to one modular exponentiation if $p \equiv 3 \pmod{4}$. Furthermore, two scalar multiplications are done and some symmetric decryption. All computations only require minimal temporary space.

4.2 Anonymous Medical Databases

Centralized medical databases are quite commonly used to provide some anonymized data for statistical investigations of certain diseases like cancer, influenza, or tuberculosis. Here, data like date of diagnosis, applied treatment, medication, chronology, and mortality are of special interest. The privacy problem with these medical databases is how to provide the anonymity of the patient. Of course, the server storing the database may be trusted, but especially in the scope of medical records, the sensitivity for privacy-endangering technologies is very high. So the best method is to remove the identifying information as soon as possible. Unfortunately, the medical data concerning the disease of a specific patient is most commonly generated over a larger period of time and hence we need some identifier to link the separate data records. Here again, we can use collision-free number generators. The authentication of the data storing party (cf. section 4.1.2) will be moved to the application level. The smartcard of the patient provides only the unique identifier of the medical record, which now depends on the patient's identifier (UI) and the identifier of the database ($DBID$).

4.2.1 Initialization

The smartcard issuer has to generate appropriate system parameters, which are stored in the smartcard. Every time a new medical database is connected (cf. Figure 7), the smartcard receives the database identifier $DBID$ and generates the corresponding pseudonym. This pseudonym will be stored in the smartcard for later usage.

4.2.2 Data Storage

During the authentication the smartcard retrieves the session key from the database server. In order to store data, all personal data of the current medical record is removed and the remaining data is encrypted by use of the session key (established during the authentication process). The resulting ciphertext is sent to the database server. The message is decrypted and the medical data is stored in the record identified by the pseudonym.

4.2.3 Data Retrieval

This application scenario does not provide data retrieval for the individual patient. Only medical institutions may retrieve data from the centralized medical register. The process to achieve this in an authentic and confidential goes beyond the scope of this paper.

4.2.4 Security & Efficiency

Security and efficiency can be considered analogously to the previous application scenario. Unlinkability is again achieved through the privacy protection that holds through the design of the used ECC-based collision-free number generator.

5 CONCLUSIONS

In this paper, we introduced a replacement for linkable unique health identifiers: locally generated but nevertheless system-wide unique digital pseudonyms. To achieve this replacement, we used so called collision-free number generators, which have been briefly discussed in the introductory part of this article. Here we presented two variants: one based on the RSA-Problem, the other based on the Elliptic Curve Discrete Logarithm Problem. Both variants

fulfill the requirements uniqueness of the generated pseudonyms, privacy in terms of hiding the originators identifier and unlinkability of individual data sets. Since the second variant is more efficient in terms of computational costs and space it has been suggested for the use within smartcards.

Beside proposing two practical generators, a protocol has been proposed through which a smartcard can efficiently authenticate anonymously to a medical database (with respect to a pseudonym). Based on this protocol medical registers can be extended and updated anonymously by the patient. Furthermore such an authentication protocol can be used to make entries in several databases. For each database, a fresh pseudonym is used so that entries of different databases are mutually computationally unlinkable. Concerning collision-free number generation further information on implementation issues and extended constructions can be found in (Schaffer and Schartner, 2007) and (Schaffer, 2007) respectively.

REFERENCES

- CNN (2000). National Health Identifier: Big Help or Big Brother? <http://www.cnn.com/HEALTH/bioethics/9807/natl.medical.id>.
- D. Hankerson, A. Menezes, S. V. (2004). *Guide to Elliptic Curve Cryptography*. Springer.
- ElGamal, T. (1985). A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms. In Blakley, G. R. and Chaum, D., editors, *Advances in Cryptology – CRYPTO’84*, volume 196 of *Lecture Notes in Computer Science*, pages 10–18. Springer.
- IEEE (2000). *IEEE 1363-2000: IEEE Standard Specifications for Public-Key Cryptography*. IEEE.
- Institute for Health Freedom (2000). What’s Happening with the “Unique Health Identifier” Plan? <http://www.forhealthfreedom.org/Publications/privacy/UniqueId.html>.
- Jonsson, J. and Kaliski, B. (2002). Public-Key Cryptography Standards (PKCS) #1: RSA Cryptography Specification. <http://www.rsa.com/rsalabs/node.asp?id=2125>.
- Medical Privacy Coalition (2007). Eliminate Unique Health Identifier. <http://www.medicalprivacycoalition.org/unique-health-identifier>.
- NIST (2001). FIPS PUB 197: Specification of the Advanced Encryption Standard (National Institute of Standards and Technology). <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>.
- Pfitzmann, A. and Köhntopp, M. (2001). Anonymity, Unobservability, and Pseudonymity – A Proposal for Terminology. In Federrath, H., editor, *Proceedings of Workshop on Design Issues in Anonymity and Unobservability*, volume 2009 of *Lecture Notes in Computer Science*, pages 1–9. Springer.
- Rivest, R. L., Shamir, A., and Adleman, L. M. (1978). A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. *Commun. ACM*, 21(2):120–126.
- Schaffer, M. (2007). *Collision-Free Number Generation: Efficiency Constructions, Privacy Issues, and Cryptographic Aspects*. PhD-Thesis, Klagenfurt University.
- Schaffer, M. and Schartner, P. (2007). Implementing Collision-Free Number Generators on JavaCards. Technical Report TR-syssec-07-03, University of Klagenfurt.
- Schaffer, M., Schartner, P., and Rass, S. (2007). Universally Unique Identifiers: How to ensure Uniqueness while Preserving the Issuer’s Privacy. In Alissi, S. and Arabnia, H. R., editors, *Proceedings of the 2007 International Conference on Security & Management – SAM’07*, pages 198–204. CSREA Press.
- Schartner, P. and Schaffer, M. (2005). Unique User-Generated Digital Pseudonyms. In Gorodetsky, V., Kotenko, I. V., and Skormin, V. A., editors, *Proceedings of Mathematical Methods, Models, and Architectures for Computer Network Security – MMM-ACNS 2005*, volume 3685 of *Lecture Notes in Computer Science*, pages 194–205. Springer.

ON PSEUDONYMOUS HEALTH REGISTERS

While they Work as Intended, they are Still Controversial in Norway

Herbjørn Andresen

Norwegian Research Centre for Computers and Law, University of Oslo, Norway
herbjorn.andresen@jus.uio.no

Keywords: Confidentiality, Privacy Regulation, Data Security, Pseudonymity, Public Health Policy.

Abstract: Patient health data has a valuable potential for secondary use, such as decision support on a national level, reimbursement settlements, and research on public health or on the effects of various treatment methods. Unfortunately, extensive secondary use of data is very likely to have disproportionate negative impact on the patients' privacy. Traditionally, privacy regulations require a balancing process; the use of data should be minimized and kept within a level where proportionate privacy is maintained. An alternative strategy is to use technological remedies to enhance privacy protection. Norwegian health data processing regulation prescribes four different ways of organising health registers (anonymous, de-identified, pseudonymous or fully identified data subjects). Pseudonymity is the most innovative of these methods, and it has been available as a legitimate means to achieve extensive secondary use of accurate and detailed data since 2001. Up to now, two different national health registers have been organised this way. The evidence from these experiences should be encouraging: Pseudonymity works as intended. Yet, there is still discernible reluctance against extending the pseudonymity principle to encompass other national health registers as well.

1 INTRODUCTION

Any patient must accept some processing of his personal data, within the confines of a medical treatment. Some data are collected from the patient himself, and some data may be generated along the course of the treatment. The confidentiality will inevitably be at risk; the risks need to be monitored and handled. Privacy regulations and data security measures, along with professional ethics, safeguard against unwarranted processing and filing practises, and against deliberate misuse.

Privacy is also at stake when the patient's health data is used beyond the appointed medical treatment. Such use could be named *secondary purposes* for processing data. The subject matter of this paper is a particular variety of secondary purposes, namely a group of national health data registers. A register is a service that comprises a database, an operating organisation, and a legal authority defining responsibilities, duties to report to the register, restrictions on use and so on. In colloquial language the term register is normally taken to mean the database itself. The organisation and the legal authority are implied.

The registers have two important features, from a privacy point of view. Firstly, they are centralised systems, containing aggregated data. Personal health data are collected from different hospitals or other treatment entities. As Norway is a small country, centralised registers usually cover the entire nation. The procedure of collecting the data could either be by electronic exchange or by printed reports which are re-typed into the centralised system.

Secondly, the data is collected and processed for secondary purposes, somewhat remote from the patient's immediate needs and interests. Roughly stated, these secondary purposes are governmental administration and medical research. Governmental administration includes both macro-level decision support and reimbursement control procedures. The demand for data is, at least in principle, limited and foreseeable. Medical research will also in most cases demand a stable amount of foreseeable data, yet in some cases it could be beneficial to use excess data or to perform inventive couplings involving different data sources. The future value of ingenious data mining is by definition unknown.

Regulations, security measures and ethics are of course at least as required for the registers as they

are for data in the immediate care systems. In addition, the registers are vulnerable to expansions of their stated and legitimate purposes. Proponents of strict privacy regulations warn against “the slippery slope”. It can be increasingly difficult on each particular occasion to turn down a proposition for extended use of a register. Such propositions often serve legitimate purposes, which is to achieve new and even benign goals more easily. Consequently, the patients’ privacy is in danger of being scooped out in the long run.

1.1 The Adage of Norway’s Favourable Conditions for Health Registers

Norway introduced a national identity number quite early. Starting in 1964, Statistics Norway assigned a unique 11-digit identification number to every individual. The primary purpose of the national identity number was to produce accurate statistics. Large public agencies, such as the Tax Administration and the National Insurance Administration, soon adopted the new unique identification number. No one imagined the vast future use of this new identity number. There was no explicit legal support for it, and hence there were no expressed restrictions on its use either (Selmer, 1992).

Due to the lack of restrictions on the use of population register data in the early years, the identity number is now the key to personal data in thousands of public as well as private IT-systems throughout the country, including primary health care systems and hospital systems. Most Norwegians will have to type or pronounce his unique personal identity number to some electronic apparatus (or to its human gatekeeper) several times a week. It is “open sesame” to enter both caves of treasures and caves of dung.

Due to the widely used national identity number, Norway may have favourable conditions for national health registers. Because the identity numbers are used in systems that are vital to virtually everyone, the data quality of the primary systems reporting to the registers almost takes care of itself. The registers could technically be suitable for collecting perfectly linkable data from the cradle to the grave.

However, the early years of vastly expanding the use of national identity numbers was succeeded by almost three decades of efforts to impart restrictions on their use. Lawful use of the national identity number is subject to a “norm of necessity”. Roughly, it goes like this: When an exact identification is *necessary*, then the controller *ought to* use a unique identifier, while using the same unique identifier

would be unlawful when unique identification is *not* necessary (this is a crude simplification, on my own behalf, of the Personal Data Act section 12). The result from these efforts to limit the lawful use of unique identifiers has not been any actual decrease in the use of the national identity numbers. Yet, many researchers perceive some uncertainty on whether they can use national identity numbers lawfully in their projects. A health register may not impart non-anonymous patient health data to research projects unless the recipient can provide sufficient justification for it.

No matter how favourable the conditions for health registers may be; legal restrictions prevent them from being used to their full potential. This situation induces two parallel debates: The first one is about the balance between privacy and legitimate uses of a health register. The second debate is about the possibility to circumvent patient identification without sacrificing the benefits of a health register.

1.2 The Origins of Digital Pseudonyms

A pseudonym is, literally, a “false name”. For ages, pseudonyms have been used by authors and artists, or even in the rare event of modest researchers, to disguise their identity. The notion of a digital pseudonym first appeared in a paper by David Chaum. He invented digital pseudonyms as a means to conceal an individual’s real identity in electronic transactions (Chaum, 1981). The intended field of application in Chaum’s paper was banking and electronic commerce. A pseudonym concealed the identity of the person who actually paid the goods.

In a few consecutive papers, he developed both the methods and the rationale further. The public key distribution system provided a secure cryptographic pseudonym. For the holder of a pseudonym to be able to communicate or inspect his own personal data, a trusted third party could manage the pseudonyms. The rationale was to introduce a new paradigm for data protection; using technological means to put the individual in control of his own data (Chaum, 1984). Organisations would not be able to share data about the individual without the data subject “acting out” his consent, so to speak. No one could collect the complete history of your transactions, debts or savings. The holder of the pseudonym would also hold the key to reverse it.

As for the proposed new paradigm of privacy in banking and electronic commerce, it seems to have lost completely to the old paradigm of widespread use of fully identified data subjects. Meanwhile, the

fields of health administration and medical research have revived the idea of digital pseudonyms.

1.3 The Pseudonymisation Process

There are different ways to carry out the process of generating a digital pseudonym to conceal the data subject's real identity.

The simplest form of pseudonyms, used for decades in research projects based on samples, is to assign a sequence number to each respondent. To enhance the respondents' trust, the researcher could hire a third party to perform the assigning process. This method works well for one-time surveys. For a panel study over time, managing sequence numbers becomes increasingly more difficult. Coupling data with relevant data from other sources would require an overt process of reversing the sequence numbers. The pseudonyms would be illusory. To exchange a "real identity" with an unrelated sequence number is only trustworthy when the researcher grants the respondent permanent anonymity, without adding to the data later. It is not a viable method for a long-term and multi-purpose health register.

A digital pseudonym in a health register involves advanced cryptography. The input to the algorithm that generates the pseudonym will have to be a stable identifying number, which does not change over time for the same patient. In Norway, the national identity number provides a convenient unique input. The health register will not need to store the national identity number, the algorithm secures that the same pseudonym is assigned to the same patient when more data is added to the register.

With a reliable and stable identification, there are, conceptually, two different ways to generate a pseudonym. One way is to use an asymmetric hash function. The encryption algorithm then generates a digest that is unique to the input, but there is no way to reverse from the encrypted digest back to the input identifier. Because the same input identifier always transforms to the same digest, it is possible to add data about the same patient in the same health register. It is, however, not possible to generate data couplings between individual-level data from two different health registers. This method provides a very high degree of confidentiality, but is on the other hand inflexible. Two health registers cannot be merged, and it would not be possible to address any registered patient, for instance if a new treatment method vital to his particular disease is developed.

The alternative way to generate a pseudonym resembles the "public key" encryption technology, and is basically the same as Chaum invented (see

section 1.2 above). The input to the algorithm is the same stable and reliable patient identity number. An encryption algorithm, using the "public key" of a key pair, generates the pseudonym. The same input, and the same public key, will make it possible to add data about the same patient to the same health register. In addition, a decryption algorithm can reverse the pseudonym back to the "real identity", by using the "private key" of the same key pair that was used for encryption. A trusted third party, which is an independent pseudonym manager, carry out the encryption, and if requested, the decryption. The health register will never see the real identity of the patient. The trusted third party, who is able to decrypt the pseudonym, does not have access to any sensitive information about the patients. This process provides more flexibility, at the cost of more fragile pseudonyms. The confidentiality of the patient is to a higher degree based on trust. Violating the pseudonyms will be somewhat easier from a technical point of view.

The latter method, a trusted third party handling reversible pseudonyms, has been the method of choice for pseudonymous health registers in Norway so far. Non-reversible pseudonyms would also conform to the legislation on health registers, yet it is not very likely that a register owner voluntarily would choose this less flexible process.

2 LEGISLATIVE SUPPORT FOR PSEUDONYMS

Recent technological innovations often seem to be far ahead of developments in legislation. Society's toolbox for protecting values and for distributing rights and obligations usually adapts slowly, to fit technological changes that have already taken place.

The introduction of pseudonyms in Norwegian health registers differs from this typical path of history. The first Norwegian national register based on pseudonyms was established in 2004. By that time participants in various legislation processes had already advocated this method for more than a decade. Technologists and professional users of the registers remained sceptic. Pseudonymous health registers have not at any rate been "technology-driven" in Norway, it would be far more correct to call it a "legislation-driven" development.

Norway has had registers for specific diseases, such as The Cancer Register, for decades. They started out as paper files, and were later converted to computer databases. The specific health registers

had proved to be useful over time, and the health authorities started to nourish a desire to establish a General National Health Register, not to be limited to any particular diagnosis.

2.1 An Early, Avant-Garde Proposal

Though the advantages of a General National Health Register were convincing, The Parliament was also much concerned about the impact on the patients' privacy. In 1989, they urged The Government to appoint a committee with a mandate to examine ways and methods to establish such register "without threshing individuals' privacy" (Boe, 1994).

The appointed committee issued a report in 1993 (ONR, 1993). An "Official Norwegian Report" is in most cases the product of an appointed drafting committee, at an early stage of the legislation process. After an official hearing among relevant stakeholders, both The Government and eventually The Parliament may make changes to the original draft, or even turn down the entire proposition all together.

The drafting committee proposed, in their report, a new act to provide legal authority to the desired General National Health Register. The proposed act was very much ahead of its time. It contained regulations on cryptographic pseudonyms generated and managed by trusted third parties, along with a profound set of rules on ensuring legitimate use of the register, data quality, and the patients' right to access and so on.

However, neither the health authorities nor the research community was in favour of this avant-garde way to organise their much-desired new health register. As the main stakeholders did not support the proposition, The Government put it on hold, and it remained so for about eight years.

Instead of either a fully identified register (which was what the health authorities wanted) or a pseudonymous health register (the proposition they turned down), the health authorities established the Norwegian Patient Register (see section 3.3 below) in 1997. The Norwegian Patient register was originally established as a de-identified register (see section 2.3 below). This was acceptable under the Personal Data Filing Systems Act of that time, and it did not require The Parliament to pass any new legislation.

2.2 Specific Privacy Regulations for Health Data and Health Registers

The Parliament passed a new general Personal Data Act on April 14, 2000. The primary motivation for replacing the old act of 1978 was to comply with the European Union Directive 95/46/EC, on protection of personal data.

The Personal Data Act regulates all processing of personal data, for any legitimate purpose. Therefore, the rules are quite flexible, leaving most assessments and decisions to the discretion of the controller. For the processing of health data, The Parliament did not consider the general act sufficient. On May 18, 2001, they passed the Personal Health Data Filing Systems Act, containing rules that are somewhat more specific. The Personal Health Data Filing Systems Act too complies with the European Union Directive, and it has many important features in common with the general Personal Data Act. For instance, the information security requirements are essentially the same.

The primary guiding rule for processing health data is a requirement to obtain the patients' consent. However, the act also recognises a need in some situations to process data without consent. A typical exception to requiring consent would be the kind of health registers where complete coverage is vital to fulfil the purpose of the register.

2.3 Four Different Levels of Patient Identification in a Health Register

The key to the regulation of health registers is section 8 of the Personal Health Data Filing Systems Act. The initial position is simply that central health registers are forbidden, unless authorised by this act or by another statute.

The remainder of section 8 spans the possibilities and preconditions for establishing health registers, providing they have an adequate legal authority. The purpose of a register shall be "to perform functions pursuant to," specified health services (the relevant acts are listed in section 8). Those functions include "the general management and planning of services, quality improvement, research and statistics". In addition, the Government shall prescribe subordinate legislation for each health register, defining specific rules, responsibilities and organisation.

An interesting feature of section 8 is the way it categorises health registers into four distinct levels of patient identification. Every health register has to conform to one of these four levels. The *choice* of a

level of identification encapsulates the privacy balancing process for each register.

Table 1: Outline of levels of patient identification (adapted from L'Abée-Lund 2006, page 28).

Personal data (being subject to privacy regulation)		Not personal data	
Data refers to unambiguous individuals		Data may refer to ambiguous data subjects	
Fully identified	Pseudonyms	De-identified	Anonymous

The bottom row of the table above shows the four different levels of patient identification. Their order, from the left to the right, reflects an order from more to less strain on the patients' privacy.

The middle row of the table shows the main division of whether the data refer to unique patients or not. Fully identified patients and pseudonymous patients both have the same granularity. They will provide the same level of statistical data accuracy.

The top row of the table merely shows that only three out of the four levels of patient identification are strictly within the definition of personal data.

Generally, fully identified health registers shall only process data about patients who consent. The only exceptions are a moderate number of health registers particularly named in section 8. By the end of June 2007 there are exactly nine fully identified health registers not requiring the patients to consent (the number of such registers was six by the time the act was originally passed, in 2001). The Parliament has to pass a formal change to section 8, specifically naming the new register, before anyone can establish a new fully identified central health register with a complete coverage (i.e. not requiring consent). That is the beauty of this construct in the Personal Health Data Filing Systems Act; it ensures an overt and highly democratic legislation process to be carried out before establishing a new register.

By using pseudonyms instead of fully identified patients, the health authorities can establish a new health register by issuing subordinate legislation. This means a Parliament decision is not necessary to establish the register. It also means the register can omit the patients' consent, if it needs complete coverage of the data. The option of pseudonymous health registers thus entered Norwegian legislation in 2001, eight years after it was first proposed.

A de-identified and a pseudonymous register have the same legal status according to section 8. The health authorities may establish a de-identified register by issuing subordinate legislation. A de-

identified register means that any clear and manifest identifying information is removed. The advantage of a de-identified over a pseudonymous register is that the de-identified register is technically easier and less expensive to operate. The paramount disadvantage of a de-identified register is that the data is not on a strictly individual level. If a hospital carries out the same surgical procedure, say four times, a de-identified register cannot tell whether it involved four different individual patients or if the same patient was involved four times.

Pseudonymous and de-identified registers share the same risk of unlawful re-identification through computational analysis of the stored data elements (Malin, 2005). The uniqueness of each registered individual becomes more transparent as the number of detailed variables increase. Coping with the risk of re-identification first requires the register owner to keep his sobriety on what data is stored and processed. Second, there is still an indispensable need for rigid access control and other conventional information security measures with pseudonymous and de-identified registers.

In an anonymous register, all information that can possibly identify individual patients is removed. In addition to removing the manifest identifiers, the register also removes, or reclassifies into categories that are more general, any data suitable for re-identification by analysis. An anonymous register takes the granularity of the data into account. Making the data anonymous often means to take deliberate action to sacrifice their accuracy.

Anonymous data may be published, and they will not require extensive data protection. The downside to anonymous data, which is why they are unapt for health registers in most cases, is that it is virtually impossible to add meaningfully to the data.

2.4 The Professionals' Responsibilities, and a Democratic Safety Valve

To summarise, the Norwegian legislation allows four different methods for storing and processing personal health registers. A method granting more privacy is less effective for achieving administration and research goals. This inverse ratio is at the heart of any privacy regulation. All the four levels of identification are in use in some existing health register, and they have all proved to work as intended. Apart from the likes and the dislikes of different stakeholders: The four different nominal levels of identification themselves provide relatively objective aids for an informed policy debate. We know what the options are, and how they work. We

also know how each of these options influence the privacy of the patients, versus the accuracy of governmental decision support data and the possibilities for providing valuable research data.

The health authorities remain chiefly responsible for all aspects of the health registers. However, the very strict preconditions for establishing a fully identified register (unless requiring patients to consent) constitute a striking democratic safety valve. Any such register require The Parliament to pass a formal change to section 8 of the Health Data Filing Systems Act. Professional agenda owners and stakeholders need not, and may not, decide alone on such privacy invasive registers. Though the process may be cumbersome and time-consuming, it also secures a highly democratic participation in the balancing between privacy and the well-grounded benefits of a health register.

3 A CURRENT STATUS, AND RECENT DEVELOPMENTS

Up to the end of June 2007, the Ministry of Health has seriously considered and deliberated on the option to make a health register pseudonymous on four different occasions. On two out of these four occasions, they actually decided to establish the proposed register with pseudonyms as its level of patient identification. For the other two registers, one of them was “promoted” to be a fully identified register, while the other one was “demoted” to be a de-identified register.

3.1 The Norwegian Prescription Database

The first pseudonymous national health register in Norway is called *The Norwegian Prescription Database* (“Reseptregisteret”, in Norwegian). In October 2003 The Ministry of Health issued the subordinate legislation providing legal authority for the register, as required by section 8 of the Health Data Filing Systems Act. The register was actually established in the beginning of 2004.

Before the Prescription Database was established, the medicine statistics were based on sales figures reported from wholesale dealers. Unquestionably, the data was insufficient both for straightforward knowledge about use of medicine, and for research on effects thereof. Various stakeholders demanded statistics based on prescriptions and actual dispatch from pharmacies to individual patients. The intended

purposes were neither to control any patient’s catch at the pharmacy nor to supervise how named doctors carried out the business of prescribing. Pseudonyms both ensure the demanded capacities of the register, and safeguard against undesirable infringements of privacy.

All pharmacies report the prescription data electronically every month. A central data collecting point transfers the data to a trusted third party. The trusted pseudonym manager is in this case Statistics Norway. They transfer the pseudonymous data to the register owner, which is the Norwegian Institute of Public Health. Both the patient’s identity (his national identity number) and the doctor’s identity (his authorised licence identifier) are replaced with pseudonyms. The pharmacies are fully identified, on an enterprise level; their licence identifier is not pseudonymised (Strøm, 2004).

The Prescription Database was in many ways a “quiet reform”. The changes have been virtually invisible to the patients. They still pick up their prescriptions and carry them to the pharmacies the same way they did before. They are not asked for consent. The existence of the Prescription Database is not a secret in any way, but neither is it of much concern to the patients. They only know about the pseudonymous data if they take a particular interest in detailed level health politics.

3.2 National Statistics Linked to Individual Needs for Care (IPLOS)

The second pseudonym-based health register is named the *National statistics linked to individual needs for care* (its Norwegian acronym is “IPLOS”, which is derived from “Individbasert pleie- og omsorgsstatistikk”). The Ministry of Health issued the subordinate legislation providing legal authority for the register in February 2006. The first mandatory reporting term to the register was February 2007, collecting data from health care services throughout all Norwegian municipalities. The register owner in this case is the Directorate for Health and Social Affairs. The Tax Administration is the trusted pseudonym manager, which illustrates the point that the main feature of a pseudonym manager is its institutional independence.

Contrary to the Prescription Database, the IPLOS has not been a “quiet reform”. The information about individual needs for care was not readily available from any existing process. Even though the patients can trust the confidentiality of the central register, they had to answer a new set of questions. Someone would type their answers into a local

database before they were sent electronically to the pseudonymous register. The crucial question was not anymore whether the pseudonym provided sufficient privacy. Many patients felt offended by some of the most invasive questions in the form. The forms were changed as a result from complaints about some of the questions, such as whether a handicapped patient needs help after going to the toilet or needs help with handling the menstrual period.

Pseudonyms only remedy privacy issues that become present after the patients have left off their participation. An important lesson is that health registers mainly deal with data that the patients hardly are aware of. In many cases, the limits to a health register are with the processes of eliciting and collecting data, and not with the confidentiality of the register itself.

3.3 The Norwegian Patient Register

The Norwegian Patient Register is a hospital and outpatient clinic discharge register. Data on each patient is collected from every hospital in Norway. The acronym NPR is used both in English and in Norwegian when this register is referred to.

The history of the NPR is complicated, and truly interesting from a privacy point of view. NPR is the actual instantiation of the “General National Health Register” which initiated the committee back in 1989, who proposed a pseudonym-based solution register in their 1993 report.

After the proposed pseudonymous register was put on hold, it was revived in 1997 as a de-identified register. The NPR was established in March 1997. It receives reports on operative procedures extracted from the patient administrative systems at all hospitals. Age, sex, place of residence, hospital and department, diagnosis, surgical procedure, and dates of admission and discharge are included in the register (Bakken et al, 2004). The name and national ID number of the patients are not included.

The NPR has proved to be a valuable register, providing much demanded data for both research and administration purposes. Yet, as a de-identified register it does have obvious shortcomings. The data do not refer to strictly unambiguous individuals.

Over the last few years, the health authorities have made efforts to “promote” the NPR into a fully identified register. Proponents of privacy argued that promoting it into a pseudonymous register would be sufficient for all purposes of the register. The health authorities and the research community argued that a pseudonymous register might not provide adequate data quality. After a heated debate, The Parliament

finally passed the necessary change to section 8 of the Personal Health Data Filing Systems Act on February 1, 2007, and included the NPR to be a fully identified health register. The subordinate legislation regulating “the new” NPR is expected anytime soon.

3.4 The Abortion Register

The latest example so far, of the Ministry of Health having seriously considered pseudonyms, is The Abortion Register. They made a proposition, intent to establish this as a pseudonymous register. The proposition went to a formal hearing; the closing date for the hearing was January 13, 2006. A large number of the bodies entitled to comment on the hearing were sceptic to a register containing as sensitive information as abortions.

The proposal was met with reluctance from different sides. Many answers to the hearing pointed out the particular strain on some of the women who decide to go through an abortion. Induced abortions are legal in Norway, yet there is a risk of social or religious condemnation from parts of the society, making the burden heavier. The Data Inspectorate, for instance, argued that the knowledge of an abortion register might influence on actual decisions on whether to have an abortion or not. Thus the register could affect, and not merely reflect, the health care activities. Recently, on June 21, 2007, The Government decided to make The Abortion Register de-identified, and not pseudonymous.

The policy debate on The Abortion Register shows an interesting limit to people’s trust in a pseudonymous register. The confidentiality is based on trust in society as we know it today. The possibility of reversing a pseudonym could be exploited sometime in the future, when privacy values may be if worse off.

4 CONCLUSIONS

4.1 Pseudonymous Identities Work

Pseudonyms are a legal and a viable means for protecting personal data in Norwegian Registers. It has been one out of four lawful levels of patient identification since 2001. There are only two health registers based on pseudonyms so far. The numbers of both fully identified registers and de-identified registers are much higher.

A case study on the Prescription Database shows that the vital functions of a pseudonymous register work as intended (L’Abée-Lund, 2006). Neither the

trusted pseudonym manager nor the register owner – nor anybody else for that matter – gets to see both the “real” identification numbers and unencrypted health data relating to the individuals at the same time. There were initial problems with the data quality for some time, but more recently, the rate of errors have been approximately the same as they are for fully identified registers.

4.2 Pseudonyms are Still Controversial

This brief report on pseudonymous health registers in Norway reveals some broad categories of arguments for and against pseudonyms.

The pro arguments are primarily a lower risk of disclosing information about patients, to people who do not need it, and for purposes where identification is unnecessary. A pseudonym increases privacy, while maintaining the statistical accuracy of the data.

The contra arguments are increased expenses due to the third party process, a risk of re-identification by analysis of the non-identifying data, and a danger of unforeseeable decrease in privacy if privacy policies change in the future. Finally, the argument most often raised against pseudonyms, is the data quality issue. The register owners will have reduced opportunities for discovering and fixing errors on their own. However, the third party may assist in structured “data laundering”-procedures.

Looking at the pros and cons of pseudonymous health registers, the amplitude of the controversies seems somewhat exaggerated. The pseudonymous registers are plainly an in-between solution. Moving either one step to the left or one step to the right, referring to the outline of levels of identifications in table 1, is merely a change in the balance of the arguments for and against pseudonyms. Choosing another level of patient identification will neither release all the advantages nor solve all the problems that may occur to a pseudonymous register. A proposal to make a particular health register pseudonymous can expect attacks from both sides; some will say identifying unambiguous individuals pseudonymously infringes privacy anyway, others will say pseudonyms place too heavy restrictions on the register.

In my opinion, though, the fact that the health authorities have only established two pseudonymous registers since that option was legislated for in 2001, while The Parliament has accepted three new fully identified registers during the same period, sadly indicates that proponents of pseudonymous registers in Norway are perhaps fighting a losing battle.

To the credit of the Personal Health Data Filing Systems Act, the construct of choosing between only four lawful levels of identification is in my opinion both clever and successful. It ensures a broad and overt democratic process, calling the attention of all stakeholders to voice their opinion.

REFERENCES

- Bakken I J, Nyland K, Halsteinli V, Kvam U H, Skjeldestad F E. The Norwegian Patient Registry. In *Norwegian Journal of Epidemiology* 2004; vol. 14 no. 1, 65-69.
- Boe, E., 1994. Pseudo-identities in Health Registers: Information technology as a vehicle for privacy protection. In *The International Privacy Bulletin* 1994; Vol. 2 no. 3, 8-13.
- Chaum, D. L., 1981. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM* 24, 2 (Feb. 1981), 84-90.
- Chaum, D., 1984. A New Paradigm for Individuals in the Information Age; In *1984 IEEE Symposium on Security and Privacy*, IEEE Computer Society Press, Washington 1984, 99-103.
- L'Abée-Lund, Å., 2006. Pseudonymisering av personopplysninger i sentrale helseregistre [Title translated: Pseudonymising personal data in central health registers]. Master thesis, *Section for Information Technology and Administrative Systems, Faculty of Law, University of Oslo*
- Malin, B., 2005. Betrayed by My Shadow: Learning Data Identity via Trail Matching. *Journal of Privacy Technology*. 2005; 20050609001. Pittsburg, PA, USA.
- ONR 1993. Pseudonyme helseregistre [Title translated: Pseudonymous Health Registers]. *Official Norwegian Report*. NOU 1993:22
- Selmer, K., 1992. Hvem er du: Om systemer for registrering og identifikasjon av personer. [Title translated: Who are you: On systems for registering and for identifying individuals] In *Lov og Rett*, 1992, 311-334
- Strøm, H., 2004. Reseptbasert legemiddelregister [Title translated: A prescription-based medicine register]. *Norsk Farmaceutisk Tidsskrift*. Nr. 1, 2004, 7-9

AUTHENTICITY AND INTEGRITY OF PORTABLE ELECTRONIC HEALTH RECORDS

Chung-Yueh Lien

*Institute of Biomedical Engineering, National Yang-Ming University, No. 155, Sec. 2, Linong st.
Beitou District, Taipei 112, Taiwan
lian@bme.ym.edu.tw*

Chia-Hung Hsiao

Department of Medical Informatics, Tzu Chi University, Hualien City, Taiwan

Lu-Chou Huang, Tsair Kao

*Institute of Biomedical Engineering, National Yang-Ming University, Taipei, Taiwan
tskao@ym.edu.tw*

Keywords: Security, Portable electronic health record, Digital signature, Integrity, Authenticity.

Abstract: In this paper, we proposed a method to secure an electronic health record stored on a portable data storage media (CDs/DVDs, diskettes, flash drives). We applied cryptography to realize the authenticity and integrity of the portable health record. A manifest signature mechanism was used to reduce the computation time of the signing and verifying processes. A DICOM DIR consists of 166 DICOM MR images was tested as an example of a portable medical record. The performance of this method is faster than the regular digital signature mechanism.

1 INTRODUCTION

Electronic health records (EHRs) may be generated by hospitals, examination laboratories, insurance companies or patients themselves (Bates et al. 2003; Wang et al. 2004; Pharow & Bolbel 2005; Tang et al. 2006). To realize clinical data exchange between healthcare providers, a trusted conduit is needed for the EHR systems and users. The integrity and authenticity of EHRs can be validated using a digital signature mechanism (Ruotsalainen & Manning 2007; Schütze et al. 2006). A digest of the digital document is calculated from an irreversible one-way hash function. The hash check of digital data is commonly used on the Internet to prevent unauthorized modification. The digital signature can be implemented by a combination of the hash algorithm and public-key cryptography such as the RSA algorithm. When the RSA algorithm is used to calculate a digital signature, the signer encrypts the digest of the digital document with his/her own private key. The recipient, with access to the

signer's public key then verifies the digital signature.

The implementation of EHRs has to conform to security regulations, laws and standards, such as Digital Imaging and Communication in Medicine (DICOM), Health Level 7 (HL7), World Wide Web Consortium (W3C), Health Insurance Portability Accountability Act (HIPAA) and ISO/TS 17090. According to the healthcare standards, a legal signed EHR must contain one digest, one digital signature and one timestamp signature. Under public-key infrastructure (PKI), the use of the RSA algorithm makes it possible to work with the certificate and trusted third party (TTP) to process inter-institutional applications such as the verification of an EHR and referral information (Lekkas & Gritzalis 2007). An EHR may contain hundreds of digital files, and then require the same number of digital signatures. However, it is impractical to implement these lengthy signing and verifying processes in the real world due to the high computation time.

In this paper, a manifest signature mechanism is proposed to reduce the computation time of the signing and verifying processes used when dealing

with a portable EHR. A DICOM DIR consists of 166 DICOM MR images was tested as an example of a portable EHR.

2 METHODS

We used a smart card system that supports the Microsoft Cryptography Service Provider (CSP) as the digital signature module. The use of the health professional card (HPC) with a smart card-based certificate is a good example and can be found in the healthcare environments of Taiwan (Yang et al. 2006), Germany (Schütze et al. 2006; Schurig, Heuser & Wedekind 2001), Belgium (France, Bangels & De Clercq 2007), etc. A health organization certificate card (HOC) holds the digital official seal of every health organization and can be used for EHRs exchange among organizations. The Health Certificate Authority Timestamping Authority (HCA-TSA) provides the timestamp service as the TTP.

2.1 The Characteristics of a Portable EHR

The flowchart of the security protection process is shown in Fig. 1. The clinical documents, integrated from various systems in a hospital stored in the portable storage media. The signed list has to be managed through the metadata that define these data in a portable EHR. The metadata has to identify the information, structures and formats to meet the needs of multimedia data exchange. The signed list is used to sign the clinical documents from the signing hospital; the documents described in the signed list will be signed using the hospital certificate with a digital time signature. According to the signed list, the digest values of the clinical documents are calculated and packaged as the digest of these documents. The hospital and TTP generate the digital signature and digital time signature to verify the authenticity of the clinical documents. The hospital certificate is used as identification on the exchanged clinical documents and to authenticate the source site, and the site receiving the data can then verify whether or not the exchanged clinical documents are valid. Table 1 summarizes the use of the characteristics of a portable EHR.

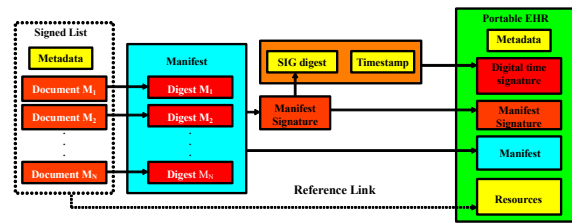


Figure 1: The characteristics of a portable EHR architecture.

Table 1: The use of the characteristics of a portable EHR.

Characteristics	Use
Metadata	Definition of the portable EHR
Signed List	Information of signed EHR
Digital Time Signature	Authenticity
Manifest	Integrity
Manifest Signature	Authenticity
Resource	The physical EHR structure

2.2 The Manifest Signature Mechanism

The manifest signature process is shown in Fig. 2. In the signed list, there are N documents M_1, M_2, \dots, M_N need to be signed. For each file $Ref(M_n) = ID(M_n) \wedge H(M_n)$ is calculated. The symbol ' \wedge ' means cascade. $Ref(M_n)$ represents the metadata of M_n , and some information related to M_n can be defined in $ID(M_n)$, such as data type, creation time, purpose, etc., for exchange among various EHR systems. $H(M_n)$ is the digest of M_n , where H is the hash function, such as MD5, SHA1, etc. We can reconstruct $Ref(M_n)$ as the manifest of N documents' digest MD_D defined as:

$$MD_D = Ref(M_1) \wedge Ref(M_2) \wedge \dots \wedge Ref(M_N)$$

Based on the RSA algorithm, using the signer's private key P_r to encrypt $H(MD_D)$, the digest value MD_D , as the digital signature of MD_D , we then define the digital signature process as follows:

$$SIG(MD_D) = RSA_{Enc}(P_r, H(MD_D)),$$

Where RSA_{Enc} is the RSA encryption function. $SIG(MD_D)$ is the manifest signature. We send the digest $H(SIG(MD_D))$ to a TSA to obtain a qualified digital time signature TSA_{SIG} , defined as:

$$TSA_{SIG} = RSA_{Enc}((P_{rTSA}, H(SIG(MD_D)) || T_{SS}))$$

The symbol ' $||$ ' represents the concatenation, P_{rTSA} is the TSA's private key and T_{SS} is the Greenwich Mean Time (GMT) timestamp.

This is an efficient mechanism to ensure the integrity and authenticity of a portable EHR. From the aspect of data integrity, $Ref(M_n)$ ensures the

integrity of a document M_n . The manifest signature $SIG(MD_D)$ and the digital time signature TSA_{SIG} provide verification of the authenticity for the documents in the signed list and the T_{SS} confirms the synchronized time.




Figure 2: Manifest of the documents' digest process.

2.3 Performance Analysis of the Manifest Signature Mechanism

Assuming the time required for hashing, signing and digital time signature retrieval are T_n , T_r and S_n , respectively, where $1 \leq n \leq N$. The timestamp retrieval time depends on the network status; here S_n is just for reference. If we create digital signature for each document one by one, the total calculation time is:

$$T_{total} = \sum_{n=1}^N T_n + N \times T_r + \sum_{n=1}^N S_n$$

However, using the method proposed in this study, the total time will be reduced to:

$$T_{total} = \sum_{n=1}^N T_n + S_l + T_r + T_M$$

T_M is the computation time of $H(MD_D)$. The number of calculations needed for verification process is still the same. Table 2 shows the number of calculations in hashing, signing and digital time signature retrieval for different methods.

Table 2: Comparison of the number of calculations between the proposed method and the one-by-one method.

Operations	Calculation time	
	Proposed method (sec)	One-by-one method (sec)
Hashing	N+1	N
Signing	1	N
Digital time signature retrieval	1	N

*N is the number of documents in the signed list.

3 RESULTS

3.1 The Manifest Signature Architecture Implemented using XML

We tested an example of the manifest signatures of 166 DICOM MR images (packaged as DICOM DIR) to be signed using both our method and the one-by-one method regulated in the DICOM standards.

First, system will search all of files in selected folder and mark the URI of files to generate the signed list. According to the URI of the signed list, system will find the direction pathway of file and calculate the digest value of each file. The $ID(M_n)$ is created from the DICOM head tags such as Transfer syntax UID (0020, 0010), SOP Instance UID (0008, 0018)...etc, and we use XML encoding to present $Ref(M_n) = ID(M_n) \wedge H(M_n)$.

Fig. 3 shows the $Ref(M_n)$ structure presented as XML. In Fig. 3, a unit of $Ref(M_n)$ is represented by the tag name "Reference". The attribute "URI" of <Reference> is the related directory pathway in addition to an identification of the resource file. $H(M_n)$ is represented by <DigestValue> and $ID(M_n)$ is represented by <DigestMethod Algorithm> and <Transforms>. Transforms means the namespace of this referenced document, which identifies the data format. In this example, the value in <Transform Algorithm> is urn:oid:1.2.840.10008.1.2.1, which expresses the explicit little endian coding for DICOM. This attribute can deal with different cases of different data formats for EHR exchange between hospitals; in addition, it can be extended for multiple data formats, which are defined by the user.

We reconstruct $Ref(M_n)$ as the manifest digest MD_D , and put MD_D into <Object> tag as the signing range of manifest signature and the attribute "Id" of <Manifest> is represent the identifier of MD_D . The cascade of the element in XML as the manifest of these DICOM files is shown in Fig. 4. We calculate $H(MD_D)$ and using HOC's private key to encrypt $H(MD_D)$ and generate the manifest signature $SIG(MD_D)$ complies with the W3C XML enveloped signature standard is shown in fig. 5. All value is encode by Base64, the value "Object" in attribute "URI" of <Reference> means sign the <Object> described above, $H(SIG(MD_D))$ is represented by <DigestValue>; the manifest signature is represented by <SignatureValue> the signing certificate is put in <X509Certificate> tag.

```
<Reference URI="SDY00000\PRS00000">
  <DigestValue>lujg6oNg+2lq+l7Gn+HTAw==</DigestValue>
  <DigestMethod Algorithm="http://www.w3.org/2000/09/xmldsig#sha1" />
- <Transforms>
  <Transform Algorithm="urn:oid:1.2.840.10008.1.2.1" />
</Transforms>
</Reference>
```

Figure 3: A unit of Ref(Mn) of one image in DICOM DIR.

```
- <Object>
- <Manifest Id="DICOMDIRManifest">
+ <Reference URI="DICOMDIR">
+ <Reference URI="SDY00000\PRS00000">
+ <Reference URI="SDY00000\SRS00000\IMG00000">
+ <Reference URI="SDY00000\SRS00000\IMG00001">
+ <Reference URI="SDY00000\SRS00000\IMG00002">
+ <Reference URI="SDY00000\SRS00000\IMG00003">
+ <Reference URI="SDY00001\SRS00000\IMG00000">
+ <Reference URI="SDY00001\SRS00001\IMG00000">
+ <Reference URI="SDY00001\SRS00002\IMG00000">
+ <Reference URI="SDY00001\SRS00003\IMG00000">
+ <Reference URI="SDY00001\SRS00004\IMG00000">
+ <Reference URI="SDY00001\SRS00005\IMG00000">
+ <Reference URI="SDY00001\SRS00006\IMG00000">
  :
  :
  :
</Manifest>
</Object>
```

Figure 4: The manifest structure of DICOM DIR presented as XML.

```
- <SignedInfo>
- <Reference URI="Object">
  <DigestMethod Algorithm="http://www.w3.org/2000/09/xmldsig#sha1" />
- <Transforms>
  <Transform Algorithm="http://www.w3.org/2000/09/xmldsig#base64" />
</Transforms>
  <DigestValue>JpJ9G0PZxRHPtXbCvWgGw==</DigestValue>
</Reference>
  <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#rsa-sha1" />
  <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-xm1-c14n-20010315" />
</SignedInfo>
- <KeyInfo>
- <X509Data>
  <X509Certificate>MIIB/TCCAWqgAwIBAgIFyqAAABQwDQYJKoZIhvcNAQEBBQAwVTE...
</X509Data>
</KeyInfo>
  <SignatureValue>NhsWRLMWRFBF0dsB1qvRymPQEUVx1+XMaRBHxGappJbLMEAQY+Iker9...
```

Figure 5: The manifest digital signature presented as W3C XML enveloped signature standard.

3.2 Performance Test

Table 3 shows a comparison of the operation time for test files by different methods. In the one-by-one method, each image has to be hashed, RSA encrypted and timestamp retrieved once, for which the total time needed was about 167.89 seconds. In the proposed method, however, the total time was reduced to 2.81 seconds. The demonstrating PC is a Pentium D 2.8 GHz, with 4 GB memory; the size of each image is about 516KB~1.58 MB; the total data size of the images is 82.3 MB; the type of the smart card is e-gate. The test of digital time signature time is retrieved from the Taiwan HCA-TSA.

Table 3: Comparison of the time needed for signing 166 DICOM MR images between the proposed method and the one-by-one method.

Operations	Calculation time	
	Proposed method (sec)	One-by-one method (sec)
Hashing	1.34	1.34
Signing	1.27	148
Digital time signature retrieval	0.2	18.55
Total time	2.81	167.89

4 DISCUSSION

If a portable EHR is to be exchanged with other organizations, the presence of a removable storage media is also required. To verify the authenticity of the EHR, the sender site creates the digital signature of the EHR as the organizational stamp. The receiver can then verify whether the EHR is valid by Certificate Authority (CA) and TTP. The digital signature created by the hospital and the digital time signature created by the TSA can record the authenticity of the EHR for inter-organization exchange. In terms of cryptography, the digital timestamp mechanism is not used to provide a qualified digital signature, but to certify a qualified time signature. Some infrastructures, such as patient identification, certificate management, and standards should be established as well. And some security issues should be noticed in implementation: e.g., data backup, audit trail, register loss, maintenance, recovery, etc.

A portable EHR contains clinical data and related setting data, the combination of these data can reconstruct representation of EHR. The signing range should contain all of the data related to clinical data. It is very important to ensure the integrity of representation of EHR. If only signing clinical data and related setting data is not signed, it could be happened in inconsistency of representation of EHR while setting data had been modified. The security protection of EHR should include all of data in portable storage media.

In general, most of the medical information standards and national regulations regulating the legal EHR do not use the manifest signature. If existing medical information digital signature rules are followed, as the practice is not feasible due to the high computational time of the signing and verifying processes. However, the signing time and timestamp retrieval time need to be reduced because a portable EHR may contain many clinical documents and

images, all of which need to be protected. For example, a study may contain hundreds of DICOM images, and following the DICOM standards in these cases is impracticable in real-world clinical operations. The manifest signature can be used not only for the exchange of EHR, but also for EHR long-term storage in hospitals.

Fast and reliable proof of authenticity and integrity is needed for security considerations when an EHR became portable. It is common that patients collect their own health records from different hospitals and manage the process by themselves. The implementation of a centralized health record containing personal health records is very difficult when taking into considerations of the physicians' intellectual property rights and patient privacy.

5 CONCLUSIONS

The computational time of this prototype is much lower than that of the one by one digital signature method. Following the existing medical information digital signature rules, the practice is not feasible. Using the proposed method, the computational time is reduced. In addition, this method can be used not only for the exchange of EHRs, but also for their long-term storage.

ACKNOWLEDGEMENTS

This work was supported by the National Science Council of Taiwan under Grant NSC 95-2221-E010-003.

REFERENCES

- Bates, DW, Ebell M, Gotlieb, E, Zapp, J, & Mullins, HC 2003, 'A Proposal for Electronic Medical Records in U.S. Primary Care', *Journal of the American Medical Informatics Association*, vol. 10, no. 1, pp. 1-10.
- France, FHR, Bangels, M & De Clercq, E 2007, 'Purposes of health identification cards and role of a secure access platform (Be-Health) in Belgium', *International Journal of Medical Informatics*, vol. 76, no. 2-3, pp. 84-88.
- Lekkas, D & Gritzalis, D 2007, 'Long-term verifiability of the electronic healthcare records' authenticity', *International Journal of Medical Informatics*, vol. 76, no. 5-6, pp. 442-448.
- Lekkas, D, Gritzalis, S & Katsikas, S 2002, 'Quality assured trusted third parties for deploying secure internet-based healthcare applications', *International Journal of Medical Informatics*, vol. 65, no. 2, pp. 79-96.
- Makoul, G, Curry, RH & Tang, PC 2001, 'The Use of Electronic Medical Records: Communication Patterns in Outpatient Encounters', *Journal of the American Medical Informatics Association*, vol. 8, no. 6, pp. 610-615.
- Pharow, P & Blobel, B 2005, 'Electronic signatures for long-lasting storage purposes in electronic archives', *International Journal of Medical Informatics*, vol. 74, no. 2-4, pp. 279-287.
- Schurig, A, Heuser, H & Wedekind, R 2001, 'Introduction of the health professional card into the SAXTELEMED-Project', *International Congress Series*, vol. 1230, pp. 867-871.
- Tang, PC, Ash, JS, Bates, DW, Overhage, JMS & Sands, DZ 2006, 'Personal Health Records: Definitions, Benefits, and Strategies for Overcoming Barriers to Adoption', *Journal of the American Medical Informatics Association*, vol. 13, no. 2, pp. 121-126.
- Wang, M, Lau, C, Matsen, FAIII & Kim, Y 2004, 'Personal health information management system and its application in referral management', *IEEE Transactions on Information Technology in Biomedicine*, vol. 8, no. 3, pp. 287-297.
- Yang, CM, Lin, HC, Chang P & Jian, WS 2006, 'Taiwan's perspective on electronic medical records' security and privacy protection: Lessons learned from HIPAA', *Computer Methods and Programs in Biomedicine*, vol. 82, no. 3, pp. 277-282.
- Zhou, XQ, Huang, HK & Lou, SL 2001, 'Authenticity and integrity of digital mammography images', *IEEE Transactions on Medical Imaging*, vol. 20, no. 8, pp. 784-791.

AUTHENTICATION OF PROFESSIONALS IN THE RTS E-HEALTH SYSTEM

André Zúquete

*IEETA / Univ. of Aveiro, Campus Univ. de Santiago, 3810-193 Aveiro, Portugal
andre.zuquete@ua.pt*

Helder Gomes

*IEETA / ESTGA / Univ. of Aveiro, Campus Univ. de Santiago, 3810-193 Aveiro, Portugal
helder.gomes@ua.pt*

João Paulo Silva Cunha

*IEETA / Univ. of Aveiro, Campus Univ. de Santiago, 3810-193 Aveiro, Portugal
jcunha@ua.pt*

Keywords: e-Health, authentication, public key certificates, PKI, smartcards, SSL/TLS, roles, RBAC.

Abstract: This paper describes the design and implementation of a PKI-based e-Health authentication architecture. This architecture was developed to authenticate e-Health Professionals accessing RTS (Rede Telemática da Saúde), a regional platform for sharing clinical data among a set of affiliated health institutions. The architecture had to accommodate specific RTS requirements, namely the security of Professionals' credentials, the mobility of Professionals, and the scalability to accommodate new health institutions. The adopted solution uses short lived certificates and cross-certification agreements between RTS and e-Health institutions for authenticating Professionals accessing the RTS. These certificates carry as well the Professional's role at their home institution for role-based authorization. Trust agreements between health institutions and RTS are necessary in order to make the certificates recognized by the RTS. As a proof of concept, a prototype was implemented with Windows technology. The presented authentication architecture is intended to be applied to other medical telematic systems.

1 INTRODUCTION

RTS (Rede Telemática da Saúde (?; ?)) is a regional health information network (RHIN) providing an aggregated view of clinical records provided by a set of affiliated health institutions (HIs). Each HI uses its own system to produce and manage clinical records, which can be browsed and presented in different ways by RTS. The goal of RTS is not to replace the systems used by the affiliated HIs, but to provide a mediated, global view of patient's clinical records independently of the HIs holding their records.

RTS provides Portals for accessing clinical records. Two Portals were foreseen: the Patients Portal and the Professionals Portal. The first is to be used by Patients to communicate with their family doctor and other health system issues such as renovation of prescription and schedule appointments. The second is to be used by healthcare Professionals for accessing clinical records required for their normal, daily work.

The RTS Professionals' Portal is a web server ac-

cessible through RIS (Rede Informática da Saúde¹), a nation-wide, private network, interconnecting all HIs, including the ones affiliated with RTS. Professionals access data provided by RTS using a normal web browser running on a computer connected to the RIS.

This paper describes an authentication architecture providing strong authentication for Professionals accessing the Professionals' Portal. Strong authentication is provided by using a two-factor approach: possession of a security token and knowledge of a secret. For the security token we chose a smartcard. Smart cards are tamperproof devices with security-related computing capabilities which are very convenient for running computations using private keys of asymmetric key pairs.

This paper is organised as follows. Section 2 overviews the authentication architecture and some RTS requirements. Section 3 presents some design goals. Section ?? presents the proposed authentica-

¹Health Computer Network

tion architecture. Section ?? presents some related work. Section ?? presents our prototype implementation. Section ?? evaluates the architecture and the implementation. Section ?? concludes the paper.

2 OVERVIEW

This paper describes an authentication architecture providing strong authentication for Professionals accessing the Professionals' Portal. Strong authentication is highly recommended in this case, as Professionals can access sensitive data — the patients' health records. The architecture allows Professionals to roam between computers of their HI or other HIs.

Our authentication architecture had also to deal with authorization issues. In fact, the RTS Portal uses a role-based access control (RBAC) policy for deriving the Professionals' authorizations to access clinical data. Therefore, each time a Professional accesses the RTS Portal, the later must learn a role that the former may legitimately play for deriving authorisations.

The proposed architecture uses public key cryptography as the basis for its operation. Each Professional is given a smartcard for storing and using personal credentials for accessing the RTS Portal. The Professionals' authentication process uses a facility provided by web browsers, the SSL/TLS client authentication with asymmetric keys and X.509 public key certificates (PKCs), to prove the authenticity of the Professional to the RTS Portal (?; ?).

Furthermore, the PKCs used by Professionals in the SSL authentication process provide extra information to the RTS Portal, besides the identity of the Professionals, such as the HIs they are affiliated to and the role they are playing. As a Professional may play several roles simultaneously (e.g. Doctor and Chief Doctor), the PKC must contain all the roles we can play, being up to the RTS Portal to chose the role to play, from the possible ones, in each session.

Since a Professional's PKC carries roles the owner can play, a mechanism must be provided to deal with role changes. A possibility was to use certification revocation for outdating given roles. However, revocation validation requires online communication between the PKC validator and the PKC issuer, which may not be possible or convenient. Furthermore, some roles are very short in time, for example, vacation substitutions, and these dynamics can be more easily managed by short lived certificates than by Certificate Revocation Lists (CRL).

Alternatively, we chose to use short-term validity periods for Professionals' PKCs, as in (?). This way, Professionals' PKC get automatically invalid after a

short period of time after their issuing and Professionals must apply for new ones. A simple and secure enrolment process was also conceived for getting new PKCs.

The public key infrastructure (PKI) for managing Professionals' credentials for accessing RTS uses a flexible, scalable grassroots approach. Each HI and the RTS have their own PKI, including root and issuing certification authorities (CA). The issuing CA of each HI is responsible for issuing RTS credentials for local Professionals. The issuing CA of RTS is responsible for issuing credentials for the RTS Portal. The validation of certificates issued by separate PKIs is enabled by cross-certification agreements. This means that the RTS Portal will only be able to validate Professionals' credentials issued by HI CAs cross-certified by RTS; other people, including Professionals from other HIs, cannot be authenticated by the RTS Portal, therefore cannot access protected clinical data.

In this paper we mainly describe our architecture for using smartcards for authenticating Professionals and the RTS Portal when interacting with each other and for providing Professionals' roles to RTS. However, the architecture was designed taking into consideration future enhancements and synergies, such as:

- Enable Professionals to use the same smartcard for producing signed data as input for health information systems.
- Enable Professionals to give signed consents regarding accesses to the clinical data.
- Adoption of a similar authentication model for authenticating Patients, possibly using the new, smartcard-enabled Citizens Card.
- Usage of PKIs deployed for managing smartcards to generate credentials for mutual authentication within secure communications between hosts or servers used in the RTS and in HIs (e.g. with IPSec or SSL/TLS (?; ?)).
- Usage of PKI deployed in each HI for managing the local authentication of Professionals accessing local services (e.g. secure wireless network access).

3 DESIGN GOALS

A set of design goals were defined at start. Those goals derived both from RTS requirements and from previous experiences with informatics services in healthcare environments.

The first goal was Professionals' mobility. The authentication architecture should not restrict the mo-

bility of Professionals; at the end it could be possible to use any computer, belonging to the RIS, to access RTS services. Naturally, this goal depends on software and hardware installed in client computers accessing Professionals' authentication tokens. Nevertheless, we tried to facilitate the widespread use of those tokens by using common hardware (e.g. USB ports) and free software packages (e.g. software packages already provided by operating system vendors).

The second goal was to be pragmatic regarding the implementation of a PKI for managing asymmetric keys and PKCs. Nation-wide PKIs do not exist for this purpose. And, though they could be advantageous, they are also difficult to deploy and to manage. Thus, we chose to start from a sort of minimalist, ad-hoc scenario, with no global PKI on top of the RTS and all the HIs, but instead with isolated, standalone PKIs on each entity, RTS and HI.

The third goal was RTS independency regarding the management of personnel in affiliated HIs. Each HI is an independent organization, with its own Professionals, human resources management department and some kind of directory service to store the Professionals' information. Independent of RTS, they will continue to manage their Professionals because of their own, internal systems. It thus makes sense to reuse HI Professionals information and let each HI to manage the access of its own Professionals to the RTS. This way, we avoid replication of information and a centralized enrolment of Professionals in RTS.

The fourth goal was to minimize communication overheads related to the authentication of Professionals and fetching/validation of role membership. Namely, we tried not to use on RTS any online services from HIs to deal with details regarding the identification, authentication and role membership of Professionals. Since Professionals' information is managed solely by their home HI (our previous goal), this means that Professionals' identification and authentication credentials should convey RTS as much information as possible, to avoid contacting online services at Professionals' home HIs.

The fifth and final goal was browser compatibility. To avoid the requirement of using a specific browser, no client-side active code (ActiveX and Java Applets) is used in RTS. Therefore, we could not use any special code for managing the authentication of Professionals using a browser to access RTS. In other words, the authentication mechanism using a two factor approach should be already available within the basic functionality of all browsers. As we will see, although the basic functionality exists in all popular browsers (support of SSL client-side authentication), the exact mechanisms and policies used to handle such support

are different and raise some problems.

4 AUTHENTICATION ARCHITECTURE

The authentication architecture is resumed in Figure ???. The Professional uses a web browser to access the web server that implements the RTS Portal, and uses an SSL secure channel for protecting the communication from eavesdropping. Furthermore, mutual authentication is required in the establishment of the SSL secure session, thus the browser authenticates the RTS Portal and the RTS Portal authenticates the Professional using the browser. Similarly, the Professional uses a web browser a mutually authenticated SSL session to access the HI Issuing CA web server for requesting fresh RTS credentials.

4.1 The Professionals' Smartcard

A Professional's smartcard carries two types of asymmetric key pair and corresponding PKCs. One type we call **RTS credentials**, which are to be used to authenticate himself when accessing the RTS Portal. The other type we call **HI credentials**, which are to be used to authenticate himself when accessing his HI issuing CA for getting new RTS credentials.

Smartcards are initialised and provided by HIs to their own Professionals. At start they only carry the HI credentials. When required, the owner uses them for requiring RTS credentials. These credentials can then be used to access the RTS Portal.

RTS credentials are short lived, lasting for one or a few days. The RTS Portal doesn't use remote HI services for checking for their validity. Instead, it assumes that a Professional's role revocation will naturally be enforced by not being able to get a new RTS credential including the revoked role. On the contrary, HI credentials are long lived, because they are used for long periods of time for getting new RTS credentials.

4.2 Professional Authentication

The Professional authentication is requested by the SSL server-side of web servers and conducted by the SSL client-side running on browsers. The SSL client-side authentication uses the Professional's smartcard for his authentication. The browser is configured to use smartcard services and when client-side authentication is required it will prompt the Professional for the right credentials, including the ones inside the smartcard, he intends to use. The Professional chose

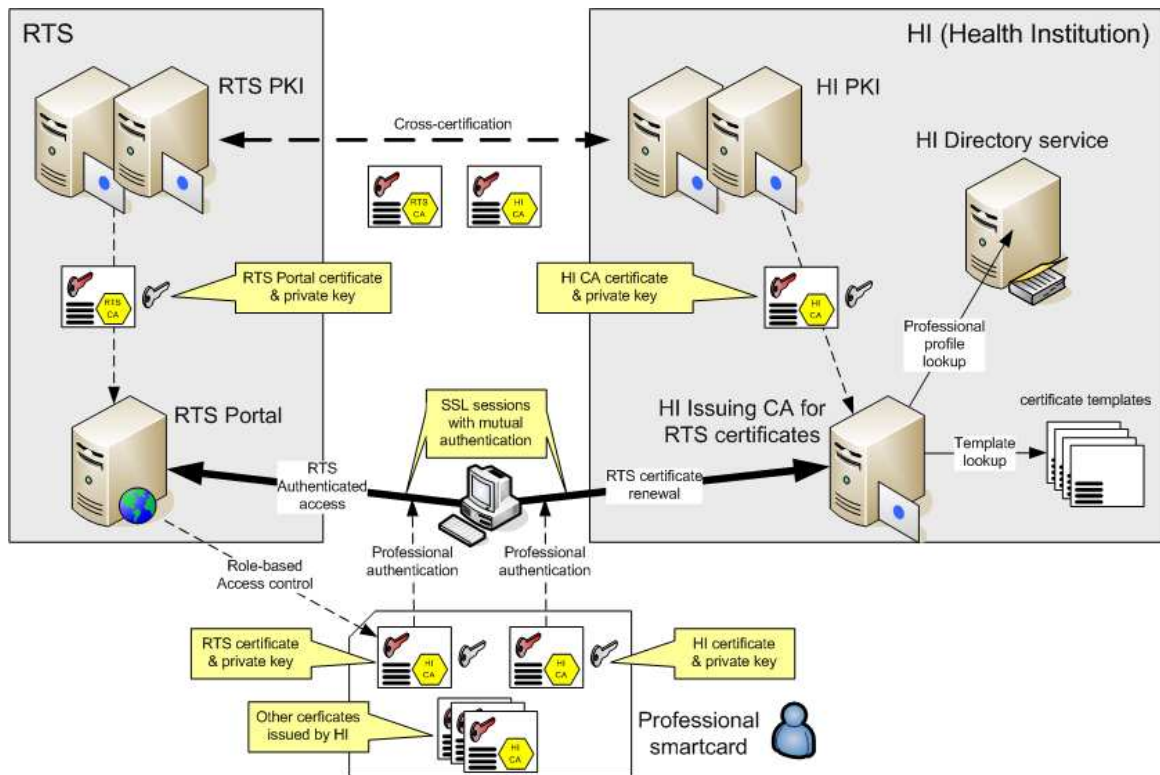


Figure 1: Overview of the authentication architecture for HI Professionals whiling to access the RTS Portal.

the right pair of asymmetric keys from the smartcard, and the PKC of the public key, and the browser uses them to provide client-side authentication.

This client-side approach is the same for accessing the RTS Portal or the HI Issuing CA. It is up to the Professional to choose the right set of credentials, from the smartcard, to get authenticated. And in all cases it needs to introduce a PIN to unblock the smart card for producing digital signatures required by the SSL authentication protocol.

The web servers used by the RTS Portal and the HI Issuing CA perform the following actions: (i) validate the PKC chose and presented by the Professional, (ii) use the certified public key to validate the SSL secure channel establishment and (iii) enable the service, RTS or CA, to access the Professional's PKC. The RTS learns from the PKC the Professional's identity, his home HI and his roles; the CA learns only the Professional's identity.

4.3 Role Assignment and Selection

The roles of each Professional are embedded in the PKC of his RTS credentials. These roles are stored in extension fields, namely the Extended Key Usage (EKU) field. Each role was given a numerical tag, an

ASN.1 Object Identifier, reserved at IANA² for RTS.

Each time a Professional requests RTS credentials, he gets, after proper authentication at the HI Issuing CA, a new PKC with the current roles he can play. This PKC is communicated to the RTS web server during SSL authentication and, if successful, the PKC is made available for consulting by the RTS Portal during the SSL session. This way, when a Professional initiates an SSL-protected session with the RTS Portal, it can easily learn the set of roles the Professional can play and prompt the Professional for selecting the right role for the current session.

4.4 Trust Relationships

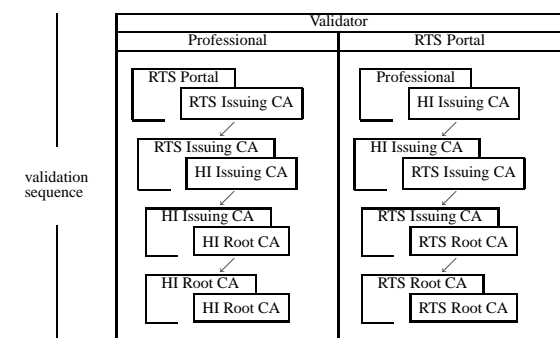
Each entity, RTS and HI, uses an independent PKI for managing RTS and HI authentication credentials used by Professionals. The RTS is not meant to serve as a CA for all HIs; it only deploys a PKI mainly for managing its own certificates. HI certification hierarchies may be isolated or integrated in wider hierarchies providing large-scale validation of certificates. For the RTS that is irrelevant, all it requires is an Issuing CA for issuing RTS credentials for local HI Professionals.

²<http://www.iana.org>

Since RTS and HI certification hierarchies are isolated from each other at the beginning, some mechanism is required to enable the RTS Portal to validate Professionals' RTS credentials, issued by HIs. Similarly, some mechanism is required to enable professionals to validate the credentials of the RTS Portal, issued by RTS. This mechanism is cross-certification. When an HI gets affiliated to the RTS, the RTS Issuing CA issues a certificate for the public key of that HI Issuing CA. With this certificate, the RTS is able to validate all the PKCs of RTS credentials issued for the Professionals of that HI. Similarly, the HI Issuing CA issues a certificate for the RTS Issuing CA, enabling local Professionals to validate the credentials of the RTS Portal.

4.5 Validation of Certificates

With this cross-certification in place, the validating of certificates' certification chains works as follows. The RTS Portal trusts only on the (self-signed) certificate of the RTS Root CA. Similarly, the Professional trusts only on the (self-signed) certificate of his HI Root CA. Since certificate chain validations progress recursively until finding an error or a trusted certificate, the validation chains are the following:



where $\begin{matrix} X \\ \hline Y \end{matrix}$ represents the PKC of X issued by Y.

Besides cross-certification for certificate chain validation, trust relationships between the RTS, HIs and their Professionals must be complemented by common certification procedures. Namely, all HIs affiliated to the RTS should follow similar procedures for issuing RTS credentials. For instance, smartcards with HI credentials should be initialised by HIs and personally delivered to Professionals.

4.6 Validity of Authentication Credentials

The authentication credentials stored inside a Professional's smartcard are the HI credentials and the RTS credentials. The first ones are used to establish an authenticated session to get the second ones.

HI credentials are to be used frequently, for instance, once per day or once for a couple of days, to fetch new RTS credentials. Therefore, they should have long validity periods and CRLs must be published to prevent unwanted use of them after a given instant. For instance, if a Professional moves from one HI to another one, his smartcard from the former HI must be returned and a CRL should be issued to invalidate the public key of the HI credentials inside the smartcard. Note, however, that CRL issuing and validation are all executed within the same HI, and not by external clients.

RTS credentials are valid only during short periods of time, one or two days. Therefore, no CRLs are used to validate them, since the error window is too narrow to allow a Professional to play a role he is no longer allowed to. Consequently, by default the HI Issuing CA doesn't publish CRLs for RTS certificates.

In special cases, such as disciplinary processes and legal inquiries, it should be possible for the HIs to provide to the RTS Portal, just in time, a list of RTS certificates that should no longer be accepted while in validity period. But since such cases should be rare, it is preferable to deal them as the exception to the general rule above stated: no CRLs exist and are checked for RTS certificates.

All the certificates used in both HI and RTS credentials do not need to be published by Issuing CAs. In fact, these certificates are used solely in the context of SSL mutual authentication, and are communicated to the interacting peers within the SSL protocol. Therefore, they need not be published in some public directory System, as other certificates do, because no one needs them for other purposes.

5 RELATED WORK

The following e-Health Systems were analysed: HYGEIAnet and MedCom/Sundhed.dk.

5.1 Hygeyanet

HYGEIAnet is the RHIN of Crete, Greece. It was developed by the Institute of Computer Science (ICS) of Foundation for Research and Technology – Hel-

las (FORTH) to provide an integrated environment for delivery of health care services in Crete Island (?; ?).

Similarly to RTS, HYGEIAnet is formed by several HIs, namely Hospitals and Primary Care Units, each with its own health data, information services and human resources. HYGEIAnet operates above these independent healthcare units, providing an infrastructure for sharing clinical information. Also, the Integrated Electronic Health Record (I-EHR) is a key element as it aggregates the patient health information in all participating healthcare units.

The trust and security frameworks are implemented in HYGEIAnet with VPNs, SSL, smartcards, PKI, security certificates and digital signatures. A Regional certification authority issues the certificates for users and applications. These certificates can be used to authentication and digital signing of documents and in case of user certificates they can be stored in smartcards (?).

Authentication is a centralized process in HYGEIAnet. All applications and services are registered in the Health Resource Service (HRS) and issued a unique ID. Each HYGEIAnet user also must register in HRS to be able to use HYGEIAnet services, and a unique user name and password is provided. The username and password are communicated to an authentication server (AS) and a certificate is issued from the regional CA.

In terms of authorization it follows a decentralized approach, where each individual service maintains and manages roles (groups) and role based permissions. The user must be assigned to the proper role in each service he is to have access.

When accessing a service, the user is authenticated through the Authentication Service and gets his individual access rights validated through the individual service.

The RTS and HYGEIAnet approaches for authentication differ: HYGEIAnet has a centralised management of resources (users and services), with certificates issued by a regional CA, and requires an online AS for user authentication. On the contrary, RTS has a decentralized management of resources, reusing the management services belonging to the affiliated HIs, with certificates issued by HI CAs for their own users and services and not requiring any online user authentication service to be used by RTS.

5.2 The Health Portal (Sundhed.dk) and the Health Data Network of Denmark (MedCom)

MedCom is the National Health Data Network of Denmark. It is working since 1994, and it connects

more than 2000 Hospitals, Pharmacies, General Practitioners (GPs) and Specialists. It started has a VAN network exchanging EDIFACT messages (?). In 2004 it started the process of migration to the Internet and EDIFACT messages were translated into XML messages. Today, both message formats are used (?).

Network security is implemented at three levels (?; ?).

At the first level are VPN connections connecting healthcare networks to a central hub in a star topology. This solution allows the reuse of Internet connections that all the health care units already have.

At a second level there is an agreement system that controls the data flows from and to any of the local healthcare networks. When a connection between two healthcare networks is needed, a previous access to the agreement system is required to establish the connection between the two networks.

The third level of security is user authentication, made locally through his username and password, or his asymmetric key pair and PKC.

The Health Portal started in December 2003. It works on top of the Health Data Network and reuses its infrastructure and services. Unlike the Health Data Network, that only provides services for Professionals, the Health Portal provides services for both Citizens and Professionals (?).

User authentication, for both Citizens and Professionals, is made using OCES certificates³ issued by the national PKI that can be used in several national public services. Professionals can use several OCES certificates: (i) Administrative digital signature for region, hospital or GP, (ii) health professional's digital signature based on personal identifier and (iii) authorization for treating patients (?).

Comparing with RTS, the Danish system extensively uses asymmetric keys and PKCs, benefiting from a nation-wide PKI. However, many of the Danish system security requirements, such as Professionals' digital signatures, are currently not required by RTS, since it is not used for entering signed data into the health information system. We believe, however, that our architecture can evolve, but keeping its basic structure, for provide security services similar to the ones provided by Danish system. Furthermore, our PKI may coexist with a national-wide PKI encompassing all HIs, though not necessarily using it.

³OCES certificate: Public Certificate for Electronic Service

6 PROTOTYPE IMPLEMENTATION

As a proof of concept, a prototype of the authentication architecture was implemented. The prototype extensively used available products for Windows operating systems, because of its dominance in the computer desktops of the HI currently affiliated to RTS.

The prototype included an RTS service, with a two-level PKI and a web Server (Professionals' Portal), one HI instance, with a two level PKI, an Active Directory Server and one registered Professional (one smartcard). CAs were implemented with Windows Certification Services available in Windows 2003 Server Enterprise Edition. When installed in Enterprise mode, this CA interacts with AD to obtain user information for certificate issuing, and uses certificate templates, stored in AD and subject to AD access control rules, for certificate issuance management.

The key aspects to test in the prototype were (i) the impact of different middleware software in smartcard deployment, (ii) the deployment of an HI PKI for the management of RTS credentials for local Professionals, and (iii) the use of short lived RTS credentials to access RTS services.

6.1 Smartcard Deployment

Since smartcards are portable devices, in theory they may be used to authenticate Professionals accessing the RTS from different computers. However, this requires some software installed in those computers: (i) the card reader driver and (ii) middleware to fill the gap between applications and smartcard services.

There are different trends in this specific middleware area. Windows applications, such as the Internet Explorer browser, use the CryptoAPI (CAPI), which can use several Cryptographic Service Providers (CSP) for interacting with different smartcards. Another approach is to use PKCS#11 (?), a standard interface for cryptographic tokens. This interface is used by Netscape and Firefox browsers.

Since middleware modules are usually specific for smartcard manufacturers and some manufacturers impose limits on the number of computers where they can be installed or do not provide similar modules for all operating systems, the following approaches were foreseen: (i) the use of smartcards with native support from the operating system, (ii) the use of open source software or free binaries and (iii) the use of non-free software providing support for multiple smartcards.

In our prototype we used only Windows XP systems for the Professional computers and two smart-

card tokens: Rainbow iKey 3000 and Schlumberger⁴ Cyberflex e-gate 32k. None of them was natively supported by Windows. Also we were not able to get an open source solution (openSC/CSP#11) working reliably. For the non-free solution we used SafeSign Standard 2.0.3 software, and both smartcards worked properly after their first initialization was made by SafeSign. If this first initialization is not made by Safe Sign, chances are that smartcards are not recognized, as it happened with Cyberflex Card.

After low-level initialization (e.g. personalization), smartcards were incepted with the Professional HI credentials, allowing the owner to enrolment for RTS certificates. The HI credentials cannot be renewed and the smartcard becomes useless when the HI certificate validation period expires.

6.2 HI PKI Deployment

The HI PKI was implemented with an offline Root CA and an online Issuer CA. The latter interfaces with an AD and with an IIS 6.0 server with a web interface for certificate enrolment, CRL Distribution Point (CDP) and Authority Information Access (AIA) functionalities.

Some groups were defined in AD, one for each identified professional role, and Professionals assigned to them. They provide access control for certificate enrolment.

Certificates issued by the Issuing CA are tailored using certificate templates. These templates allow the definition of certificate characteristics and access control rules. Certificate templates were created, one for each professional role that only differ in the specification of certificate extensions and certificate security. An application policy was defined for each Professional role, to be included in RTS certificates issued for the role; application policies are simple ASN.1 OIDs used by RTS after reservation at IANA. The application policy OID is stored in the certificate EKU (Extended Key Usage) field. Also a certificate template was defined for the HI certificate, with an application policy for RTS certificate renewal.

Note that RTS certificates should contain all the current roles of a Professional, and not only one. However, that is simply not possible with certificate templates and access control rules. Therefore, in the prototype a Professional may have several RTS credentials, one for each role, and chose the proper one when starting a session with the RTS Portal.

The customization of Windows certificate templates has also some limitations. Namely, certificate templates do not allow for validity periods shorter

⁴Now Gemalto, after being Axalto

than two days. This may be problematic if two days is considered a large risk window for RTS credentials. But in our opinion two days is perfectly reasonable.

The Issuer CAs also issues cross-certificates for the public key of the RTS Issuing CA. Name constraints were used to define the name space for the accepted certificates. Certificate (issuance) constraints and application constraints were not used because they are not interpreted by browsers, since they require some application context (?).

A web interface, adapted from the Microsoft Certificate Services web interface, was deployed for enrollment for RTS certificates. After the Professional authentication using its HI credentials, RTS certificates are immediately issued and installed in the Professional's smartcard. Both Internet Explorer and Netscape can be used for RTS credential renewal.

6.3 Usage of RTS Credentials for Accessing the RTS Portal

The validation of Professionals' RTS credentials by the RTS Portal, an IIS 6.0 web server, was performed at two different levels. At the IIS level, validation follows SSL rules and certification chains. At the application level, validation includes checking RTS OID values placed in EKU field of the received RTS certificate. The Portal only initiates a session with a Professional if his certificate is considered valid at both levels.

Finally, Professional can use both Internet Explorer and Mozilla Firefox to access the RTS Portal. Tests were made in order to determine if the number PKCs from the HI PKI hierarchy in the smartcard could be reduced, but due to different approaches between browsers for building and validation of certificate chain, we conclude that all certificates must be present in order to allow both browsers to be used.

7 EVALUATION

In this section we evaluate the architecture and implementation of our authentication system taking into consideration the design goals presented in Section 3.

Concerning the first goal, a pragmatic PKI implementation, it was achieved, since no specific, large-scale PKI is required. On the contrary, the PKI is build on top of independent PKIs and cross-certification agreements. Trust relationships between RTS and affiliated HIs are reflected in such cross-certification and on common policies for issuing RTS certificates for Professionals.

Concerning the second goal, Professionals' mobility, smartcards embedded in USB tokens are the most promising solution nowadays but still raise some problems. For instance, they (still) cannot be used with PDAs and smartphones. Furthermore, and more problematic, the usage of smartcards in USB-enabled computers still raises the problem of software installation for dealing with them. As we saw in Section ??, it is not simple to find a ubiquitous, free solution for the middleware required by different applications (browsers) to interact with many smartcards.

Concerning the third goal, leaving RTS out of the management of Professionals working at the HIs, it was totally attained. The RTS Portal only requires Professionals to have a valid certificate issued by their HI and containing a set of role on it. HIs have full control on the management of local Professionals and their role, enabling RTS access by issuing RTS certificates with the proper contents, namely Professional identity, HI affiliation and possible roles.

The fourth goal, to minimize communication overheads between RTS and HIs for authenticating Professionals and getting their role, was also fully attained. The RTS Portal, by itself, is capable of authenticating Professionals just by validating their certificate, without checking CRLs remotely, and capable of learning their role also from the certificate. No on-line communication between RTS and HIs is required in this process.

The fifth and final goal was browser compatibility. In this case we must say that it may be difficult to provide the same set of functionalities with all the browsers, because of the differences between the existing middleware solutions for bridging the gap between applications and smartcards (CAPI, PKCS#11, etc.). Furthermore, some smartcard management activities, such as garbage collection of useless credentials inside the smartcard, may require the deployment of active code for running within Professionals' browsers.

8 CONCLUSIONS

In this paper we described the design and implementation of an authentication architecture for Professionals working within the RTS e-Health environment. Since Professionals access RTS services using a browser and an RTS Portal, the authentication of Professionals was mapped on top of SSL client-side authentication. The credentials used in this authentication are provided by their HIs and formed by a private key and a short-lived X.509 PKC, both stored inside a smartcard. The short lifetime of these certifi-

cates allows issuing CAs to simplify their PKI: they are not published and they are not listed in CRLs.

The key characteristics of the authentication architecture are (i) the use of smartcards for strong authentication, to store Professional credentials and to improve their mobility, (ii) the use of short-lived RTS certificates carrying Professional identification and roles for authentication on the RTS Portal and authorization of operations required to the RTS, (iii) the use of “normal”-lived HI certificates for Professional enrolment for RTS certificates, (iv) a PKI where the RTS and each HI run their own, private PKI with (v) cross-certification for the establishment of trust relations required to validate Professionals credentials and RTS credentials within SSL sessions. This authentication architecture is highly scalable and is prepared to be applied to other medical telematic projects such as the Brain Imaging Network Grid (BING) (?) and the Grid-Enabled REpoSitories for medicine (GERESmed), two medical networks now under development an IEETA/University of Aveiro.

A prototype was implemented as proof of concept and based exclusively in technology provided by Windows systems or developed for Windows systems. Regarding the browsers used by Professionals, we tested two of the most popular ones also on Windows systems: Internet Explorer and Mozilla Firefox.

The major source of problems that we found for implementing the prototype was the use and management of smart cards by Professionals’ systems and browsers. The variety of middleware existing for managing smart cards and the different approaches followed by different applications (browsers) regarding the middleware make it very hard to provide a clean, ubiquitous interface for Professionals. Furthermore, this is a critical issue in the deployment of this authentication architecture along many different systems and computers.

ACKNOWLEDGEMENTS

This work was financed by the Aveiro Digital Programme 2003-2006 of the Portugal Digital Initiative, through the POSI programme of the Portuguese Government, and by the FCT (Portuguese R&D agency) through the programs INGrid 2007 (grant GRID/GRI/81819/2006) and FEDER.

REFERENCES

- Cunha, J. P. (2007). RTS Network: Improving Regional Health Services through Clinical Telematic Web-based Communication System. In *eHealth Conference 2007*, Berlin, Germany.
- Cunha, J. P. S., Cruz, I., Oliveira, I., Pereira, A. S., Costa, C. T., Oliveira, A. M., and Pereira, A. (2006). The RTS Project: Promoting secure and effective clinical telematic communication within the Aveiro region. In *eHealth 2006 High Level Conference*, Malaga, Spain.
- Cunha, J. P. S., Oliveira, I., Fernandes, J. M., Campilho, A., Castelo-Branco, M., Sousa, N., and Pereira, A. S. (2007). BING: The Portuguese Brain Imaging Network GRID. In *IberGRID 2007*, pages 268–276, Santiago de Compostela, Spain.
- Dierks, T. and Rescorla, E. (2006). The TLS Protocol Version 1.1. RFC 4346, IETF.
- Housley, R., Ford, W., Polk, W., and Solo, D. (1999). Internet X.509 Public Key Infrastructure Certificate and CRL Profile. RFC 2459, IETF.
- ISO 9735 (1988). Electronic data interchange for administration, commerce and transport (EDIFACT). <http://www.iso.org>.
- Katehakis, D. G., Sfakianakis, S. G., Anthoulakis, D., Kavlentakis, G., Tzelepis, T. Z., Orphanoudakis, S. C., and Tsiknakis, M. (2005). A Holistic Approach for the Delivery of the Integrated Electronic Health Record within a Regional Health Information Network. Technical Report 350 (FORTH-ICS/TR-350), Foundation for Research and Technology - Hellas, Institute of Computer Science, Heraklion, Crete, Greece.
- Kent, S. and Atkinson, R. (1998). Security Architecture for the Internet Protocol. RFC 2401, IETF.
- Lloyd, S. (2001). CA-CA Interoperability. PKI Forum.
- MedCom IV (2003). MedCom – the Danish Healthcare Data Network. MedCom IV, Status Plans and Projects. <http://www.medcom.dk/dwn396>.
- Pedersen, C. D. (2005). An baltic healthcare network and interoperability challenges. Cisco eHealth think tank meeting.
- PKCS#11 (2004). PKCS #11 v2.20: Cryptographic Token Interface Standard. RSA Laboratories.
- Ribeiro, C., Silva, F., and Zúquete, A. (2004). A Roaming Authentication Solution for WiFi using IPSec VPNs with Client Certificates. In *TERENA Networking Conference 2004*, Rhodes, Greece.
- Rossing, N. (2005). The Health Portal (www.sundhed.dk) And The Health Data Network Of Denmark. Executive Summary of Presentaion in eHealth Athens 2005. <http://www.ehealthathens2005.gr>.
- Sundhed.dk (2006). The Danish eHealth experience: One Portal for Citizens and Professionals. <http://dialog.sundhed.dk>.
- Tsiknakis, M., Katehakis, D. G., Sfakianakis, S., Kavlentakis, G., and Orphanoudakis, S. C. (2005). An Architecture for Regional Health Information Networks Addressing Issues of Modularity and Interoperability. *Journal of Telecommunications and Information Technology (JTIT)*, 4:26–39.
- Voss, H., Heimly, V., and Sjögren, L. H. (2005). The Baltic ehealth Network – taking secure, Internet-based healthcare networks to the next level. Norwegian Centre for Informatics in Health and Social Care.

COMMUNICATION OF MEDICAL INFORMATION USING AGENTS

John McGrory, Frank Clarke
Dublin Institute of Technology, Dublin 8, Ireland
john.mcgrory@dit.ie, frank.clarke@dit.ie

Jane Grimson
School of Computing, Trinity College Dublin, Dublin 2, Ireland
jane.grimson@tcd.ie

Peter Gaffney
Adelaide & Meath Hospital, Incorporating the National Children's Hospital (AMNCH), Tallaght, Dublin 24, Ireland
peter.gaffney@amnch.ie

Keywords: HL7, *OpenEHR*, CEN-ENV13606, agents, computer, guidelines and protocols.

Abstract: Agents are self-contained software entities which act faithfully and autonomously on behalf of a body of knowledge. They can operate in a standalone capacity, or as part of a social group collaborating and coordinating activities with other software agents. To access their knowledge, agents are interfaced with using message passing communication. The principle behind medical communications is to provide a means for exchanging information and knowledge from one computerised location to another, whilst preserving its true meaning and understanding between the listener and sender. Agent communication is similar to medical communications, but must provide an additional framework element to allow agents to interact at a social and operational level. Social aspects relate to agents collaborating on shared objectives, and operational aspects relate to coordination of tasks between the loosely coupled agents working as part of a group. Medical communications focus on data exchanges specific to the medical domain, while agent communication was designed for a much broader audience. Therefore, it is essential to verify if agent communications can support standard medical data exchanges. This paper investigates current forms of agent based communications and demonstrates they can support medical communication, yet retain their social and interaction information exchange functionality.

1 INTRODUCTION

An agent is a standalone, self-contained software application. It contains its own inference engine and is encoded using goals, plans and beliefs (Rao et al., 1995). The goals are used to describe its motivation, the plans are used to describe its intention (e.g. different types of workflow activities) and its beliefs are used to describe facts (e.g. weight, height or gender). Each plan and goal has some triggering condition using beliefs which must be satisfied before the plan or goal can be executed. The inference engine selects a goal and then chooses plans which can satisfy that goal whose triggering condition is also satisfied. Using this execution dynamic the agent reacts to known beliefs and

events by selecting plans and goals. This permits the agent to be autonomous and self-contained. Using this principle the authors have developed an agent to act on behalf of a clinical and laboratory guideline. However, guidelines are rarely used in isolation and it is sometimes necessary for clinicians to review a number of these documents so a customised healthcare plan can be made for their specific patient. For example, if the patient was obese with early signs of renal failure. A guideline focusing on the obesity may indicate reducing the carbohydrates and increasing the protein to lower the weight. However, another guideline focusing on the renal failure would indicate reduce the protein in order to preserve the renal state of the patient. Technically

both guidelines are correct, but provide conflicting information affecting patient healthcare.

Applying the agency concept to guidelines, it was stated that an agent can be encoded with all the goals, plans and beliefs related to a particular guideline. Therefore, if the two guidelines referred to above were encoded into separate autonomous agents, they could communicate, work as a group and collaborate to provide a solution. Supporting knowledge from one guideline could be sent as a message to the other guideline. When this knowledge is received by the other it would update its beliefs (say proteins must be lowered) and this information changes the goals and plans that can be selected. This is because the goals and plans are selected in relation to the agent's beliefs.

The principle aim of medical data communication is to provide a means for exchanging health information and knowledge from one computerised location to another in a complete and context rich format. Its goal is to preserve the true meaning and understanding of information when communicated to another application. To realise this different medical communication standards exist, but none contain any facility for communication to take place at a social or collaborative level.

The thrust of this paper is to illustrate software agent communications, particularly the Foundation for Intelligent Physical Agents (FIPA) message standard is capable of providing a context rich data transmission, similar to that offered by existing medical communication standards, and yet retain its agent social and interaction information exchange functionality. This would permit the agent communication approach to provide equivalent data exchanges as that provided through the medical communication standards, but yet allow the agent act as a socialite with a group of loosely coupled agents which can collaborate and coordinate to solve problems.

2 MESSAGE CONSTRUCTS

From a medical perspective there are three main standards used for constructing messages:

1. CEN-ENV13606-4:1999(CEN, 1999)
(currently under revision prEN13606:2004(E)),
2. HL7, Release 2,
3. *OpenEHR*, Release 1,

In the current pre-standard documentation release of CEN prEN13606:2004 (only certain

parts available at this time) there is a synergy between the three standards, and although not officially a standard yet, the CEN standard will be the focus of this paper.

From the agent perspective there are two main standards used for constructing agent messages:

1. FIPA Agent Communication Language (ACL) Message Structure.
2. Knowledge Query Manipulation Language (KQML)

The FIPA standard incorporates and extends many aspects of KQML, and has been the most widely used agent communication standard to date (Luck et al., 2004). Thus for the purposes of this paper the agent communication will focus on the FIPA standard.

Both the CEN medical and FIPA agent communication standards achieve their exchanges by message passing where they simultaneously integrate two types of information within a message (FIPA, 2002) (CEN, 1999):

1. Lower-level information(message payload)
2. Meta-information about the content of the message (message envelope).

The message payload is the structure of the message contents, such as XML, schema, or objects. The message envelope component relates to how the message is seen at a network level between the message sender and receiver. Both the message payload and envelope have different impacts on message passing communication. In the following subsections the function, use and meaning behind the message envelopes and payloads are discussed.

2.1 Message Payload

The CEN pre-standard prEN13606:2004 does not insist on a particular message payload type, and accepts formats such as an XML document, schema, or an object. All of these formats preserve the message data relationship model when communicating information from one system to another. XML permits developers to organise the structure and ordering of information in a document, whilst isolating it from the actual technical content. XML, in combination with other standards, makes it possible to define the content of a document separately from its formatting, making it easier to reuse content in other applications, or for other presentation environments. Most importantly XML provides a basic syntax used to share information between different kinds of computers, different

applications, and different organisations without needing to pass through many layers of conversion. In this case the XML file itself is the message payload. The schema method is like a map or plan, where the information elements in the message are stored in a rigid format, and have a relationship to each other by virtue of their position in the schema. The receiving party is aware of the schema structure and can access the information slots to retrieve the required information. In this case the schema data file itself is the message payload. The Object method is where information is transmitted as a software object. The receiving party can interface with the object to retrieve the desired information. In this case the Object itself is the message payload. FIPA standard allows any Java object to be sent as part of a message payload. This object can be an XML, Schema or other type object which can be handled by the Java object class.

The primary difficulty in exchanging messages is to ensure the message being sent is understood and has the same meaning both by the sender and receiver. One simple natural language expression, such as “The woman is on the bus” can be used to illustrate the complexities associated with communications between different systems. This statement can be interpreted in several ways e.g., “the woman is travelling in the bus”, or “the woman is painted onto the side of the bus”, or “the woman is travelling on top of the bus”. This ambiguity is in addition to the assumption that the observer (the listener) receiving the message knows what the “bus” object is (and this interpretation is the same as the sender) and “on” is a relationship description used when discussing the object “bus”. The confusion associated with this bus example stems from an overlapping of ontologies. Language and ontologies are two interconnected components which are used to formalise the meaning of data, and preserve that meaning when sending and receiving messages (FIPA_a, 2002) (Noy et al., 2001). An ontology is a data model which represents language of a domain and is used to reason about these described objects and relationships between objects. Most people possess the capability to handle more than one ontology, such as a domestic and a work-related ontology. Therefore, different ontologies can co-exist in the one entity, but care must be made to ensure the message exchanges are filtered to match the ontology of the other party. Difficulties in communicating and sharing medical information between institutes, individuals or groups has generated a multitude of ontology and language implementations for example Galen (Rector et al.,

2005) (Stuckenschmidt et al., 2004), Tambis (Baker et al., 1999), UMLS (Unified Medical Language System) (NLM, 2006), ONIONS (Gangemi et al., 1999), HL7 RIM (Beeler, 2001), GENE (Egana, 2005). These ontologies and language implementations specify various medical domains through an abstract conceptualised model of the real world environment. This demonstrates that no unique “one-stop-shop” ontology for the medical domain exists. The FIPA message structure recognises that in the real world different ontologies are present, and instead of forcing a single ontology, it allows many exist in the same environment and includes a framework to define, describe and manage them.

The FIPA ontology is composed of two parts, a vocabulary which describes the terminology of concepts used by agents in their realm of communication (e.g. dietitian or renal), and the classification of the relationships between these concepts, semantics and structures (FIPA_a, 2002). Exchanging messages using a specific ontology provides a richer contextual environment in which to share information between separate software entities.

In summary, a payload holds (or contains) the actual context rich medical information to be exchanged between two or more systems. The message is formed using specific ontologies to ensure the message is understood and has the same meaning between the sender and receiver. Both the medical CEN and agent FIPA standards allow similar types of message payloads to be transmitted. But for communications to work effectively it is vital that the message gets to the correct destination.

2.2 Message Envelope

To deliver a message payload to a specific destination it is necessary to wrap or encapsulate the payload using a message *envelope*. A message *envelope* consists of a number of key parameters which allows the message sender, receiver and content to be clearly identified during message transmission. Agents not only use messages for communicating information in a context rich form, but also for social and collaborative interaction so they transmit more envelope parameters, and messages in general. It is therefore imperative to compare the parameters used by medical and agent message standards to ensure the agent system can support them. This will identify what parameters (if any) would have to be added to the agent

communication model in order to support medical transmissions.

By analysing parameters used by the ENV 13606-4:1999 (prEN13606:2004-Part 5: Exchange models was not available) medical communication standard and comparing them to the FIPA messaging standard it can be shown that six of the twelve CEN parameters have similar technical meanings. A list of these parameters is detailed in Table 1.

Table 1: CEN to FIPA message envelope parameters comparison.

Item	CEN ENV 13606-4:1999	FIPA
1	identification of message by originator	sender, conversation-id
2	EHCR source	sender
3	EHCR destination	receiver
4	EHCR message related agent	reply-to
5	language	language
6	message reference	conversation-id

The main purpose of the FIPA Agent Communication Language (ACL) messaging structure is to allow agents to communicate effectively when being utilised by a wide audience base, and were not designed specifically for a medical application. However, FIPA implementations are free to include customised user-defined message parameters other than the items specified within the standard itself. The semantics of these user-defined parameters is not defined by FIPA, and FIPA compliance does not require any particular interpretation of these parameters (FIPA, 2002). The prefix "X-" must be used for the names of these non-FIPA additional items. By reusing the overlapping parameters as detailed in Table 1, and adding the remaining six parameters using the prefix "X-", the fixed size of the message envelope is 65kbytes. The agent parameters now include the ontology parameter, so a message's ontology can be clearly identified, in addition to language before the message payload is accessed.

In summary, the FIPA messaging standard can be adapted with the addition of these user-defined parameters, to provide a similar message model to that detailed in ENV 13606-4:1999, and yet retain its agent communication functionality.

3 UTILISING AGENT COMMUNICATION

To compare the agent communications to an existing approach consider an example were four clinical

guidelines are used together. The implementation chosen for illustration purposes was the evaluation of a set of Liver Function Tests to determine the cause of a chronic anaemia in patients.

One approach to managing the activities of the four guidelines is to decompose the guidelines into workflow activities and management rules. The management rules from each guideline are linked together centrally using an inference engine. This inference engine is constructed using rules that logically link the various workflow activities together. These management rules provide the motivation for a centralised inference engine to choose particular workflow activities depending on the patients known characteristics (e.g. weight, gender, height). The patient data is retrieved from the LIS using an XML message payload coupled with a standard CEN message envelope. As more guidelines are decomposed and added, the number of workflow activities and the size and complexity of the centralised inference engine increases. But all the decisions on choosing a particular workflow activity are performed centrally by the inference engine as no other separate modules exist in this system. The guidelines no longer exist as separate entities.

An alternative solution to capturing the knowledge of the four guidelines is to encode each as a separate autonomous agent, one agent for each guideline. This is achieved by encoding the guidelines beliefs, goals and plans together in the agent and using each agent's inference engine to interpret them. In this example there is no centralised management resource, therefore each agent must establish links to the other agents. To achieve this agents use message passing to share supportive information and coordinate activities. The patient data is retrieved from the LIS by each agent separately using an XML message payload coupled with a standard CEN message envelope. When an agent received a message from another guideline it altered (if necessary) its own execution based on this belief. If an agent wished to forward supporting information to the other agents it used message passing. Therefore, the only communication between the separate agent modules to coordinate activities together was via message passing, not direct linking as in the centralised solution. In this implementation no centralised resource was used to coordinate activities between the separate agents. After the agents completed deliberations the outcomes were compared to the documented case histories from which the evaluation data was taken. Although outcomes

derived from the agent approach matched that provided by the case histories, it highlighted that the agent approach transmitted more messages than the former centralised approach. On average the group of four agents transmitted a total of 12 messages per second between them. This is because the separate agents relied solely on message passing in which to share information. During this evaluation it was also found that although the message envelope was a fixed size the message payload size could vary dramatically. This was because different types of information were being sent between the agents. Small messages were in the form of short supportive information comments. A longer message is where a more detailed packet of information was sent. Depending on the data being transmitted varying sizes of message payload ranging from 2kbytes to 360kbytes could occur. So how does this affect the computer system using the agents? If the agents, because they can be distributed were located on separate machines and a total of 12 messages was transmitted between them the network bandwidth would range from 8kbps for 2kbyte payloads, to 51kbps for the 360kbyte payloads (assuming the network had an efficiency of 10bits/byte). This is a substantial network overhead. However, this network overhead could be reduced if the agents with high message coupling were located on the same machine. The main reason for this message overhead is that the separate agents need to transmit data, social interaction data and collaboration data in order to operate. Whereas the centralised approach used fixed links between the workflow activities within the one application and so no network traffic was needed. However the agent approach was capable of providing distributed processing of the activities and only relationships between guidelines needing to collaborate needed to be established, where as the centralised approach needed to be run on a single machine.

4 CONCLUSIONS

This research demonstrates the agent communication approach offers the capability to transmit medical information, in an equivalent fashion to that provided by the existing medical communication standards, yet retains its agent social and interaction information exchange functionality. The FIPA ACL standard also allows for the ontology of a message to be identified so a richer form of data interaction can occur between agents. However, the concept of using the agents to represent guidelines highlighted

that as the agents communicated data using messages, they also used messages to socialise and coordinate activities which could have a substantial impact on the network overhead. A summary of the difference between the two approaches is shown in Table 2.

Table 2: Summary Agent and centralised approaches.

Aspect	Agents	Centralised
<i>Workflow activity links between guidelines.</i>	Achieved using coordination message passing.	Achieved within the centralised inference engine.
<i>Sharing of information.</i>	Achieved using message passing.	Achieved by triggering rules within the inference engine.
<i>Access to LIS.</i>	Achieved using message passing. But each agent accesses it separately.	Achieved using message passing. But all required information accessed via one message.
<i>Processing of guideline knowledge.</i>	Distributed.	Centralised.
<i>Guideline knowledge file size.</i>	Small as each agent is self-contained.	Large as inference engine must cover all guidelines.
<i>Adding, altering or deleting of guideline knowledge.</i>	As each agent is independent there is no fixed link between them. So guidelines can be added, altered or deleted without impacting any other resources.	Adding, altering or deleting requires the links in the centralised inference engine to be modified. Therefore, all other resources affected.

REFERENCES

- Baker P.G., Brass A., Bechhofer S., Goble C., Paton N., Stevens R., 1999. "TAMBIS-Transparent Access to Multiple Bioinformatics Information Sources", Scientific and Statistical Database Management, 1999. Eleventh International Conference on 28-30 July 1999 Page: 280.
- Beeler G.W., 2001. *Reference Information Model (RIM) HL7 Version 3*, tutorial in two parts Part I & Part II accessed via <http://www.hl7.org>, 2001.
- CEN/TC 251 Health Informatics, 1999. "ENV13606-4:1999 Health informatics – Electronic Healthcare record communication - Part 4: messages for the exchange of information", accessed via www.centc251.org.
- Egana A., "Ontology Design Patterns for the Formalisation of Biological Ontologies", Degree of Masters of Philosophy, University of Manchester, 2005 accessed via <http://gong.man.ac.uk/docs/MPhilThesis.pdf>

- FIPA, 2002(Foundation for Intelligent Physical Agents). *FIPA Agent Message Transport Service Specification*, Document Number SC00067F, Dated 2002/12/03, accessed via <http://www.fipa.org>
- FIPA_a, 2002(Foundation for Intelligent Physical Agents). "*FIPA Device Ontology Specification*", Document Number SC00091E, Dated 2002/12/03, accessed via <http://www.fipa.org>
- Gangemi A., Pisanelli D.M., Steve G., 1999. "*An Overview of the ONIONS Project: Applying Ontologies to the Integration of Medical Terminologies*", Data and Knowledge Engineering, vol. 31, 1999, Pages 183-220, published by Elsevier.
- Luck M., McBurney P., Preist C., 2004. *Agent Technology: Enabling Next Generation Computing: A Roadmap for Agent-Based Computing*, accessed via AgentLink web site <http://www.agentlink.org/roadmap/> 2004 ISBN 0854 327886.
- NLM 2006, National Library of Medicine USA *Unified Medical Language System* accessed via <http://www.nlm.nih.gov/pubs/factsheets/umls.html>, accessed October 2006.
- Noy N.F., McGuinness D.L., 2001. *Ontology Development 101: A Guide to Creating Your First Ontology*. http://protege.stanford.edu/publications/ontology_development/ontology101.pdf. Stanford Knowledge Systems Laboratory Technical Report KSL-01-05, March 2001.
- Rao A., Georgeff M., 1995. *BDI Agents: From Theory to Practice*, proceedings of the first international conference on multi agent systems (ICMMAS-95) San Francisco USA 1995.
- Rector A., Roger J., 2005. *Ontological & Practical Issues in using a Description Logic to Represent Medical Concepts: Experience from GALEN*, School of Computer Science, The University of Manchester, Preprint Series, CSPP-35.
- Stuckenschmidt H., van Harmelen F., Serafini L., Bouquet P., Giunchiglia F., 2004. *Using C-OWL for the Alignment and Merging of Medical Ontologies*, Proceedings of the First International Workshop on Formal Biomedical Knowledge Representation KRMed'04", Whistler, BC, Canada, 1 June 2004.

AN ONTOLOGICAL APPROACH TO REPRESENTING AND REASONING WITH TEMPORAL CONSTRAINTS IN CLINICAL TRIAL PROTOCOLS

Ravi D. Shankar, Susana B. Martins, Martin J. O'Connor and Amar K. Das
Stanford Medical Informatics, Stanford University
251 Campus Drive, MSOB X215, Stanford, California, 94305, USA
{ravi.shankar, amar.das}@stanford.edu

Keywords: Ontology, temporal reasoning, clinical trials, biomedical informatics, Semantic Web, OWL.

Abstract: Temporal constraints play an important role in the specification and implementation of clinical trial protocols, and subsequently, in the querying of the generated trial data. Protocols specify a temporal schedule of clinical trial activities such as tests, procedures, and medications. The schedule includes temporal constraints on the sequence of these activities, on their duration, and on potential cycles. In this paper, we present our approach to formally represent temporal constraints found in clinical trials. We have identified a representative set of temporal constraints found in protocols to study immune tolerance. Our research group has developed a temporal constraint ontology that allows us to formulate the temporal constraints to the extent required to support clinical trials management. We use this ontology to provide temporal annotation of clinical activities in an encoded clinical trial protocol. We have developed a temporal model that represents time-stamped data and facilitates interval-based temporal operations on the data. Using semantic web technologies, we are building a knowledge-based framework that integrates the temporal constraint ontology with the temporal model to support queries on clinical trial data. Using our approach, we can formally specify temporal constraints, and reason with the temporal knowledge to support management of clinical trials.

1 INTRODUCTION

Clinical trials are formal studies on participants to systematically evaluate the safety and efficacy of new or unproven approaches in the prevention and treatment of medical conditions in humans. A clinical trial protocol is a document that includes study objectives, study design, participant eligibility criteria, enrollment schedule, and study plan. It specifies a temporal schedule of clinical trial activities such as tests, procedures, and medications. The schedule includes temporal constraints on the sequence of these activities, on their duration, and on potential cycles. A temporal constraint is defined as an interval-based temporal annotation on a domain entity in relationship with other entities. Temporal constraints are fundamental to the descriptions of protocol entities, such as the following specifications: *Participants will be enrolled at least two days apart; Participant is ineligible if he/she had vaccination with a live virus within the last 6 weeks before enrollment; The first*

dose will be infused over a minimum of 12 hours; Visit 10 for the participant occurs 3 weeks \pm 2 days from the day of transplant. There is an enormous requirement on the execution of a clinical trial to conform to the temporal constraints found in the protocol. Studies need to be tracked for the purposes of general planning, gauging progression, monitoring patient safety, and managing personnel and clinical resources. The tracking effort is compounded by the fact that a trial often is carried out at multiple sites, geographically distributed, sometimes across the world. The validity of the findings of the clinical trial depends on the clinical trial personnel and the participants performing clinical trial activities as planned in the protocol. More importantly, the treatment and assessment schedules should be strictly followed to ensure the safety of participants.

We have developed an ontological framework that we call Epoch (Shankar et al., 2006), to support the management of clinical trials at the Immune Tolerance Network, or ITN (Rotrosen et al., 2002)

(<http://www.immunetolerance.org/>). As part of this effort, we have developed a suite of ontologies that, along with semantic inferences and rules, provide a formal protocol definition for clinical trial applications. We use the OWL Web Ontology language (<http://www.w3.org/2004/OWL/>), which is a W3C standard language for use in Semantic Web where machines can provide enhanced services by reasoning with facts and definitions expressed in OWL. Central to our ontological effort is the modeling of temporal constraints that we identified in clinical trial protocols. We have created a temporal constraint ontology to formally represent temporal constraints. The ontological representation can then be used to construct rules that can be used in turn, for reasoning with temporal constraints. Thus, at protocol specification phase, a domain expert can capture the essence of temporal constraints using higher-level ontological constructs. At a later time, a software developer can fully encode the constraints by creating rules in terms of temporal patterns and other protocol entities in the ontologies. We are using SWRL, the Semantic Web Rule Language (<http://www.w3.org/Submission/SWRL/>) to write the rules. At execution time of the protocol, the rule elements use the protocol knowledge specified in the Epoch ontologies, and the clinical trial data collected in the clinical trial databases to reason with the temporal constraints. In this paper, we discuss our work in identifying temporal constraints found in ITN's clinical trial protocols. We then discuss our temporal constraint ontology using some patterns that we found in the temporal constraints. We then show how we use the temporal constraint ontology along with other Epoch ontologies to create rules that are then executed at runtime to support clinical trial management.

2 TRIAL CONSTRAINTS

A clinical trial protocol defines a protocol schema that divides the temporal span of the study into phases such as the treatment phase and follow-up phase, and specifies the temporal sequence of the phases. It also includes a schedule of activities that enumerates a sequence of protocol visits that are planned at each phase, and, for each visit, specifies the time window when the visit should happen and a list of protocol activities (assessments, procedures and tests) that are planned at that visit. Activities such as medication need not be confined to visits and can be planned to occur in a time window within a protocol phase. An activity can have sub activities

that impose additional temporal constraints. For example, an assessment activity can include collection and processing of biological specimens with associated temporal constraints.

Here is a representative set of temporal constraints that we found in the ITN protocols that we are encoding:

1. Visit 17 must occur at least 1 week but no later than 4 weeks after the end of 2003 ragweed season.
2. Administer Rapamune 1 week from Visit 0 daily for 84 days.
3. Visit 1 should occur 2 weeks \pm 3 days after transplant.
4. Screening visit evaluations must occur between 30 days prior to Visit -1 and 45 days prior to Visit 0.
5. The vital signs of the participant should be obtained at routine time points starting at 10 minutes post infusion, then at 20-minute intervals until the participant is discharged.
6. Administer study medication at weekly intervals for 3 months.
7. Clinical assessments are required twice a week until Day 28 or discharge from hospital.
8. The first and second blood draws are 10 days apart, and the third draw is 11-14 days after the second.
9. On days that both IT and omalizumab are administered, omalizumab will be injected 60 minutes after the IT.
10. Monitor cyclosporine levels 3 times per week while in-patient, then weekly as out-patient.

As evident in the constraints, clinical activities—we are using the terms *activity* and *event* interchangeably—are temporally dependent on each other. The temporal annotations in the constraints are specified in relative terms typically with reference to one or more clinical events. At the protocol execution time, the actual times of these events found in the clinical data will be used to reason with the constraints. There can also be fuzziness in the relative start and end times as well as in the duration of the activity. An activity can be repeated at a periodic interval for a specific number of times or until a condition is satisfied. The periodic interval can be a single offset or a set of offsets. The temporal annotation of an activity or the temporal ordering of activities can be conditional on other events.

3 TEMPORAL REPRESENTATION

We have developed a temporal constraint ontology that can be used to formally specify the temporal constraints found in the clinical trial protocols. We briefly describe the core entities of the underlying temporal representation below:

Anchor defines an unbound time point that can be used to specify temporal relations among activities. It can be used as a reference point to define the start of another event before or after the anchor. In example 1 (of the constraints listed earlier), *end of 2003 ragweed season* is an anchor used to define the start of *Visit 17*. During the execution of the protocol, an anchor is bound to the absolute time of the anchor as recorded in the clinical trial data.

Duration is the difference between two time points. It is used typically to specify how long an activity lasts. In example 2, *84 days* is a duration.

Anchored Duration relates two activities with a temporal offset. In example 2, the activity *administer Rapamune* is offset from the anchor *Visit 0* by *1 week*.

Varying Duration is defined as duration with a high variance and a low variance. In example 3, *2 weeks ± 3 days* specifies a varying offset between *transplant* and *Visit 1*.

Start and End Expression constrains the start and the end of an activity and is expressed as offsets before or after one or more reference events. In example 4, the start of the activity *Screening visit evaluations* is 30 days before the anchor *Visit -1* and the end is 45 days before another anchor *Visit 0*.

Cyclical Plan Expression formulates events that are repeated at periodic intervals. The repetition ends typically when a specific number of cycles is reached or until a specific condition is satisfied. There are two types of cyclical plans with subtle differences. The first type has a single anchor point with potentially multiple intervals. In example 5, the *vital signs* assessments are planned at 10, 30, 60, 90, 120, and 180 minutes after *infusion*. If the participant gets off schedule because the assessment is made at minute 35 instead of minute 30, then the participant gets back on schedule with the next assessment at minute 60. This type of cyclical plan is used generally with assessments and tests where evaluations need to be made at specific intervals after a clinical intervention. The second type of cyclical plan can potentially have multiple anchors with a single offset. In example 6, the plan is to administer medication at weekly intervals for 3 months. The initial anchor is the event of administering the first dose. According to the

schedule, the second dose will be 1 week later, and the third 1 week later from the second dose. If the participant gets off schedule because the drug was administered 5 days after first dose and not 7 days, then the participant gets back on schedule with the next dose at 7 days from the last dose. This type of cyclical plan is used typically with drug administration where fixed intervals between dosages need to be maintained for safety and efficacy purposes.

Conditional Expression allows associating different temporal annotations with a single activity based on a condition. There are three patterns of conditional expressions – *if-then*, *if-then-else* and *until-then* patterns. Example 9 illustrates the *if-then* pattern – the temporal constraint between the administrations of two drugs is dependent on the condition that the two drugs are administered on the same day. Example 10 illustrates the *until-then* pattern – the monitoring activity is performed 3 times a week until the participant is in in-patient status, and when the status changes to out-patient then the activity is performed weekly.

4 EPOCH ONTOLOGIES

In order to support clinical trial management activities, the Epoch knowledge-based approach provides three methods: 1. knowledge acquisition methods that allow users to encode protocols, 2. ontology-database mapping methods that integrate the protocol and biomedical knowledge with clinical trial data including clinical results and operational data stored in the ITN data repository, and 3. concept-driven querying methods that support integrated data management, and that can be used to create high-level abstractions of clinical data during analysis of clinical results. At the center of all these methods is the suite of Epoch ontologies that provide a common nomenclature and semantics of clinical trial protocol elements. We shall describe each of the core ontologies.

4.1 Protocol Ontology

The *protocol ontology* is a knowledge model of the clinical trial protocol. It simplifies the complexity inherent in the full structure of the protocol by focusing only on concepts required to support clinical trial management. Other concepts are either ignored or partially represented. The main concepts represented in the protocol ontology are the protocol schema and the schedule of activities.

4.2 Temporal Constraints Ontology

The *temporal constraints ontology* models the class of temporal constraints found in clinical trial protocols (see Section 3).

4.3 Virtual Trial Data Ontology

The *virtual trial data ontology* encapsulates the study data that is being collected, such as participant clinical record, specimen workflow logs, and site related data. A mapping component can then map clinical trial data (found in a relational database) to these virtual data records using a mapping ontology. The data model concept is similar to the Virtual Medical Record (Johnson et al., 2001) specification promoted in the clinical guideline modeling efforts.

4.4 Temporal Model

The *temporal model* provides a valid-time model of the temporal component of clinical trial data. In this model, all facts have temporal extent and are associated with instants or intervals denoting the times that they are held to be true. The core concept in the model is the *extended proposition* class that represents information that extends over time. There are two types of extended propositions in the model: 1. *extended primitive propositions* that represent data derived directly from secondary storage, and 2. *extended abstract propositions* that are abstracted from other propositions. These extended propositions can be used to consistently represent temporal information in ontologies. For example, a set of participant visits in a clinical trial data can be represented by defining a class called *VisitRecord* that inherits the *valid time* property from *extended proposition* class. The *valid time* property will then hold a visit's actual occurrence time. Similarly, an extended primitive proposition can be used to represent a drug regimen, with a value of type string to hold the drug name and a set of periods in the valid time property to hold drug delivery times. A more detailed discussion of the temporal model can be found elsewhere in the literature (O'Connor et al., 2006).

5 OWL IMPLEMENTATION

We have developed these ontologies in OWL by building hierarchies of *classes* describing concepts in the ontologies and relating the classes to each other using *properties*. OWL can also represent data

as instances of OWL classes —referred to as *individuals*— and also provides mechanisms for reasoning with the data and manipulating it. OWL also provides a powerful constraint language for precisely defining how concepts in ontology should be interpreted. The Semantic Web Rule Language (SWRL) allows users to write Horn-like rules that can be expressed in terms of OWL concepts and that can reason about OWL individuals. SWRL provides deductive reasoning capabilities that can infer new knowledge from an existing OWL knowledge base. We use SWRL to specify temporal constraints. Once all temporal information is represented consistently using the temporal model, then SWRL rules can be written in terms of this model and the temporal constraint ontology. However, the core SWRL language has limited temporal reasoning capabilities. A few temporal predicates called *built-ins* are included in the set of standard predicates, but they have limited expressive power. SWRL provides an extension mechanism to add user-defined predicates. We used this mechanism to define a set of temporal predicates to operate on temporal values. These predicates support the standard Allen temporal operators (Allen, 1993). Using these built-in operators in conjunction with the temporal model, we can express complex temporal rules. Here is an example SWRL rule to check if participants conform to the visit schedule specified in the protocol:

```
Participant(?p) ^
hasVisitRecord(?p, ?vr) ^
hasVisitId(?vr, ?vid1) ^
hasValidTime(?vr, ?vt) ^
Visit(?v) ^
hasVisitId(?vr, ?vid2) ^
hasStartExpression(?vr, ?se) ^
swrlb:equal(?vid1, ?vid2) ^
temporal:inside(?vt, ?se) ^
-> ConformingParticipant(?p)
```

This rule uses concepts such as *Participant* and *Visit* from the *protocol ontology* and the concept of *Start Expression* from the *temporal constraint ontology*. The class of actual visits undertaken by a participant is the *VisitRecord* in the *virtual trial data ontology*, and is modeled as an extended proposition. The rule uses two built-ins – *equal*, that checks if two strings are equal, and *inside*, which is a built-in that we developed to check if an absolute time is within an anchored varying duration (see Section 3). Protégé (Knublauch, 2004) (<http://protege.stanford.edu/>) is a

software tool that supports the specification and maintenance of terminologies, ontologies and knowledge-bases in OWL. It has a plug-in called SWRL Tab (O'Connor et al., 2005), an editor for SWRL rules. We used Protégé to create the ontologies in OWL and SWRL. We then encoded specific protocols using Protégé's knowledge-acquisition facilities. The data generated from the implementation and execution of clinical trials is stored in a relational database. The types of data include participant enrollment data, specimen shipping and receiving logs, participant visits and activities, and clinical results. We have implemented a dynamic OWL-to-relational mapping method and have used SWRL to provide a high-level query language that uses this mapping methodology. A *schema ontology* describes the schema of an arbitrary relational database. A *mapping ontology* describes the mapping of data stored in tables in a relational database to entities in an OWL ontology. A *mapping software* uses the data source and mapping ontologies to dynamically map data to entities in the *clinical trial data ontology*. A detailed description of the mapping techniques can be found elsewhere in the literature (O'Connor et al., 2007). We are currently using JESS (<http://www.jessrules.com/>), a production rule-engine, to selectively execute the SWRL rules based on the context. For example, the rule that specifies the constraint on a visit time window will alone need to be executed when checking if a specific participant's visit satisfied the constraint. Thus, a temporal constraint is defined first using the temporal constraint ontology, then is formulated as a rule, finally, is reasoned with real clinical data using dynamic mappings.

6 RELATED WORK

Over the years, many expressive models have been developed to represent temporal constraints (Bettini et al., 2002), (Combi, 2004), (Terenziani, 2002), (Duftschmid, 2002). Shahar's approach (Shahar, 1996) identifies temporal abstractions of data and properties using interpolation-based techniques and knowledge-based reasoning. In recent years, there have been a number of initiatives to create clinical trial protocol models that encapsulate clinical trial activities and associated temporal constraints found in a protocol. These ontologies are then used to automate different clinical trial management activities such as eligibility determination, participant tracking, and site management. The

ontologies can also be used when subsequently analyzing the clinical trial data.

In the past few years, we have seen considerable interest in building knowledge-based systems that automate clinical trial protocols and clinical practice guidelines. Our Epoch framework employs a task-based paradigm that combines an explicit representation of the clinical trial domain with rules that capture the logical conditions and temporal constraints found in the trial management process. There have been a number of proposals on task-based clinical guideline representation formats – EON (Musen et al., 1996), PROforma (Fox et al., 1996), GLIF (Boxwala et al., 2004), etc. that deal with temporal constraints on patient data and on activities found in clinical guidelines.

In the area of clinical trials, several modelling efforts have addressed different requirements of trial management activities. An ontology to represent temporal information and cyclical event patterns in clinical trial protocols has been proposed by (Weng et al., 2002). The Trial Bank Project (Sim et al., 2003) is a trial registry that uses a protocol ontology to capture information on randomized clinical trials such as intervention, outcomes, and eligibility criteria. The underlying knowledge base can support systematic reviewing and evidence-based practice.

There is an ongoing effort by CDISC (<http://www.cdisc.org/>), an industry-lead, multidisciplinary organization, to develop and support the electronic acquisition, exchange, submission and archiving of clinical trials data. As part of this effort, CDISC is developing the Structured Protocol Representation that identifies standard elements of a clinical trial protocol that can be codified to facilitate the data interchange among systems and stakeholders including regulatory authorities, biopharmaceutical industry, statisticians, project managers, etc. A parallel effort is the BRIDG (Weng et al., 2007) project, a partnership of several organizations including CDISC, the HL7 (<http://www.hl7.org/>) standards body, the National Cancer Institute and the Federal Drug Administration, that consumes the Trial Design Model work to build a comprehensive domain analysis model representing protocol-driven biomedical/clinical research. The BRIDG model is a work in progress to elaborately define functions and behaviors throughout clinical trials, and uses the Unified Modeling Language (UML) for representation. The model, in its current state, lacks formalization of and reasoning with temporal constraints, and thus, cannot fully support the requirements of ITN's clinical trial management.

7 DISCUSSION

The increasing complexity of clinical trials has generated an enormous requirement for knowledge and information management at all stages of the trials – planning, specification, implementation, and analysis. Our focus is currently on two application areas: 1. tracking participants of the trial as they advance through the studies, and 2. tracking clinical specimens as they are processed at the trial laboratories. The core of the Epoch framework is a suite of ontologies that encodes knowledge about the clinical trial domain that is relevant to trial management activities. This focus on just supporting trial management activities is also reflected in our approach to temporal constraint reasoning. Thus, in the temporal constraint ontology and in our reasoning approach with rules, we have limited ourselves to the types of temporal constraints, to the complexity of formalism and to the levels of reasoning to just support the clinical trial management activities. For example, we do not support checking temporal constraints for consistency. We continue to work on the *temporal constraints ontology* to support newer and more complex constraints. With any complex constraint, one concern is the power, or lack thereof, of our reasoning approach with SWRL rules,

Since we use OWL ontologies and SWRL rules, native RDF Store (storing data as RDF triples) would have been a natural solution for storing clinical trial data, and then seamlessly operate on the data using our ontologies and rules. ITN uses a legacy relational database system to store clinical trial data, and therefore, prevents us from using native RDF Stores as our backend. We have built techniques to map the database tables to our *virtual trial data ontology* OWL classes. With these solutions, our data model remains flexible and independent of the structure of the data sources. We are yet to undertake a thorough evaluation of our dynamic mapping methodology especially in the area of scalability

An often over-looked aspect of knowledge-based reasoning approaches is the task of knowledge-acquisition. Currently, we use the Protégé-OWL editor to build the Epoch models. Based on the class and property definitions, Protégé automatically generates graphical user interface (GUI) forms that can be used to create instances of these classes (OWL *individuals*). Thus, domain specialists can use to enter a specification of a protocol, say for a transplant clinical trial, using these Protégé-generated forms. Unfortunately, domain specialists

find it cumbersome and non-intuitive to use the generic user interfaces as they are exposed to the complexities of the Epoch ontologies, the OWL expressions and the SWRL rules. We are building custom graphical user interfaces that hide the complexities of the knowledge models, and that facilitate guided knowledge-acquisition. Providing a friendly user interface to enter SWRL rules can be challenging.

The knowledge requirements borne out of the need for managing clinical trials align well with the touted strengths of semantic web technologies – uniform domain-specific semantics, flexible information models, and inference technology. Using these technologies, we have built a knowledge-based framework for temporal constraints reasoning that is, above all, practical.

ACKNOWLEDGEMENTS

This work was supported in part by the Immune Tolerance Network, which is funded by the National Institutes of Health under Grant NO1-AI-15416.

REFERENCES

- Allen, J.F., 1993. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11): 832-843.
- Bettini, C., Jajodia, S., Wang, X., 2002. Solving multi-granularity constraint networks. *Artificial Intelligence*, 140(1-2):107-152.
- Boxwala, A.A., Peleg, M., Tu, S. W., Ogunyemi, O., Zeng, Q. T., Wang, D., Patel, V. L., Greenes, R. A., Shortliffe, E. H. 2004. GLIF3: A Representation Format for Sharable Computer-Interpretable Clinical Practice. *Journal of Biomedical Informatics*, 37(3):147-161.
- Combi, C., Franceschet, M., and Peron, A, 2004. Representing and Reasoning about Temporal Granularities. *Journal of Logic and Computation*, 14(1):51-77.
- Duftschnid, G., Miksch, S., Gall, W. 2002. Verification of temporal scheduling constraints in clinical practice guidelines. *Artificial Intelligence in Medicine* 25(2): 93-121.
- Fox, J., Johns, N., Rahmzadeh, A., Thomson, R. 1996. PROfarma: A method and language for specifying clinical guidelines and protocols. *Proceedings of Medical Informatics Europe*.
- Johnson, P.D., Tu, S. W., Musen, M. A., Purves, I., 2001. A Virtual Medical Record for Guideline-Based Decision Support. *Proceedings of the 2001 AMIA Annual Symposium*, 294-298.

- Knublauch, H. Fergerson, R.W., Noy, N.F. and Musen, M.A., 2004. The Protégé OWL Plugin: An Open Development Environment for Semantic Web applications. *Proceedings of the Third International Semantic Web Conference*. 229-243.
- Musen, M.A., Tu, S.W., Das, A.K., Shahar, Y. 1996. EON: A component-based approach to automation of protocol-directed therapy, *Journal of the American Medical Informatics Association*, 3(6): 367–388.
- O'Connor, M.J., Knublauch, H., Tu, S.W., Grossof, B., Dean, M., Grosso, W.E., Musen, M.A., 2005, Supporting Rule System Interoperability on the Semantic Web with SWRL. *Proceedings of the Fourth International Semantic Web Conference*. 974-986.
- O'Connor, M.J., Shankar, R.D. Das, A.K., 2006. An Ontology-Driven Mediator for Querying Time-Oriented Biomedical Data. *Proceedings of the 19th IEEE International Symposium on Computer-Based Medical Systems*. 264-269.
- O'Connor, M.J., Shankar, R.D. Tu, S.W., Nyulas, C., Musen, M.A., Das, A.K., 2007. Using Semantic Web Technologies for Knowledge-Driven Queries in Clinical Trials. *Proceedings of the 11th Conference on Artificial Intelligence in Medicine*.
- Rotrosen, D., Matthews, J.B., Bluestone, J.A., 2002. The Immune Tolerance Network: a New Paradigm for Developing Tolerance-Inducing Therapies. *Journal of Allergy and Clinical Immunology*, 110(1):17-23.
- Shahar Y., Musen, M.A.. 1996. Knowledge-Based Temporal Abstraction in Clinical Domains. *Artificial Intelligence in Medicine*, 8:267-298.
- Shankar, R.D., Martins, S.B., O'Connor, M.J., Parrish, D.B., Das, A.K. 2006. Epoch: an ontological framework to support clinical trials management. *Proceedings of the International Workshop on Health Information and Knowledge Management*, 25–32.
- Shankar, R.D., Martins, S.B., O'Connor, M.J., Parrish, D.B., Das, A.K., 2006. Towards Semantic Interoperability in a Clinical Trials Management System. *Proceedings of the Fifth International Semantic Web Conference*. 901-912.
- Sim, I., Olasov, B., Carini, S. 2003. The Trial Bank system: capturing randomized trials for evidence-based medicine. *Proceedings of the American Medical Informatics Association Fall Symposium*. 1076.
- Terenziani, P., 2002. Toward a Unifying Ontology Dealing with Both User-Defined Periodicity and Temporal Constraints About Repeated Events, *Computational Intelligence* 18(3):336-385.
- Weng, C., Kahn, M., Gennari, J.H. 2002. Temporal Knowledge Representation for Scheduling Tasks in Clinical Trial Protocols. *Proceedings of the American Medical Informatics Association Fall Symposium*. 879 – 883.
- Weng, C., Gennari, J.H., Fridsma, D.B. 2007. User-centered semantic harmonization; A case study. *Journal of Biomedical Informatics* 40: 353–364.

MEDLINE ABSTRACTS CLASSIFICATION

Average-based Discrimination for Noun Phrases Selection and Weighting Applied to Categorization of MEDLINE Abstracts

Fernando Ruiz-Rico, Jose Luis Vicedo and María-Consuelo Rubio-Sánchez

University of Alicante, Spain

frr@alu.ua.es, vicedo@dlsi.ua.es, mcrcs7@alu.ua.es

Keywords: Text classification, MEDLINE abstracts.

Abstract: Many algorithms have come up in the last years to tackle automated text categorization. They have been exhaustively studied, leading to several variants and combinations not only in the particular procedures but also in the treatment of the input data. A widely used approach is representing documents as Bag-Of-Words (BOW) and weighting tokens with the TFIDF schema. Many researchers have thrown into precision and recall improvements and classification time reduction enriching BOW with stemming, n-grams, feature selection, noun phrases, metadata, weight normalization, etc. We contribute to this field with a novel combination of these techniques. For evaluation purposes, we provide comparisons to previous works with SVM against the simple BOW. The well known OHSUMED corpus is exploited and different sets of categories are selected, as previously done in the literature. The conclusion is that the proposed method can be successfully applied to existing binary classifiers such as SVM outperforming the mixture of BOW and TFIDF approaches.

1 INTRODUCTION

In order to arrange all data in MEDLINE database, each time a new document is added, it must be assigned to one or several MESH¹ terms. More than 100,000 citations are inserted every year, leading to a tedious task, hard to be completed. During the last decades, an important effort has been focused on developing systems to automate the categorization process. In this context, several statistical and machine learning techniques have been extensively studied. We can emphasize Rocchio's based approaches, Bayesian classifiers, Support Vector Machines, Decision Trees and k-Nearest Neighbours among others (Sebastiani, 1999; Aas and Eikvil, 1999; Yang and Liu, 1999). Most of them treat the classified items as feature vectors, where documents are transformed into vectors using the Bag-Of-Words (BOW) representation, where commonly each feature corresponds to a single word or token.

At a first sight, some problems may arise from using the simple BOW. First, a lot of linguistic information is lost, such as word sequence. Also different terms have different importance in a text, so we should think about how to quantify the relevance of a feature so that we have a valid indicative of the de-

gree of the information represented. From an intuitive point of view, a simple consideration of phrases as features may increment the quality and quantity of information contained by feature vectors. For example, the expression "heart diseases" loses its meaning if both words are treated separately. Moreover, we can associate to each phrase sophisticated weights containing some statistical information such as the number of occurrences in a document, or within the whole training set or even how the phrase is distributed among different categories.

The paper is organized as follows. First, we have a look at previous efforts on the same matter by reviewing the literature and pointing out some relevant techniques for feature selection and weighting. Second, we try to remark the most important characteristics of our algorithm by explaining the intuitions that took us to carry out our experiments. Third, the details of the investigation are given by providing a full description of the algorithm and the evaluation procedure. Finally, several results and comparisons are presented and discussed.

2 RELATED WORK

The above observations have led numerous researchers to focus on enriching the BOW model for

¹Medical Subject Headings. More information in www.nlm.nih.gov/mesh/meshhome.html

many years. Most of them have experimented with n-grams (n consequent words) (Scott and Matwin, 1999; Tan et al., 2002; Tesar et al., 2006) and others with itemsets (n words occurring together in a document) (Antonie and Zaane, 2002; Z. Yang and Zhanhuai, 2003; Tesar et al., 2006). In some cases, a significant increment in the performance was reported, but many times only marginal improvement or even a certain decrease was given.

This work proposes a new automatic feature selection and weighting schema. Some characteristics of this approach are based on ideas (noun phrases, meta-information, PoS tagging, stopwords, dimensionality reduction, etc.) that have been successfully tried out in the past (Basili et al., 2000; Granitzer, 2003; Moschitti and Basili, 2004). However, they have never been combined altogether in the way we propose.

For concept detection and isolation we use especial n-grams as features, also known as noun phrases (Scott and Matwin, 1999). The ones we propose are exclusively made of nouns that may or may not be preceded by other nouns or adjectives.

For selecting the relevant expressions of each category, a lot of approaches use only the TF (*term frequency*) and DF (*document frequency*) measures calculated over the whole corpus. This way, most valuable terms occur frequently and have a discriminative nature for their occurrence in only a few documents. We propose using this concept along with the use of TF and DF as individual category membership indicatives, since the greater they are for a particular category corpus, the more related category and term are. Additionally, we define and use *category frequency* (CF) as discriminator among categories. Each of these measures (TF, DF, CF, ...) isolates a set of relevant expressions for a category using its average as discrimination threshold (*average-based discrimination*). The final representative expressions for a category will be obtained by intersecting the sets of relevant expressions for each of the proposed measures.

Finally, relevant expressions are weighted according to a new schema that aggregates all these measures into a single weight.

Normalization over the feature weights has also been proved to be effective (Buckley, 1993). For example, it solves the problem of differences in document sizes: long documents usually use the same terms repeatedly and they have also numerous different terms, increasing the average contribution of their features towards the query document similarity in preference over shorter documents (Singhal et al., 1996).

Once each document in the corpus is represented as features, an algorithm must be provided to get

the final classification. Although quite efficient category ranking approaches can be found in the literature (Ruiz-Rico et al., 2006), they are not suitable for MEDLINE abstracts indexation, where a document must be classified as relevant or not relevant to every particular MESH topic, by taking binary decisions over each of them. For this purpose, SVM is known to be a very accurate binary classifier. Since its complexity grows considerably with the number of features, it is often used together with some techniques for dimensionality reduction.

3 SPECIFIC CONSIDERATIONS

This section highlights some concepts whose analysis is considered important before describing in detail the process of feature selection and weighting.

Parts of Speech and Roots. There are several tools to identify the part of speech of each word and to get its root (word stemming). To achieve the best performance, this paper proposes using both a PoS tagger² and a dictionary³ working together.

Category Descriptors. Training document collections used for classification purposes are usually built in a manual way. That is, human beings assign documents to one or more categories depending on the classification they are dealing with. To help this process, and to be sure that different people use a similar criteria, each class is represented by a set of keywords (*category descriptors*) which identifies the subject of the documents that belong to that category. A document containing some of these keywords should reinforce its relation with particular categories.

Nouns and Adjectives. There are types of words whose contribution is not important for classification tasks (e.g. articles or prepositions) because concepts are typically symbolized in texts within noun phrases (Scott and Matwin, 1999). Also, if we have a look at the category descriptors, we can observe that almost all the words are nouns and adjectives. So, it makes sense to think that the word types used to describe the subject of each category should be also the word types to be extracted from the training documents to identify the category they belong to.

We must assume that it is almost impossible to detect every noun phrase. Moreover, technical corpus are continuously being updated with new words

²SVMTool (Màrquez and Giménez, 2004)

³www-formal.stanford.edu/jsierra/cs193l-project/morphological-db.lisp

and abbreviations. We propose considering these unknown terms as nouns because they are implicitly uncommon and discriminative.

Words and Expressions. When a word along with its adjoining words (a phrase) is considered towards building a category profile, it could be a good discriminator. This tight packaging of words could bring in some semantic value, and it could also filter out words occurring frequently in isolation that do not bear much weight towards characterizing that category (Kongovi et al., 2002). Moreover, it may be useful to group the words so that the number of terms in an expression (TL or text length) can be taken as a new relevance measure.

Non-descriptive expressions. The presence of neutral or void expressions can be avoided by using a fixed list of stopwords (Granitzer, 2003). However, if we only have a general list, some terms may be left out of it. We show that building this list automatically is not only possible but convenient.

Document's Title. The documents to be categorized have a title which briefly summarizes the contents of the full document in only one sentence. Some algorithms would discard an expression that only appears once or twice in a couple of titles because its relevance cannot be confirmed⁴. This paper proposes not only not discarding it, but giving it more importance.

4 FEATURE SELECTION AND WEIGHTING SCHEMA

There are two main processes involved in the task of building the category prototype vector for each category. First, the training data is analyzed in order to detect and extract the most relevant expressions (*expression selection*). These expressions will be used as dimensions of the category prototype vectors. Second, the category prototypes are weighted according to the training set (*expression weighting*).

4.1 Average-based Discrimination

An average or central tendency of a set (list) of values refers to a measure of the "middle value" of the data set. In our case, having a set E of n expressions where each expression is weighted according to a measure

⁴A threshold of 3 is usually chosen (Granitzer, 2003), which means that terms not occurring at least within 3 documents are discarded before learning

W ($\{w_1 \dots w_n\}$), the average of the set E for the measure W (that we denote as \bar{W}) is defined as the arithmetic mean for the values w_i as follows:

$$\bar{W} = \frac{\sum_{i=1}^n w_i}{n}$$

Average discrimination uses the average of a measure W over a set E as threshold for discarding those elements of E whose weight w_i is higher (*H-Average discrimination*) or lower (*L-Average discrimination*) than \bar{W} depending on the selected criteria. In the context of this work, this technique will be applied on different measures for selecting the most representative or discriminative expressions from the training data.

4.2 Expression Selection

The most relevant expressions are selected from the training data by using the *average discrimination measure* of different characteristics as cutting threshold.

4.2.1 Selecting Valid Terms

This process detects and extracts relevant expressions from each document as follows:

1. The words are reduced to their roots.
2. Only nouns and adjectives are taken into consideration. Any other part of speech (verbs, prepositions, conjunctions, etc.) is discarded. For this purpose, a PoS tagger and a dictionary are used. The words which are not found in the dictionary are considered to be nouns.
3. Sentences are divided into expressions: sequences of nouns or adjectives terminating in a noun. In regular expression form this is represented as "{Adjective, Noun}* Noun". For instance, the expressions extracted from "Ultrasound examinations detect cardiac abnormalities" are:

ultrasound	cardiac abnormality
ultrasound examination	abnormality
examination	

This process will give us the set of valid terms (VT) in the whole collection. From now on the words 'term' and 'expression' are used interchangeably.

4.2.2 Computing Term, Document and Category Frequencies

Our starting point is m training collections, each one containing the training documents belonging to each category. Every subset is processed separately to compute the frequencies for each expression in all the

categories. This way, the values required by the algorithm to get the final weights could be put in a matrix where columns correspond to expressions and rows correspond to categories:

	e_1	e_2	...	e_n	
c_1	TF_{11}, DF_{11}	TF_{12}, DF_{12}	...	TF_{1n}, DF_{1n}	N_1
c_2	TF_{21}, DF_{21}	TF_{22}, DF_{22}	...	TF_{2n}, DF_{2n}	N_2
...
c_i	TF_{i1}, DF_{i1}	TF_{i2}, DF_{i2}	...	TF_{in}, DF_{in}	N_i
...
c_m	TF_{m1}, DF_{m1}	TF_{m2}, DF_{m2}	...	TF_{mn}, DF_{mn}	N_m
	CF_1	CF_2	...	CF_n	

where:

- c_i = category i .
- n = number of expressions extracted from all the training documents.
- m = number of categories.
- TF_{ij} = Term Frequency of the expression e_j , that is, number of times that the expression e_j appears in all the training documents for the category c_i .
- DF_{ij} = Document Frequency of the expression e_j , that is, number of training documents for the category c_i in which the expression e_j appears.
- CF_j = Category Frequency of the expression e_j , that is, number of categories in which the expression e_j appears.
- N_i = number of expressions extracted from the training documents of the category c_i .

Every CF_j and N_i can be easily calculated from TF_{ij} by:

$$CF_j = \sum_{i=1}^m x_{ij}; N_i = \sum_{j=1}^n x_{ij}; x_{ij} = \begin{cases} 1 & \text{if } TF_{ij} \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

Across this paper, some examples are shown with the expressions obtained from the training process. The TF and DF corresponding to each expression are put together between brackets, i.e. (TF, DF).

Frequencies for Expressions in the Titles. Some experiments have been performed to get an appropriate factor which increases the weight of the expressions that appear in document's titles (Ruiz-Rico et al., 2006). Doubling TF and DF is proved to be a consideration which optimizes the performance.

4.2.3 Getting the Most Representative Terms (MRT)

The expressions obtained from documents are associated to the categories each document belongs to in the training collection. As a result, we will get m sets of expressions, each one representing a specific category.

For example, after analysing every document associated to the category "Carcinoid Heart Disease", some of the representative expressions are:

carcinoid disease (1,1)	tricuspid stenosis (1,1)
carcinoid heart (26,8)	ventric. enlargement (1,1)
carcin. heart disease (26,8)	ventricular failure (4,1)
carcinoid syndrome (8,3)	ventricular volume (1,1)
carcinoid tumour (3,2)	ventric. vol. overload (1,1)

For each category, we have to select the terms which best identify each category. Three criteria are used to carry out this selection:

- *Predominance inside the whole corpus.* The more times a term occurs inside the full training collection, the more important it is. L-Average discrimination using TF and DF over all the expressions in the corpus (\overline{TF} , \overline{DF}) is used to identify and select the best terms (BT) across the whole corpus.
- *Discrimination among categories.* The more categories a term represents, the less discriminative it is. Expressions appearing in more than half of the categories are not considered discriminative enough. Some authors use fixed stopword lists (Granitzer, 2003) for discarding expressions during the learning and classification processes. Our approach produces this list automatically so that it is adjusted to the number of categories, documents and vocabulary of the training collection.

In this case, the set of category discriminative terms (CDT) for a category is obtained by removing expressions that are representative in more than half of the categories. That is, for every category, an expression e_j will be removed if:

$$CF_j > (m/2 + 1)$$

where m stands for the number of categories.

- *Predominance inside a specific category.* The more times a term occurs inside a category, the more representative it is for that particular category. L-Average discrimination using TF and DF values over all the expressions in each category i (\overline{TF}_i , \overline{DF}_i) are used to identify the best terms in a category (BTC_i).

So, we propose using these TF, DF and CF measures for dimensionality reduction as follows. For each category i :

1. Select the set of terms that are predominant inside the corpus (BT).
2. Select the set of terms that are discriminant among categories (CDT).
3. Select the set of terms that are predominant into this category (BTC_i).

The most representative terms of the category ($MRTC_i$) are obtained from the intersection of the three enumerated sets of terms:

$$\{MRTC\}_i = \{BT\} \cap \{CDT\} \cap \{BTC\}_i$$

As a result, we will get a subset of expressions for each category. For example, the category ‘‘Carcinoid Heart Disease’’ is identified by the following expressions:

carcinoid (44,8)	heart disease (40,9)
carcinoid heart (26,8)	tricuspid valve (11,5)
carcinoid heart disease (26,8)	

4.3 Expression Weighting

At this point, we have the most relevant terms for each of the m categories in the training set. These expressions are now weighted in order to measure their respective importance in a category. This process is accomplished as follows.

4.3.1 Normalization

The corpus of each category has its own characteristics (e.g. different number of training documents, longer or shorter expressions). So, we should not use the TF, DF and TL values directly obtained from the corpus. They can be normalized so that the final weights do not depend on the size of each category’s training set neither on the differences on the averaged length over the representative expressions.

As also stated in (Singhal et al., 1996), we consider that expressions whose frequencies and lengths are very close to the average, are the most appropriate, and their weights should remain unchanged, i.e. they should get unit or no normalization. By selecting an average normalization factor as the pivot, normalized values for TF, DF and TL (TFn, DFn and TLn) are calculated in terms of proportion between the total values and the average over all the expressions in the category:

$$TFn_{ij} = \frac{TF_{ij}}{TF_i} ; DFn_{ij} = \frac{DF_{ij}}{DF_i} ; TLn_{ij} = \frac{TL_j}{TL_i}$$

where i stands for the category c_i and j for the expression e_j respectively.

Normalized values higher than 1 indicate relevance higher than the average, therefore they point to quite significant expressions.

4.3.2 Expressions Matching Category Descriptors

Normalized values measure how much a term stands out over the average. Since category descriptors are

special expressions which can be considered more important than the average, if an expression e_j contains some of the category descriptors of c_i , its normalized frequencies and length should be 1 or higher. To assure this, TFn, DFn and TLn are set to 1 for category descriptors with normalized weights lower than 1.

4.3.3 Weighting

All the proposed values are put together to get a single relevance measure (weight). The proposed weighting schema contains much more information than the common TFIDF approach. Usually, a single set of features is extracted from the training data, and each feature is assigned a single weight. We extract an individual set of features per category, obtaining also different weights for the same expression in different categories.

Every expression e_j is weighted for each of the categories c_i according the following formula:

$$w_{ij} = \frac{(TFn_{ij} + DFn_{ij}) \cdot TLn_{ij} \cdot TFnew_j}{CF_j}$$

where $TFnew_j$ stands for the single number of times that the expression e_j appears in the current document which is being represented as a vector. The greater w_{ij} becomes, the more representative e_j is for c_i . By following the intuitions explained in section 3, this equation makes the weight grow proportionally to the term length and frequencies and makes it lower when the term is more distributed among the different categories.

Since the goal is building a binary classifier, we must have a class c_p representing the positive samples in the training set. Intuitively, to get an even more separable case, the weights of the expressions representing c_p should be calculated differently from the ones representing other categories. For the latter case, we propose accumulating the weight of the negative classes. More formally, we obtain the weight w_j of the expression e_j as following:

$$w_j = \begin{cases} w_{pj} & \text{if } e_j \in c_p \\ \sum_{i=1}^m w_{ij} \forall i \neq p & \text{otherwise} \end{cases}$$

where w_{pj} stands for weight calculated from the positive samples, and $\sum_{i=1}^m w_{ij} \forall i \neq p$ represents the weight calculated from negative samples.

5 EVALUATION

Comparison to previous works is proposed using SVM against the simple BOW. We have represented

the input data as feature vectors under the proposed schema (noun phrases) to make comparisons using a well-known training corpus such as the OHSUMED collection. Results will show that our method increases substantially the classification performance.

5.1 Classification Algorithm

For more accurate comparisons against previous works (Granitzer, 2003), *SVM^{light}* software (Joachims, 1999) has been used for evaluation purposes as a baseline classifier. All default parameters are selected except the cost-factor (“-j”) (Joachims, 2003), which controls the relative weighting of positive to negative examples, and thus provides a way to compensate for unbalanced classes. Leave-one-out cross-validation (LOO) (turned on by “-x 1” parameter) is used to compute a training set contingency table corresponding to each setting of “-j”. *SVM^{light}* is run multiple times for each category, once for each of the resulting values from 0.25 to 4.0 with 0.25 increments (e.g. 0.25, 0.5, 0.75 ... 3.75, 4.0).

Since SVM yields better error bounds by using euclidean norm (Zu et al., 2003), all feature vectors (both in the training and test set) are normalized to euclidean length 1.

5.2 Data Sets

The OHSUMED collection consists of 348,566 citations from medical journals published from 1987 to 1991. Only 233,445 documents contain a title and an abstract. Each document was manually assigned to one or several topics, selected from a list of 14,321 MESH terms. Since automating this process leads to a quite difficult classification problem, most of the authors use smaller data sets. We have chosen the diseases (Joachims, 1998) and heart diseases sub-trees (Granitzer, 2003).

For the diseases hierarchy, MESH terms below the same root node are grouped, leading to 23 categories. The first 10,000 documents in 1991 which have abstracts are used for training, and the second 10,000 are used for testing.

For the heart diseases sub-tree, the categories which have no training documents are discarded, leaving only 16,592 documents and 102 possible categories. The documents from 1987 to 1990 are used as the training set, and the 1991 ones are used as the test set.

5.3 Evaluation Measures

The algorithm performance has been evaluated through the standard BEP and F1 measures (Joachims, 2000):

$$F_1 = \frac{2 \cdot \text{Recall} \cdot \text{Precision}}{\text{Recall} + \text{Precision}}$$

where recall is defined to be the ratio of correct assignments by the system divided by the total number of correct assignments, and precision is the ratio of correct assignments by the system divided by the total number of the system’s assignments. The precision and recall are necessarily equal (BEP) when the number of test examples predicted to be in the positive class equals the true number of positive test examples.

The relevant list of topics for each category is evaluated first, and the average performance score is calculated for all documents (micro-averaged) and for all categories (macro-averaged).

5.4 Relevance of the Parameters for the Classification Task

To compute the relevance of the parameters, the micro-averaged F1 performance is obtained from the original algorithm. After removing each parameter individually, the evaluation is performed again and the percentage of deterioration from the original algorithm is calculated.

Figure 1 reflects the influence of the main characteristics for the final categorization results. The following points describe the conditions applied for the different evaluations along with their associated labels in this figure:

- **Phrases:** Expressions are made of single words.
- **Phr.PoS:** Expressions are made of one single word of any part of speech (no dictionary nor PoS tagger are used).
- **Titles:** Expressions in the titles have the same weight as the other ones.
- **Cat.Des.:** Category descriptors do not have any influence for weighting the expressions.
- **TF-DF, TFnew, CF, TL:** TFn_{ij} , DFn_{ij} , $TFnew_j$, CF_j and TLn_j respectively do not have any effect during the weighting process. This is achieved by modifying the equation given in section 4.3.3 to omit in each case the indicated value:

$$TF - DF \Rightarrow w_{ij} = \frac{TLn_{ij} \cdot TFnew_j}{CF_j}$$

$$TFnew \Rightarrow w_{ij} = \frac{(TFn_{ij} + DFn_{ij}) \cdot TLn_{ij}}{CF_j}$$




Figure 1: Deterioration in the performance after removing each parameter. Parameters are put from left to right in increasing order, from the least to the most relevant one. Tests performed over the diseases sub-tree data set.

Table 1: Number of features in relation with micro-averaged F1 performance. Results obtained over the 23 diseases categories.

	# features	F1
Noun words	2055	63.9
Any words	2221	66.0
Noun phrases	24823	68.6

$$CF \Rightarrow w_{ij} = (TFn_{ij} + DFn_{ij}) \cdot TLn_{ij} \cdot TFnew_j$$

$$TL \Rightarrow w_{ij} = \frac{(TFn_{ij} + DFn_{ij}) \cdot TFnew_j}{CF_j}$$

Figure 1 indicates how relevant each parameter is for the whole categorization process. It shows the percentage of deterioration in the performance when each parameter is removed from the algorithm.

The TF-DF measures lead the graph, meaning that term frequencies have a crucial significance as known from many other previous works. The use of phrases or noun phrases instead of single words of any type is the second most important parameter. The increment of the weights for those expressions in titles also improves significantly the performance. Category Frequency is the fourth most important parameter, which confirms that the more categories an expression represents, the less discriminative it is.

Documents from other corpora may be better represented by taking only single words as features and using simple weighting schemas (Dumais et al., 1998). However, it is not the same for OHSUMED. As far as we know, the results here presented are the best ever achieved, leading us to the conclusion that for some type of data such as MEDLINE documents, we should try to increment the number of features and the amount of information they contain, as confirmed in table 1.

6 RESULTS

Next tables show the results obtained in previous works followed by the results achieved by applying the new proposed algorithm for feature selection and weighting over the same training and test sets. The best values are in boldface.

Table 2: Break even point on 5 most frequent categories and micro-averaged performance over all 23 diseases categories (Joachims, 1998).

	SVM (words)	SVM (noun phrases)
Pathology	58.1	52.7
Cardiovascular	77.6	80.9
Immunologic	73.5	77.1
Neoplasms	70.7	81.5
Digestive system	73.8	77.5
Micro avg (23 cat.)	66.1	68.6

Table 3: Averaged F1 performance over 102 heart diseases categories (Granitzer, 2003).

	SVM (words)	SVM (noun phrases)
Micro avg	63.2	69.9
Macro avg	50.3	55.5

Table 2 shows the micro-averaged BEP performance calculated over the 23 diseases categories. Noun phrases gets a global 3.8% improvement, also outperforming almost all categories individually.

Table 3 contains both the micro and macro averaged F1 performance over the 102 categories of the heart diseases sub-tree. For this corpus we have achieved more than 10% improvements (10.6% for micro and 10.3% for macro measures respectively).

7 CONCLUSIONS

Using a proper feature selection and weighting schema is known to be decisive. This work proposes a particular way to choose, extract and weight special n-grams from documents in plain text format so that we get a high performing representation. Moreover, the new algorithm is fast, easy to implement and it contains some necessary adjustments to automatically fit both existing and incoming MEDLINE documents.

REFERENCES

- Aas, K. and Eikvil, L. (1999). Text categorisation: A survey. Technical report, Norwegian Computer Center.
- Antonie, M. and Zaane, O. (2002). Text document categorization by term association. *IEEE International Conference on Data Mining (ICDM)*, pages 19–26.
- Basili, R., Moschitti, A., and Pazienza, M. T. (2000). Language-sensitive text classification. In *Proceeding of RIAO-00, 6th International Conference "Recherche d'Information Assistee par Ordinateur"*, pages 331–343, Paris, FR.
- Buckley, C. (1993). The importance of proper weighting methods. In M. Bates, editor, *Human Language Technology Morgan Kaufman*.
- Dumais, S. T., Platt, J., Heckerman, D., and Sahami, M. (1998). Inductive learning algorithms and representations for text categorization. In Gardarin, G., French, J. C., Pissinou, N., Makki, K., and Bouganim, L., editors, *Proceedings of CIKM-98, 7th ACM International Conference on Information and Knowledge Management*, pages 148–155, Bethesda, US. ACM Press, New York, US.
- Granitzer, M. (2003). Hierarchical text classification using methods from machine learning. Master's thesis, Graz University of Technology.
- Joachims, T. (1998). Text categorization with support vector machines: learning with many relevant features. In Nédellec, C. and Rouveirol, C., editors, *Proceedings of ECML-98, 10th European Conference on Machine Learning*, pages 137–142, Chemnitz, DE. Springer Verlag, Heidelberg, DE. "Lecture Notes in Computer Science" series, number 1398.
- Joachims, T. (1999). *Making large-Scale SVM Learning Practical. Advances in Kernel Methods - Support Vector Learning*. Software available at <http://svmlight.joachims.org/>.
- Joachims, T. (2000). Estimating the generalization performance of a svm efficiently. In Langley, P., editor, *Proceedings of ICML-00, 17th International Conference on Machine Learning*, pages 431–438, Stanford, US. Morgan Kaufmann Publishers, San Francisco, US.
- Joachims, T. (2003). *Support Vector and Kernel Methods*. SIGIR 2003 Tutorial.
- Kongovi, M., Guzman, J. C., and Dasigi, V. (2002). Text categorization: An experiment using phrases. In *Proceedings of ECIR-02, 24th European Colloquium on Information Retrieval Research*, pages 213–228, Glasgow, UK.
- Màrquez, L. and Giménez, J. (2004). A general pos tagger generator based on support vector machines. *Journal of Machine Learning Research*. Software available at www.lsi.upc.edu/nlp/SVMTool.
- Moschitti, A. and Basili, R. (2004). Complex linguistic features for text classification: A comprehensive study. In *Proceedings of ECIR-04, 26th European Conference on Information Retrieval Research*, pages 181–196. Springer Verlag, Heidelberg, DE.
- Ruiz-Rico, F., Vicedo, J. L., and Rubio-Sánchez, M.-C. (2006). Newpar: an automatic feature selection and weighting schema for category ranking. In *Proceedings of DocEng-06, 6th ACM symposium on Document engineering*, pages 128–137.
- Scott, S. and Matwin, S. (1999). Feature engineering for text classification. In Bratko, I. and Dzeroski, S., editors, *Proceedings of ICML-99, 16th International Conference on Machine Learning*, pages 379–388, Bled, SL. Morgan Kaufmann Publishers, San Francisco, US.
- Sebastiani, F. (1999). A tutorial on automated text categorisation. In Amandi, A. and Zunino, R., editors, *Proceedings of ASAI-99, 1st Argentinian Symposium on Artificial Intelligence*, pages 7–35, Buenos Aires, AR. An extended version appears as (?).
- Singhal, A., Buckley, C., and Mitra, M. (1996). Pivoted document length normalization. *Department of Computer Science, Cornell University, Ithaca, NY 14853*.
- Tan, C.-M., Wang, Y.-F., and Lee, C.-D. (2002). The use of bigrams to enhance text categorization. *Information Processing and Management*, 38(4):529–546.
- Tesar, R., Strnad, V., Jezek, K., and Poesio, M. (2006). Extending the single words-based document model: a comparison of bigrams and 2-itemsets. In *DocEng '06: Proceedings of the 2006 ACM symposium on Document engineering*, pages 138–146, New York, NY, USA. ACM Press.
- Yang, Y. and Liu, X. (1999). A re-examination of text categorization methods. In Hearst, M. A., Gey, F., and Tong, R., editors, *Proceedings of SIGIR-99, 22nd ACM International Conference on Research and Development in Information Retrieval*, pages 42–49, Berkeley, US. ACM Press, New York, US.
- Z. Yang, Z. Lijun, Y. J. and Zhanhuai, L. (2003). Using association features to enhance the performance of naive bayes text classifier. *Fifth International Conference on Computational Intelligence and Multimedia Applications, ICCIMA '03*, pages 336–341.
- Zu, G., Ohyama, W., Wakabayashi, T., and Kimura, F. (2003). Accuracy improvement of automatic text classification based on feature transformation. In *Proceedings of DOCENG-03, ACM Symposium on Document engineering*, pages 118–120, Grenoble, FR. ACM Press, New York, US.

TOWARDS ON-DEMAND BIOMEDICAL KNOWLEDGE EXTRACTION

Vincenzo Lanza

Anesthesia Department, Buccheri La Ferla Hospital, Fatebenefratelli, Palermo, Italy
lanza@unipa.it

M. Ignazia Cascio

Sicilian Center for Training and Research in Public Health (CEFPAS), Caltanissetta, Italy
lanza@unipa.it

Chun-Hsi Huang

Dept. of Computer Science and Engineering, University of Connecticut, Storrs, CT 06269, USA
huang@cse.uconn.edu

Keywords: Knowledge Grids, Biomedical Computing.

Abstract: This paper outlines a UMLS-compatible distributed genomic semantic network. The system aims at providing cooperative reasoning on distributed genomic information, complying with the UMLS concept representation, from distributed repositories. The distributed semantic network has currently incorporated most of the 871,584 concepts (named by 2.1 million terms) of the 2002 version UMLS Metathesaurus, with inter-concept relationships across multiple vocabularies and concept categorization supported. Modern information and compute infrastructure is incorporated to allow seamless access to geographically dispersed users.

1 INTRODUCTION

The complete sequencing of numerous genomes has stimulated new cross-domain and cross-discipline research topics. Computationally, researchers have been exploring the massive genomic and proteomic information, attempting to generate new hypotheses for gene/protein functions, as well as novel targets for the development of insecticides, antibiotics, antiviral drugs, and health related drugs. *Semantically*, researchers study biological information from individual (clinical practice) to the population level (social health-care), as well as the infrastructure for high-performance, automated integration and analysis of these information, in an attempt to better individual and public health care.

It is crucial in most of these novel researches that the massive genomic data produced are well represented so that useful biological information may be efficiently extracted. A useful tool for effective knowledge representation is the *semantic network* system (Lee et al., 2003). A semantic network is a conceptual model for knowledge representation, in which the knowledge entities are represented by nodes (or vertices), while the edges (or arcs) are the relations between entities (Cerccone, 1992; Fahlman,

1982; Brachman and Schmolze, 1985; Shapiro and The SNePS Implementation Group, 1998; Chung and Moldovan, 1993; Surdeanu et al., 2002; Moldovan et al., 1992; Evett et al., 1991; Stoffel et al., 1996). A semantic network is an effective tool, serving as the backbone knowledge representation system for genomic, clinical and medical data. Usually these knowledge bases are stored at locations geographically distributed. This highlights the importance of an efficient distributed semantic network system enabling distributed knowledge integration and inference.

The semantic network is a key component of the *Unified Medical Language System* (UMLS) project initiated in 1986 by the U.S. National Library of Medicine (NLM). The goal of the UMLS is to facilitate associative retrieval and integration of biomedical information so researchers and health professionals can use such information from different (readable) sources (Lindberg et al., 1993). The UMLS project consists of three core components: (1) the **Metathesaurus**, providing a common structure for more than 95 source biomedical vocabularies. It is organized by concept, which is a cluster of terms, *e.g.*, synonyms, lexical variants, and translations, with the same meaning. (2) the **Semantic Network**, categorizing these concepts

by semantic types and relationships, and (3) the **SPECIALIST lexicon** and associated lexical tools, containing over 30,000 English words, including various biomedical terminologies. Information for each entry, including base form, spelling variants, syntactic category, inflectional variation of nouns and conjugation of verbs, is used by the lexical tools. The 2002 version of the Metathesaurus contains 871,584 concepts named by 2.1 million terms. It also includes inter-concept relationships across multiple vocabularies, concept categorization, and information on concept co-occurrence in MEDLINE.

2 THE PILOT SYSTEM

We are currently developing a UMLS-compatible distributed genomic semantic network. This system aims at providing cooperative reasoning on distributed genomic information, complying with the UMLS concept representation, from distributed repositories. Representative inference rules (path-based) and commands (SNePS-like (Moldovan et al., 2003)) are briefed in Appendices A and B.

The infrastructure of the cooperative software components is extended from the TROJAN system (Lee et al., 2004; Lee and Huang, 2004). The pilot system emphasizes the task-based and message-driven model to exploit parallelism at both task and data levels. The system also features *multi-threading* and *task migration* to support communication latency hiding and load balancing, respectively. In the task model, queries are decomposed into tasks and distributed among processors for execution. When a task is completed, a message is generated to either spawn new tasks or trigger further processing, depending on the property and current status of the task. This process is carried out by two collaborating components: the *host system* and the *slave system*. The host system interacts with users and processes information for the slave system, while the slave system executes compute tasks.

The host system is composed of the following major components. The *language front-end* interacts with the user and decomposes the commands into either knowledge or tasks. All the preprocessing and distributing are carried out in the *command processing module*. The *object-oriented packing module* is the communication channel between processors. When the slave module finishes a query, the answer messages are then sent back to the *host answer processing module* of the host system to be merged into a final inference conclusion. Some knowledge is kept in the *host knowledge base* for simple queries. Fig. 5

illustrates the host system.

The major components comprising the slave system are as follows. The *shared knowledge management module* stores and exchanges knowledge in the *shared knowledge base*. The *task execution module* is the kernel of task execution. Several sub-modules are embedded in the task execution module, including the *kernel message module*, the *task execution engine*, and the *load balancing module*, etc. The *duplicate checking module* records the answers that have been reached to save repeated executions. The *slave scheduler* schedules task execution and swapping. The *object-oriented packing system* is similar to that of the host. The slave system is depicted in Fig. 6.

Commands in our semantic network system are generally categorized into three groups: (1) network building (e.g. build and assert, etc.), (2) inferencing (e.g. find, findassert, etc.) and (3) others (e.g. nodeset operation commands, etc.). Commands in groups (1) and (2) usually need to communicate with slave PEs, while those in (3) are answered directly inside the host module. Our system provides three commands, build, assert and add, to construct the semantic network. The syntax of these commands are listed below:

- build: (build {*relation nodeset*}*)
- assert: (assert {*relation nodeset*}* *context-specifier*)
- add: (add {*relation nodeset*}* *context-specifier*)

For example, the command

```
(assert member Saccharomyces-cerevisiae
      class yeast)
```

defines the concept “Saccharomyces-cerevisiae is yeast”. In the system, two base nodes Saccharomyces-cerevisiae and yeast are generated by the command. The molecular node M1 (index depending on the current knowledge base state) is generated by the system, where “!” stands for the “assertion” concept. Two forward links member and class are defined by the user, two reverse links member- and class-, indicated by dash lines, are generated automatically. Hierarchical concepts can be constructed similarly by following the links of “subclass-” and “supclass”.

To sum up, the network building commands put a node into the network with an arc labeled *relation* to each node in the following *nodeset*, and returns the newly built node. An attempt to build a currently existing node will immediately return such an existing node. build creates an unasserted node unless an asserted node exists in the network with a superset of the relations of the new node, in which case the new node

is also asserted. *assert* is just like *build*, but creates the node with assertion. *add* acts like *assert*, but in addition triggers forward inference. *relation* has to be a unit-path and *non converse*. The converse relation *relation*- that connects each node of the nodeset to the built node is constructed implicitly by the system.

A heuristic approach is used to partition the semantic network to get around the *NP*-hardness of optimal partitioning (Cormen et al., 2000; Garey et al., 1976). Starting from PE_1 as the initial target PE, while a network construction command is issued, the host sends the newly built nodes to the target PE until a certain number of nodes have accumulated and then cyclically shifts to the next PE as the new target.

Several (path-based) inference commands are provided by our system, including the *find* family (*find*, *findassert*, *findbase*, *findconstant*, *findpattern* and *findvariable*). While a query is made, corresponding tasks are generated by the *task preprocessor* under the command of the parser in the language front end, and then stored in the host *task queue* temporarily while waiting to be dispatched by the *task distributor*. These query tasks are split according to the implied parallelism of the command. For example, when a path-based query is made, the command is usually in the format of (*find* {*path nodeset*}*), which is equivalent to $\bigcap_{i=1}^n \bigcup_{j=1}^{f(i)} (\text{find } path_i \text{ node}_{i,j})$, where the \cap and \cup are the nodeset intersection and union, respectively. These (sub)query tasks are formed and later dispatched.

Path-based inference is the fundamental inference mechanism of all semantic networks. By tracing the arcs between nodes, new knowledge can be derived. In our system, the relation between two nodes can be either explicit (direct arc between two nodes), or implicit (an arc across several intermediate nodes). The implicit relation is defined by the command *define-path*.

The command *find*, designed for path-based inference queries, has the following syntax:

(*find* {*path nodeset*}*).

This command returns a set of nodes such that each node in the set has every specified *path* going from it to at least one node in the accompanying *nodeset*. When the command

(*find subclass* (human animal))

is issued, the system answers (M2 M3) since M2 and M3 each has an edge *subclass* to *either* node human *or* animal.

The distributed semantic network has currently incorporated most of the 871,584 concepts (named by 2.1 million terms) of the 2002 version

UMLS Metathesaurus, with inter-concept relationships across multiple vocabularies and concept categorization supported.

3 BASIC INFRASTRUCTURE

Modern Grid technology represents an emerging and expanding instrumentation, computing, information and storage platform that allows geographically distributed resources, which are under distinct control, to be linked together in a transparent fashion (Berman et al., 2003; Foster and Kesselman, 1999). The power of the Grids lays not only in the aggregate computing ability, data storage, and network bandwidth that can readily be brought to bear on a particular problem, but also on its ease of use. After a decade's research effort, Grids are moving out of research laboratories into early-adopter production systems, such as the Computational Grid for certain computation-intensive applications, the Data Grid for distributed and optimized storage of large amounts of accessible data, as well as the Knowledge Grid for intelligent use of the Data Grid for knowledge creation and tools to all users.

Here we refer to the Cross-Campus (or Continent) Computational Grid as the C^3 -Grid; and the Cross-Campus (or Continent) Data Grid as the C^2D -Grid.

The development of the C^3 -Grid portal focuses on the establishment of a robust set of APIs (Application Programming Interfaces). The implementation of C^3 -Grid is largely based on the Globus Toolkit middleware (2.2.4) The web portal is served by an Apache HTTP Server located at the University of Connecticut. The C^3 -Grid database regularly aggregates compute platform statistics such as job status, backfill availability, queue schedule, as well as production rates. The job monitoring system provides real-time snapshots of critical computational job metrics stored in a database and presented to the user via dynamic web pages. Computation jobs are classified into a few classes, each with a pre-specified priority. Statistics for each job class are created in a real-time manner so as to provide intelligent management of resources.

The C^2D -Grid adds another dimension of functionality to the C^3 -Grid in terms of efficient management of the often-curated biomedical knowledge-base. Our goal is to transparently and efficiently manage the biomedical knowledge-base distributed across the participating campuses, providing access via a uniform interface (web-portal). Basic file management functions are available via a user-friendly and platform-independent interface. Basic file transfer, editing and search capabilities are available via a uni-

form interface. The logical display of files for a given (local or remote) user is also available. Data/file migration is implemented to minimize the bandwidth consumption and to maximize the storage utilization rate on a per user basis.

Additional technical and configuration details in regards to the compute and data grid infrastructure are elided, according to reviewers' comments and the page limit.

4 WORKFLOW CONTROL

The design of the our workflow control toolkit over the C^3 -Grid is largely based on the Genome Analysis and Database Update system (GADU) (Pearson, 1994; Shpaer et al., 1996; Mulder, 2003; Bateman et al., 2002; Henikoff et al., 1999; Pearl et al., 2003; Sulakhe et al.,). GADU has successfully used Grid resources with different architectures and software environments like the 64-bit processors in TeraGrid and 32-bit processors in the Open Science Grid or DOE Science Grid¹.

The opportunistic availability and the different architectures and environments of these resources make it extremely difficult to use them simultaneously through a single common system. GADU addresses these issues by providing a resource-independent system that can execute the bioinformatics applications as workflows simultaneously on these heterogeneous Grid resources. and is easily scalable to add new Grid resources or individual clusters into its pool of resources, thus providing more high-throughput computational power to its scientific applications. The workflow control toolkit has wide applications in genomics as the interpretation of every newly sequenced genome involves the analysis of sequence data by a variety of computationally intensive bioinformatics tools, the execution of result and annotation parsers, and other intermediate data-transforming scripts.

Our toolkit will act as a gateway to the C^3 -Grid and the C^2D -Grid, handling all the high-throughput computations necessary for knowledge inference and extraction from our semantic network. Analogous to GADU, our workflow control toolkit will be implemented in two modules, an *analysis* server and an *update* server. The analysis server automatically creates workflows in the abstract Virtual Data Language, based on predefined templates that it executes on distributed Grid resources. The update server updates the integrated knowledge-base with recently changed data from participating sites

The toolkit will execute its parallel jobs simultaneously on different Grid resources. It expresses the workflows in the form of a directed acyclic graph (DAG) and executes it on a specified Grid site using Condor-G (Frey et al., 2002). The toolkit will use the GriPhyN Virtual Data System (Foster et al., 2002) to express, execute, and track the results of the workflows that help in using the grid resources.

To sum up, this workflow controller will provide a resource-independent configuration to execute the workflows over the C^3 -Grid. It can submit jobs remotely to a resource, as long as the resource provides a Globus GRAM interface (*e.g.*, the Jazz cluster). All the transformations of a workflow are expressed as Condor submit files and a DAG using Pegasus. The Condor-G submits the workflow to a remote resource using the GRAM interface and also monitors the workflow. The toolkit will also automatically manage the dynamic changes in the state of the Grid resources using monitoring and information services along with the authentication and access models used at different Grids.

5 CONCLUDING REMARKS

Biomedical research increasingly relies on globally distributed information and knowledge repositories. The quality and performance of future computing and storage infrastructure in support of such research depends heavily on the ability to exploit these repositories, to integrate these resources with local information processing environments in a flexible and intuitive way, and to support information extraction and analysis in a timely and on-demand manner.

This paper outlines a UMLS-compatible distributed genomic semantic network. The system provides cooperative reasoning on distributed genomic information, complying with the UMLS concept representation, from distributed repositories. The distributed semantic network has currently incorporated most of the 871,584 concepts (named by 2.1 million terms) of the 2002 version UMLS Metathesaurus, with inter-concept relationships across multiple vocabularies and concept categorization supported.

The knowledge database and semantic network are to be installed within a cross-campus data grid framework. The knowledge inference will be decomposed into sub-tasks and distributed across the participating compute nodes for computation.

¹<http://www.doesciencegrid.org>

ACKNOWLEDGEMENTS

The workflow control for the C^3 -Grid job execution was based on an earlier collaborative project and discussions on GADU with Dr. N. Maltsev at the Argonne National Lab, US.

REFERENCES

- Bateman, A., Birney, E., Cerruti, L., Durbin, R., Etwiller, L., Eddy, S., Griffiths Jones, S., Howe, K., Marshall, M., and Sonnhammer, E. (2002). The Pfam protein families database. *Nucleic Acids Res.*, 30:276–280.
- Berman, F., Fox, G., and Hey, T. (2003). Grid Computing: Making the Global Infrastructure a Reality. *John Wiley & Sons*.
- Brachman, R. J. and Schmolze, J. G. (1985). An Overview of the KL-ONE Knowledge Representation System. *Cognitive Sci.*, 9:171–216.
- Cercone, N. (1992). The ECO Family. *Computers Math. Applic.*, 23(2-5):95–131.
- Chung, S. and Moldovan, D. I. (1993). Modeling Semantic Networks on the Connection Machine. *Journal of Parallel and Distributed Computing*, 17:152–163.
- Cormen, T. H., Leiserson, C. E., and Rivest, R. L. (2000). *Introduction to Algorithms*. McGraw-Hill.
- Evett, M. P., Hendler, J. A., and Spector, L. (1991). Parallel Knowledge Representation on the Connection Machine. *Journal of Parallel and Distributed Computing*, 22:168–184.
- Fahlman, S. E. (1982). *NETL: A System for Representing and Using Real-World Knowledge*. The MIT Press.
- Foster, I. and Kesselman, C. (1999). *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, San Francisco.
- Foster, I., Voekler, J., Wilde, M., and Zhou, Y. (2002). Chimera: A virtual data system for representing, querying, and automating data derivation. In *Proceedings of the 14-th Conference on Scientific and Statistical Database*.
- Frey, J., Tannenbaum, T., Foster, I., Livny, M., and Tuecke, S. (2002). Condor-G: a computation management agent for multi-institutional Grids. *Cluster Computing*, 5:237–246.
- Garey, M. R., Johnson, D. J., and Stockmeyer, L. (1976). Some Simplified NP-Complete Graph Problems. *Theoret. Comput. Sci.*, 1:237–267.
- Henikoff, S., Henikoff, J., and Pietrovski, S. (1999). Blocks+: a non-redundant database of protein alignment blocks derived from multiple compilations. *Bioinformatics*, 15:471–479.
- Lee, C.-W. and Huang, C.-H. (2004). Toward Cooperative Genomic Knowledge Inference. *Parallel Computing Journal*, 30(9-10):1127–1135.
- Lee, C.-W., Huang, C.-H., and Rajasekaran, S. (2003). TROJAN: A Scalable Parallel Semantic Network System. In *Proceedings of the 15th IEEE International Conference on Tools with Artificial Intelligence*, pages 219–223.
- Lee, C.-W., Huang, C.-H., Yang, L., and Rajasekaran, S. (2004). Path-Based Distributed Knowledge Inference in Semantic Networks. *Journal of Supercomputing*, 29(2):211–227.
- Lindberg, D., Humphreys, B., and McCray, A. (1993). The Unified Medical Language System. *Methods Inf. Med.*, 32(4):281–291.
- Moldovan, D., Lee, W., Lin, C., and Chung, M. (1992). SNAP, Parallel Processing Applied to AI. *IEEE Computer*, pages 39–49.
- Moldovan, D. I., Pasca, M., Harabagiu, S., and Surdeanu, M. (2003). Performance Issues and Error Analysis in an Open-Domain Question Answering System. *ACM Tran. on Information Systems*, 21(2):133–154.
- Mulder, N.J., e. a. (2003). The InterPro Database, 2003 brings increased coverage and new features. *Nucleic Acids Res.*, 31:315–318.
- Pearl, F., Bennett, C., Bray, J., Harrison, A., Martin, N., Shepherd, A., Sillitoe, I., Thornton, J., and Orengo, C. (2003). The CATH database: an extended protein family resource for structural and functional genomics. *Nucleic Acids Res.*, 31:452–455.
- Pearson, W. (1994). Using the FASTA program to search protein and DNA sequence databases. *Methods Mol. Biol.*, 24:307–331.
- Shapiro, S. C. and The SNePS Implementation Group (1998). *SNePS-2.4 User's Manual*. Department of Computer Science at SUNY Buffalo.
- Shpaer, E., M., R., Yee, D., Candlin, J., Mines, R., and Hunkapiller, T. (1996). Sensitivity and selectivity in protein similarity searches: a comparison of Smith-Waterman in hardware to BLAST and FASTA. *Genomics*, 38:179–191.
- Stoffel, K., Hendler, J., Saltz, J., and Anderson, B. (1996). Parka on MIMD-Supercomputers. Technical Report CS-TR-3672, Computer Science Dept., UM Institute for Advanced Computer Studies, University of Maryland, College Park.
- Sulakhe, D., Rodriguez, A., D'Souza, M., Wilde, M., Nefedova, V., Foster, I., and Maltsev, N. G. (2002). Gnare: automated system for high-throughput genome analysis with grid computational backend. *J. Clin. Monit. and Comput.*, 19(4).
- Surdeanu, M., Moldovan, D. I., and Harabagiu, S. M. (2002). Performance Analysis of a Distributed Question/Answering System. *IEEE Trans. on Parallel and Distributed Systems*, 13(6):579–596.

APPENDIX

Path Definition Primitives

- $unitpath ::= relation$.
- $unitpath ::= relation^-$: A unit path can be either a relation or a converse of a relation.
- $path ::= unitpath$: A path can be either a unitpath or the composition of various pathes defined by the following definition.
- $path ::= (COMPOSE \{path\}^*)$
If x_1, \dots, x_n are nodes and P_i is a path from x_1 to x_{i+1} , then $(COMPOSE P_1 \dots P_{n-1})$ is a path from x_1 to x_n .
- $path ::= (KSTAR path)$
If path P is composed with itself zero or more times from node x to node y , then $(KSTAR P)$ is a path from x to y .
- $path ::= (KPLUS path)$
If path P is a composed with itself one or more times from node x to node y , then $(KPLUS P)$ is a path from x to y .
- $path ::= (OR \{path\}^*)$
If P_1 is path from node x to node y or P_2 is a path from x to y or \dots or P_n is a path from x to y then $(OR P_1, P_2 \dots P_n)$ is a path from x to y .
- $path ::= (AND \{path\}^*)$
If P_1 is path from node x to node y and P_2 is a path from x to y and \dots and P_n is a path from x to y then $(AND P_1, P_2 \dots P_n)$ is a path from x to y .
- $path ::= (NOT path)$
If there is no path P from node x to node y , then $(NOT P)$ is a path from x to y .
- $path ::= (RELATIVE-COMPLEMENT path path)$
If P is a path from node x to node y and there is no path Q from x to y , then $(RELATIVE-COMPLEMENT P Q)$ is a path from x to y . The situation can be seen in Fig. 1.
- $path ::= (IRREFLEXIVE-RESTRICT path)$
If P is a path from node x to node y , and $x \neq y$, then $(IRREFLEXIVE-RESTRICT P)$ is a path from x to y . The situation can be seen in Fig. 2.
- $path ::= (DOMAIN-RESTRICT (path node) path)$
If P is a path from node x to node y and Q is a path from x to node z , then $(DOMAIN-RESTRICT (Q z) P)$ is a path from x to y . The situation can be seen in Fig. 3.
- $path ::= (RANGE-RESTRICT path (node path))$
If P is a path from node x to node y and Q is a path from y to node z , then $(RANGE-RESTRICT P (Q z))$ is a path from x to y . The situation can be seen in Fig. 4.

- $path ::= (path^*)$
The definition is the same as $(COMPOSE \{path\}^*)$




Figure 1: (RELATIVE-COMPLEMENT P Q).




Figure 2: (IRREFLEXIVE-RESTRICT P).




Figure 3: (DOMAIN-RESTRICT (Q z) P).




Figure 4: (RANGE-RESTRICT P (Q z)).

Grammars

```

⟨sneps_command⟩ ::= ( ⟨path_defn_command⟩ )
                  | ( ⟨file_command⟩ )
                  | ( ⟨delete_command⟩ )
                  | ( ⟨misc_command⟩ )
                  | ⟨multi_node_command⟩
                  | ⟨snepsul_var⟩ ■
⟨path_defn_command⟩ ::= ⇒ DEFINE ⟨unitpath⟩+
                       | UNDEFINE ⟨unitpath⟩+ ■
⟨file_command⟩ ::= ⇒ INNET “⟨string⟩”
                  | OUTNET “⟨string⟩” ■
⟨delete_command⟩ ::= ⇒ RESETNET 'T
                   | RESETNET 'NIL
                   | RESETNET ■
⟨misc_command⟩ ::= ⇒ LISP
                  | ⟨dup_check_command⟩
                  | ⟨load_bal_command⟩

```

<pre> <knowledge_base_command> ■ <dup_check_command> ::= DUPCHECK 'T DUPCHECK 'NIL DUPCHECK ■ <load_bal_command> ::= LOADBAL 'T LOADBAL 'NIL LOADBAL ■ <knowledge_base_command> ::= PRIVATE <integer> PRIVATE ALL CACHE <integer> CACHE ALL ■ <multi_node_command> ::= (<multiple_nodes>) <node_command> ■ <multiple_nodes> ::= <node_command>+ ■ <node_command> ::= (<inference_command>) (<display_command>) (<net_build_command>) (<nodeset_op_command>) <snepsul_var> <multi_dollar_node> <multi_hash_node> <node_name> ■ <inference_command> ::= FIND <multi_path_nodeset> FINDASSERT <multi_path_nodeset> FINDCONSTANT <multi_path_nodeset> FINDBASE <multi_path_nodeset> FINDVARIABLE <multi_path_nodeset> FINDPATTERN <multi_path_nodeset> ■ <display_command> ::= DUMP <multi_node_command> DESCRIBE <multi_node_command> ■ <net_build_command> ::= ASSERT <multi_relation_nodeset> BUILD <multi_relation_nodeset> ■ <nodeset_op_command> ::= & <multi_node_command> <multi_node_command> + <multi_node_command> <multi_node_command> - <multi_node_command> <multi_node_command> = <multi_node_command> <symbol> - <multi_node_command> <unitpathset> </pre>	<pre> > <unitpathset> <symbol> ■ <snepsul_var> ::= * <symbol> ■ <multi_dollar_node> ::= <dollar_node> (<dollar_nodeset>) ■ <dollar_nodeset> ::= <dollar_node>+ ■ <dollar_node> ::= \$ <symbol> ■ <unitpathset> ::= <unitpath> (<multi_unitpath>) ■ <multi_unitpath> ::= <unitpath>+ ■ <node_name> ::= <symbol> <integer> ■ <multi_hash_node> ::= <hash_node> (<hash_nodeset>) ■ <hash_nodeset> ::= <hash_node>+ ■ <hash_node> ::= # <symbol> ■ <question_node> ::= ? <symbol> ■ <multi_relation_nodeset> ::= <relation_nodeset>+ ■ <relation_nodeset> ::= <relation> <multi_node_command> ■ <multi_path_nodeset> ::= <path_nodeset>+ <path_question_node> <multi_path_nodeset> <path_question_node> ■ <path_question_node> ::= <path> <question_node> ■ <path_nodeset> ::= <path> <multi_node_command> ■ <path> ::= <unitpath> (COMPOSE <multi_path>) (KSTAR <path>) (KPLUS <path>) (OR <multi_path>) (AND <multi_path>) (NOT <path>) (RELATIVE_COMPLEMENT <path> <path>) (IRREFLEXIVE_RESTRICT <path>) (DOMAIN_RESTRICT (<path> <node_name>) <path>) (RANGE_RESTRICT <path> (<path> <node_name>)) (<multi_path>) ■ <multi_path> ::= <path>+ ■ <relation> ::= <unitpath> <unitpath>- ■ <unitpath> ::= <symbol> ■ <digit> ::= [0-9] ■ <non_neg_char> ::= [0-9a-zA-Z_] ■ <integer> ::= <digit>+ ■ <symbol> ::= <non_neg_char>+ ■ <string> ::= <char>+ ■ </pre>
--	---

Software Architectures

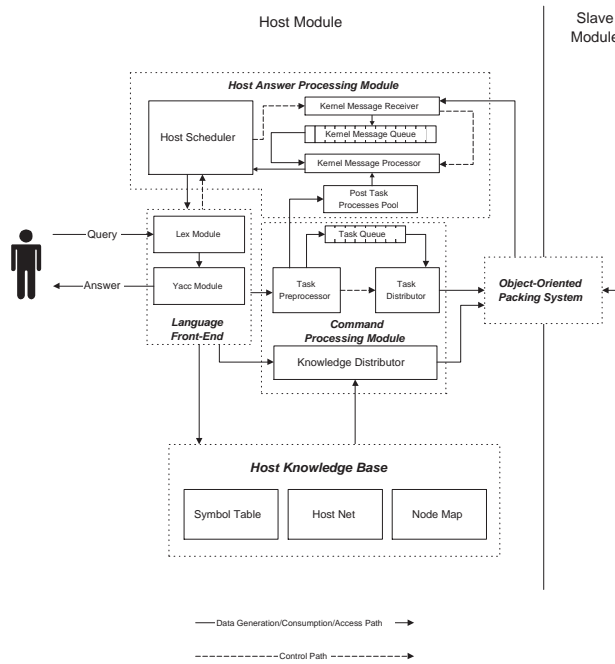


Figure 5: Host System Software Architecture.

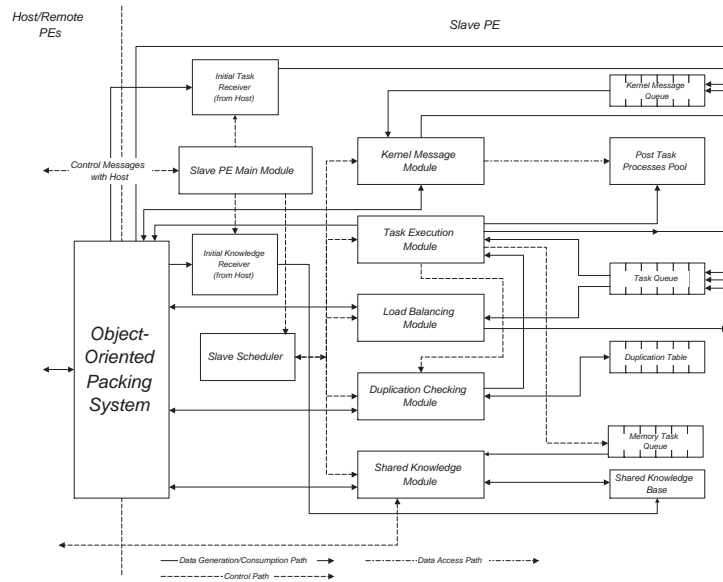


Figure 6: Slave System Software Architecture.

A BPM-BASED MOBILE U-HEALTH SERVICE FRAMEWORK

Dongsoo Han, In-Young Ko, Sungjoon Park, Minkyu Lee

*School of Engineering, Information and Communications University, 119, Munjiro, Yuseong-gu, Daejeon, Korea
dshan@icu.ac.kr, iko@icu.ac.kr, sungjoon@icu.ac.kr, niklaus@icu.ac.kr,*

Suntae Jung

*Samsung Electronics, 416, Maetan-dong, Youngtong-gu, Suwon, Kyunggi-do, Korea
suntae1.jung@samsung.com*

Keywords: e-Health, u-Health service platform, ontology, mobile, BPM.

Abstract: The integration of mobile bio-sensors and cellular phones opens a new horizon for healthcare service. Mobile u-health service, which usually incorporates mobile bio-sensor attached cellular phones, provides users real time healthcare services at the right time in the right manner. While the u-health services may look different from the service point of view, they often share many common features at various levels such as the service structure, unit service, and data levels. Thus, it is necessary to have a common platform on which various u-health services can be developed by effectively sharing and reusing the common features and services rather than developing the services independently from the scratch. In this paper, we propose a mobile u-health platform that provides core functions and facilities to develop mobile u-health services. Main elements of the platform include the u-health ontology and common data structures, and Business Process Management (BPM) based service integration framework. The platform provides commonly reusable features and functions in developing u-health services. According to the early evaluation, our platform turned out to have strength in terms of service flexibility, accessibility, evolvability, reusability, adaptability, interoperability and guideline provision for developing u-health services.

1 INTRODUCTION

The integration of mobile bio-sensors and cellular phones opens a new horizon for healthcare service. It enables us to provide various healthcare services in mobile environment. Actually, couple of mobile u-health services are developed and announced to the public. But, currently available mobile u-health services are very limited in terms of scopes and types of the services. This situation will be changed as more mobile bio-sensors become available and the number of mobile u-healthcare users increases.

Meanwhile, due to the lack of commonly available mobile u-health service development and running platform, current mobile u-health services are usually developed in ad hoc and separate manner. Although u-health services may look different from the service point of view, they often share many common features at various levels such as the service structure, unit service, and data levels. Thus, it is necessary to have a common platform on which various u-health services can be developed by effectively sharing and reusing the common features and services rather than developing the services independently from the scratch.

In this paper, we propose a mobile u-health platform that provides core functions and facilities to develop mobile u-health services. Main elements of the platform include the u-health ontology and common data structures, and Business Process Management (BPM) based service integration framework. The platform provides commonly reusable features and functions in developing u-health services.

The platform has several unique features. First, the platform interprets and treats mobile u-health service as service process and it extends BPMS to healthcare service. As a consequence, the platform can provide functions and facilities of general Business Process Management System (BPMS). In regard with the u-health service process, we define a typical mobile u-health service process template and deploy the process template on the platform. Second, the platform is equipped with a very unique matrix based patient group identification method. The method is quite useful in mobile environment in which less precise bio-signals data is frequently gathered from a large number of users. Third, the platform provides several advanced features and functions by incorporating ontology technologies in

defining mobile u-health services and disease inferring processes.

We implemented the prototype of the proposed platform and evaluated the prototype based on the various software design criteria. According to the early evaluation, our platform turned out to have strength in the aspects of service flexibility, accessibility, evolvability, reusability, adaptability, interoperability and guideline provision for developing u-health services.

2 RELATED WORK

Konstantas et al. (Konstantas et al., 2006) have presented a Java-based service platform for remote monitoring of patients. They collect patient data using medical sensors integrated into the system which captures six different bio-signals: ECG, HRV, pulse oximetry, temperature, marking, respiration, and motion/activity. The captured bio-signals are delivered to remote healthcare experts for analysis. However, the system does not provide any means to manage the health care process between the experts and the patients.

There were some attempts to use a workflow management system or Business Process Management System (BPMS) for the integration of the Patient Record Manager (PRM) module and the Personal Organizer (PO) system in healthcare support (Pappas et al., 2002). Healthcare and emergency response organizations are the users of such solutions. One of the drawbacks of a healthcare system without BPM is insufficient handling of incomplete task definitions (Broens et al., 2005).

Oliver et al. (Oliver and Flores-Mangas, 2006) have designed and developed MPTrain, a mobile phone based system that utilizes the effects of music in physical exercises. MPTrain enables users to more easily achieve their exercise goals by constantly monitoring their physiology such as HRV, and by selecting/playing music with specific features that make the users speed up, slow down or keep the pace to maximize the exercise results.

The work done by Jovanov et al. (Jovanov et al., 2003) is closely related to the healthcare monitoring. They capture HRV to quantify the levels of stress for a group of individuals undergoing military training. They collect HRV data of military personnel prior to, during and after training and assess their stress levels and predict their stress resistance.

However, none of the work support a systematic approach in developing mobile u-health services by using a common u-health service platform like the one described in this paper.

3 MOBILE U-HEALTH SERVICES

A mobile u-health service system is a total system that enables u-health services to be coherently provided to the users. The system integrates various components such as bio-sensors, cellular phones, associated software, and devices that are essential for u-health services. In this section, we describe the core components that constitute mobile u-health services.

We have identified the main categories of the core components that play essential roles in providing u-health services. Bio-data gathering and management, bio-data analysis, and knowledge extraction and decision support are the service categories that we identified based on the steps or roles involved in the u-health service process.

3.1 Bio-data Gathering and Management

A mobile u-health service starts its function with periodically or randomly gathering input data i.e. capturing the bio-signal data from users. We use some bio-sensors that are wearable by the users or imbedded into cellular phones. Thus, the bio-sensor devices and cellular phones, which act as a gateway between the bio-sensors and the u-health server, are the essential components for gathering input data.

Besides, questionnaires that can be provided via cellular phones are necessary to obtain information that cannot be gathered from bio-sensors. The physical symptoms that a user experiences, the location of the user, and the weather information are examples of such information that need to be obtained directly from the users by using the questionnaires. We have developed a generic questionnaire composer to accommodate various symptoms and environmental information in making questionnaires.

Our mobile u-health service platform provides the sensing modules and questionnaire interfaces independently from the core functionality of u-health services. This improves the reusability of the sensing modules and questionnaires for different healthcare services.

3.2 Bio-data Analysis

Since huge amount of bio-data need to be gathered and analyzed in mobile u-health services, it is essential to have an efficient data structure that allows the system to effectively store and manage bio-data. We investigated that a matrix is one of the good candidates for storing and analyzing a large

quantity of data. We devised a bio-signal and symptom combination matrix, in which the appearances of bio-signal and symptom combination pairs of a normal user group and those of a patient user group are registered. We call the matrix as the Disease Combination Appearance Probability (DCAP) matrix. Two DCAP matrices are created, one for normal groups, and another for patient groups. The accumulated data in DCAP matrices becomes the basis of the next step of the u-health service process, the knowledge extraction and decision support.

3.3 Knowledge Extraction and Decision Support

Once a large amount of bio-data is accumulated, knowledge extraction and decision support for diagnosis can be provided by using data mining technology. Sometimes, we can use well established health indices to diagnose certain diseases to support some u-health services. However, in many cases, new health indices may need to be developed by performing machine learning or pattern recognition based on the accumulated data.

For the accumulation and management of bio-data and information obtained from questionnaires, we have developed data ontology. Based on this data ontology, symptom, disease and bio-signal data can be archived in a structured manner. Some additional information such as a weight value to represent the degree of association between each input data and a certain disease type is assigned to the data ontology. Note that the weight values are defined not by users but by domain experts. The data ontology grows as new services and/or disease information are incorporated, and this contributes to make our platform evolvable.

One assumption that we made in identifying health indices is that diseases can be diagnosed based on the combination of bio-signals and symptom information. Another assumption is that bio-signals and symptom data of normal and patient groups can be obtained in some way. If the difference between the data gathered from the normal group and the patient group is obvious, the interrelation patterns between bio-signals and symptoms provide good criteria to classify users into the two groups. A statistical equation is devised based on the DCAP matrices to discern a normal group from a patient group. Once bio-signals and symptoms of a disease or a service are stored in DCAP matrices, (Eq. 1) is used to compute the probability for a person who has the CP(p) bio-signals and symptom pairs to belong to the patient

group of the disease. X(p) and Y(P) denote weight vectors of bio-signals and symptoms respectively. We leave the detailed explanation to our previous work and we do not delve into the details in this paper.

$$P(p \in \text{Patient Group} \mid \text{CP}(p)) = \frac{\sum_{i=1}^m \left[\frac{X(p)_i}{\sum_{j=1}^m X(p)_j} \cdot \left(\frac{\sum_{u=1}^n Y(p)_u \cdot \text{DCAP}(p)_{i,u}}{\sum_{v=1}^n Y(p)_v} \right) \right]}{\quad} \quad (\text{Eq. 1})$$

4 MOBILE U-HEALTH SERVICE SCENARIO

Fig. 1 illustrates the generic service scenario for mobile u-health services. In the first step of the service scenario, users fill out questionnaires to provide information about physical symptoms and their environments, which cannot be obtained from bio-sensors. The sensors embedded in cellular phones capture necessary bio-signal and relay the data to the u-health server.

After gathering bio-data from questionnaires and sensors, a mobile u-health service process is initiated by an associated BPM engine. The next activity of the scenario process is to store the relayed data to a database so that the history of bio-data of a person can be persistently kept for further analyses.

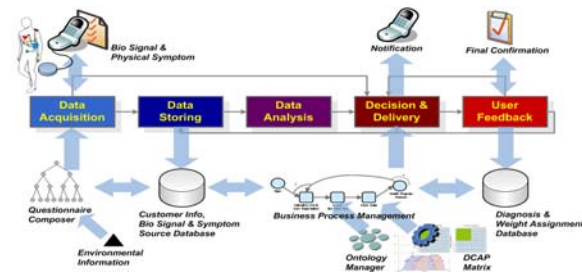


Figure 1: Mobile u-health service scenario process.

In the data analysis and decision support steps, the data ontology manager, which keeps semantically structured data for a specific set of diseases, plays a key role in identifying potential diseases based on the symptoms and bio-signals. For each symptom and bio-signal, the data ontology manager assigns a weight value that represents the degree of association between the symptom or bio-signal and a certain disease.

The final diagnosis decision is made by (Eq. 1). The computed probability is delivered to the user to show whether the user belongs to the patient group or not. One of the unique features of our approach is the user feedback mechanism. Using the feed back

mechanism, users can send their feed back on analysis results (the diagnosed diseases) as the response to the system. Based on this feedback data, the system adjusts weight values of corresponding symptoms and bio-signals to provide more personalized services.

The u-health service scenario itself, and the data ontology and the DCAP matrix components are independent from any specific u-health services. Therefore, the platform we designed is considered to be a general platform. In other words, the u-health service scenario and its supporting components can be used for a variety of u-health services. A mobile u-health service for specific disease can be implemented by extending the service scenario. Please see Section 6 for more details.

5 PLATFORM ARCHITECTURE

The mobile u-health service platform is a middleware that enables the integration of diverse services by using BPM. It serves as the hub to integrate techniques and functions that are associated with mobile u-health services, and provides an environment to develop and run the u-health services. Fig. 2 illustrates the architecture of the platform that shows the major components and the connections to the surrounding elements in the u-health service framework.

Since cellular phones play the role of a gateway between bio-sensors and servers, mobile message handling is essential for the platform. The mobile message handling module relays all the messages from bio-sensors to the server. Not only the bio signal data but also service request messages are delivered by the message handling module. Sometimes, it contributes to filter out some noise signals from the received messages.

The bio-data delivered to a server is stored and managed by the huge temporal data management module. The bio-data is stored in diverse forms so that various services can utilize the data to perform their functions.

Data mining and pattern recognition techniques are used to identify health index from the accumulated bio-data. Also, an external or internal expert system may refer to the data as feedback information. In order to support this, the database schema of the temporal database needs to be designed to satisfy the requirements of data mining or pattern recognition techniques and expert systems.

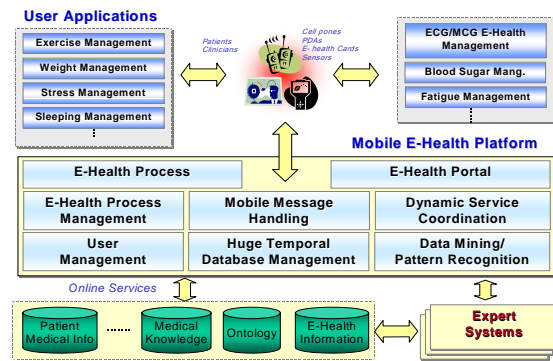


Figure 2: Mobile u-health service platform architecture.

As explained in Section 4, in order to develop a mobile u-health service on the service platform, the u-health service should be defined in a form of process containing steps like bio-data gathering, storing, analysis, and result reporting. The u-health process block in Fig. 2 denotes a set of mobile u-health processes derivable from the u-health service scenario process explained in Section 4.

The u-health process definition tool allows developers to easily define new mobile u-health services. Mobile u-health services defined in a form of process are controlled and enacted by the u-health process management module. The u-health process management module provides not only the enactment function but also monitoring and administration functions for u-health processes.

The mobile u-health management module and u-health process definition tool play a key role in making the platform evolvable. When a new mobile u-health service process is defined, reusable process templates, steps, and data structures are identified and registered to the server for later use. The management module ties together a set of u-health services into a group by specifying execution dependencies and the dataflow between the services. There may be also some constraints that prevent a particular service to be initiated until some other services finish their functions.

The user management module is essential for providing personalized service to individual users. The user management module stores user profile information such as age, gender, and occupation. Since a user is a participant in a mobile u-health process as well, this module is closely connected to the participant information in the u-health process management module.

The dynamic service coordination module executes a service process by initiating and synchronizing services during runtime. It also ensures reliable u-health service execution by

replacing one service with an alternative one when a fault occurs in the service or the user's requirements are changed. Data mining and pattern recognition functions may need to be developed for specific u-health services. Whenever such functions become available, they are placed on the data mining/pattern recognition module. An expert system engaged with a u-health service may need to access the functions in the data mining/pattern recognition module to get a decision making support.

Mobile u-health services are accessible not only from mobile devices, but also through a Web browser. Web users can connect to a healthcare Web portal, and access their health information. The healthcare portal also provides useful services such as registering user information, composing questionnaires, and browsing expert advices. Especially, the Web portal allows users to access services that cannot be provided through mobile devices due to the limitations of the devices such as the small screen size, and memory and processor constraints.

6 IMPLEMENTATION

6.1 Mobile u-Health Modeling Tool

We have implemented a prototype of the u-health process modeling tool that allows application developers to semi-automatically compose mobile u-health services.

Fig. 3 shows the prototype of our mobile u-health process modeling tool. The tool is implemented in Java™ on Eclipse Foundation's Eclipse™ platform. Part (a) of Fig. 3 shows the canvas of the visual business process editing tool. The tool graphically shows the basic u-health service process template, and allows application developers to specialize and extend the process to meet their requirements. In the course of defining the process, the process modeling tool recommends a set of candidate services that can be used to implement each step of the process, and are interoperable with other services that are already put in the process. It also automatically inserts adapters and converters when it finds some syntactic and semantic mismatches between services in the process. Part (b) of the modeling tool is the service recommendation tool that shows a list of candidate services for each step in a process. When a user marks a service in a service process, the tool lists the services that are interoperable with the marked service. Part (c) of Fig. 3 is the ontology-hierarchy browser that allows developers to browse through the service hierarchy.

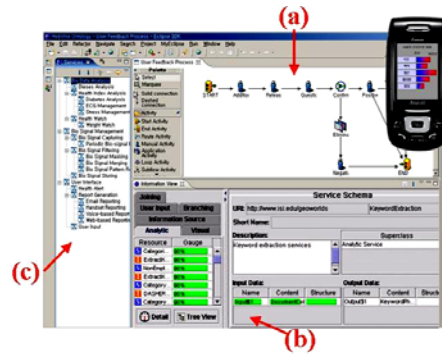


Figure 3: Prototype of the service composition tool - (a) Visual business process editing tool; (b) Service recommendation tool; (c) Ontology-hierarchy browser.

6.2 Mobile u-Health Ontology Manager

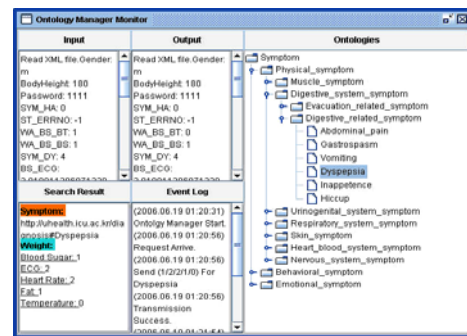


Figure 4: Ontology manager interface.

As explained earlier, in our platform, u-health data such as symptoms, diseases, and bio-signals, are systematically managed by constructing their ontology. In addition, services are classified and managed based on their functional ontology. The ontology manager allows experts to easily manage the u-health ontology. Fig. 4 shows the user interface of the ontology manager. By using the tool, experts can browse, add, remove, and search u-health ontology.

6.3 Mobile u-Health DCAP Matrix

In our u-health service platform, the interrelation patterns are identified in the learning and diagnosis stage based on the patterns that is constructed in the prediction stage. In the learning stage, interrelation patterns of bio signal data and physical symptoms of the normal and patient groups are identified. We assume that enough learning data for the identification of interrelations is available for this stage.

In the learning stage, interaction frequencies of bio-signals and physical symptoms from the normal and patient groups are counted and registered in the Combination Interaction (CI) matrices. Each CI matrix contains combination patterns of specific bio-signals and physical symptoms of the normal or patient group with frequency scores. We integrated both CI matrices into the DCAP matrix. In the DCAP matrix, each element has a probability which means how likely the combination pattern of the element can be shown in a patient group. In the prediction stage, bio signals and physical symptoms of an unidentified person are submitted to the prediction system. Then, the prediction system determines how likely the person can be categorized into the patient group. Since the DCAP matrix is constructed based on the CI matrix, we can easily figure out which bio-signal and physical-symptom combinations contribute to a certain diagnosis. In Fig. 5-(b), shaded areas are corresponding bio-signal and symptom combination elements, and a probability, ranging from 0 to 1, is calculated by using (Eq. 1).

The DCAP matrix enables diverse forms of bio-signals and physical symptoms to be used for analyses. As shown in Fig. 5-(a), users can compare their position in the matrix with those of others. In addition, when we use the DCAP matrix, finding primary bio-signal and physical symptom pairs that contribute for the user to be classified into the patient group is relatively easy. As a result, the prediction system can show the causes of a certain disease to users in a comprehensive manner. The more clinical data is accumulated in the platform, the more services the platform can provide in the future.



Figure 5: Disease combination appearance probability matrix – (a) Disease distribution chart for a learning set; (b) Test result for a user health condition by the DCAP matrix.

6.4 Mobile Client

Due to the diverse types and characteristics of bio-sensors and various application scenarios, there are many different requirements imposed on mobile u-health systems. As a mobile gateway, a mobile device is responsible for collecting data from sensors by periodically sensing the human body.

Once a communication link is established, bio-data and functions are delivered to the u-health server using the wireless interface. The interface sends and receives XML messages to and from BPM. An XML message can a flexible structure to facilitate transmission of different data types and formats generated based on various bio-signals and physical symptoms. By parsing the XML message, the server classifies bio-data into several types and then the server responds with the results of classified bio-data in an XML message.

After forwarding bio-data to the server, the application changes its state to the stand-by state until it receives a response message from the server. No event or signal can be pushed to the client asynchronously. A process polling mechanism is used to continuously pull the event of clients' requests to the decision support step of the u-health service process. When the server process is completed, the result is delivered to the client.



Figure 6: WUPI Application Mounted on KTF Emulator & Samsung Anycall SPH-V8900.

Once the client receives the probability of a disease, the result is graphically displayed on the mobile device. Figure 6 shows a graph displayed on a cellular phone after receiving the results from the server. If the user does not agree with the results, he/she can send a feedback to the server to revise his/her own weight vectors. The DCAP matrices are updated if the feedback data is sufficiently accumulated.

7 EVALUATION

We considered some quality attributes in designing the mobile u-health service platform, which is a large-scale service platform. We think that flexibility, accessibility, evolvability, reusability, adaptability and interoperability are the essential requirements for large-scale service platforms.

Table 1: Evaluation of the Mobile u-Health Service Platform.

Core Components Design Goals	Bio-data gathering and management		Bio-data analysis		Decision support
	WI PI	Web Portal	Ontology	DCAP Matrix	BPM
Flexibility	○	△	○	○	○
Accessibility	○	○	×	×	△
Evolvability	N/A	N/A	○	○	△
Reusability	△	○	○	△	○
Adaptability	△	△	△	○	○
Interoperability	○	○	△	△	○

Remarks: ○: Supported, ×: Not Supported, △: Insufficient, N/A: Not Applicable

Table 1 summarize these quality attributes and explain how the core components of our service platform contribute to satisfy those requirements. The following sub-sections discuss the design rationales of our platform to make it satisfy each of the quality attributes.

7.1 Flexibility

There are some reasons why the mobile u-health service platform needs to be flexible. Different sets of bio-sensors such as ECG, HRV, and temperature sensors need to be supported for different u-health services. Furthermore, some of the sensors may need to be replaced with other types depending on the mobile-device types and some other constraints. In addition, the functionality provided by the platform or the user interface on a mobile device may need to be customized for a specific user group. Also, the bio-signals and symptoms that the healthcare server handles may need to be changed based on the disease types to cover by specific u-health services.

In order to support flexibility, we have developed a Web-based questionnaire composer. It can easily accommodate new questionnaires or updates of questionnaires for new u-health scenarios. In addition, since we use the BPM process modeling tool for the definition of u-health service processes,

we can easily update the defined process when a change is required. This enables the platform to be used for a variety of u-health services with different scenarios. The ontology-based data management and the DCAP matrix are the features that improve flexibility of our platform as well. The structure of ontology and DCAP matrix can be dynamically changed based on users' feedback information.

7.2 Accessibility

The multi-channel approach, which makes users' health information accessible via both mobile and Web-based interfaces, supports a good access environment that enables the benefits of hybrid accessibility. In the healthcare service environment, a real-time and mobile disease-checking function can be supported using mobile devices whereas the in-depth analysis of the disease can be provided on the Web. This makes the users to access their health information in various degrees of detail.

7.3 Evolvability

As discussed in Section 3, the data ontology evolves as new data and feedback information from users are accumulated. In addition, the DCAP matrices used in our u-health service platform are also evolvable. According to the feedback on the analysis results, the DCAP matrices appropriately change their values for more precise and personalized analyses.

7.4 Reusability

The BPMS-based service development mechanism makes service components reusable for various types of u-health services. In other words, when a new mobile u-health service is defined with the u-health process definition tool, developers can reuse various service components, and other assets in different abstraction levels. Process templates, activities, applications pertaining to an activity, and data structures are the reusable assets in our platform. When we consider that most u-health service processes shares the basic u-health service scenario described in Section 4, many components of a u-health service process can be reusable by other u-health processes. Since our platform is based on BPMS, it is competitive in supporting reusability.

7.5 Adaptability

The circumstances and environment surrounding mobile u-health services are ever changing. New bio-sensors, which produce new forms of signals or

signals never handled before, may become available, and new service scenarios may be developed. For example, a mobile u-health service may be developed to handle only a single type of sensor. But sometime it may need to handle multiple sensors in a bundle. Then, the implementation of the u-health service needs to be modified to cope with the changes. Since our platform maintains a property file that keeps the descriptions about the sensor types, the platform can adapt to the change more efficiently. This property file enables multiple sensors to be linked to a mobile device and to transmit the sensor data over a single wireless link to the u-health server.

When hospitals and insurance companies are engaged with the mobile u-health service scenario described in Section 4, the u-health service scenario needs to be extended. Such kind of environmental change circumventing u-health services may be frequent in real situations. Since u-health services are defined in the form of processes using the BPMT tool in our platform, they can adapt to such changes more efficiently by adding new service components and/or by replacing some of the service components with alternative ones.

7.6 Interoperability

One of important criteria for designing the mobile u-health service platform is interoperability. BPMS is usually equipped with facilities for supporting interoperability. Since our platform is based on BPMS, the facilities of BPMS for interoperability can be shared in the mobile u-health service platform. The remote applications and systems can be integrated with u-health services in our platform without any extra work. The Web Services API for BPMS and the Web portal provide an easy way of exchanging information between processes.

8 CONCLUSIONS

In this paper, we have identified several core components of the u-health service scenario process, and described the overall architecture of the mobile u-health platform. We have also developed essential functions and facilities that allow service developers to effectively use the core components to develop various mobile u-health services with less effort.

Our mobile u-health service platform is very unique in that it adopts BPMS as its underlying platform, and u-health services are designed as service processes. The standard u-health service process provides a guideline to u-health service

developers. They can specialize and/or extend the service process template to develop their own u-health services. Six design goals are identified for the mobile u-health platform, and we have tried to achieve them in our BPMS-based u-health service platform.

One of the lessons that we learned while working with the mobile environment is that different trade-offs need to be made for different circumstances. For example, we need to consider resource limitations on mobile devices, communication bandwidth, and constraints on user interfaces to make a trade-off between different quality of services for different environments and different user groups.

There are also some pending works to be done to make the platform more secure and reliable. We are currently investigating how to support the features of managing personal healthcare data securely, and enhancing the performance by compressing messages exchanged between mobile devices and the u-health server.

REFERENCES

- D. Konstantas, R. Bults, A. Van Halteren, K. Wac, V. Jones, I. Widya, R. Herzog, B. Streimelweger, "Mobile Health Care: Towards a commercialization of research results", In proceedings of 1st European Conference on eHealth - ECEH06 -Fribourg, Switzerland, pp.12 – 13, October 2006.
- Pappas, C. Coscia, E. Doderio, G. Gianuzzi, V. Earney, M., "A Mobile E-Health System Based on Workflow Automation Tools", Computer-Based Medical Systems. Proceedings of the 15th IEEE Symposium on pp.271 – 276, June 2002.
- Broens, T.H.F., van Halteren, A.T., van Sinderen, M.J., Wac, K.E. "Towards an application framework for context-aware m-health applications", In: EUNICE 2005 "Networked Applications". 11th open European Summer School, pp.6-8 July 2005.
- N. Oliver, F. Flores-Mangas, "MPTrain: a mobile, music and physiology-based personal trainer", ACM International Conference Proceeding Series; Vol.159, pp.21 – 28, 2006.
- E. Jovanov, D. Raskovicl, A. O. Lords, P. Cox, R. Adharni, F. Andrasik, "Stress monitoring Using a Distributed Wireless Intelligent Sensor System", in IEEE Engineering in Medicine and Biology Magazine, pp. 49-55, May/June 2003.

PROCESS MINING IN HEALTHCARE

A Case Study

R. S. Mans, M. H. Schonenberg, M. Song, W. M. P. van der Aalst
Eindhoven University of Technology, P.O. Box 513, NL-5600 MB, Eindhoven, The Netherlands
{r.s.mans, m.h.schonenberg, m.s.song, w.m.p.v.d.aalst}@tue.nl

P. J. M. Bakker
Academic Medical Center, University of Amsterdam, Department of Innovation and Process Management
Amsterdam, The Netherlands
p.j.bakker@amc.uva.nl

Keywords: Process mining, healthcare processes.

Abstract: To gain competitive advantage, hospitals try to streamline their processes. In order to do so, it is essential to have an accurate view of the “careflows” under consideration. In this paper, we apply process mining techniques to obtain meaningful knowledge about these flows, e.g., to discover typical paths followed by particular groups of patients. This is a non-trivial task given the dynamic nature of healthcare processes. The paper demonstrates the applicability of process mining using a real case of a gynecological oncology process in a Dutch hospital. Using a variety of process mining techniques, we analyzed the healthcare process from three different perspectives: (1) the control flow perspective, (2) the organizational perspective and (3) the performance perspective. In order to do so we extracted relevant event logs from the hospitals information system and analyzed these logs using the ProM framework. The results show that process mining can be used to provide new insights that facilitate the improvement of existing careflows.

1 INTRODUCTION

In a competitive health-care market, hospitals have to focus on ways to streamline their processes in order to deliver high quality care while at the same time reducing costs (Anyanwu et al., 2003). Furthermore, also on the governmental side and on the side of the health insurance companies, more and more pressure is put on hospitals to work in the most efficient way as possible, whereas in the future, an increase in the demand for care is expected.

A complicating factor is that healthcare is characterized by highly *complex* and extremely *flexible* patient care processes, also referred to as “careflows”. Moreover, many disciplines are involved for which it is found that they are working in isolation and hardly have any idea about what happens within other disciplines. Another issue is that within healthcare many autonomous, independently developed applications are found (Lenz et al., 2002). A consequence of this all is that *it is not known what happens in a healthcare process for a group of patients with the same diagnosis*.

The concept of process mining provides an in-

teresting opportunity for providing a solution to this problem. Process mining (van der Aalst et al., 2003) aims at extracting process knowledge from so-called “event logs” which may originate from all kinds of systems, like enterprise information systems or hospital information systems. Typically, these event logs contain information about the start/completion of process steps together with related context data (e.g. actors and resources). Furthermore, process mining is a very broad area both in terms of (1) applications (from banks to embedded systems) and (2) techniques.

This paper focusses on the *applicability* of process mining in the healthcare domain. Process mining has already been successfully applied in the service industry (van der Aalst et al., 2007a). In this paper, we demonstrate the applicability of process mining to the healthcare domain. We will show how process mining can be used for obtaining insights related to careflows, e.g., the identification of care paths and (strong) collaboration between departments. To this end, in Section 3, we will use several mining techniques which will also show the diversity of process mining techniques available, like control flow discovery but also the discovery of organizational aspects.

In this paper, we present a case study where we use raw data of the AMC hospital in Amsterdam, a large academic hospital in the Netherlands. This raw data contains data about a group of 627 gynecological oncology patients treated in 2005 and 2006 and for which all diagnostic and treatment activities have been recorded for financial purposes. Note that we did not use any a-priori knowledge about the care process of this group of patients and that we also did not have any process model at hand.

Today's Business Intelligence (BI) tools used in the healthcare domain, like Cognos, Business Objects, or SAP BI, typically look at aggregate data seen from an external perspective (frequencies, averages, utilization, service levels, etc.). These BI tools focus on performance indicators such as the number of knee operations, the length of waiting lists, and the success rate of surgery. Process mining looks "inside the process" at different abstraction levels. So, in the context of a hospital, unlike BI tools, we are more concerned with the care paths followed by individual patients and whether certain procedures are followed or not.

This paper is structured as follows: Section 2 provides an overview of process mining. In Section 3 we will show the applicability of process mining in the healthcare domain using data obtained for a group of 627 gynecological oncology patients. Section 4 concludes the paper.

2 PROCESS MINING

Process mining is applicable to a wide range of systems. These systems may be pure information systems (e.g., ERP systems) or systems where the hardware plays a more prominent role (e.g., embedded systems). The only requirement is that the system produces *event logs*, thus recording (parts of) the actual behavior.

An interesting class of information systems that produce event logs are the so-called *Process-Aware Information Systems* (PAISs) (Dumas et al., 2005). Examples are classical workflow management systems (e.g. Staffware), ERP systems (e.g. SAP), case handling systems (e.g. FLOWer), PDM systems (e.g. Windchill), CRM systems (e.g. Microsoft Dynamics CRM), middleware (e.g., IBM's WebSphere), hospital information systems (e.g., Chipsoft), etc. These systems provide very detailed information about the activities that have been executed.

However, not only PAISs are recording events. Also, in a typical hospital there is a wide variety of systems that record events. For example, in an inten-

sive care unit, a system can record which examinations or treatments a patient undergoes and also it can record occurring complications for a patient. For a radiology department the whole process of admittance of a patient till the archival of the photograph can be recorded. However, frequently these systems are limited to one department only. However, systems used for billing purposes have to ensure that all services delivered to the patient will be paid. In order for these systems to work properly, information from different systems needs to be collected so that it is clear which activities have been performed in the care process of a patient. In this way, these systems within the hospital can contain information about processes *within* one department but also *across* departments. This information can be used for improving processes within departments itself or improving the services offered to patients.

The goal of process mining is to extract information (e.g., process models) from these logs, i.e., process mining describes a family of *a-posteriori* analysis techniques exploiting the information recorded in the event logs. Typically, these approaches assume that it is possible to sequentially record events such that each event refers to an activity (i.e., a well-defined step in the process) and is related to a particular case (i.e., a process instance). Furthermore, some mining techniques use additional information such as the performer or originator of the event (i.e., the person/resource executing or initiating the activity), the timestamp of the event, or data elements recorded with the event (e.g., the size of an order).

Process mining addresses the problem that most "process/system owners" have limited information about what is actually happening. In practice, there is often a significant gap between what is prescribed or supposed to happen, and what *actually* happens. Only a concise assessment of reality, which process mining strives to deliver, can help in verifying process models, and ultimately be used in system or process redesign efforts.

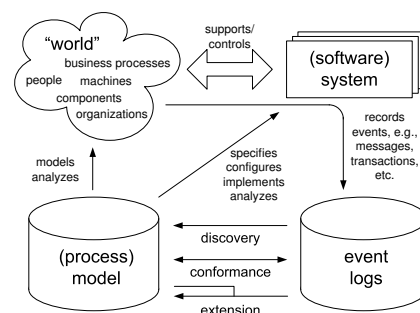


Figure 1: Three types of process mining: (1) Discovery, (2) Conformance, and (3) Extension.

The idea of process mining is to discover, monitor and improve real processes (i.e., not assumed processes) by extracting knowledge from event logs. We consider three basic types of process mining (Figure 1): (1) *discovery*, (2) *conformance*, and (3) *extension*.

Discovery: Traditionally, process mining has been focusing on *discovery*, i.e., deriving information about the original process model, the organizational context, and execution properties from enactment logs. An example of a technique addressing the control flow perspective is the α -algorithm (van der Aalst et al., 2004), which constructs a Petri net model describing the behavior observed in the event log. It is important to mention that there is no a-priori model, i.e., based on an event log some model is constructed. However, process mining is not limited to process models (i.e., control flow) and recent process mining techniques are more and more focusing on other perspectives, e.g., the organizational perspective, performance perspective or the data perspective. For example, there are approaches to extract social networks from event logs and analyze them using social network analysis (van der Aalst et al., 2005). This allows organizations to monitor how people, groups, or software/system components are working together. Also, there are approaches to visualize performance related information, e.g. there are approaches which graphically shows the bottlenecks and all kinds of performance indicators, e.g., average/variance of the total flow time or the time spent between two activities.

Conformance: There is an a-priori model. This model is used to check if reality conforms to the model. For example, there may be a process model indicating that purchase orders of more than one million Euro require two checks. Another example is the checking of the so-called “four-eyes” principle. Conformance checking may be used to detect deviations, to locate and explain these deviations, and to measure the severity of these deviations.

Extension: There is an a-priori model. This model is extended with a new aspect or perspective, i.e., the goal is not to check conformance but to enrich the model with the data in the event log. An example is the extension of a process model with performance data, i.e., some a-priori process model is used on which bottlenecks are projected.

At this point in time there are mature tools such as the ProM framework (van der Aalst et al., 2007b), featuring an extensive set of analysis techniques which can be applied to real-life logs while supporting the whole spectrum depicted in Figure 1.

3 HEALTHCARE PROCESS

In this section, we want to show the *applicability* of process mining in healthcare. However, as healthcare processes are characterized by the fact that *several organizational units* can be involved in the treatment process of patients and that these organizational units often have their own specific IT applications, it becomes clear that getting data, which is related to healthcare processes, is not an easy task. In spite of this, systems used in hospitals need to provide an integrated view on all these IT applications as it needs to be guaranteed that the hospital gets paid for every service delivered to a patient. Consequently, these kind of systems contain process-related information about healthcare processes and are therefore an interesting candidate for providing the data needed for process mining.

To this end, as case study for showing the applicability of process mining in health care, we use raw data collected by the billing system of the AMC hospital. This raw data contains information about a group of 627 gynecological oncology patients treated in 2005 and 2006 and for which all diagnostic and treatment activities have been recorded. The process for gynecological oncology patients is supported by several different departments, e.g. gynecology, radiology and several labs.

For this data set, we have extracted event logs from the AMC’s databases where each event refers to a service delivered to a patient. As the data is coming from a billing system, we have to face the interesting problem that for each service delivered for a patient it is only known on which *day* the service has been delivered. In other words, we do not have any information about the actual timestamps of the start and completion of the service delivered. Consequently, the ordering of events which happen on the same day do not necessarily conform with the order in which events of that day were executed.

Nevertheless, as the log contains *real* data about the services delivered to gynecological oncology patients it is still an interesting and representative data set for showing the applicability of process mining in healthcare as still many techniques can be applied. Note that the log contains 376 different event names which indicates that we are dealing with a non-trivial careflow process.

In the remainder of this section we will focus on obtaining, in an explorative way, *insights into the gynecological oncology healthcare process*. So, we will only focus on the *discovery* part of process mining, instead of the *conformance* and *extension* part. Furthermore, obtaining these insights should not be

limited to one perspective only. Therefore, in sections 3.2.1, 3.2.2 and 3.2.3, we focus on the discovery of *care paths followed by patients*, the discovery of *organizational aspects* and the discovery of *performance related information*, respectively. This also demonstrates the diversity of process mining techniques available. However, as will be discussed in Section 3.1, we first need to perform some preprocessing before being able to present information on the right level of detail.

3.1 Preprocessing of Logs

The log of the AMC hospital contains a huge amount of distinct activities, of which many are rather low level activities, i.e., events at a low abstraction level. For example, for our purpose, the logged lab activities are at a too low abstraction level, e.g. determination of chloride, lactic acid and erythrocyte sedimentation rate (ESR). We would like to consider all these low level lab tests as a single lab test. Mining a log that contains many distinct activities would result in a too detailed spaghetti-like model, that is difficult to understand. Hence, we first apply some preprocessing on the logs to obtain interpretable results during mining. During preprocessing we want to “simplify” the log by removing the excess of low level activities. In addition, our goal is to consider only events at the department level. In this way, we can, for example, focus on care paths and interactions between departments. We applied two different approaches to do this.

Our first approach is to detect a *representative* for the lower level activities. In our logs, this approach can be applied to the before mentioned lab activities. In the logs we can find an activity that can serve as representative for the lab activities, namely the activity that is always executed when samples are offered to the lab. All other (low level) lab activities in the log are simply discarded.

The log also contains groups of low level activities for which there is no representative. For instance at the radiology department many activities can occur (e.g., echo abdomen, thorax and CT brain), but the logs do not contain a single event that occurs for every visit to this department, like a registration event for example. We apply *aggregation* for low level activities in groups without a representative by (1) defining a representative, (2) mapping all activities from the group to this representative and (3) removing repetitions of events from the log. For example, for the radiology department we define “radiology” as representative. A log that originally contains “...ultrasound scan abdomen,

chest X-ray, CT scan brain,...”, would contain “...radiology,...”, after mapping low level radiology activities to this representative and removing any duplicates.

3.2 Mining

In this section, we present some results obtained through a detailed analysis of the ACM’s event log for the gynecological oncology process. We concentrate on the discovery part to show actual situations (e.g. control flows, organizational interactions) in the healthcare process. More specifically, we elaborate on mining results based on three major perspectives (i.e. control flow, organizational, performance perspectives) in process mining.

3.2.1 Control Flow Perspective

One of the most promising mining techniques is control flow mining which automatically derives process models from process logs. The generated process model reflects the actual process as observed through real process executions. If we generate process models from healthcare process logs, they give insight into care paths for patients. Until now, there are several process mining algorithms such as the α -mining algorithm, heuristic mining algorithm, region mining algorithm, etc (van der Aalst et al., 2004; Weijters and van der Aalst, 2003; van Dongen et al., 2007). In this paper, we use the heuristic mining algorithm, since it can deal with noise and exceptions, and enables users to focus on the main process flow instead of on every detail of the behavior appearing in the process log (Weijters and van der Aalst, 2003). Figure 2 shows the process model for all cases obtained using the Heuristics Miner. Despite its ability to focus on the most frequent paths, the process, depicted in Figure 2, is still spaghetti-like and too complex to understand easily.

Since processes in the healthcare domain do not have a single kind of flow but a lot of variants based on patients and diseases, it is not surprising that the derived process model is spaghetti-like and convoluted.

One of the methods for handling this problem is breaking down a log into two or more sub-logs until these become simple enough to be analyzed clearly. We apply clustering techniques to divide a process log into several groups (i.e. clusters), where the cases in the same cluster have similar properties. Clustering is a very useful technique for logs which contain many cases following different procedures, as is the usual case in healthcare systems. Depending on the interest (e.g., exceptional or frequent procedures), a cluster can be selected. There are several clustering




Figure 2: Derived process model for all cases.

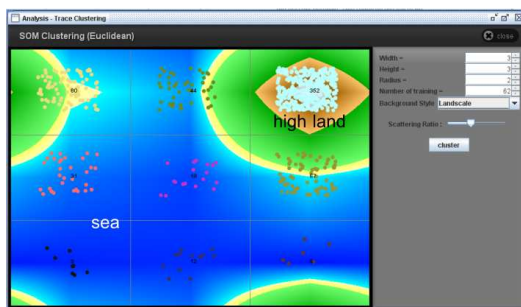


Figure 3: Log clustering result.

techniques available. Among these, we use the SOM (Self Organizing Map) algorithm to cluster the log because of its performance (i.e., speed). Figure 3 shows the clustering result obtained by applying the Trace Clustering plug-in. Nine clusters are obtained from the log. In the figure, the instances in the same cell belong to the same cluster. The figure also shows a contour map based on the number of instances in each cell. It is very useful to take a quick glance at the clusters – are there clusters with many similarities (high land), or are there many clusters with exceptional cases (sea).

By using this approach, we obtained several clusters of reasonable size. In this paper we show only the result for the biggest cluster, containing 352 cases all with similar properties. Figure 4 shows the heuristic net derived from the biggest cluster. The result is much simpler than the model in Figure 2. Furthermore, the fitness of this model is “good”. The model represents the procedure for most cases in the cluster,

i.e., these cases “fit” in the generated process model. A closer inspection of this main cluster by domain experts confirmed that this is indeed main stream followed by most gynecological oncology patients.

When discussing the result with the people involved in the process, it was noted that patients, referred to the AMC by another hospital, only visit the outpatient clinic once or twice. These patients are already diagnosed, and afterwards they are referred to another department, like radiotherapy, for treatment and which is then responsible for the treatment process. Also, very ill patients are immediately referred to another department for treatment after their first visit.

3.2.2 Organizational Perspective

There are several process mining techniques that address organizational perspective, e.g., organizational mining, social network mining, mining staff assignment rules, etc. (van der Aalst et al., 2005). In this paper, we elaborate on social network mining to provide insights into the collaboration between departments in the hospital. The Social Network Miner allows for the discovery of social networks from process logs. Since there are several social network analysis techniques and research results available, the generated social network allows for analysis of social relations between originators involving process executions. Figure 5 shows the derived social network. To derive the network, we used the *Handover of Work* metric (van der Aalst et al., 2005) that measures the frequency of transfers of work among departments.

all other departments. For instance there is no relationship (within this threshold) between this OC and the vascular lab. This means that there is no, or not much, interaction between these two departments.

When this result was presented to the people involved in the process, they confirmed the strong collaboration with the departments shown in Figure 5. However, they were surprised about the rather strong collaboration with the dietics department. Nevertheless, this can be explained by the fact that, when a patient has to go to several chemotherapy sessions, then a visit to the dietician is also often needed.

Moreover, they also noted that the many interactions between the lab and other departments is misleading as all the examinations are requested by gynecological oncology and not by the lab. This can be explained by the many lab tests and resulting interactions between the lab and other departments.

3.2.3 Performance Perspective

Process mining provides several performance analysis techniques. Among these, the dotted chart is a method suitable for case handling processes which are flexible and knowledge intensive business processes and focus not on the routing of work or the activities but on the case (e.g. careflows). In this paper, we use the dotted chart to show overall events and performance information of the log. Figure 6 shows the dotted chart. In the chart, events are displayed as dots, and the time is measured along the horizontal axis of the chart. The vertical axis represents case IDs and events are colored according to their task IDs. It supports several time options such as actual, relative, logical, etc. In the diagram, we use relative time which shows the duration from the beginning of an instance to a certain event. Thus it indicates the case duration of each instance. It also provides performance metrics such as the time of the first and of the last events, case durations, the number of events in an instance, etc. For example, in the figure (top right, average spread in seconds), the average case duration is about 49 days.

Users can obtain useful insights from the chart, e.g., it is easy to find interesting patterns by looking at the dotted chart. In Figure 6, the density of events on the left side of the diagram is higher than the density of those on the right side. This shows that initially patients have more diagnosis and treatment events than in the later parts of the process. When we focus on the long duration instances (i.e. the instances having events in the right side of the diagram), it can be observed that they mainly consist of regular consultation (red dot), consultation by phone (red dot), and lab test (violet dot) activities. It reflects the situation that patients have regular consultation by visiting or being

phoned by the hospital and sometimes have a test after or before the consultation. It is also easy to discover patterns in the occurrences of activities. For example, seven instances have the pattern that consists of a lab test and an admittance to the nursing ward activities.

When the results were presented to the people involved in the process, they confirmed the patterns that we found. Furthermore, for the last pattern they indicated that the pattern deals about patients who get a chemotherapy regularly. The day before, they come for a lab test and when the result is good, they get the next chemotherapy.

4 CONCLUSIONS

In this paper, we have focussed on the applicability of process mining in the healthcare domain. For our case study, we have used data coming from a non-trivial care process of the AMC hospital. We focussed on obtaining insights into the careflow by looking at the control-flow, organizational and performance perspective. For these three perspectives, we presented some initial results. We have shown that it is possible to mine complex hospital processes giving insights into the process. In addition, with existing techniques we were able to derive *understandable* models for large groups of patients. This was also confirmed by people of the AMC hospital.

Furthermore, we compared our results with a flowchart for the diagnostic trajectory of the gynecological oncology healthcare process, and where a top-down approach had been used for creating the flowchart and obtaining the logistical data (Elhuizen et al., 2007). With regard to the flowchart, comparable results have been obtained. However, a lot of effort was needed for creating the flowchart and obtaining the logistical data, where with process mining there is the opportunity to obtain these kind of data in a semi-automatic way.

Unfortunately, traditional process mining approaches have problems dealing with unstructured processes as, for example, can be found in a hospital environment. Future work will focus on both developing *new* mining techniques and on using *existing* techniques in an innovative way to obtain understandable, high-level information instead of “spaghetti-like” models showing all details. Obviously, we plan to evaluate these results in healthcare organizations such as the AMC.




Figure 6: Dotted Chart.

ACKNOWLEDGEMENTS

This research is supported by EIT, NWO-EW, the Technology Foundation STW, and the SUPER project (FP6). Moreover, we would like to thank the many people involved in the development of ProM.

REFERENCES

- Anyanwu, K., Sheth, A., Cardoso, J., Miller, J., and Kochut, K. (2003). Healthcare Enterprise Process Development and Integration. *Journal of Research and Practice in Information Technology*, 35(2):83–98.
- Dumas, M., van der Aalst, W., and ter Hofstede, A. (2005). *Process-Aware Information Systems: Bridging People and Software through Process Technology*. Wiley & Sons.
- Elhuizen, S., Burger, M., Jonkers, R., Limburg, M., Klazinga, N., and Bakker, P. (2007). Using Business Process Redesign to Reduce Wait Times at a University Hospital in the Netherlands. *The Joint Commission Journal on Quality and Patient Safety*, 33(6):332–341.
- Lenz, R., Elstner, T., Siegele, H., and Kuhn, K. (2002). A Practical Approach to Process Support in Health Information Systems. *Journal of the American Medical Informatics Association*, 9(6):571–585.
- van der Aalst, W., Reijers, H., and Song, M. (2005). Discovering Social Networks from Event Logs. *Computer Supported Cooperative Work*, 14(6):549–593.
- van der Aalst, W., Reijers, H., Weijters, A., van Dongen, B., de Medeiros, A. A., Song, M., and Verbeek, H. (2007a). Business process mining : an industrial application. *Information Systems*, 32(5).
- van der Aalst, W., van Dongen, B., Günther, C., Mans, R., de Medeiros, A. A., Rozinat, A., Rubin, V., Song, M., Verbeek, H., and Weijters, A. (2007b). ProM 4.0: Comprehensive Support for Real Process Analysis. In Kleijn, J. and Yakovlev, A., editors, *Application and Theory of Petri Nets and Other Models of Concurrency (ICATPN 2007)*, volume 4546 of *Lecture Notes in Computer Science*, pages 484–494. Springer-Verlag, Berlin.
- van der Aalst, W., van Dongen, B., Herbst, J., Maruster, L., Schimm, G., and Weijters, A. (2003). Workflow Mining: A survey of Issues and Approaches. *Data and Knowledge Engineering*, 47(2).
- van der Aalst, W., Weijters, A., and Maruster, L. (2004). Workflow Mining: Discovering Process Models from Event Logs. *IEEE Transactions on Knowledge and Data Engineering*, 16(9):1128–1142.
- van Dongen, B., Busi, N., Pinnaand, G., and van der Aalst, W. (2007). An Iterative Algorithm for Applying the Theory of Regions in Process Mining. BETA Working Paper Series, WP 195, Eindhoven University of Technology, Eindhoven.
- Weijters, A. and van der Aalst, W. (2003). Rediscovering Workflow Models from Event-Based Data using Little Thumb. *Integrated Computer-Aided Engineering*, 10(2):151–162.

HEALTHCARE RISK MODELING FOR MEDICAID PATIENTS

The Impact of Sampling on the Prediction of High-Cost Patients

Sai T. Moturu, Huan Liu

*Department of Computer Science and Engineering, School of Computing and Informatics
Arizona State University, Tempe, AZ 85287
smoturu@asu.edu, hliu@asu.edu*

William G. Johnson

*Center for Health Information & Research (CHiR), School of Biomedical Informatics
Arizona State University, Tempe, AZ 85287
william.g.johnson@asu.edu*

Keywords: Predictive risk modeling, healthcare costs, high-cost patients, high-risk patients, non-random sampling, over-sampling, under-sampling, imbalanced data, skewed data, Medicaid, data mining, classification.

Abstract: Healthcare data from the Arizona Health Care Cost Containment System, Arizona's Medicaid program provides a unique opportunity to exploit state-of-the-art data processing and analysis algorithms to mine data and provide actionable findings that can aid cost containment. Our work addresses specific challenges in this real-life healthcare application to build predictive risk models for forecasting future high-cost patients. We survey the literature and propose novel data mining approaches customized for this compelling application with specific focus on non-random sampling. Our empirical study indicates that the proposed approach is highly effective and can benefit further research on cost containment in the healthcare industry.

1 INTRODUCTION

The Center for Health Information and Research (CHiR) at Arizona State University houses a community health data system called Arizona HealthQuery (AZHQ). AZHQ contains comprehensive health records of patients from the state of Arizona linked across systems and time. The data, which include more than six million persons, offer the opportunity for research that can impact on the health of the community by delivering actionable results for healthcare researchers and policy makers.

One of the primary issues plaguing the healthcare system is the problem of rapidly rising costs. Many reasons have been put forward for the consistent growth in health care expenditures ranging from the lack of a free market and the development of innovative technologies to external factors like economy and population growth (Bodenheimer, 2005). A first step to tackle these issues is to devise effective cost containment measures. One efficient approach to cost containment is to focus on high-cost patients responsible for these expenditures and undertake

measures to reduce these costs. Predictive risk modeling is a relatively recent attempt at proactively identifying prospective high-cost patients to reduce costs. We embark on the challenging task of building predictive risk models using real-life data from the Arizona Health Care Cost Containment System (AHCCCS), Arizona's Medicaid program, available in AZHQ. The AHCCCS data was selected because it contains a large number of patients who can be tracked over multiple years and it contains many features needed for the analysis in this study.

Apart from data analysis challenges due to the voluminous amount of patient records and the considerable amount of variation for similarly grouped patients, such cost data provides a bigger challenge. It has been commonly observed that a small proportion of the patients are responsible for a large share of the total healthcare expenditures. This skewed pattern has remained constant over many decades. Previous studies show that more than two-thirds of the health costs are from the top ten percent of the population (Berk & Monheit, 2001). Similar patterns are observed in our empirical study.

Since a tiny percentage of patients create a large portion of the impact, identifying these patients beforehand would allow for designing better cost containment measures. Early identification could help design targeted interventions for these higher risk patients who could then be part of more effective, specially designed disease or case management programs. Early identification could help defer or mitigate extremely negative outcomes.

This approach also ensures that the different players shaping the healthcare market be satisfied. Insurers and employers who pay for the healthcare costs would stand to gain considerably from reduced costs. Employers in particular have an added incentive as this would reduce other “indirect costs” incurred due to the time taken by the patient to return to work and the resulting loss of productivity. Additional benefits for these players include better return on investment due to an improvement in the allocation of available resources and a basis for the establishment of capitation reimbursements. On the other hand, such an approach does not directly impact providers and suppliers who provide services to the patients. However, before achieving such gains, the imbalanced nature of the data provides a considerable challenge for accurate prediction.

As a part of this study, we propose a predictive risk modeling approach to identify high-risk patients. We use data mining and machine learning techniques to design such an approach as they are known to work well with large data and in particular when the data collection has been automated and performance takes precedence over interpretability (Scheffer, 2002). Data mining has been successfully used in the past for financial applications like credit card fraud detection, stock market prediction, and bankruptcy prediction (Zhang and Zhou, 2004).

Healthcare data provides a unique opportunity for knowledge discovery using data mining while also presenting considerable challenges. Despite the success of data mining in various areas, it hasn't been regularly used to tackle these challenges though limited examples exist (Anderson, Balkrishnan, & Camacho, 2004; Cios & Moore, 2002; Li et al., 2005). We study the possibility of applying data mining techniques to aid in healthcare risk modeling, where we aim to forecast whether a patient would be of high cost for the next year based on data from the current year.

2 RELATED WORK

2.1 Learning from Imbalanced Data

Due to the existence of high-risk, high-cost patients healthcare expenditure data is highly skewed. As a result, it is essential to pay attention to the data imbalance when dealing with such data. This is not uncommon and has been observed in applications like credit card fraud detection, network intrusion detection, insurance risk management, text classification, and medical diagnosis. The problems of dealing with imbalanced data for classification have been widely studied by the data mining and machine learning community (Chawla, Japkowicz, & Kolcz, 2004). Most classification algorithms assume that the class distribution in the data is uniform. Since the metric of classification accuracy is based on this assumption, the algorithms often try to improve this faulty metric while learning.

The two most common solutions to this problem include non-random sampling (under-sampling or down-sampling, over-sampling or up-sampling and a combination of both) and cost-sensitive learning. Both solutions have a few drawbacks (most importantly, under-sampling might neglect few key instances while over-sampling might cause overfitting) but they have shown improvement over conventional techniques (McCarthy, Zabar, & Weiss, 2005; Weiss & Provost, 2001).

Various studies have compared over-sampling, under-sampling and cost-sensitive learning. While some found that there was little difference in the results from these methods, others found one among them to be the best. Results from different studies are inconclusive in selecting the best among them (Batista, Prati, & Monard, 2004; Drummond & Holte, 2003; Maloof, 2003; McCarthy et al., 2005). The use of a combination of under-sampling and over-sampling has also been found to provide improved results over the individual use of these techniques. Additionally, it has been found using varying ratios of the minority and majority classes that the best results were generally obtained when the minority class was overrepresented in the training data (Estabrooks, Jo, & Japkowicz, 2004; Weiss & Provost, 2001). The use of synthetically generated instances for the minority class has also been proposed (Chawla, Bowyer, Hall, & Kegelmeyer, 2002) but the prudence of using this technique for highly varied instances in healthcare data needs to be evaluated.

Despite the reported success of these techniques in other domains, none have been applied with

respect to healthcare expenditure data in the past. In this study, we explore the possibility of using non-random sampling as a key element in creating predictive models for identifying high-risk patients. Preliminary work has confirmed the usefulness of this approach (Moturu et al., 2007).

2.2 Techniques and Predictors

Healthcare data sets have been used in the past to predict future healthcare utilization of patients where the goal varied from being able to predict individual expenditures to the prediction of total healthcare expenditures. Typically, various regression techniques have been employed in the past with varying success for these tasks but the assumptions of independence, normality and homoscedasticity are not satisfied by the skewed distribution of the costs. Regression techniques generally tend to predict the average cost for a group of patients satisfactorily but on an individual basis, the predictions aren't too accurate. Other approaches include the transformation of the distribution to match the assumptions of the analysis technique and the use of the Cox proportional hazards model (Diehr, Yanez, Ash, Hornbrook, & Lin, 1999).

Apart from these statistical methods, multiple risk-adjustment models that can forecast individual annual healthcare expenses are currently available. These can be used to predict high-cost patients by setting a cost threshold. Popular models including Adjusted Clinical Groups (ACG), Diagnostic Cost Groups (DCG), Global Risk-Adjustment Model (GRAM), RxRisk, and Prior Expense show comparable performance (Meenan et al., 2003).

The performance of predictive modeling techniques is highly dependent on the data and features used. Different sources have provided data for the prediction of future utilization. Self-reported health status information gathered from patients using surveys has been used to predict medical expenditures (Fleishman, Cohen, Manning, & Kosinski, 2006) and group patients into cost categories (Anderson et al., 2004). Unlike these studies, our work employs administrative claims-based data. For such data both demographic and disease-related features have proven to be useful in the past. Demographic variables like age have been known to work well as predictors for expenditure. Disease-related information in the form of comorbidity indices has been used in the past as predictors of healthcare costs and the use of both inpatient and outpatient information was found to be useful (Perkins et al., 2004). However, simple count

measures like number of prescriptions and number of claims were found to be better predictors of healthcare costs than comorbidity indices (Farley, Harrdley, & Devine, 2006). Though the performance of comorbidity indices might vary, disease-related information is still a key predictor. Such information from various utilization classes such as inpatient, outpatient and pharmacy information has been used in the past, either separately or together to predict cost outcomes. Combining information from different utilization classes has been found to be useful (Zhao et al., 2005). In this study we use a set of features similar to those that have proven useful in the past together with data mining techniques that haven't been explored with respect to this area.

3 PREDICTIVE RISK MODELING

3.1 Data and Features

The substantially large amount of data in AZHQ necessitates the selection of a specific subset for analysis. The requirement for a multi-year claims-based data set representing patients of varied demographics and containing disease-related information from various utilization classes, AHCCCS data is well-suited for risk modeling. Despite being only a small part of AZHQ, AHCCCS data provides a large sample size of 139039 patients.




Figure 1: Illustration of Training and Test Data Sets.

Four hundred and thirty seven demographic and disease-related features, either categorical or binary, were extracted from the original AHCCCS data. The patients were categorized into the minority or rare class (high-cost) and the majority class based on the paid amount. Figure 1 depicts the structure of the data and its division into training and test data. Since the goal is to predict future healthcare costs, features from one year and class from the following year have been used together. Training data was constructed with features from 2002 and class from 2003 while test data was constructed with features from 2003 and class from 2004.

The demographic variables employed include age category (ages in groups of five), gender, race (Asian, Black, Hispanic, Native American, White and Other), marital status (single, married, divorced, separated and widowed) and county. Age and gender have been included due to previous success while race, location and marital status have been added as they could impact both financial and health aspects.

We avoid comorbidity or multimorbidity indices due to lack of flexibility. To allow the inclusion of inpatient, outpatient and emergency department information, International Classification of Diseases (ICD) procedure codes have been further grouped into twenty major diagnostic categories (MDC). For pharmacy data, the classification has been derived from the National Drug Code (NDC) classification with 136 categories. The practice of discounting billed charges in the healthcare industry requires that the amounts paid for the services are used as measures of costs rather than the amounts charged. Payments are used in this study and we select two different thresholds for the separation of high-cost patients. These thresholds of \$50,000 (954 or 0.69 % high-cost patients) and \$25,000 (3028 or 2.18% high-cost patients) ensure that the resultant data is sufficiently highly skewed.

3.2 Analysis

Knowledge discovery using data mining requires clear understanding of the problem domain and the nuances of the data. These are achieved in the previous sections. Further, the analysis consists of three major steps. The first step is data preprocessing and is considered one of the most important parts of data mining. This is followed by the application of data mining techniques on training data to learn an appropriate model. Finally, this model is evaluated on test data using suitable evaluation metrics.

Training and test data are created in the data preprocessing step with required features being extracted from the data. The creation of a training data set provides a major challenge. The large size of the data makes the learning task tedious and necessitates the sampling of instances to reduce size. The nature of imbalanced data sets, which invariably result in poor performance while using conventional analysis techniques, needs to be taken into consideration for the selection of appropriate training instances. To address this challenge, non-random sampling has been employed as a combination of over-sampling the minority class and under-sampling the majority class to create a training sample. This approach is reasonable as it

has been employed successfully with such data in the past. Though the use of an equal number of training instances from both classes seems intuitive, it has been suggested that a higher number of instances from the minority class might improve sensitivity (Weiss & Provost, 2001). We evaluate this suggestion using multiple training samples with varying proportions of the two classes.

The next step is the creation of predictive models. We have preliminarily tested a variety of popular classification algorithms to focus on the challenge of learning from the training data. Out of the algorithms tested, five have worked considerably better. These include AdaBoost (with 250 iterations of a Decision Stump classifier), LogitBoost (also with 250 iterations of a Decision Stump classifier), Logistic Regression, Logistic Model Trees, and the Support Vector Machine (SVM) classifier.

Performance evaluation provides the final challenge in our analysis. Since the data is highly skewed, traditional measures like accuracy aren't particularly useful. We propose the following four evaluation metrics to gauge performance:

- Sensitivity: Sensitivity corresponds to the proportion of correctly predicted instances of the minority class with respect to all such instances of that class. It is equal to the number of true positives over the sum of true positives and false negatives.

$$S_T = \frac{N_{TP}}{N_{TP} + N_{FN}}$$

- Specificity: Specificity corresponds to the proportion of correctly predicted instances of the majority class with respect to all such instances of that class. It is equal to the number of true negatives over the sum of true negative and false positives.

$$S_P = \frac{N_{TN}}{N_{TN} + N_{FP}}$$

- F-measure: F-measure is typically used as a single performance measure that combines precision and recall and is defined as the harmonic mean of the two. Here we use it as a combination of sensitivity and specificity.

$$F_M = \frac{2 * S_T * S_P}{S_T + S_P}$$

- G-mean: G-mean typically refers to geometric mean and in this study it is the geometric mean of sensitivity and specificity.

$$G_M = \sqrt{S_T * S_P}$$

To evaluate the performance of predictive risk models, it is necessary to understand the relevance of their predictions. The identification of high-cost patients allows for targeted interventions and better case management. Therefore, identifying most of these patients would prove useful. Such high sensitivity is achieved with a corresponding decrease in specificity, which is acceptable due to the cost benefits from identifying a large percentage of the high-cost patients. Consider the following example of two predictive models created using non-random and random sampling whose predictions are depicted through a confusion matrix in table 1. Identifying a limited number of high-cost patients (32 as opposed to 675) with greater prediction accuracy means that a large percentage of high-cost patients are unidentified and therefore a considerable portion of the health and cost benefits are unattainable. Alternatively holding targeted interventions and providing effective disease management for 22487 patients (675 correct and 21812 incorrect) could result in health benefits for the actual high-risk patients and cost benefits for the employers and insurers. This example indicates the need for high sensitivity along with an acceptable trade-off between specificity and sensitivity.

Table 1: Random vs. Non-random Sampling.

	Non-Random Sample		Random Sample	
	Positive	Negative	Positive	Negative
Predicted Positive	675	21812	32	82
Predicted Negative	279	116273	922	138003

3.3 Predictive Modeling

Recall that our preliminary results indicate the usefulness of non-random sampling for predictive modeling. Further, we identified five classification algorithms that show promise and delineated four measures for performance evaluation considering the imbalance of data. These elements set the stage for an empirical study designed to markedly indicate the usefulness of non-random sampling to our approach for predictive modeling. Further, this sampling technique is applied on suitably varied training data samples. Additionally, two class thresholds are used to check for the robustness of our approach to differently skewed data sets. These experiments help to provide a comparative outlook of our approach and also indicate its benefits and flexibility.

4 EMPIRICAL STUDY

We first provide details of our experimental design along with the software environment, algorithms and then discuss experimental results.

4.1 Experimental Design

Employing the AHCCCS data as depicted in Section 3.1, we evaluate the predictions across an extensive range of experiments. All experiments have been performed using the Weka software (Witten & Frank, 2005). Training data is created from the data set with features from 2002 and class from 2003. The model learned from this training data is used to predict on the test data set with features from 2003 and class from 2004. Non-random sampling was used to create training data as a default. The default class threshold used was \$50,000. For each experiment, the five algorithms listed previously have been used to create predictive models with a goal of identifying the best one. The following dimensions were used for comparison.

4.1.1 Random Versus Non-Random Sampling

Experiments across this dimension were designed to depict the differences in performance between the sampling techniques. One set of experiments used random sampling where 50% of the data was randomly selected for training. Another set of experiments used non-random sampling where the minority class was over-sampled and the majority class was under-sampled. Twenty different random samples were obtained for both classes, with every sample containing 1,000 instances. The resulting training data sample contained 40,000 instances.

4.1.2 Varying Proportions of the Minority Class Instances in the Training Data

These experiments were designed to evaluate the differences of learning using non-randomly sampled data with varied proportions of rare class instances. Multiple training data sets were created with proportions of instances from the minority class being 10%, 25%, 40%, 60%, 75% and 90%. Random samples of 1000 instances each were drawn both classes according to the appropriate proportion for that training data set. However, the total number of instances was maintained at 40,000. For example, the training set with 40% rare class instances had 16 random samples from that class resulting in 16,000 instances. Six different non-randomly sampled

training data sets were obtained in addition to the existing one with equal instances from both classes.

4.1.3 Varying the Class Threshold

Two different thresholds (\$50,000 and \$25,000) for the differentiation of high-cost patients have been used for the various training data samples described in Section 4.1.2 to assess whether our approach is robust to variations along this boundary.

4.2 Results and Discussion

4.2.1 Importance of Non-Random Sampling

Both random and non-random samples are drawn from the same data set to form training data in order to build predictive models. The purpose of this experiment is twofold: (1) to verify whether non-random sampling is indeed necessary as suggested in our preliminary analysis, and (2) to use a baseline to compare predictions from the two techniques. It is apparent from Table 2 that random sampling provides very poor sensitivity with less than ten percent of the high-cost patients identified correctly. We can also consider a baseline model where patients are predicted to be in the same class as they were in the previous year. Such a model performs better with a sensitivity of 0.276 and a specificity of 0.993 for this data set resulting in an F-measure of 0.432 and a G-mean of 0.524. The low sensitivity indicates that not many high-cost patients remain in that category the following year making predictive modeling more difficult. Non-random sampling shows a marked improvement but as one would expect, this comes with a loss in specificity. Nevertheless, the F-measure and G-mean are much higher indicating that the trade-off between sensitivity and specificity is better than the baseline. These results clearly indicate the effectiveness of non-random sampling for predictive modeling.

4.2.2 Classification Algorithm Performance

Five different classification algorithms were used to learn predictive models across the experiments with the purpose of identifying the best among them. Recall that these algorithms were selected over many other algorithms based on our preliminary analysis. Results from Table 2 (and similar comparisons in Section 4.2.3 as shown in Figure 2) clearly indicate that these five algorithms perform consistently well with very similar sensitivity and specificity making it difficult to select the best one. One can only conclude that any of these algorithms

could be used to learn a suitable predictive model from a non-randomly sampled training data set. Combining results in Section 4.2.1, we conclude that all classification models perform similarly poorly or well with random or non-random sampling. Hence, non-random sampling plays an instrumental role in significantly boosting performance.

Table 2: Random vs. Non-random Sampling.

Algorithm	Comparison	S _T	Sp	F _M	G _M
AdaBoost	Random	0.019	1	0.037	0.138
	Non-Random	0.668	0.85	0.748	0.754
LogitBoost	Random	0.063	0.999	0.118	0.251
	Non-Random	0.646	0.894	0.75	0.760
Logistic Regression	Random	0.058	0.999	0.109	0.241
	Non-Random	0.646	0.899	0.752	0.762
Logistic Model Trees	Random	0	1	0	0.000
	Non-Random	0.632	0.902	0.743	0.755
SVM	Random	0.004	1	0.008	0.063
	Non-Random	0.594	0.919	0.722	0.739

Table 3: Varying class proportions in training data.

Rare Class Percentage		S _T	Sp	F _M	G _M
10	Threshold: \$25000	0.324	0.981	0.487	0.564
	Threshold: \$50000	0.289	0.985	0.447	0.534
25	Threshold: \$25000	0.534	0.945	0.682	0.710
	Threshold: \$50000	0.463	0.958	0.625	0.666
40	Threshold: \$25000	0.637	0.898	0.746	0.757
	Threshold: \$50000	0.602	0.919	0.727	0.744
50	Threshold: \$25000	0.637	0.863	0.733	0.742
	Threshold: \$50000	0.646	0.894	0.750	0.760
60	Threshold: \$25000	0.764	0.799	0.781	0.782
	Threshold: \$50000	0.731	0.843	0.783	0.785
75	Threshold: \$25000	0.895	0.682	0.774	0.782
	Threshold: \$50000	0.847	0.742	0.791	0.792
90	Threshold: \$25000	0.979	0.475	0.640	0.682
	Threshold: \$50000	0.945	0.553	0.698	0.723

4.2.3 Using Varied Class Proportions

Using a higher proportion of minority class instances in the training data sample is expected to improve results (Weiss & Provost, 2001). Experiments were designed to evaluate this expectation and this trend is observed with our data as well. Table 3 depicts the results for this comparison using the LogitBoost algorithm. Using a higher proportion of minority class instances in the sample (60% and 75%) performs better than an equal proportion as indicated by both the F-measure and the G-mean. A receiver operating characteristics (ROC) curve can be generated from these different proportions. Figure 2 depicts such a curve that provides a better visual representation of the improvement in results. It has to be noted that the two cases with improved results

(60% and 75%) show a very different trade-off between sensitivity and specificity despite similar values for the F-measure and G-mean. Such an observation indicates a unique opportunity to deal with differences across the differently proportioned samples. It is difficult to identify a suitable trade-off without the availability of data that can establish the cost benefits to be gained from a particular trade-off. In such a scenario, such experiments can be invaluable as they provide multiple trade-offs to choose from. Upon the availability of information about the cost benefits, the suitably proportioned training data sample can be selected for analysis.




Figure 2: ROC Curve.

4.2.4 Varying the Class Threshold

Two thresholds for the differentiation of cost categories have been used to indicate the robustness of our approach to changes in class threshold. We observe from Table 3 that results for both the thresholds are comparable with the higher threshold proving slightly better as indicated by F-measure and G-mean. Since the training data is balanced by non-random sampling, the slight underperformance from the data with lower threshold could be due to the fact that there are more patients closer to the lower threshold, increasing the chance of an error in prediction. This particular comparison serves to indicate the adaptability of our approach while using differently skewed data sets for predictive modeling.

5 CONCLUSIONS

Predictive risk modeling for forecasting high-cost patients is an important area of research and this study provides a look at a beneficial new technique using a real-world data set. Results indicate that creating training data using non-random sampling helps balance the challenges resulting from the

skewed nature of healthcare cost data sets. Further, over-representing the minority class in the training data helps improve performance. Our study manifests the significance of sampling in building predictive risk models. However, it is hard to judge the best trade-off between specificity and sensitivity when there is no available data on the cost benefits. In this sense, using varied proportions of instances from the two classes in the training data can work as a boon in disguise. When data on cost benefits is available, one can test the use of different proportions of instances from the two classes to select the case with the best cost benefit. This makes our approach for predictive modeling much more adaptable.

Our comparison of classification algorithms for this task indicates that all of the selected ones work almost equally well. Though we find that it is hard to choose between these algorithms, results indicate to future users a handful of appropriate classification techniques to be used along with non-random sampling for predictive modeling. Our proposed approach creates a model by learning from the data and is therefore not restricted to the use of a specific type of data or features. Further, the threshold for high-cost patients is tunable and can be varied depending on the goals of a particular study. All these taken together signify the flexibility of predictive risk modeling for future high-cost patients using classification techniques to learn from non-randomly sampled training data and the benefits that can be obtained from such analyses.

Considering the variation in data, predictors and evaluation metrics, comparison with previous studies is improper. Nevertheless, the ROC curve in Figure 2 is similar (the performance of the best model is comparable) to that obtained for existing risk-adjustment models (Meenan et al., 2003). The numbers are also better (our results double the sensitivity at about the same level of specificity) than a decision-tree based predictive modeling technique (Anderson et al., 2005). This validates the usefulness of this technique that is further enhanced by its flexibility. As can be observed, sampling is the most important component of this technique and is very beneficial for predictive modeling.

Predictive risk modeling is a useful technique with practical application for numerous employers and insurers in the goal to contain costs. We provide a promising approach that is valuable, flexible and proven to be successful on real-world data. Nevertheless, there is further scope to improve the interpretation of these results. It is commonly observed that a considerable percentage of high-cost

patients do not remain that way every year. Also, two patients could share very similar profiles with only one of them being high-cost. Studying these seemingly anomalous patients could provide a better understanding of how a high-cost patient is different from other patients. In addition, the current sampling approach and available classification techniques could be further tuned to improve results.

Apart from these possibilities, the most promising future direction is in working with key data partners. This avenue provides the opportunity to obtain information on the cost containment methods used and their efficiency as well as real data on the cost benefits obtained from previous predictive models. Working with such partners, we endeavor to provide a reasonable, patient-specific answer to this question that would significantly impact cost containment in the healthcare industry.

REFERENCES

- Anderson, R.T., Balkrishnan, R., & Camacho, F. (2004). Risk Classification of Medicare HMO Enrollee Cost Levels using a Decision-Tree Approach. *Am J Managed Care*, 10(2), 89-98.
- Batista, G. E. A. P. A., Prati, R. C., & Monard, M. C. (2004). A Study of the Behavior of Several Methods for Balancing Machine Learning Training Data. *ACM SIGKDD Explorations Newsletter*, 6(1), 20-29.
- Bodenheimer, T. (2005). High and Rising Health Care Costs. Part 1: Seeking an Explanation. *Ann Intern Med*, 142, 847-854.
- Berk, M. L., & Monheit, A. C. (2001). The Concentration of Health Care Expenditures, Revisited. *Health Affairs*, 20 (2), 9-18.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16, 321-357.
- Chawla, N. V., Japkowicz, N., & Kolecz, A. (2004). Editorial: Special Issue on Learning from Imbalanced Data Sets, *ACM SIGKDD Explorations Newsletter*, 6(1), 1-6.
- Cios, K. J., & Moore, G. W. (2002). Uniqueness of Medical Data Mining. *Artificial Intelligence in Medicine*, 26(1-2), 1-24.
- Diehr, P., Yanez, D., Ash, A., Hornbrook, M., & Lin, D. Y. (1999). Methods For Analysing Health Care Utilization and Costs. *Ann Rev Public Health*, 20, 125-144.
- Drummond, C., & Holte, R. C. (2003). C4.5, Class Imbalance, and Cost Sensitivity: Why Under-Sampling beats Over-Sampling. *ICML Workshop Learning From Imbalanced Datasets II*, 2003.
- Estabrooks, A., Jo, T., & Japkowicz, N. (2004). A Multiple Resampling Method For Learning From Imbalanced Data Sets. *Computational Intelligence*, 20(1), 18-36.
- Farley, J. F., Harrdley, C. R., & Devine, J. W. (2006). A Comparison of Comorbidity Measurements to Predict Health care Expenditures. *Am J Manag Care*, 12, 110-117.
- Fleishman, J. A., Cohen, J. W., Manning, W.G., & Kosinski, M. (2006). Using the SF-12 Health Status Measure to Improve Predictions of Medical Expenditures. *Med Care*, 44(5S), I-54-I-66.
- Li, J., Fu, A. W., He, H., Chen, J., Jin, H., McAullay, D. et al. (2005). Mining Risk Patterns in Medical Data. *Proc 11th ACM SIGKDD Int'l Conf. Knowledge Discovery in Data Mining (KDD'05)*, 770-775.
- Maloof, M. (2003). Learning When Data Sets are Imbalanced and When Costs are Unequal and Unknown. *ICML Workshop Learning From Imbalanced Datasets II*, 2003.
- McCarthy, K., Zabar, B., & Weiss, G. Does cost-sensitive learning beat sampling for classifying rare classes? *Proc 1st Int'l Workshop on Utility-based data mining (UBDM '05)*, 69-77.
- Meenan, R. T., Goodman, M. J., Fishman, P. A., Hornbrook, M. C., O'Keeffe-Rosetti, M. C., & Bachman, D. J. (2003). Using Risk-Adjustment Models to Identify High-Cost Risks. *Med Care*, 41(11), 1301-1312.
- Moturu, S.T., Johnson, W.G., & Liu, H. (2007). Predicting Future High-Cost Patients: A Real-World Risk Modeling Application. *Proc IEEE International Conference on Bioinformatics and Biomedicine 2007*, Accepted.
- Perkins, A. J., Kroenke, K., Unutzer, J., Katon, W., Williams Jr., J. W., Hope, C. et al. (2004). Common comorbidity scales were similar in their ability to predict health care costs and mortality. *J Clin Epidemiology*, 57, 1040-1048.
- Scheffer, J. (2002). Data Mining in the Survey Setting: Why do Children go off the Rails? *Res. Lett. Inf. Math. Sci.*, 3, 161-189.
- Weiss, G.M., & Provost, F. (2001). The Effect of Class Distribution on Classifier Learning: An Empirical Study, (Dept. Computer Science, Rutgers University, 2001), tech report ML-TR-44.
- Witten, I. H., & Frank, E. (2005). *Data Mining: Practical machine learning tools and techniques, 2nd Edition*, San Francisco: Morgan Kaufmann.
- Zhang, D., & Zhou, L. (2004). Discovering Golden Nuggets: Data Mining in Financial Application. *IEEE Trans. Sys. Man Cybernet.*, 34(4), 513-522.
- Zhao, Y., Ash, A. S., Ellis, R. P., Ayanian, J. Z., Pope, G. C., Bowen, B. et al. (2005). Predicting Pharmacy Costs and Other Medical Costs Using Diagnoses and Drug Claims. *Med Care*, 43(1), 34-43.
- Zweifel, P., Felder, S., & Meiers, M. (1999). Ageing of Population and Health Care Expenditure: A Red Herring. *Health Econ*, 8, 485-496.

MEDICAL IMAGE UNDERSTANDING THROUGH THE INTEGRATION OF CROSS-MODAL OBJECT RECOGNITION WITH FORMAL DOMAIN KNOWLEDGE

Manuel Möller, Michael Sintek, Paul Buitelaar

*German Research Center for Artificial Intelligence, Trippstadter Str. 122, 67663 Kaiserslautern, Germany
manuel.moeller@dfki.de, michael.sintek@dfki.de, paul.buitelaar@dfki.de*

Saikat Mukherjee

*Siemens Corporate Research, 755 College Road East, Princeton, USA
saikat.mukherjee@siemens.com*

Xiang Sean Zhou

*Siemens Medical Solutions, 51 Valley Stream Parkway, Malvern PA, USA
xiang.zhou@siemens.com*

Jörg Freund

*Siemens Medical Solutions, Hartmannstr. 16, 91052 Erlangen, Germany
joerg.freund@siemens.com*

Keywords: Semantic Web, Ontologies, NLP, Medical Imaging, Image Retrieval.

Abstract: Rapid advances in medical imaging scanner technology have increased dramatically in the last decade the amount of medical image data generated every day. By contrast, the software technology that would allow the efficient exploitation of the highly informational content of medical images has evolved much slower. Despite the research outcomes in image understanding and semantic modeling, current image databases are still indexed by keywords assigned by humans and not by the image content. The reason for this slow progress is the lack of scalable and generic information representations capable of overcoming the high-dimensional nature of image data. Indeed, most of the current content-based image search applications are focused on the indexing of certain image features that do not generalize well and use inflexible queries. We propose a system combining medical imaging information with semantic background knowledge from formalized ontologies, that provides a basis for building universal knowledge repositories, giving clinicians a fully cross-lingual and cross-modal access to biomedical information of all forms.

1 INTRODUCTION

Rapid advances in imaging technology have dramatically increased the amount of medical image data generated daily by hospitals, pharmaceutical companies, and academic medical research¹. Technologies such as 4D 64-slice CT, whole-body MR, 4D Ultrasound, and PET/CT can provide incredible detail and a wealth of information with respect to the human body anatomy, function and disease associations. This increase in the volume of data has brought about significant advances in techniques for analyzing such

¹For example, University Hospital of Erlangen, Germany, has a total of about 50 TB of medical images. Currently they have approx. 150.000 medical examinations producing 13 TB per year.

data. The precision and sophistication of different image understanding techniques, such as object recognition and image segmentation, have also improved to cope with the increasing complexity of the data.

However, these improvements in analysis have not resulted in more flexible or generic image understanding techniques. Instead, the analysis techniques are very object specific and not generic enough to be applied across different applications. Throughout this paper we will address this fact as “lack of scalability”. Consequently, current image search techniques, whether for Web sources or for medical PACS (Picture Archiving and Communications System), are still dependent on the subjective association of keywords to images for retrieval.

One of the important reasons behind this lack

of scalability in image understanding techniques has been the absence of *generic* information representation structures capable of overcoming the feature-space complexity of image data. Indeed, most current content-based image search applications are focused on indexing syntactic image features that do not generalize well across domains. As a result, current image search technology does not operate at the *semantic* level and, hence, is not scalable.

We propose to use hierarchical information representation structures, which integrate state-of-the-art object recognition algorithms with generic domain semantics, for a more scalable approach to image understanding. Such a system will be able to provide direct and seamless access to the informational content of image databases.

Our approach is based on the following main techniques:

- Integrate the state-of-the-art in semantics and image understanding to build a sound bridge between the symbolic and the subsymbolic world. This cross-layer research approach defines our road-map to quasi-generic image search.
- Integrate higher level knowledge represented by formal ontologies that will help explain different semantic views on the same medical image: structure, function, and disease. These different semantic views will be coupled to a backbone ontology of the human body.
- Exploit the intrinsic constraints of the medical imaging domain to define a rich set of queries for concepts in the human body ontology. The ontology not only provides a natural abstraction over these queries but also statistical image algorithms could be associated to semantic concepts for answering these queries.

Our focus is on filling the gap between what is current practice in image search (*i. e.*, indexing by keywords) and the needs of modern health provision and research. The overall goal is to empower the medical imaging content-stakeholders (clinicians, pharmaceutical specialists, patients, citizens, and policy makers) by providing flexible and scalable semantic access to medical image databases. Our short term goal is to develop a basic image search engine and prove its functionality in various medical applications.

In 2001 Berners-Lee and others published a visionary article on the Semantic Web (Berners-Lee et al., 2001). The use-case they described was about the use of meta-knowledge by computers. For our goals we propose to build a system on existing Semantic Web technologies like RDF (Brickley and Guha, 2004) and OWL (McGuinness and van Harme-

len, 2004) which were developed to lay the foundations of Berners-Lee's vision. From this point of view it is also a Semantic Web project.

Therefore we propose a system that combines medical imaging information with semantic background knowledge from formalized ontologies and provides a basis for building universal knowledge repositories, giving clinicians *cross-modal* (independent from different modalities like PET, CT, ultrasound) as well as *cross-lingual* (independent of particular languages like English and German) access to various forms of biomedical information.

2 GENERAL IDEA

There are numerous advanced object recognition algorithms for the detection of particular objects on medical images: (Hong et al., 2006) at anatomical level, (Tu et al., 2006) at disease level and (Comaniciu et al., 2004) at functional level. Their specificity is also their limitation: Existing object recognition algorithms are not at all generic. Given an arbitrary image it still needs human intelligence to select the right object recognizers to apply to an image. Aiming to gain a pseudo-general object recognition one can try to apply the whole spectrum of available object recognition algorithms. But it turns out that in generic scenarios even with state-of-the-art object recognition tools the accuracy is below 50 percent (Chan et al., 2006; Müller et al., 2006).

In automatic image understanding there is a semantic gap between low-level image features and techniques for complex pattern recognition. Existing work aims to bridge this gap by ad-hoc and application specific knowledge. In contrast our objective is to create a formal fusion of semantic knowledge and image understanding to bridge this gap to support more flexible and scalable queries.

For instance, human anatomical knowledge tells us that it is almost impossible to find a heart valve next to a knee joint. Only in cases of very severe injuries these two objects might be found next to each other. But in most cases the anatomical intuition is correct and, hence, the background knowledge precludes the recognition of certain anatomical parts given the presence of other parts. It is in this use of formalized knowledge that ontologies² come into play within our framework.

In the context of medical imaging it is necessary to define image semantics for parts of human anatomy.

²According to Gruber (Gruber, 1995), an ontology is a specification of a (shared) conceptualization.

In this domain the expert’s knowledge is already formalized in comprehensive ontologies like the *Foundational Model of Anatomy* (Rosse and Mejino, 2003) for human anatomy or the *International Statistical Classification of Diseases and Related Health Problems* (ICD-10)³ of World Health Organization for a classification of human diseases. These ontologies represent a rich medical knowledge in a standardized and machine interpretable format.

In contrast to current work which defines ad-hoc semantics, we take the novel view that within a constrained domain the semantics of a concept is defined by the queries associated with it. We will investigate which types of queries are asked by medical experts to ensure that the necessary concepts are integrated into the knowledge base. We believe that in IR applications this view will allow a number of advances which will be described in the following sections.

We chose the medical domain as our area of application. Unlike common language and many other scientific areas the medical domain has clear definitions for its technical terms. Ambiguities are rare which eases the task of finding a semantic abstraction for a particular text or image. However, our framework is generic and can be applied to other domains with well-defined semantics.

3 ASPECTS OF USING ONTOLOGIES

Ontologies (usually) define the semantics for a *set of objects* in the world using a *set of classes*, each of which may be identified by a particular symbol (either linguistic, as image, or otherwise). In this way, ontologies cover all three sides of the "semiotic triangle" that includes *object*, *referent*, and *sign* (see Fig. 1). *I.e.*, an *object* in the world is defined by its *referent* and represented by a *symbol* (Ogden and Richards, 1923—based on Peirce, de Saussure and others). Currently, ontology development and the Se-




Figure 1: Semiotic Triangle.

semantic Web effort in general have been mostly di-

³<http://www.who.int/classifications/icd/en/>

rected at the *referent* side of the triangle, and much less at the *symbol* side. To allow for automatic multilingual and multimedia knowledge markup a richer representation is needed of the linguistic and image-based symbols for the object classes that are defined by the ontology (Buitelaar et al., 2006).

From our point of view a semantic representation should not be encapsulated into a single module. Instead we think that a layered approach as shown in Fig. 2 has a number of advantages.




Figure 2: Interacting Layers in Feature Extraction and Representation.

Once there is a representation established at the semantic level there are a number of benefits compared to conventional IR systems. For a more detailed description of the abstraction process see Sect. 4.

Cross-Modal Image Retrieval. Current systems for medical IR depend on the modality of the stored images. But in medical diagnosis very different imaging techniques are used such as PET, CT, ultrasonography, or time series data from 4D CT *etc.* Each technique produces images with characteristic appearance. For tumor detection, for example, often PET (to identify the tumor) and CT (to have a view on the anatomy) are combined, to formulate a precise diagnosis with a proper localization of the tumor. The proposed system will allow to answer queries based on semantic similarity and not only visual similarity.

Full Modality Independence. Cross-Modality even can be driven another step forward by integrating documents of any format into one single database. We plan to also include text documents like medical reports and diagnoses. On the level of semantic representation they will be treated like the images. Accordingly, the system will be able to answer queries not only with images but also with

text documents including similar concepts as in the retrieved images.

Improved Relevance of Results. Current search engines retrieve documents which contain the keyword from the query. The documents in the result set are ranked by various techniques using information such as their inter-connectivity, statistical language models, or the like. For huge datasets search by keyword often returns very large result sets. Ranking by relevance is hard.

This holds for low-level image retrieval as well. Here only two similarity measures are applicable: through visual similarity which can be completely independent from the object and context and via a comparison between keywords and image annotations. With current IR systems the user is forced to use pure keyword-based search as a detour while in fact he or she is searching for documents and/or images including certain concepts.

Our notion of keyword-based querying goes beyond current search engines by mapping keywords from the query to ontological concepts. Our system provides the user with a semantic representation. That allows the user to ask for a concept or a set of concepts with particular relationships. This allows far more precise queries than a simple keyword-based retrieval mechanism and likewise better matching between query and result set.

Inferencing of Hidden Knowledge. By mapping the keywords from a text-based query to ontological concepts and the use of semantics the system is able to infer⁴ implicit results. This allows us to retrieve images which are not annotated explicitly with the query concepts but with concepts related to them through the ontology.

To represent the complex knowledge of the medical domain and allow a maximum of flexibility in the queries we will have to enrich the ontology by rules and allow to use rules in the queries. Another point will be an integration of spatial representation of anatomical relations as well as an efficient implementation of spatial reasoning.

4 LEVELS OF SEMANTIC ABSTRACTION

Our notion of semantic imaging is to ground the semantics of a human anatomical concept on a set of

⁴We aim at using standard OWL reasoners like Racer, FaCT++ or Pellet.

queries associated with it. The constrained domain of a human body enables us to have a rich coverage of these queries and, consequently, define image semantics at various levels of the hierarchy of the human anatomy.

Fig. 3 gives an overview of the different abstraction levels in the intended system. For the proposed system we want to take a step beyond the simple dichotomy between a symbolic and subsymbolic representation of images. Instead, from our perspective there is a spectrum ranging from regarding the images as simple bit vectors over color histograms, shapes and objects to a fully semantic description of what is depicted. The most formal and generic level of representation is in form of an ontology (*formal ontological modeling*). The ontology holds information about the general structure of things. Concrete entities are to be represented as *semantic instances* according to the schema formalized in the ontology.

To emphasize the difference to the dichotomic view, we call the lower end of this spectrum *informal* and the upper end *formal* representation. From our perspective the abstraction has to be modeled as a multi-step process across several sublevels of abstraction. This makes it easier to close the gap between the symbolic and subsymbolic levels from the classic perspective. Depending on the similarity measure that is to be applied for a concrete task different levels of the abstraction process will be accessed.

If a medical expert searches for images that are *looking* similar to the one he or she recently got from an examination, the system will use low-level features like histograms or the bit vector representation. In another case the expert might search for information about a particular syndrome. In that case the system will use features from higher abstraction levels like the semantic description of images and texts in the database to be able to return results from completely different modalities.

We believe that text documents have to be understood in a similar way. Per se, a text document is just a string of characters. This is similar to regarding an image as a sheer bit vector. Starting from the string of characters, in a first step relatively simple methods can be used to identify terms. In further steps technologies from concept based Cross Language Information Retrieval (CLIR)⁵ are applied to map terms in the documents to concepts in the ontology (Volk et al., 2003; Vintar et al., 2003; Carbonell et al., 1997; Eichmann et al., 1998). CLIR currently can be divided into three different methods: approaches based on bilingual dictionary lookup or

⁵The research project *MUCHMORE* (<http://muchmore.dfki.de/>) was focused on this aspect.

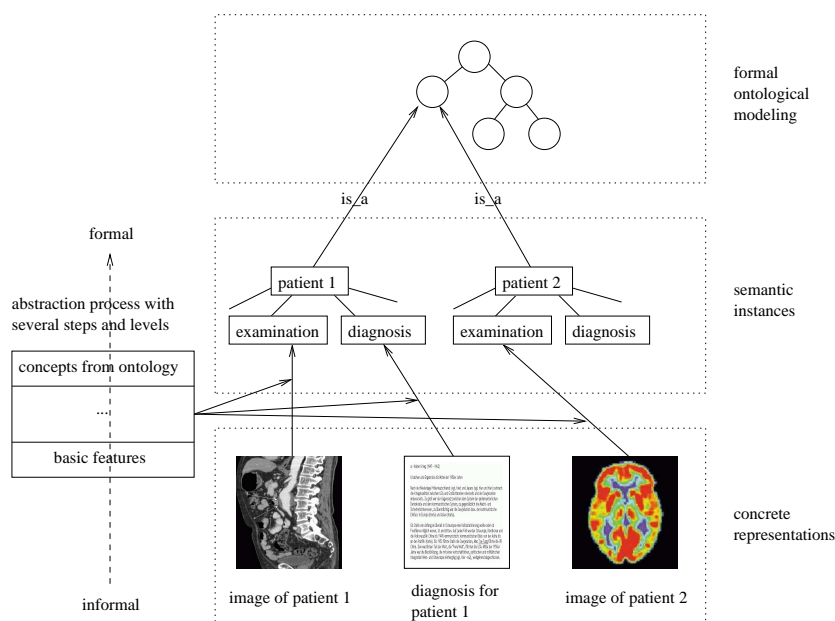


Figure 3: Abstraction Model.

Machine Translation (MT); corpus based approaches utilizing a range of IR-specific statistical measures; and concept-driven approaches, which exploit semantic information to bridge the gap between surface linguistic form and meaning. The latter allows not only to cross borders between modalities but also between languages to make the system *cross-modal* and *cross-lingual*.

Modern hospitals often have tens of terabytes of concrete data from medical cases. In many cases this data is already stored in digital format. In our abstraction model it belongs to the level of *concrete representations* (cf. Fig. 3). It will be stored in a conventional database. The entities populating the database can be images from all available medical imaging modalities. Since the system is designed to be completely independent from the modality it can be also texts describing examination results, diagnoses, medical publications, etc.

Recall that, in our framework, we take the novel view that the semantics of a concept is defined by the queries associated with it. These queries can either be defined by domain experts, such as physicians, radiologists, nurses, or they can be mined from text corpora. Medical knowledge repositories, such as clinical books, journals, etc., contain information on image-centric questions on different body parts which are of interest to physicians. For instance, typical image queries of the heart could be about detection of the ventricles, etc. Similar information also exists in physician reports, laboratory notes, etc.

We will make use of methods for extracting infor-

mation from unstructured text for automatically discovering these queries. There is a significant research on information extraction (Laender et al., 2002) from natural language text as well as learning patterns of free-text questions from examples (Ravichandran and Hovy, 2002; Ling et al., 2002). We will apply these techniques for identifying relevant questions and their answers in natural language text. This will help us in collecting a broad coverage of possible questions associated with concepts.

5 SUPPORTED QUERY SCENARIOS

Iterative Retrieval Process. In most present systems the retrieval is started by a query sent by the user which is subsequently answered by a result set. We think that the retrieval should be merely understood as an iterative process. The user starts the process of information retrieval by submitting a query. In many cases this query will be either too general resulting in an result set which is too huge. The other extreme is a query which is too specific leading to an empty result set. To support the user with the navigation through the available information we aim to have a close interaction between user and system. Step by step the user can refine the search query using aid which is given proactively from the domain knowledge of the system.

We envision three primary ways in which users will query the semantic imaging platform. Users can

query either through sample images, or pose structured queries using conceptual descriptors, or use natural language to describe queries. In the following, we explain each of these different methods.

Query by Sample Image. Basically there are two different approaches to image based queries. The *first approach* retrieves images from the database which are *looking* similar. Only low-level image features are used to select results for this type of query. The ability to match the image of a current patient to similar images from a database of former medical cases can be of great help in assisting the medical professional in his diagnosis (see Sect 4) we believe that image understanding has to be an abstraction over several levels. To answer queries by sample image we will make use of the more informal features extracted from the images. The support for these queries is based on state-of-the-art similarity-based image retrieval techniques (Deselaers et al., 2005).

Today there are various image modalities in modern medicine. Many diseases like cancer require to look at images from different modalities to formulate a reliable diagnosis (see example in Sect. 3). The *second approach* therefore takes the image from the query and extracts the semantics of what can be seen on the image. Through mapping the concrete image to concepts in the ontology, an abstract representation is generated. This representation can be used for a matching on the level of image semantics with other images in the database. Applying this method makes it possible to use a CT image of the brain to search for images from all available modalities in the database (see Fig. 4–6).

Query by Conceptual Descriptions. Similar to the use of SQL for querying structured relational databases, special purpose languages are also required for querying semantic metadata. Relying on well-established standards we propose using a language on top of RDF, such as SPARQL, for supporting generic structured semantic queries.

Query by Natural Language. From the point of the medical expert having a natural language interface is very important. Through a textual interface the user directly enters keywords which are mapped to ontology concepts. Current systems like the IRMA-Project (see Sect. 6) only allow to search for keywords which are extracted offline and stored as annotations. Since our system has to compose a semantic representation of each query, the ontological background knowledge can be used in an iterative process of query refine-

ment. Additionally, it will be possible to use complete text documents as queries.

In cases where the system cannot generate a semantic representation—due to missing knowledge about a knew syndrome, therapy, drug or the like—it will fall back to a normal full text search. If the same keyword is used frequently this can be used as evidence that the foundational ontology has to be extended to cover the corresponding concept(s).

6 RELATED WORK

Most current work in content based image retrieval models object recognition as a probabilistic inferencing problem and use various mathematical methods to cope with the problems of image understanding. These techniques use image features which are tied to particular applications and, hence, suffer from a lack of scalability.

Among extant work in fusing semantics with image understanding, (Hunter, 2001) describes a technique for modeling the MPEG-7 standard, which is a set of structured schemas for describing multimedia content, as an ontology in RDFS. There has been some research (Barnard et al., 2003; Lavrenko et al., 2003; Lim, 1999; Carneiro and Vasconcelos, 2005; Town, 2006; Mezaris et al., 2003; Mojsilovic et al., 2004) on semantic imaging relying primarily on associating word distributions to image features. However, these works used hierarchies of words for semantic interpretation and did not attempt to model image features themselves in levels of abstraction. Furthermore, the lack of formal modeling made these techniques susceptible to subjective interpretations of the word hierarchies and, hence, were not truly scalable. Especially in the context of medical imaging, our notion of semantics is tied to information gathered from physics, biology, anatomy, *etc.* This is in contrast to perception-based subjective semantics in these works.

The goal of the project “Image Retrieval in Medical Applications” (IRMA) (Lehmann et al., 2003) was an automated classification of radiographs based on global features with respect to imaging modality, direction, body region examined and the biological system under investigation. The aim was to identify image features that are relevant for medical diagnosis. These features were derived from a database of 10.000 a-priori classified and registered images. By means of prototypes in each category, identification of objects and their geometrical or temporal relationships are handled in the object and the knowledge layer, respectively. Compared to the system proposed



Figure 4: CT scan.




Figure 5: PET scan.




Figure 6: 3D model.

in this paper the IRMA project only used a taxonomy without a formal representation of anatomical, functional, or pathological relations. Therefore it lacks any inference of implicit statements in the query as well as using the background knowledge for a relevance test during the automatic annotation of new images and to score results.

Mechouche and colleagues did research in the area of integrating ontologies enriched by rules with low level object recognizers for semantic annotation of brain MRI images (Mechouche et al., 2007). While they are only focusing on one organ and only one modality they prove that the combination of formal high-level knowledge and low-level feature extraction can be beneficial.

7 CONCLUSION AND FUTURE WORK

In this paper we proposed a close integration of sub-symbolic pattern recognition algorithms and semantic domain knowledge represented in formal ontologies. The vision is to combine the techniques from both fields to bridge the gap between a symbolic and sub-symbolic world for a generic understanding of medical images and text. We take the novel view that within a constrained domain the semantics of a concept, as described in a physics-based ontology of human anatomy, is defined by typical queries associated with it. Thus, our framework is different from research which fuses image understanding with subjective semantics.

The use of formal ontologies, and their reasoning capabilities, forms the essence behind better information retrieval. By abstracting from the syntactic content representation, it is possible to perform semantic matching between queries and the content. Additionally, the user is provided with an extremely flexible interface which allows cross-modal as well as cross-lingual queries. By matching at the level of semantic

concepts, abstracting from syntactic representations where necessary, and using low-level features where necessary, our framework enables scalable querying on images and text across different anatomical concepts.

Extensive research has been done on the extraction of semantics from text documents. Therefore this component of our system can rely on an existing state-of-the-art. The next research task in implementing the proposed system will be the integration of formalized background knowledge with low-level object recognition algorithms. The section Related Work shows that this is currently an area of intensive work. All existing approaches lack the generality which we aim at for the proposed system. Therefore developing a truly generic and scalable integration will be our next step.

ACKNOWLEDGEMENTS

This research has been supported in part by the THESEUS Program in the MEDICO Project, which is funded by the German Federal Ministry of Economics and Technology under the grant number 01MQ07016. The responsibility for this publication lies with the authors.

REFERENCES

- Barnard, K., Duygulu, P., Forsyth, D., de Freitas, N., Blei, D., and Jordan, M. (2003). Matching words and pictures. *Journal of Machine Learning Research*, 3:1107–1135.
- Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The semantic web – a new form of web content that is meaningful to computers will unleash a revolution of new possibilities. *Scientific American*.
- Brickley, D. and Guha, R. (2004). RDF vocabulary description language 1.0: RDF Schema.

- Buitelaar, P., Sintek, M., and Kiesel, M. (2006). A lexicon model for multilingual/multimedia ontologies. *Proceedings of the 3rd European Semantic Web Conference (ESWC06)*.
- Carbonell, J. G., Yang, Y., Frederking, R. E., Brown, R. D., Geng, Y., and Lee, D. (1997). Translingual information retrieval: A comparative evaluation. *International Journal of Computational Intelligence and Applications*, 1:708–715.
- Carneiro, G. and Vasconcelos, N. (2005). A database centric view of semantic image annotation and retrieval. *SIGIR '05: Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 559–566.
- Chan, A. B., Moreno, P. J., and Vasconcelos, N. (2006). Using statistics to search and annotate pictures: an evaluation of semantic image annotation and retrieval on large databases. *Proceedings of Joint Statistical Meetings (JSM)*.
- Comaniciu, D., Zhou, X., and Krishnan, S. (2004). Robust real-time myocardial border tracking for echocardiography: an information fusion approach. *IEEE Transactions in Medical Imaging*, 23 (7):849–860.
- Deselaers, T., Weyand, T., Keyzers, D., Macherey, W., and Ney, H. (2005). FIRE in ImageCLEF 2005: Combining content-based image retrieval with textual information retrieval. *Working Notes of the CLEF Workshop*.
- Eichmann, D., Ruiz, M. E., and Srinivasan, P. (1998). Cross-language information retrieval with the UMLS metathesaurus. *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 72–80.
- Gruber, T. R. (1995). Toward principles for the design of ontologies used for knowledge sharing. In *International Journal of Human-Computer Studies*, volume 43, pages 907–928.
- Hong, W., Georgescu, B., Zhou, X. S., Krishnan, S., Ma, Y., and Comaniciu, D. (2006). Database-guided simultaneous multi-slice 3D segmentation for volumetric data. In Leonardis, A., Bischof, H., and Prinz, A., editors, *Journal of the European Conference on Computer Vision (ECCV 2006)*, volume 3954, pages 397–409. Springer-Verlag.
- Hunter, J. (2001). Adding multimedia to the semantic web - building an MPEG-7 ontology. *International Semantic Web Working Symposium (SWWS)*.
- Laender, A., Ribeiro-Neto, B., Silva, A., and Teixeira, J. (2002). A brief survey of web data extraction tools. In *Special Interest Group on Management of Data (SIGMOD) Record*, volume 31 (2).
- Lavrenko, V., Manmatha, R., and Jeon, J. (2003). A model for learning the semantics of pictures. *Proceedings of the Neural Information Processing Systems Conference*.
- Lehmann, T., Güld, M., Thies, C., Fischer, B., Spitzer, K., Keyzers, D., Ney, H., Kohnen, M., Schubert, H., and Wein, B. (2003). The IRMA project. A state of the art report on content-based image retrieval in medical applications. *Proceedings 7th Korea-Germany Joint Workshop on Advanced Medical Image Processing*, pages 161–171.
- Lim, J.-H. (1999). Learnable visual keywords for image classification. In *DL '99: Proceedings of the fourth ACM conference on Digital libraries*, pages 139–145, New York, NY, USA. ACM Press.
- Ling, C., Gao, J., Zhang, H., Qian, W., and Zhang, H. (2002). Improving encarta search engine performance by mining user logs. *International Journal of Pattern Recognition and Artificial Intelligence*, 16 (8).
- McGuinness, D. L. and van Harmelen, F. (2004). OWL web ontology language overview.
- Mechouche, A., Golbreich, C., and Gibaud, B. (2007). Towards a hybrid system using an ontology enriched by rules for the semantic annotation of brain MRI images. In Marchiori, M., Pan, J., and de Sainte Marie, C., editors, *Lecture Notes in Computer Science*, volume 4524, pages 219–228.
- Mezaris, V., Kompatsiaris, I., and Strintzis, M. (2003). An ontology approach to object-based image retrieval. *International Conference on Image Processing (ICIP)*.
- Müller, H., Deselaers, T., Lehmann, T., Clough, P., Kim, E., and Hersh, W. (2006). Overview of the ImageCLEFmed 2006 medical retrieval and annotation tasks. In *Accessing Multilingual Information Repositories*, volume 4022 of *Lecture Notes in Computer Science*. Springer-Verlag.
- Mojsilovic, A., Gomes, J., and Rogowitz, B. (2004). Semantic-friendly indexing and querying of images based on the extraction of the objective semantic cues. *International Journal of Computer Vision*, 56:79–107.
- Ravichandran, D. and Hovy, E. (2002). Learning surface text for a question answering system. *Proceedings of the Association of Computational Linguistics Conference (ACL)*.
- Rosse, C. and Mejino, R. L. V. (2003). A reference ontology for bioinformatics: The Foundational Model of Anatomy. In *Journal of Biomedical Informatics*, volume 36, pages 478–500.
- Town, C. (2006). Ontological inference for image and video analysis. *International Journal of Machine Vision and Applications*, 17(2):94–115.
- Tu, Z., Zhou, X. S., Bogoni, L., Barbu, A., and Comaniciu, D. (2006). Probabilistic 3D polyp detection in CT images: The role of sample alignment. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2:1544–1551.
- Vintar, S., Buitelaar, P., and Volk, M. (2003). Semantic relations in concept-based cross-language medical information retrieval. In *Proceedings of the ECML/PKDD Workshop on Adaptive Text Extraction and Mining (ATEM)*.
- Volk, M., Ripplinger, B., Vintar, S., Buitelaar, P., Raileanu, D., and Sacaleanu, B. (2003). Semantic annotation for concept-based cross-language medical information retrieval. *International Journal of Medical Informatics*.

SOFTWARE AGENTS REPRESENTING MEDICAL GUIDELINES

John McGrory, Frank Clarke
Dublin Institute of Technology, Dublin 8, Ireland
john.mcgrory@dit.ie, frank.clarke@dit.ie

Jane Grimson
School of Computing, Trinity College Dublin, Dublin 2, Ireland
jane.grimson@tcd.ie

Peter Gaffney
Adelaide & Meath Hospital, Incorporating the National Children's Hospital (AMNCH), Tallaght, Dublin 24, Ireland
peter.gaffney@amnch.ie

Keywords: Clinical laboratory validation, patient-centred, patient-focused, agents, computer-interpretable-guideline (CIG), guidelines and protocols. Software Agents representing medical guidelines.

Abstract: Guidelines are self-contained documents which healthcare professionals reference to obtain specific disease or medical condition knowledge for a particular population cohort. They view these documents and apply known facts about their patients to access useful supportive information to aid in developing a diagnosis or manage a condition. Traditional CIG models decompose these guidelines into workflow plans, which are then called using certain motivational trigger conditions controlled by a centralised management engine. Therefore, CIG guidelines are not self-contained documents, which specialise in a particular condition or disease, but are effectively a list of workflow plans, which are called and used when the patient information is available. The software BDI agent offers an alternative approach which more closely matches the modus operandi of narrative based medical guidelines. An agent's beliefs capture information attributes, plans capture the deliberative and action attributes, and desire captures the motivational attributes of the guideline in a self-contained autonomous software module. This synergy between the narrative guideline and the BDI agent offers an improved solution for computerising medical guidelines when compared to the CIG approach.

1 INTRODUCTION

Clinical guidelines are condition focused documents through which domain specific aims, goals, procedures, plans and normal reference ranges are disseminated to healthcare professionals (Browne, 2005) (Oosterhuis, 2004). To maximise influence on a market audience these documents are written in domain specific languages and ontologies (e.g. cardiology, neurology and paediatrics). The purpose of these documents are to guide the reader, and streamline activities around a particular medical condition, organ or disease using evidence based supportive information. When a clinical or laboratory guideline is developed by an expert group they focus on best practice for the specific disease, condition or organ. They include all relevant knowledge, logic and motivational aspects they

deem necessary to adequately describe the domain. Therefore, guidelines are autonomous self-contained documents, which can be used in whole or in part, to provide supporting information for healthcare decision-making, once their meaning is not taken out of context.

Clinicians and laboratory technologists care for patients so it is their responsibility to filter through these guidelines acting on a patient's behalf. They must make use of the maximum decision-making support offered by these resources based on the known facts about the individual. This is termed *patient-centred validation*. However, the enormous quantity and presentation style of these documents makes it difficult for professionals to quickly identify relevant guidelines, and extract information and intended logic contained within them, and apply

them usefully in a patient-centred fashion (Peleg et al., 2003).

One approach to overcome these issues is to use Computer-Interpretable-Guidelines (CIG). The underlying operation of CIG models is to decompose the narrative guideline into separate workflow activities and management rules. The management rules from each guideline are linked together centrally using an inference engine. This inference engine is constructed by the CIG developers using rules that link the various workflow activities together. These management rules provide the motivation for a centralised inference engine to choose particular workflow activities depending on the patients known characteristics (e.g. weight, gender, height). As more guidelines are decomposed and added, the number of workflow activities and the size and complexity of the centralised inference engine increases. Using this technique the original self-contained knowledge developed by the guideline author's is fragmented and the true context is lost.

Agent-oriented software operates on the notion that human intelligence and decisions can be synthesised by managing tangible elements, such as beliefs, goals, sub-goals, plans and intentions. Each agent is an autonomous software module that is self-contained, has its own inference engine and a set of beliefs, goals and plans. The goals, beliefs and plans are encoded into the agent and the inference engine interprets them. When a goal is chosen the inference engine searches plans that match currently known beliefs and prioritises them. The most appropriate plan is selected and executed to perform some activity or choose another goal. In addition to its inference engine the agent approach also permits separate agents to socialise, work in groups and collaborate to solve common goals.

There are characteristic similarities between guidelines and a particular type of agent, namely the Belief-Desire-Intention (BDI) agent as shown in Table 1.

Table 1: BDI to Guideline mapping.

BDI Structure		Characteristics		Guideline Structure
Belief	↔	Informational attributes	↔	Facts
Desire	↔	Motivational attributes	↔	Goals and aims
Intension	↔	Deliberative attributes	↔	Selection of Actions

By encoding the agent with the knowledge, logic and motivational components of a guideline, it creates an opportunity for an individual agent to

represent a single guideline. As additional guidelines are added, new agents are created to represent them. This technique means that the original self-contained knowledge developed by author's remains as an integrated software module and the context is retained.

The thrust of this paper is to demonstrate that a single software agent possesses the ability to reproduce the function and content of a clinical guideline as a 1:1 ratio. This is because a synergy exists between guideline and agent characteristics.

2 COMPARISON OF GUIDELINES, CIG'S AND AGENTS

Guidelines play a crucial supporting role in the detection, diagnosis, treatment, and supervision of diseases in patients, in a modern healthcare setting. In some cases these documents aid in the planning of treatments, quantifying of medication amounts, monitoring of patient responses, which could all have serious patient safety issues if inaccurate information was used. Therefore, it is vital the knowledge, logic and motivation contained within the guideline are clearly understood with reference to a specific patient. But how do CIG's and agents reproduce the knowledge, logic, action and more specifically motivation of guidelines? In the following subsections the function, use and methodology behind guidelines, CIGs and agents are discussed.

2.1 Guideline Infrastructure

Guidelines are autonomous literature sources which seldom rely on any other knowledge, except some base-level understanding. They provide information, logic and motivation in order to describe aspects of a particular domain. The guideline document itself is made up of a series of aims, goals, procedures and plans. It was never expected guidelines would be used exclusively in a form of cookbook medicine, but more in a complementary and supportive role by providing domain specific information. Therefore, these documents are rarely used in whole, but more in part, to aid in healthcare delivery. This means that the guideline should be viewed holistically, and not just as a collection of separate workflow elements taken out of context.

In summary, a clinical guideline is a self-contained document which specialises in a specific

condition, organ or disease. Medical professionals interface with this document in light of the known patient information to support their patient-centred healthcare delivery.

2.2 Computer-Interpretable-Guidelines

The underlying principle behind the CIG approach is to disassemble the guideline into separate workflow activities, and then orchestrate rules to link these activities together centrally based on the presented patient data. The majority of CIG's decompose the medical guideline into an Arden Syntax Medical Logic Module (MLM) or similar by dividing the narrative guideline into 'evoked', 'data', 'logic' and 'action' slots. These slots are used to develop a software workflow activity plan, complete with a triggering condition, but contain no motivation or goal. As more guidelines are added, rules controlling the selection, merging and division of workflow activities increases. The centralised CIG control engine which manages these activities provides the motivation or goal for the software to operate. As all guidelines are coupled together via the centralised set of management rules it is not possible to distribute this activity among a number of computer systems. This is fundamentally different to the true operation of a medical guideline where many copies of the same guideline exist.

In summary, a CIG is a list of workflow plans which are called and used when the relevant patient information is available. It is centralised and the original knowledge is fragmented and cannot be truly interfaced with.

2.3 Software Agents

An agent is an autonomous self-contained software module which is programmed using belief, goal and plan attributes. Each agent has its own inference engine which interprets these attributes in order to perform some activity. The principle of a BDI agent's operation is based on a belief capturing the informational attributes, the desire capturing motivational attributes and the intention capturing the deliberative attributes of an agent (Rao et al., 1995). An agent shell is a generalised version of an agent which can be adapted for a wide variety of different applications. There are a number of BDI agent shells available such as Jason, 3APL and Jadex. Jadex 0.932 was the agent shell used in this research. Although there are characteristic similarities between agents and guidelines the raw Jadex shell is not capable of capturing guideline

functionality without some modification. To this end the Jadex agent shell was adapted and this modified version of the Jadex agent was titled Autonomous Socialising Knowledge agent (ASK-agent).

Using the ASK-agent model a guideline can be decomposed into workflow activities, but instead of having a centralised inference engine to manage all the guidelines, each guideline has its own inference engine. The guidelines have the ability to communicate with other resources using message passing. The ASK-agent encodes the separate workflow paths and motivational management rules of the guideline within a single autonomous software module. There is no centralised engine managing the separate ASK-agents, therefore the approach is distributed. Each agent registers itself, complete with the services it provides, language and ontology it uses with a Directory Facilitator (DF). The DF acts as a *goldenpages* for agents, allowing them to be looked-up, accessed and used as independent resources.

Table 2: Adapted MLM to Agent map.

Adapted MLM Slot	ASK-Agent Component
<i>Evoke</i>	The ASK-agent's action trigger to perform some task, or perform goal.
<i>Data</i>	Belief, the facts the agent uses to trigger logic. A belief can be any Java object.
<i>Logic</i>	Condition, precondition or trigger based on beliefs used for the selection of an appropriate workflow activity plan or goal.
<i>Action</i>	Execution of the workflow activity.
NEW <i>Achieve Goal</i>	Motivates the ASK-agent to achieve a specific goal to reach some desired state, such as determine patient gender, age, PatientID.
NEW <i>Maintain Goal</i>	Motivates the ASK-agent to maintain a specific condition (e.g. maintain the plausibility value above 60%)
NEW <i>Query Goals</i>	Motivates the ASK-agent to seek alternative avenues on an IF basis without committing to the workflow activity (e.g. test alternative paths before committing to the path, such as would knowing the gender of the patient alter the outcome?).
NEW <i>Meta-level reasoning</i>	If more than one workflow activity is triggered simultaneously, but only one should be chosen the ASK-agent must choose the most appropriate course of action to take.
NEW <i>Modal reasoning</i>	If data stored in the beliefs or received in a message has a level of truth, but cannot be established as 100% true or false, then the ASK-agent can weight its selection of an appropriate action (e.g. the patient could have kidney failure, probability of 60%, but the patient could also have liver failure, probability of 55%)

The agent's beliefs capture and encode the information attributes of the guideline, the agent's

plans capture and encode the deliberative and action attributes of the guideline, and the desire captures and encodes the motivational attributes of the guideline. This can be accurately and efficiently completed by using an adapted MLM to agent map as shown in Table 2, which was used in a proof of concept agent application developed by the authors.

The extended MLM structure provides a means to extract the guideline motivational components in the form of achieve, maintain and query goals. These motivational components provide the driving force behind the agent's activity. Using this approach an agent can act faithfully and autonomously on behalf of the guideline in a self-contained capacity. Thus when patient specific information is presented to the individual agents, via message passing, they have the ability to apply their encoded knowledge and logic, and provide a supportive response based solely on that information. Using this operation dynamic an agent module can also make use of the maximum supportive response from the other separate agent's (which represent guidelines) based on the known facts about the individual patient. By providing a framework which allows separate agents broadcast supportive communications to each other, the agent approach offers the opportunity for the data to be validated in a patient-centred fashion.

In summary, an ASK-agent is a self-contained software module which specialises in a specific condition, organ or disease affecting a particular population cohort. Other agents (software or human) interface with this software knowledge in light of the known patient information, and receive supportive information from the agent to aid in their patient specific healthcare delivery.

3 PROCEDURE TO CONVERT A GUIDELINE TO ASK-AGENT

To provide a consistent technique for developing ASK-agent modules, which can act faithfully and accurately on behalf of a guideline, it is important to document, and explain the conversion process. The development of an ASK-agent module from a specific narrative guideline is based on the cycle shown in Figure 1. This procedure uses the expanded MLM mapping to the Agent Definition File (ADF) Extended Mark-up Language (XML) file and Java plans. The procedure is divided into six steps each of which is described in the following subsections.

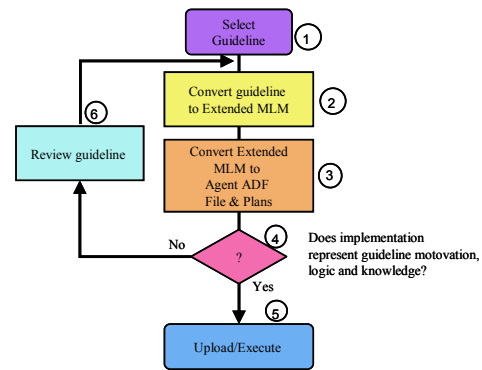


Figure 1: Encoding of expert agent procedure.

3.1 Select Guideline

The medical institution chooses a guideline that matches and complements their core activities. The types of guidelines being discussed are text based documents with conclusions being derived using rule-based or statistic-based decisions or numerical analysis. Neither the proposed ASK-agent nor existing CIG approaches can manage guidelines which have been developed using graphical imagery and charts.

3.2 Convert Guideline to Extended MLM

The guideline document is an amalgamation of a number of workflow activities which are triggered using motivational goals. Consider for illustration purposes the example of a generalised guideline extract shown in Table 3.

Table 3: Generalised guideline extract.

Evolve:

Arrival of laboratory results for analyte results X, Y and Z.

Data:

PatientAge
PatientGender

Logic:

IF (25Years < PatientAge >50Years) AND (40U/L < X > 130U/L) THEN perform *ActivityA*.

IF (PatientGender = Male AND (8U/L < Y > 40U/L)) THEN perform *ActivityB*.

IF (PatientGender = Female AND (6U/L < Y > 35U/L)) THEN perform *ActivityB*.

IF (PatientGender = Female AND (50U/L > Z)) THEN perform *ActivityC*.

Action:

ActivityA
ActivityB
ActivityC

The *evoked* slot relates to the triggering condition which is used to start this specific workflow path. The *data* information is retrieved from the Laboratory Information System (LIS) by completing an SQL query, or by executing an archetype query. An archetype is a recently employed term used in medical informatics. It is described in the CEN pre-standard prEN13606:2004(E) as a reusable, template model which presents a specific viewpoint of a domain reference model. The *logic* relates to a rule for selecting an activity. Finally the action is the activity to be performed, such as “print message to screen” or perform another activity.

3.3 Convert MLM into Agent ADF File and Java Plans

The agent is encoded by directly mapping the extended MLM components to the ASK-agent. Code 1 illustrates how the Jadex ADF is encoded using XML. This file contains MLM slot components such as *data*, *evoked*, *action*, *logic*, *query goal*, *achieve goal*, *maintain goal* and *meta-level priority*. Meta-level reasoning is where two or more plans or goals can be triggered at the same time and a choice is needed on which one takes priority. For example consider the generalised guideline extract shown in Table 2. The first logical statement may be triggered calling on *activityA* to be executed. But any of the other three logic statements could also be triggered simultaneously calling on *activityB* or *activityC*. Do all need to be performed together, or should one activity be completed first? These decisions are traditionally realised centrally within the CIG engine, and are not part of the distributed MLM modules themselves. This motivation must be extracted from the original narrative guideline document and added to the BDI, providing the motivating force behind the agent’s activity. Using Jadex this ordering of activities can be completed using priority attributes. Other components utilised by the expert agent model but not explicitly declared in either the narrative guideline or MLM are language, ontology and service descriptions. If the guideline was developed using a specific language or ontology, it is important that the sender of messages to it are aware of this, so a richer form of communication can be established. To this end the Jadex agent permits the declaration of an agent’s language and ontology. The service description is used to register services the guideline can perform. In practice a guideline can perform many services, for example validate analyte Y and Z results if the *PatientAge* is known, or validate analyte X and Z

results if *PatientGender* is known. These are services the ASK-agent is able to provide and are described using the Service Description component. For each service described using the ASK-agent eight parameters are used to describe it. These parameters are *name*, *type*, *ownership*, *GuidelineReference*, *InformationNeeded*, *ValidationType*, *ontology* and *protocol*. These services are registered with the DF *goldenpages* agent. *Type* is used to identify the field of expertise to which the ASK-agent belongs (e.g. liver, kidney, and haematology). *Ownership* is used to identify the guideline implementation owner (e.g. department). *GuidelineReference* is used to identify the guideline that the ASK-agent represents (e.g. ISBN). *InformationNeeded* is a list of the data the ASK-agent requires in order to process the validation request (e.g. SerumALK_P, SerumALT, SerumGGT). The *InformationNeeded* titles can be developed to match the medical domain titles. Both ontology and protocol where described earlier. Their presence in the service description is so other ASK-agents can identify the ontology and protocol used by this particular agent.

```

<!-- First Heading -->
<!-- <H3> Second Heading</H3>
      <H4> Third Heading</H4> -->
<agent xmlns:xsi="http..."
<imports>...</imports>
<capabilities>...</capabilities>
<beliefs>Data_MLM_Slot </beliefs>
  <goals>Evoke_MLM_Slot, Query
Goals, Achieve Goals, Maintain
Goals and Meta-level priority
</goals>
  precondition, conditions and
  triggers activating the goal are
  represented by the Logic_MLM
  Slot
  <plans>Action_MLM_Slot, Meta-
  level priority </plans>
  precondition, conditions and
  triggers activating the plans are
  represented by the Logic_MLM
  Slot
<events>...</events>
<languages>...</languages>
<ontologies>...</ontologies>
<expression>...</expression>
<servicedescription>...
</servicedescription>
</agent>

```

Code 1: Jadex agent ADF XML constructor.

3.4 Expert Agent Testing

Once the Agent ADF XML file and associated plans are developed the implementation needs to be thoroughly tested to ensure it represents the original guideline's motivations, logic and knowledge. Guidelines are autonomous self-contained entities; therefore the agent must act in a similar fashion. If the implementation of the guideline achieves this then it is uploaded into the main agent platform and can be used as part of the ASK-agent group.

3.5 Upload and Execute Agent Module

The developed ADF XML file and Java plans which represent the guideline are loaded into the JADE Remote Agent Management. Using the JADE Remote Agent Management GUI each agent is selected in turn using the "Start New Agent" button and loaded into the manager. Once installed the agent is able to act as part of the social group. The uploaded agents can be switched on, off or suspended by using the JADE Remote Agent Management GUI.

3.6 Review Guideline Conversion to ASK-Agent

If the agent module autonomous action does not represent the intended operation and understanding of the original narrative guideline, its encoding should be reviewed. This review continues until the agent represents the true operation of the guideline.

4 CONCLUSIONS

The purpose of this research is not to dismiss the CIG approach, but re-examine it from a processing point of view. This research demonstrates the agency approach offers a solution to the management of medical guidelines electronically, in a manner similar to that provided by the original narrative guidelines themselves. This is because of the synergy between the knowledge base, plans, decisions, action, goals and the self-contained components between guidelines and agents. It illustrates that this method permits the guideline content to be stored within one standalone agent module, and the entire knowledge, logic and data can be interfaced with by other agents when necessary. The ASK-agent approach permits a single guideline to be expressed within a standalone agent, but for true patient-centred validation to occur these separate agents must be able to communicate with

other modules and collaborate. These last two aspects are covered in more detail in separate papers presented at this conference. So in summary the key differences between the CIG and ASK-agent approaches are detailed in Table 4.

Table 4: Generalised guideline extract.

<i>Aspect</i>	<i>CIG</i>	<i>ASK-Agent</i>
<i>Processing</i>	Centralised	Distributed
<i>Guideline content</i>	Fragmented within CIG engine.	Integrated as a standalone module.
<i>Method of accessing information</i>	Data entry and the execution of CIG rules.	Message passing.
<i>Encoding of guideline content</i>	Menu or/and graphical.	XML and Java.
<i>Size of file(s) created</i>	For a single guideline the CIG engine is small. However, as more MLM's are added the CIG engine application file becomes very large. Processing cannot be distributed.	For a single guideline the ASK-agent is larger than the CIG. However, as model is distributed the system file size is not a limitation

ACKNOWLEDGEMENTS

The authors wish to acknowledge and thank the work undertaken by developers of the Jadex agent platform at Distributed Systems and Information Systems Group, Computer Science Department, University of Hamburg, <http://vsis-www.informatik.uni-hamburg.de/projects/jadex/>

REFERENCES

- Browne E.D., 2005. *Workflow Modelling of Coordinated Inter-Health-Provider Care Plans*, School of Computer and Information Science, University of South Australia, Adelaide, PhD Thesis January 2005.
- Oosterhuis W.P., Bruns D.E., Watine J., Sandberg S., Horvath A. 2004. *Evidence-Based Guidelines in Laboratory Medicine: Principles and Methods*, published in Clin Chem, year 2004, Vol50 pgs 806-818.
- Peleg M., Tu S., Bury J., Ciccarse P., Fox J., Greenes R., Hall R., Johnson P., Jones N., Kumar A., Miksch S., Quaglini, S., Shortliffe E., Stefanelli M. 2003. *Comparing Computer-interpretable Guideline Models: A Case-study Approach*, accessed via Journal of the American Medical Informatics Ass. Volume 10 Number 1 Jan / Feb 2003.
- Rao A., Georgeff M., 1995. *BDI Agents: From Theory to Practice*, proceedings of the first international conference on multi agent systems (ICMMAS-95) San Francisco USA 1995.

FORMAL ANALYSIS OF INTELLIGENT AGENTS FOR MODEL-BASED MEDICINE USAGE MANAGEMENT

Mark Hoogendoorn, Michel Klein, Zulfiqar Memon and Jan Treur

Vrije Universiteit Amsterdam, Department of Artificial Intelligence

De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands

{mhoogen, m.klein, za.memon, treur}@few.vu.nl

Keywords: Model-based, agent, medicine usage management, formal analysis.

Abstract: A model-based agent system model for medicine usage management is presented and formally analysed. The model incorporates an intelligent ambient agent model that has an explicit representation of a dynamical system model to estimate the medicine level in the patient's body by simulation, is able to analyse whether the patient intends to take the medicine too early or too late, and can take measures to prevent this.

1 INTRODUCTION

A challenge for medicine usage management is to achieve in a non-intrusive manner that patients for whom it is crucial that they take medicine regularly, indeed do so. Examples of specific relevant groups include independently living elderly people, psychiatric patients or HIV-infected persons. One of the earlier solutions reported in the literature provides the sending of automatically generated SMS reminder messages to a patient's cell phone at the relevant times; e.g., (Safren, Hendriksen, Desousa, Boswell, and Mayer, 2003). A disadvantage of this approach is that patients are disturbed often, even if they do take the medicine at the right times themselves, and that due to this after some time a number of patients start to ignore the messages.

A more sophisticated approach can be based on a recently developed automated medicine box that has a sensor that detects whether a medicine is taken from the box, and can communicate this to a server; cf. SIMpill (Green, 2005). This enables to send SMS messages only when at a relevant point in time no medicine intake is detected. A next step is to let a computing device find out more precisely what relevant times for medicine intake are. One way is to base this on prespecified prescription schemes that indicate at what time points medicine should be taken. However, this may be inflexible in cases that a patient did not follow the scheme precisely. To obtain a more robust and flexible approach, this paper explores and analyses possibilities to use an

automated medicine box in combination with model-based intelligent agents to dynamically determine the (estimated) medicine level over time.

The agent-based model for medicine usage management discussed was formally specified in an executable manner and formally analysed using dedicated tools. The system incorporates a model-based intelligent agent that includes an explicitly represented dynamic system model to estimate the medicine level in the patient's body by simulation. Based on this it is able to dynamically determine at what point in time the patient should take medicine, and given that, to analyse whether the patient intends to take medicine too early or too late, and to take measures to prevent this.

In this paper, Section 2 describes the multi-agent system introduced, whereas Section 3 present detailed information about the specific agents. Furthermore, Section 4 presents simulation results, and Section 5 formal analysis of these results. Finally, Section 6 is a discussion.

2 OVERVIEW OF THE SYSTEM

Figure 1 presents an overview of the entire system as considered. The top right corner shows the patient, who interacts with the medicine box, and communicates with the patient cell phone. The Medicine Box detects whether medicine is taken out of the medicine box. The Medicine Box Agent (MBA) observes this medicine box. In case, for example, the patient intends to take the medicine too

soon after the previous dose, it finds out that the medicine should not be taken at the moment (i.e., the sum of the estimated current medicine level plus a new dose is too high), and communicates a warning to the patient by a beep. Furthermore, all information obtained by this agent is passed on to the Usage Support Agent (USA). All information about medicine usage is stored in the patient database by this agent. If the patient tries to take the medicine too early, a warning SMS with a short explanation is communicated to the cell phone of the patient, in addition to the beep sound already communicated by the Medicine Box Agent.

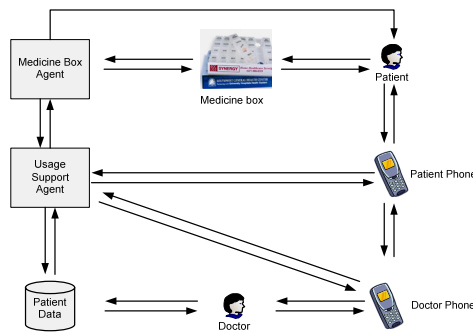


Figure 1: Multi-Agent System: Overview.

On the other hand, in case the Usage Support Agent finds out that the medicine is not taken early enough (i.e., the medicine concentration is estimated too low for the patient and no medicine was taken yet), it can take measures as well. First of all, it can warn the patient by communicating an SMS to the patient cell phone. This is done soon after the patient should have taken the medicine. In case that after some time the patient still does not take medicine, the agent can communicate an SMS to cell phone of the appropriate doctor. The doctor can look into the patient database to see the medicine usage, and in case the doctor feels it is necessary to discuss the state of affairs with the patient, he or she can contact the patient via a call using the doctor cell phone to the patient cell phone.

3 AGENT PROPERTIES

The model used for the Usage Support Agent (USA) makes (re)use of elements of the Generic Agent Model GAM described in (Brazier, Jonker, and Treur, 2000). In addition, it makes use of an explicitly represented dynamical model to for the medicine level over time within the patient. Moreover, the model for the Usage Support Agent includes a reasoning method (based on simulation) to estimate the current medicine level based on the

dynamical model and information on medicine taking in the past.

To express the agent’s internal states and processes, a state ontology partly shown in Table 1 was specified. An example of an expression that can be formed by combining elements from this ontology is

$\text{belief}(\text{leads_to_after}(I:\text{INFO_EL}, J:\text{INFO_EL}, D:\text{REAL}))$

which expresses that the agent has the knowledge that state property I leads to state property J with a certain time delay specified by D. This type of expression is used to represent the agent’s knowledge of a dynamical model of a process. Using the ontology, the functionality of the agent has been specified by generic and domain-specific temporal rules.

Table 1: Ontology for the Usage Support Agent Model.

Formalisation	Description
$\text{belief}(I:\text{INFO_EL})$	information I is believed
$\text{world_fact}(I:\text{INFO_EL})$	I is a world fact
$\text{has_effect}(A:\text{ACTION}, I:\text{INFO_EL})$	action A has effect I
$\text{leads_to_after}(I:\text{INFO_EL}, J:\text{INFO_EL}, D:\text{REAL})$	state property I leads to state property J after duration D
$\text{at}(I:\text{INFO_EL}, T:\text{TIME})$	state property I holds at time T

Note the formal form of the agent properties has been omitted below for the sake of brevity. Generic rules specify that incoming information (by observation or communication) from a source that is believed to be reliable is internally stored in the form of beliefs. When the sources are assumed always reliable, the conditions on reliability can be left out:

IB(X) From Input to Beliefs

If agent X observes some world fact, then it will believe this.

If X gets information communicated, then it will believe this.

Execution of a dynamical model is specified by:

SE(X) Simulation Execution

If it is believed that I holds at time T, and it is believed that I leads to J after time duration D, then it is believed that J holds at time T+D.

This temporal rule specifies how a dynamic model that is represented as part of the agent’s knowledge can be used by the agent to extend its beliefs about the world at different points in time.

Domain-specific rules for the Usage Support Agent are shown below. The Usage Support Agent’s specific functionality is described by three sets of temporal rules. First, the agent maintains a dynamic model for the concentration of medicine in the patient over time in the form of a belief about a *leads to after* relation.

USA1: Maintain dynamic model

The Usage Support Agent believes that if the medicine level for medicine M is C, and the usage effect of the medicine is E, then

after duration D the medicine level of medicine M is $C+E$ minus $G*(C+E)*D$ with G the decay value. Formally:

In order to reason about the usage information, this information is interpreted (USA2), and stored in the database (USA3).

USA2: Prepare storage usage

If the agent has a belief concerning usage of medicine M and the current time is T , then a belief is generated that this is the last usage of medicine M , and the intention is generated to store this in the patient database.

USA3: Store usage in database

If the agent has the intention to store the medicine usage in the patient database, then the agent performs this action.

Finally, temporal rules were specified for taking the appropriate measures. Three types of measures are possible. First, in case of early intake, a warning SMS is communicated (USA4). Second, in case the patient is too late with taking medicine, a different SMS is communicated, suggesting to take the medicine (USA5). Finally, when the patient does not respond to such SMSs, the doctor is informed by SMS (USA6).

USA4: Send early warning SMS

If the agent has the belief that an intention was shown by the patient to take medicine too early, then an SMS is communicated to the patient cell phone that the medicine should be put back in the box, and the patient should wait for a new SMS before taking more medicine.

USA5: SMS to patient when medicine not taken on time

If the agent has the belief that the level of medicine M is C at the current time point, and the level is considered to be too low, and the last message has been communicated before the last usage, and at the current time point no more medicine will be absorbed by the patient due to previous intake, then an SMS is sent to the patient cell phone to take the medicine M .

USA6: SMS to doctor when no patient response to SMS

If the agent has the belief that the last SMS to the patient has been communicated at time T , and the last SMS to the doctor was communicated before this time point, and furthermore, the last recorded usage is before the time point at which the SMS has been sent to the patient, and finally, the current time is later than time T plus a certain delay parameter for informing the doctor, then an SMS is communicated to the cell phone of the doctor that the patient has not taken medicine M .

USA7: Communicate Current Concentration

If the agent has the belief that the level of medicine M is C at the current time point then the agent informs the medicine box agent about this level.

The Medicine Box Agent has functionality concerning communication to both the patient and the Usage Support Agent. Generic temporal rules are included as for the Usage Support Agent (see above). Domain-specific temporal rules are both shown below. First of all, the observed usage of medicine is communicated to the Usage Support Agent in case the medicine is not taken too early, as specified in MBA1.

MBA1: Medicine usage communication

If the Medicine Box Agent has a belief that the patient has taken medicine from a certain position in the box, and that the particular position contains a certain type of medicine M , and taking the medicine does not result in a too high medicine concentration of medicine M within the patient, then the usage of this type of medicine is communicated to the USA.

In case medicine is taken out of the box too early, a warning is communicated by a beep and the information is forwarded to the Usage Support Agent (MBA2 and MBA3).

MBA2: Too early medicine usage prevention

If the Medicine Box Agent has the belief that the patient has taken medicine from a certain position in the box, that this position contains a certain type of medicine M , and taking the medicine results in a too high medicine concentration of medicine M within the patient, then a warning beep is communicated to the patient.

MBA3: Early medicine usage communication

If the Medicine Box Agent has a belief that the patient was taking medicine from a certain position in the box, and that the particular position contains a certain type of medicine M , and taking the medicine would result in a too high concentration of medicine M within the patient, then this is communicated to the Usage Support Agent.

4 SIMULATION

In order to show how the above presented system functions, the executable properties have been implemented in a dedicated software environment that can execute such specifications (Bosse, Jonker, Meij, and Treur., 2007). To enable creation of simulations, a patient model is used that simulates the behaviour of the patient in a stochastic manner. The model specifies four possible behaviours of the patient, each with its own probability: (1) too early intake, (2) correct intake (on time), (3) responding to an SMS warning that medicine should be taken, and (4) responding to a doctor request by phone. Based upon such probabilities, the entire behaviour of the patient regarding medicine usage can be simulated. In the following simulations, values of respectively 0.1, 0.8, 0.9 and 1.0 have been used.

Figure 2 shows an example of a simulation trace whereby the medicine support system is active.

The figure indicates the medicine level over time as estimated by the agent based on its dynamic model. Here the x-axis represents time whereas the y-axis represents the medicine level.

Note that in this case, the minimum level of medicine within the patient is set to 0.35 whereas the maximum level is 1.5. These numbers are based on the medicine half-life value, that can vary per type of medicine. For more details on the formal properties behind the simulation, and a more elaborate discussion of the results, see (Hoogendoorn, Klein, and Treur, 2007).




Figure 2: Medicine level over time.

5 FORMAL ANALYSIS

When a model such as the one described above, has been specified, it is easy to produce various simulations based on different settings, initial conditions and external events offered. Moreover, it is possible to incorporate nondeterministic behaviours by temporal rules that involve probabilistic effects (cf. Bosse, *et al.*, 2007). Thus large sets of traces can be generated. When such a set is given, it is more convenient to check them on relevant properties automatically, than going through them by hand. Furthermore, it may also be useful when insight is provided how dynamic properties of the multi-agent system as a whole depend on dynamic properties of the agents within the system, and further on, how these relate to properties of specific components within the agents. This section shows how this can be achieved. To this end a number of *dynamic properties* have been specified for different aggregation levels of the multi-agent system, cf. (Jonker and Treur 2002; Bosse, Jonker, Meij, Sharpanskykh, and Treur, , 2006). The main property considered for the system as a whole is: will the medicine level in the patient be maintained between the required minimum and maximum level? This desired situation is called ‘S’. That a value V of a variable P should be within a specific range between the lower threshold TL and the upper threshold TU , is specified as follows:

$$\text{has_value}(P,V) \wedge (V > TU \vee V < TL) \quad (S)$$

This has been applied to the variable ‘medicine_level’.

GP1 At any point in time the medicine level is between TL and TU .

$\forall T:\text{TIMEPOINT}, V:\text{REAL}:$

$$\text{state}(M, T) \models \text{has_value}(P,V) \Rightarrow (V \leq TU \wedge V \geq TL)$$

Here M is a trace, and $\text{state}(M, T)$ denotes the state in this trace at time T . Moreover, $\text{state}(M, T) \models p$ denotes that state property p holds in state $\text{state}(M, T)$.

Related to this, global properties can be defined that specify that the total number of violations of the threshold values is smaller than some maximum number, or that the total duration of the violation is smaller than some maximum period. In these definitions $\Sigma \text{case}(p, 1, 0)$ is a special construct in the

language that calculates the sum of timepoints for which a state holds.

GP2 The total number of times that the medicine level falls below TL or raises above TU is smaller than MAX_OCCURANCES .

$$\forall M:\text{TRACE}: \forall T1, T2:\text{TIMEPOINT}: \\ \Sigma \text{case}(T1 \leq T \wedge T \leq T2 \wedge \text{state}(M, T) \models S \wedge \\ \text{state}(M, T+1) \models \neg S, 1, 0) < \text{MAX_OCCURANCES}$$

GP3 The total time period that the medicine level is not between TL and TU is smaller than MAX_DURATION .

$$\forall M:\text{TRACE} \forall T1, T2:\text{TIMEPOINT}: \\ \Sigma \text{case}(T1 \leq T \wedge T \leq T2 \wedge \text{state}(M, T) \models \neg S, 1, 0) < \\ \text{MAX_DURATION}$$

5.1 Evaluation of Traces

The formal properties have been used to evaluate the usefulness of the medicine usage management system. For this, 60 simulation traces of the medicine level within a patient have been generated, with a length of 36 hours. In half of the traces the medicine usage management system was supporting the patient, in the other half the system did not take any action, but was still monitoring the medicine usage. As a consequence of the stochastic patient model (the probabilities used are the ones mentioned in Section 5), this resulted in 60 different simulation traces.

For all traces it has been checked whether, how often and how long, the medicine level is between the required values of 0.5 and 1.35. It has also been checked whether this is the case for *preferred* values. It can be assumed that, in addition to the required range for the medicine level, there is also an optimal or preferred range. Table 2 lists the total number of violations in the 30 traces with support of the system and the 30 traces without support for different maximum and minimum levels. Table 3 shows the duration of the violations time in minutes for the same set of traces.

Table 2: Total number of violations of the threshold values (used in property GP2).

number of violations	with support	without
above 1.5 (required)	2	8
below 0.35 (required)	5	18
total required	7	26

Table 3: Total duration of violations of the threshold values (used in property GP3).

duration (in minutes)	with support	without
above 1.5 (required)	0.8	25.36
below 0.35 (required)	3.53	179.13
total required	4.33	204.49
above 1.2 (preferred)	161.6	328.20
below 0.5 (preferred)	218.1	327.63
total preferred	379.7	655.83

From the figures in tables it is immediately apparent that the medicine level is much more often and for a much longer time between the required or preferred values. However, it is also clear that even with support of the system the medicine level is sometimes outside the required range. In fact, property GP1 did not hold in 5 out of the 30 simulation traces in which the system was active. An analysis of these traces revealed that this is a side-effect of the specification of the simulation: as every communication between agents and between components within agents costs at least one time-step, it takes a number of time-steps before a message of the system has reached the patient (in fact 24 minutes). In between sending and receiving a message, it could happen that the medicine level has dropped below the minimum value, or that the patient has taken a pill already. For violations of the minimum level it is possible to compensate for this artificial delay by allowing an additional decrease that is equivalent to the decay of the medicine during the time of the delay. This means that medicine level should not drop below the 0,3335 if the delay is taken into account. Table 4 shows the duration of the violations for the corrected minimum level. In this case, there are no violations of the lower threshold in traces where the system is active.

Unfortunately, a similar correction for violations of the maximum level is not possible, as these violations are caused by taking two pills within a short period, which is a non-monotonic effect.

Table 4: Corrected values for duration the violations of the minimum level.

duration (in minutes)	with support	without
below 0.3335	0	164.77

5.2 Relating Global Properties to Executable Properties

Besides the verification of properties against simulation traces, the correctness of the entire model can also be proven (given certain conditions). This proves that for all possible outcomes of the model the global properties are indeed achieved under these specific conditions. Such correctness of a model can be proved using the SMV model checker

(MacMillan, 1995). In order to come to such a proof, an additional property level is introduced, namely the external behavioural properties for the components within the system. Thereafter, the relationship between the executable properties of the Medicine Box Agent, and the Usage Support Agent are related to these behavioural properties, and furthermore, these behavioural properties are shown to entail the top-level properties.

5.2.1 External Behavioural Properties

First a number of external properties for the Usage Support Agent (USA) are introduced. The first property specifies that the patient should be warned that medicine should be taken in case the medicine level is close to being too low (EUSA1). Secondly, property EUSA2 specifies that the USA should warn the doctor in case such a warning has been sent to the patient, but there has been no response. Moreover, the storage of the usage history is specified in EUSA3, and the sending of an early warning message is addressed in EUSA4. Finally, EUSA5 describes that the USA should communicate the current medicine level within the patient. Note that in all properties a parameter e is used, which specifies the maximum delay after which these communications or actions should occur. Such a parameter can vary per property, and is used throughout this section for almost all non-executable properties.

EUSA1: Usage Support Agent Behaviour External View Patient Warning Take Medicine

If the Usage Support Agent received communicated medicine intake, and based on these, the estimated accumulated concentration is C , and $C < TL$, then it communicates an SMS to the Patient Cell Phone that medicine should be taken.

$\forall t: \text{TIME}, \gamma: \text{TRACE}, M: \text{MEDICINE}, C: \text{REAL}$
 $\text{history_implied_value}(\gamma, \text{input}(\text{USA}), t, M, C) \ \& \ C < TL \Rightarrow$
 $\exists t' \ t' \leq t + e \ \& \ \text{state}(\gamma, t', \text{output}(\text{USA})) \models$

$\text{communication_from_to}(\text{sms_take_medicine}(M),$
 $\text{usage_support_agent}, \text{patient_cell_phone})$

EUSA2: Usage Support Agent Behaviour External View Doctor Warning

If the Usage Support Agent sent out a warning message to the patient, and the patient did not take medicine within X time steps after the warning, the Usage Support Agent sends a message to the Doctor Cell Phone.

$\forall t: \text{TIME}, \gamma: \text{TRACE} \ \text{state}(\gamma, t, \text{output}(\text{USA})) \models$

$\text{communication_from_to}(\text{sms_take_medicine}(M),$
 $\text{usage_support_agent}, \text{patient_cell_phone}) \ \&$
 $\neg \exists t2: \text{TIME} \ [t2 \geq t \ \& \ t2 < t + X \ \& \ \text{state}(\gamma, t2, \text{input}(\text{USA})) \models$
 $\text{communicated_from_to}(\text{medicine_used}(M),$
 $\text{medicine_box_agent}, \text{usage_support_agent})]$
 $\Rightarrow \exists t': \text{TIME} \ t + X \leq t' \leq (t + X) + e \ \text{state}(\gamma, t', \text{output}(\text{USA})) \models$

$\text{communication_from_to}(\text{sms_not_taken_medicine}(M),$
 $\text{usage_support_agent}, \text{doctor_cell_phone})$

EUSA3: Usage Support Agent Behaviour External View Store Information in Database

If the Usage Support Agent receives a communication that medicine has been taken, then the agent stores this information in the patient database.

$$\forall t:\text{TIME}, \gamma:\text{TRACE}, M:\text{MEDICINE} \text{ state}(\gamma, t, \text{input}(\text{USA})) \mid =$$

$$\text{communicated_from_to}(\text{medicine_used}(M),$$

$$\text{medicine_box_agent},$$

$$\text{usage_support_agent}) \ \&$$

$$\Rightarrow \exists t':\text{TIME} \ t \leq t' \leq t + e \ [\text{state}(\gamma, t', \text{output}(\text{USA})) \mid =$$

$$\text{performing_in}(\text{store_usage}(M, t),$$

$$\text{patient_database})]$$
EUSA4: Usage Support Agent Behaviour External View Store Send Early Warning Message to Patient

If the Usage Support Agent receives a communication that the patient attempted to take medicine too early, then the agent sends an SMS to the patient cell phone.

$$\forall t:\text{TIME}, \gamma:\text{TRACE}, M:\text{MEDICINE} \ \text{state}(\gamma, t, \text{input}(\text{USA})) \mid =$$

$$\text{communicated_from_to}(\text{too_early_intake_intention},$$

$$\text{medicine_box_agent}, \text{usage_support_agent}) \ \&$$

$$\Rightarrow \exists t':\text{TIME} \ t \leq t' \leq t + e \ [\text{state}(\gamma, t', \text{output}(\text{USA})) \mid =$$

$$\text{communication_from_to}(\text{put_medicine_back_and_wait_for_signal},$$

$$\text{usage_support_agent}, \text{patient_cell_phone})]$$
EUSA5: Usage Support Agent Behaviour External View Send Approximated Concentration to Medicine Box Agent

If the history of medicine usage implies a certain medicine level, then the Usage Support Agent communicates this value to the Medicine Box Agent.

$$\forall t:\text{TIME}, \gamma:\text{TRACE}, M:\text{MEDICINE}, C:\text{REAL}$$

$$\text{history_implied_value}(\gamma, \text{input}(\text{USA}), t, M, C) \Rightarrow$$

$$\exists t' \ t \leq t' \leq t + e \ \& \ \text{state}(\gamma, t', \text{output}(\text{USA})) \mid =$$

$$\text{communication_from_to}(\text{medicine_level}(M,$$

$$C),$$

$$\text{usage_support_agent}, \text{medicine_box_agent})$$

Besides the Usage Support Agent, the Medicine Box Agent (MBA) plays an important role within the system as well. From an external perspective, the MBA has three behavioural properties. The first (EMBA1) expresses that the MBA should communicate to the Usage Support Agent that medicine has been taken by the patient. This only occurs when the medicine is not taken too early. Properties EMBA2 and EMBA3 concern the communication in case of an early intake. First of all, the MBA should sound a beep (EMBA2), and furthermore, the MBA should communicate this information to the USA (EMBA3). Again, a parameter e is used to specify the maximum delay for these properties. Note that these properties are later referred to as EMBA, which is the conjunction of the three properties specified below.

EMBA1: Medicine Box Agent Behaviour External View Communicate Usage Non-Early Intake

When the medicine box agent observes medicine is taken from position X, Y in the box and the medicine is of type M , and the medicine level of M communicated to the agent is C , and furthermore, the medicine level plus a dose does not exceed the overall maximum, then the medicine box agent outputs medicine has been taken to the Usage Support Agent.

$$\forall \gamma:\text{TRACE}, t:\text{TIME}, X, Y:\text{INTEGER}, C:\text{REAL}, M:\text{MEDICINE}$$

$$\text{state}(\gamma, t, \text{input}(\text{medicine_box_agent})) \mid =$$

$$\text{observed_result_from}(\text{medicine_taken_from_position}(\text{x_y_coordinate}(X, Y), \text{medicine_box}) \ \&$$

$$\text{state}(\gamma, t, \text{input}(\text{medicine_box_agent})) \mid =$$

$$\text{communicated_from_to}(\text{medicine_level}(M, C),$$

$$\text{usage_support_agent}, \text{medicine_box_agent}) \ \&$$

$$C + \text{DOSE} \leq \text{MAX_MEDICINE_LEVEL} \ \&$$

$$\text{medicine_at_location}(X, Y, M)$$

$$\Rightarrow \exists t':\text{TIME} \ t \leq t' \leq t + e \ [\text{state}(\gamma, t', \text{output}(\text{medicine_box_agent})) \mid =$$

$$\text{communication_from_to}(\text{medicine_used}(M),$$

$$\text{medicine_box_agent},$$

$$\text{usage_support_agent})]$$
EMBA2: Medicine Box Agent Behaviour External View Communicate Beep when Early Intake

When the medicine box agent observes medicine is taken from position X, Y in the box and the medicine is of type M , and the medicine level of M communicated to the agent is C , and furthermore, the medicine level plus a dose exceeds the overall maximum, then the medicine box agent outputs a beep to the Patient.

$$\forall \gamma:\text{TRACE}, t:\text{TIME}, X, Y:\text{INTEGER}, C:\text{REAL}, M:\text{MEDICINE}$$

$$\text{state}(\gamma, t, \text{input}(\text{medicine_box_agent})) \mid =$$

$$\text{observed_result_from}(\text{medicine_taken_from_position}(\text{x_y_coordinate}(X, Y), \text{medicine_box}) \ \&$$

$$\text{state}(\gamma, t, \text{input}(\text{medicine_box_agent})) \mid =$$

$$\text{communicated_from_to}(\text{medicine_level}(M, C),$$

$$\text{usage_support_agent}, \text{medicine_box_agent}) \ \&$$

$$C + \text{DOSE} > \text{MAX_MEDICINE_LEVEL} \ \&$$

$$\text{medicine_at_location}(X, Y, \text{hiv_slowers})$$

$$\Rightarrow \exists t':\text{TIME} \ t \leq t' \leq t + e \ [\text{state}(\gamma, t', \text{output}(\text{medicine_box_agent})) \mid =$$

$$\text{communication_from_to}(\text{sound_beep}, \text{medicine_box_agent},$$

$$\text{patient})]$$
EMBA3: Medicine Box Agent Behaviour External View Communicate Early Intake to Usage Support Agent

When the medicine box agent observes medicine is taken from position X, Y in the box and the medicine is of type M , and the medicine level of M communicated to the agent is C , and furthermore, the medicine level plus a dose exceeds the overall maximum, then the medicine box agent outputs a communication concerning this early intake intention to the Usage Support Agent.

$$\forall \gamma:\text{TRACE}, t:\text{TIME}, X, Y:\text{INTEGER}, C:\text{REAL}, M:\text{MEDICINE}$$

$$\text{state}(\gamma, t, \text{input}(\text{medicine_box_agent})) \mid =$$

$$\text{observed_result_from}(\text{medicine_taken_from_position}(\text{x_y_coordinate}(X, Y), \text{medicine_box}) \ \&$$

$$\text{state}(\gamma, t, \text{input}(\text{medicine_box_agent})) \mid =$$

$$\text{communicated_from_to}(\text{medicine_level}(M, C),$$

$$\text{usage_support_agent}, \text{medicine_box_agent}) \ \&$$

$$C + \text{DOSE} > \text{MAX_MEDICINE_LEVEL} \ \&$$

$$\text{medicine_at_location}(X, Y, \text{hiv_slowers})$$

$$\Rightarrow \exists t':\text{TIME} \ t \leq t' \leq t + e \ [\text{state}(\gamma, t', \text{output}(\text{medicine_box_agent})) \mid =$$

$$\text{communication_from_to}(\text{too_early_intake_intention},$$

$$\text{medicine_box_agent}, \text{usage_support_agent})]$$

Furthermore, a number of properties are specified for the external behavior of the other components. These properties include basic forwarding of information by the patient cell phone (PCP), the doctor cell phone (DCP), communication between various communicating components (CP). Furthermore, it includes the specification of the observation results of performing actions in the medicine box (EMD), the storage of information in the database (PDP), and the transfer of those actions and observations (WP).

Finally, in order for such external behavioral properties to establish the global property, certain assumptions need to be made concerning the behavior of the doctor and of the patient. For the proof, the minimal behavior of the patient is used. This minimal behavior is specified by stating that the patient should at least respond to the doctor

communication. Of course, most likely would be that the patient already responds to the SMS being sent, but using the minimal behavior it can already be proven that this is not even required, as long as the patient responds to the doctor communication. This ensures, that even if the patient does not respond on an SMS his medicine level will still remain within the boundaries set. The behavior of the patient is expressed in property PB. Besides the patient, also the doctor needs to respond to the SMS of the system in a particular way, namely that he immediately contacts the patient, as expressed in property DB.

PB: Respond to Doctor request

When a patient receives a doctor warning that the patient should take medicine M, then the patient takes the medicine from the appropriate place in the box.

```

 $\forall \gamma: \text{TRACE}, t: \text{TIME}, M: \text{MEDICINE}, X, Y: \text{INTEGER}$ 
state( $\gamma, t$ , input(patient)) |=
communicated_from_to(doctor_request_take_medicine(M),
patient_cell_phone, patient)
medicine_at_location(X, Y, M)
 $\Rightarrow \exists t': \text{TIME } t \leq t' \leq t + e$ 
state( $\gamma, t'$ , output(patient)) |=
performing_in(take_medicine_from_position(x_y_coordinate(X, Y)),
medicine_box)

```

DB: Warn patient after SMS

When the doctor has received an SMS concerning a patient that has not used medicine, then the doctor uses its cell phone communicating that the patient to take medicine.

```

 $\forall \gamma: \text{TRACE}, t: \text{TIME}, M: \text{MEDICINE}$ 
state( $\gamma, t$ , input(doctor)) |=

communicated_from_to(sms_not_taken_medicine(M),
doctor_cell_phone, doctor)
 $\Rightarrow \exists t': \text{TIME } t \leq t' \leq t + e$  [ state( $\gamma, t'$ , output(doctor)) |=
communication_from_to(doctor_request_take_medicine(M),
doctor,
doctor_cell_phone) ]

```

5.2.2 Relating Properties

The relation between the external behavioral properties introduced in Section 6.2.1 and the top-level global property is shown in Figure 4. The figure also shows the relationship between the executable properties of the Usage Support Agent, and the Medicine Box Agent and the external behavioral properties thereof. Note that the external level of both the Usage Support Agent as well as the Medicine Box Agent has been represented in the figure as EUSA and UMBA respectively. This is simply the conjunction of the external properties of these agents. Furthermore, the external behavioral properties of the other components directly translate to executable properties in the specification. The following relations hold from the local to the external level.

Usage Support Agent External Behavior

```

IB & SE & USA1 & USA5  $\Rightarrow$  EUSA1
IB & USA2 & USA5 & USA6  $\Rightarrow$  EUSA2
IB & USA2 & USA3  $\Rightarrow$  EUSA3

```

```

IB & USA4  $\Rightarrow$  EUSA4
IB & SE & USA1 & USA7  $\Rightarrow$  EUSA5

```

Medicine Box Agent External Behavior

```

IB & MBA1 & MBA2 & MBA3  $\Rightarrow$  EMBA

```

Furthermore, the global property GP1 has the following relationship with the external properties of the various components.

Global Behavior

```

EMBA & EUSA & PCP & DCP & CP & EMD & PDP & WP &
PB & DB  $\Rightarrow$  GP1

```

In order to prove the above specified relationships, the relations have been checked within the SMV model checker. Using the following parameters, the inter-level relations within the tree expressed in Figure 4 are indeed shown to hold using SMV: The initial medicine level is set to 60, the reduction level per step is set to 98 (i.e. 0.98), the maximum medicine level is set to 150 and the minimum level to 30. Furthermore, the warning level is set to 60 (after which an SMS is sent to the patient). Finally, the delay for a doctor message is set to 10 time steps. Note that in the case of the SMV checks no communication delay is specified. Therefore, this formal verification shows that without this delay the model indeed entails the global property.

6 DISCUSSION

In this paper, possible support in the domain of medicine usage management was addressed. A multi-agent system model that supports the users of medicine in taking their medicine at the appropriate time was discussed and formally analysed by simulation and verification. The system has been specified using a formal modeling approach which enables the specification of both quantitative as well as qualitative aspects (Jonker and Treur 2002; Bosse, Jonker, Meij, and Treur, 2006). To specify the model, both generic and domain specific temporal rules have been used, enabling reuse of the presented model. The analysis of the model has been conducted by means of (1) generation of a variety of simulation runs using a stochastic model for patients; (2) specification of dynamic properties at different aggregation levels; (3) specification of interlevel relations between these properties; (4) automated verification of properties specified in (2) against traces generated in (1), and (5) automated verification of the interlevel relations specified in (3). The simulation execution in (1) has been achieved making use of the LEADSTO software automated verification of such properties against set of simulation traces in (4). Verification of interlevel relations in (5) has been performed using the SMV

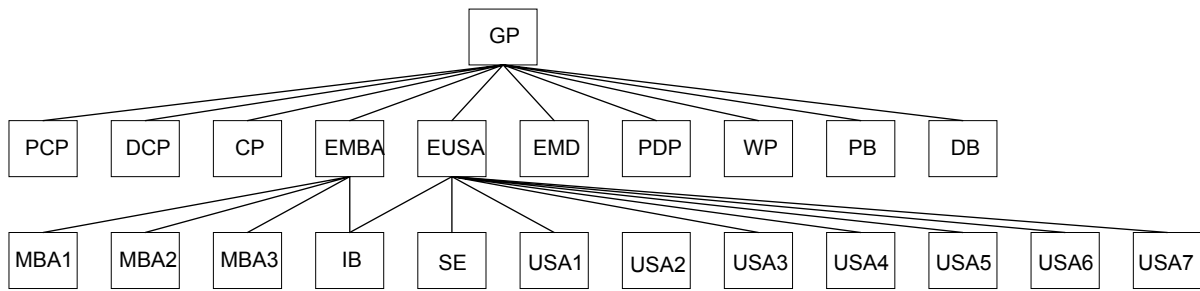


Figure 3: Property hierarchy in the form of an AND tree.

model checking environment (MacMillan, 1995). Evaluation of the system with actual users is part of future work.

The presented analysis fits well in the recent developments in Ambient Intelligence (Aarts, Collier, Loenen, Ruyter, 2001; Aarts, Harwig, and Schuurmans, 2003; Riva *et al.*, 2005). Furthermore, it also shows that multi-agent system technology can be of great benefit in health care applications, as also acknowledged in (Moreno and Nealon, 2004). More approaches to support medicine usage of patients have been developed. Both in (Greene, 2005) as well as (Floerkemeier and Siegemund, 2003) models are presented that do not simply always send an SMS that medicine should be taken such as proposed by (Safren, *et al.*, 2003). Both approaches only send SMS messages in case the patient does not adhere to the prescribed usage. The model presented in this paper however adds an additional dimension to such a support system, namely the explicit representation and simulation of the estimated medicine level inside the patient. Having such an explicit model enables the support agent to optimally support the patient; see also (Bosse, Memon, and Treur, 2007) for the use of such a model for mental processes of another agent.

REFERENCES

- Aarts, E.; Collier, R.; van Loenen, E.; Ruyter, B. de (eds.) (2003). *Ambient Intelligence. Proc. EUSAI 2003*. LNCS vol. 2875. Springer Verlag, 2003, pp. 432.
- Aarts, E., Harwig, R., and Schuurmans, M. (2001). Ambient Intelligence. In: P. Denning (ed.), *The Invisible Future*, McGraw Hill, pp. 235-250.
- Bosse, T., Jonker, C.M., Meij, L. van der, Sharpanskykh, A., and Treur, J. (2006). Specification and Verification of Dynamics in Cognitive Agent Models. In: Nishida, T. et al. (eds.), *Proceedings of IAT'06*. IEEE Computer Society Press, 2006, pp. 247-254.
- Bosse, T., Jonker, C.M., Meij, L. van der, and Treur, J. (2007). A Language and Environment for Analysis of Dynamics by Simulation. *Int. Journal of Artificial Intelligence Tools*, vol. 16, 2007, pp. 435-464.
- Bosse, T., Memon, Z.A., and Treur, J. (2007). A Two-level BDI-Agent Model for Theory of Mind and its Use in Social Manipulation. In: P. Olivier, C. Kray (eds.), *Proceedings of AISB'07*. AISB Publications, 2007, pp. 335-342.
- Brazier, F.M.T., Jonker, C.M., and Treur, J. (2000). Compositional Design and Reuse of a Generic Agent Model. *Applied Artificial Intelligence Journal*, vol. 14, 2000, pp. 491-538.
- Floerkemeier, C., Siegemund, F. (2003). Improving the Effectiveness of Medical Treatment with Pervasive Computing Technologies. *Workshop on Ubiquitous Computing for Pervasive Healthcare Applications at Ubicomp 2003*, Seattle, October 2003
- Green D.J. (2005). Realtime Compliance Management Using a Wireless Realtime Pillbottle – A Report on the Pilot Study of SIMPILL. In: *Proc. of the International Conference for eHealth, Telemedicine and Health, Med-e-Tel'05*, 2005, Luxemburg.
- Hoogendoorn, M., Klein, M., and Treur, J., Formal Design and Simulation of an Ambient Multi-Agent System Model for Medicine Usage Management. In: *Proceedings of the International Workshop on Ambient Assisted Living (AAL 2007)*, to appear, 2007.
- Jonker, C.M., and Treur, J. (2002). Compositional Verification of Multi-Agent Systems: a Formal Analysis of Pro-activeness and Reactiveness. *Intl. J. of Coop. Inf. Systems*, vol. 11, 2002, pp. 51-92.
- McMillan, K.L., (1993). *Symbolic Model Checking: An Approach to the State Explosion Problem*. PhD thesis, School of Computer Science, CMU, Pittsburgh, 1992. Published by Kluwer Academic Publishers, 1993.
- Moreno, A., Nealon, J.L. (eds.), *Applications of Software Agent Technology in Health Care Domain*. Birkhäuser Basel, 2004.
- Riva, G., F. Vatalaro, F. Davide, M. Alcañiz (eds.) (2005). *Ambient Intelligence*. IOS Press, 2001.
- Safren, S.A., Hendriksen, E.S., Desousa, N., Boswell, S.L., Mayer, K.H., (2003). Use of an on-line pager system to increase adherence to antiretroviral medications. In: *AIDS CARE*, vol. 15, pp. 787 – 793.

POSTERS

INTELLIGENT SYSTEM FOR ASSISTING ELDERLY PEOPLE AT HOME

Juan Pascual, Miguel A. Sanz-Bobi and David Contreras

*Computer Science Department, Engineering School, Pontificia Comillas University, Alberto Aguilera 23, Madrid, Spain
masanz@upcomillas.es*

Keywords: Multi-agent system, risks at the homes of elderly people, intelligent system, human activities monitoring.

Abstract: Nowadays the number of elderly people in our society is increasing thanks to continuous and important advances related to health care. The ideal situation for these elderly people is to spend as much time as possible enjoying their life within their family and social environment, without the need to abandon their homes to go live in specialized centres as long as their physical health permits. This paper describes an architecture of pro-active intelligent agents in which the main objective is to extend the amount of time as much as possible that these elderly people can reside in their own homes by means of providing continuous vigilance of certain parameters concerning daily activities which possibly could be risky by facilitating reminders to complete specific tasks, by easing communication with the exterior world, and in the case that it is necessary, by automatically calling for emergency services.

1 INTRODUCTION

At present an important demographic change is taking place regarding the mean age of our society. The increase of the population aged 60 years or older is very significant. This is a world wide phenomenon as observed by the United Nations (United Nations, 2006). In Europe the number of new born children is less than it was 20 years ago and life expectancy is much longer and in consequence the ageing effect of the population is even more important.

When elderly people are healthy, but with typical problems due to ageing, the best recommendation is that they spend as much time as possible enjoying their life within their family and social environment, without the need to abandon their home to go live in specialized centres as long as their physical health permits. In many cases, simply because there do not exist sufficient means or vigilance time by a caretaker in their home, these people are forced to leave their homes to go live in specialized centres such as nursing homes.

In order to try to solve this problem, several organizations and companies offer teleassistance services to elderly people at home. The most part of this type of services is based on the demand of the user (elderly person) by simply pushing a button on a small device that he/she carries on him/her

(Aguilera, 2003). A specialized call centre attends any request from the user at any time, and also, the call centre can contact the user periodically in order to know that everything is going well. All these services are very helpful, but they require inputs based on human decisions, the user or people attending the call centre. It seems that a further step is needed in order to try to monitor some daily activities of these elderly people. There are important advances in the use of new information technologies for monitoring some activities of elderly people at home (Jih, 2006), (Fishkin, 2005), for assistance to find the way, if one is disoriented (Liu, 2006) and to monitor some important biological parameters (Pollack, 2005). Also, some efforts have been developed in the elaboration of technological platforms able to integrate different kinds of services of remote assistance to elderly people (Robocare, 2007), (Hill, 2005), (Attentianet, 2007). This paper is in line with these examples of research but with the aim of preventing possible risks for elderly people at home carrying out their daily tasks.

This paper is organised with the following sections. The first section presents the objectives of SIAM, the next section presents its main strategy, the following two sections describe the agents and the implementation of SIAM, and finally, an example is provided.

2 SIAM OBJECTIVES

This paper describes a multi-agent system named SIAM, which in Spanish stands for intelligent system for the assistance to elderly people at home. Its main objective is to contribute to extending the amount of time as much as possible that elderly people can reside in their own homes with automatic and intelligent assistance based on new information technologies. This is reached by:

- continuous vigilance of certain variables that could be important for risk detection in daily activities at home
- facilitating reminders to complete specific tasks
- easing communication with the exterior world
- in the case that it is necessary, automatically calling for emergency services.

SIAM is based on a set of intelligent agents of a multi-agent system capable of integrating the previously mentioned pro-active assistant services for elderly people as they carry out their usual activities at home.

The scope of the current version of SIAM is limited to the agent architecture and their relationships. It was decided to not install SIAM in a real environment in the house of an elderly person without an intensive testing phase of the software developed in order to prevent possible unnecessary disturbances. Furthermore it should be noted that this version of SIAM does not include interaction with physical sensors. The information coming from the sensors is obtained from a simulation environment which represents the main rooms of a house and the sensors that are installed in each one. A direct input over a sensor in the simulation environment can be used to switch its status on/off.

3 SIAM STRATEGY

The SIAM strategy is conceived as a pro-active set of intelligent services able to help elderly people at home.

These intelligent services are the following:

- Detection of possible actions or situations in the house that could be a risk for the elderly person, and trying to prevent it. Examples of possible risk situations are: a sudden fall of the person on to the dining room floor, unattended gas open in the kitchen, unattended water running in the bath, etc.
- Prevention of other types of risks related with forgetting things such as medications, an

appointment with the doctor, a payment of some important bill, etc.

c. Facilitation of the communication between the elderly person and the external world to his/her home: contact with the caretaker, with family, with the CMD (Central Monitoring Department), etc.

These services can be accessed by several actors with different roles. A strategy based on a multi-agent system (Weiss, 2000), (Wooldridge, 2002) was chosen in order to fulfil all the requirements of these services.

The following are the four types of actors: the user or elderly person, the CMD, the caretaker and the virtual caretaker or intelligent system.

SIAM has to cover different roles through its different agents for interaction with the different actors. The main roles to be covered are the following:

- Communication to and from the elderly person at home.
- Communication to and from the caregiver of the elderly person, if such a person exists.
- Communication to and from the CMD.
- Collection of information coming from sensors.
- Intelligent analysis of the information collected in order to predict a possible risk for the person and to issue the corresponding actions.

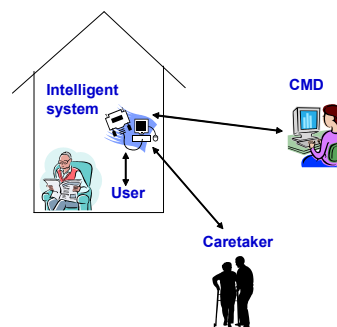


Figure 1: SIAM actors and relationships.

Figure 1 shows the design logic of SIAM based on the interaction among actors according to the roles of each one. Usually, SIAM collects information from sensors installed in the home of the elderly person in order to detect if a possible risky situation is produced by an action or event when he/she is at home carrying out his/her activities. In the case that some anomaly or risk is detected, SIAM will first try to contact the elderly person and, if an answer is not received, this will be notified to the CMD and/or to the caretaker. Also, from Figure 1, it is possible to observe that the user can activate a request to the CMD and the caretaker and vice versa.

4 SIAM AGENTS

SIAM contains the following type of agents:

- **USER.** It is in charge of all the communications from SIAM to the person and from the person to SIAM. This agent is located in a mobile device that the person carries on him or her. At the moment this agent is installed in a PDA.
- **DATA COLLECTOR.** This agent is in charge of the collection of information from key places of the elderly person's home in order to know if some particular activities are occurring which cause a certain risk for the person. These agents are located in the house being monitored.
- **HOME CONTROLLER.** It is in charge of the integration and pre-analysis of all the information collected by the Data Collector agents in order to obtain a global view of activities in the house. This agent is located in the house being monitored.
- **DIAGNOSTIC.** This agent performs a diagnostic of possible risks in the house of the elderly person according to the information collected and specialized knowledge previously stored in a knowledge base. The structure of this agent is similar to an expert system. This agent is located in the Central Monitoring Department responsible for taking care of a group of elderly people using the SIAM platform.
- **CARETAKER.** Its objective is taking care of all the communications from SIAM to the caretaker of the elderly person and from the caretaker to SIAM. This agent is located in a mobile device that the caretaker carries on him or her. At the moment this agent is installed in a PDA.

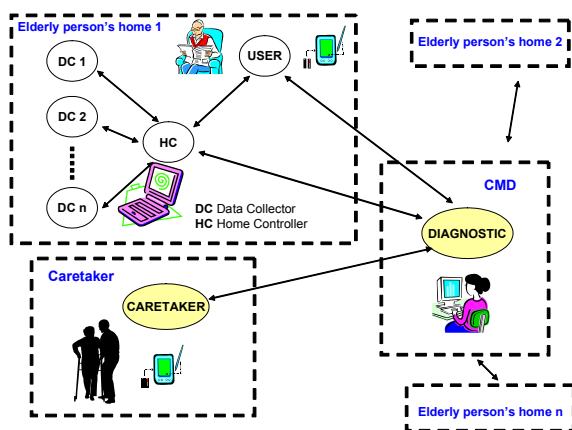


Figure 2: SIAM architecture.

Figure 2 shows the multi-agent architecture. As can be observed, the diagnostic agent is physically

located at the CMD. The CMD is conceived as a specialized centre that could be a nursery home or in general a company dedicated to taking care of elderly people. The CMD can monitor several houses, each one having its own diagnostic agent.

The data collector agents and the home controller have to be located in a computer in the house of the elderly person. This computer does not require attention by the person and no screen associated to the PC is needed if it is not required by the user. The maintenance of the SIAM agents must be done by remote control from the CMD.

The caretaker agent is in a mobile device that could be supported by another PDA in the caretaker's home or by another form of communication.

5 SIAM IMPLEMENTATION

The implementation of SIAM is based on the multi-agent system free software tool named JADE (Java Agent DEvelopment Framework) (JADE, 2007). SIAM was designed taking into account a low cost for the resulting product. The JADE architecture is based on a set of platforms that include containers. SIAM consists of a central platform that emulates the CMD. This platform includes a container with the diagnostic agent.

Each user has his/her own platform consisting of three containers. The first container includes all the agents that monitor activities at the house of the elderly person: data collector agents and home controller. The second and third containers include the user and caretaker agents respectively. These agents are supported by portable devices such as a PDA, and they have been implemented in JADE Leap, a special version for these cases, using J9 from IBM as a virtual machine. The communication of these agents with the rest is WiFi.

All the SIAM agents have been developed in JAVA and their interfaces using SWT: The Standard Widget Toolkit, another free software tool. Finally, the permanent storage of all the data collected and other data required by SIAM uses MySQL as database server.

The diagnostic agent is an intelligent agent physically running in a computer at the CMD. It has a graphical interface for each house being monitored. It receives all the relevant facts from the corresponding home controller agent of a house and uses an inference engine to reach conclusions about possible risks for elderly people in their homes or attending their demands. The architecture of the

diagnostic agent is based on an expert system including a knowledge base where the definitions are included of risky situations in a house and actions to take to prevent them. The knowledge base is particular for each user and can include special circumstances of each user. The elaboration of the knowledge base has to be done under the direction of personnel specialized in the care of elderly people. The knowledge base has a very simple architecture based on production rules including certainty factors. The inference engine that uses the knowledge is a classical forward-chaining engine.

The diagnostic agent can take autonomous decisions without waiting for a confirmation from the CMD personnel which simply will be informed. This is an important feature of SIAM because it can react very fast when a risk or demand is coming.

The diagnostic agent exchanges information with the user through the home controller that is located at the elderly person's home. This agent is in charge of the communication with all the devices that support SIAM, and also, of keeping the diagnostic agent updated about the profile of new risk situations in the house. When a new user is coming to SIAM, the first agent to start is the home controller agent. Once it is alive, it tries to make contact with the diagnostic agent, and if it accepts a new user, it is monitored by the CMD. At this moment the other agents associated to the user start to work.

In this version of SIAM four data collector agents have been developed in order to monitor some activities at the elderly person's home. Each one is in charge of activities in a room of the house. They are the following:

- data collector agent monitoring activities in the bathroom: presence of the person, closed/opened water valve in the shower and washbasin and high vibrations.
- data collector agent monitoring activities in the kitchen: presence of the person, closed/opened water valve, close/open gas and high vibrations.
- data collector agent monitoring activities in the dining room: presence of the person, and high vibrations.
- data collector agent monitoring activities in the entrance of the house: presence of the person and closed/opened door.

As was mentioned previously, this version of SIAM is focused on the development of the architecture of the whole multi-agent system. There are no sensors connected yet, however a simulation environment to simulate the operation of the real

sensors has been developed. More investigation has to be done to select and install the appropriate sensors and data collector according to expert opinion on elderly people, but this will not change the current operation of SIAM.

The user agent is running in a mobile device, in this SIAM version it is in a PDA. It has an interface which is extremely simple to use with big symbols to communicate to and from the elderly person. The user can activate one of three big icons: emergency situation, communication with the caretaker and asking some questions to the CMD. Figures 3.a and 3.b show how the user observes these icons.



Figure 3: Basic interface for the user agent.

Also, the user can receive warnings using big icons about the need to review something in the house that could be a risk for the person. Figure 3.b shows an example of a warning telling the user to turn off the gas. The structure and interfaces of the caretaker agent are similar to those of the user agent, and also, it is ready for running in a PDA.

6 EXAMPLE

This section describes a simple example of the operation of SIAM. Let us suppose that the elderly person is at home and he/she decides to go to the bathroom to wash his/her hands. In this case a signal corresponding to the presence of the person switches from off to on, and a moment later the same occurs with the signal corresponding to water running in the washbasin. This information is collected by the data collector agent that monitors the bathroom and it is sent to the home controller agent. Once this agent has pre-processed all the information in the house, it sends these two new events to the diagnostic agent in the CMD. Figure 4 shows the interface of the diagnostic agent corresponding to this user. In this figure, two new lines of information appear, representing the two new facts received. Immediately, the expert system is started and it does not reach any conclusion about risk situations for the user and so nothing happens.

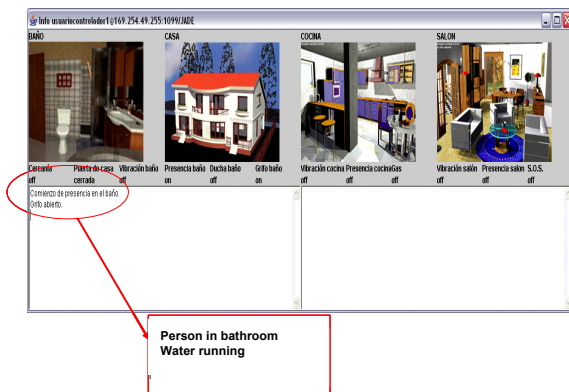


Figure 4: interface of the diagnostic agent.

However, suppose that the person leaves the bathroom and he/she is in the kitchen, but the water is still running. In this case, the data collector in charge of the bathroom updates the situation and resends the new events to the home controller agent, the data collector for the kitchen does the same, and finally the home collector sends all the information to the diagnostic agent. It starts the expert system and it concludes the detection of a risk situation taking the decision to wait some minutes more, expecting new information from the home controller. If the risk situation does not disappear, the diagnostic agent takes the decision to send a message to the home controller in order to issue a warning message to the user agent to turn the water off in the bathroom. In the case that the required action is not executed because the diagnostic agent detects that the problem persists, it automatically decides to contact the caretaker. In order to do this it sends a message to the caretaker agent similar to that received by the user. Finally, if some time passes without any action, then the CMD personnel is informed.

Several similar situations were tested and the general performance of all the processes was successful.

7 CONCLUSIONS

This paper describes the architecture of the multi-agent system named SIAM. It has been designed for an automatic detection of possible risks of elderly people at home and for assistance in these cases or on demand of the user. SIAM will contribute to extending the amount of time as much as possible that elderly people can reside in their own homes assisted continuously by an intelligent agent and using new information technologies.

The resulting application has important flexibility and incorporates new knowledge and new features in an easy manner. Simulations of several situations have been tested and the results are very promising.

At present, a next phase of development of SIAM is starting in order to include information coming from real sensors. In parallel, a deeper analysis concerning risky conditions in the elderly person's home is being developed. After these new steps are completed in SIAM, it is expected to be tested in real environments.

REFERENCES

- Aguilera, S., Colás, J., Campos, E., 2003. *Multi Functional Call Center for Disabled and Elderly People*. Assistive Technology Conference AAAT, 2003. Dublin
- Attentionet. Project web description retrieved June 20, 2007 from: <http://www.attentionet.eu>
- Fishkin, K.P., Philipose, M. Rea, A. 2005. *Hands-on RFID: Wireless wearables for detecting use of objects*. IEEE International Symposium on Wearable Computers.
- JADE. Project web page retrieved June 20, 2007 from: <http://robocare.ip.rm.cnr.it>
- Jih, W., Hsu, J. Y., Tsai, T., 2006. *Context-aware service integration for elderly care in a smart environment*. American Association for Artificial Intelligence Workshop.
- Liu, A. Hile, A., Kautz, H., Borriello, G., Brown, P., Harniss, M., Johnson, K., 2006. *Indoor wayfinding: Developing a functional interface for individuals with cognitive impairment*. Proceedings of the 8th International ACM SIGACCESS Conference on Computers and Accessibility.
- Pollack, M., 2005. *Intelligent technology for an aging population: the use of AI to assists elders with cognitive impairment*. AI Magazine, 26(2)
- Robocare. Project web page retrieved June 20, 2007 from: <http://robocare.ip.rm.cnr.it>
- United Nations, 2006. Department of Economic and Social Affairs. *Population ageing report 2006*. Retrieved June 20, 2007 from: <http://www.un.org/esa/population/publications/ageing/ageing2006.htm>.
- Weiss, G. 2000. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, The MIT Press.
- Wooldridge, M. 2002. *An Introduction to MultiAgent Systems*, John Wiley & Sons.

ASSISTING WELLBEING

The Challenges of using Technology to Improve Wellbeing in Older Adults

Douglas Millward

*School of Computing and Information Technology, University of Wolverhampton
Wulfruna Street, Wolverhampton, WV1 1SB, UK
Djm63@wlv.ac.uk*

Wendy Nicholls

*School of Applied Sciences, University of Wolverhampton, Wulfruna Street, Wolverhampton, WV1 1SB, UK
Wendy.Nicholls@wlv.ac.uk*

Keywords: Assistive technology, integration, psychosocial, telecare, user interface, wellbeing.

Abstract: Telecare is an increasingly important application of technology that is designed to increase the independence of older adults, amongst other goals. The programme of research described below aims to identify important issues with the deployment of this technology to the target group. It describes an ongoing programme of research that attempts to classify these issues, and posit solutions. It additionally proposes a new area of research into the effects of telecare and related technologies on a client's psychosocial wellbeing.

1 INTRODUCTION

The UK Government has a published plan to provide Telecare into every home that requires it by 2010 (Curry, Trejo-Tinoco, & Wardle, 2002). One of the problems faced in requirements gathering and analysis in the assistive technology (AT) field is the number of stakeholders involved. These parties include healthcare providers, social services, formal and informal carers. The needs of these groups have been studied (e.g., Lines & Hone, 2004). However, few studies have tried to gather older adults requirements and aspirations.

It is argued that the shift towards using technology to provide or even replace services that promote client independence should be accompanied by a focus on addressing more holistic needs, such as improving well being and quality of life. The purpose of the current research programme is therefore threefold:

1. Telecare is an emerging and increasingly important aspect of AT (Barlow, Bayer, & Curry, 2003). However, although there is evidence to support the view that telecare can assist in ensuring the physical well being of clients, there has been little research around the effect of such systems on the psychosocial aspects of wellbeing. Aim 1 is

therefore to test the effect of contemporary system provisions on the psychosocial well being of older adults. In order to achieve this aim the authors have created a system that provides similar outputs to those of other commercial offerings in the telecare arena.

2. It is hypothesised that telecare systems will have a greater effect on the psychological well being of users if users are provided with increased feedback on the information that these systems can provide. To this end the prototype system discussed above is provided with a user interface that is capable of providing continual feedback on the status of the system, and by inference the status of the client's home. The second aim is to investigate if the provision of an easily accessible user interface will improve on the effects measured above.

3. The third aim of this programme is to identify the needs of older adults with respect to future developments, direction, and provision of an Enhanced Electronic Assistive Technology System (EEATS). Research (Dickinson, Eisma, & Gregor, 2003) indicates that the increasingly (over)complex user interfaces associated with modern computer applications (and systems) places a large burden on a users cognitive abilities as they try to build a mental model of the operation of the system and

this, combined with lower confidence makes it difficult to master new technologies (Marquie, Jourdan-Boddaert, & Huet, 2002). Training alone will not increase confidence (Segrist, 2004), and as anxiety about using computers is positively associated with age (Ellis & Allaire, 1999) it may be that pursuing the use of computers specifically needs revision, and technologies, such as the telephone (Reed & Monk, 2004), or other familiar technologies need to be revisited as viable alternatives. White et al (White et al., 2002) recognised that a simplified interface would be of benefit to older adults. The first stage of the present study therefore sought to identify the current trends in IT and AT use in older adults, with a view to establishing the ideal mode of delivery for AT in the proposed intervention.

2 METHOD

2.1 Design

There are three stages to the present study.

Stage 1: The aim was to gauge established practice with IT, AT and familiarity with household technology using a postal survey.

Stage 2: The aim is to explore the effect of AT on mental health well being and to elicit an evaluation of the system provided. Measures of psychological functioning, quality of life, relationships with others, and well being, will be taken before and after a trial period of one week using the AT.

Stage 3: The aim is to explore possibilities for the type of applications older adults desired in order to inform development and design. Three focus groups will be conducted.

Thus far, stage 1 has been implemented and results are discussed below.

2.2 Participants

There were 59 participants of which 20 were male and 39 were female. With respect to age, 22% were under 65, 42.4% were 65-75, 25.4% were 75-84 and 10.2% were 85 years of age and over. Where n values vary in the below analyses, this is due to missing data. Thirty-one participants were recruited from a supported housing project and 28 from an IT club for older adults.

2.3 Materials

Participants were issued with a postal survey. The purpose was to get an overview of the technologies already in use by older adults and their preferences for the medium through which assistive technology should be provided. Items asked about the frequency with which household technologies are used, confidence with these technologies, and ease of use. Participants were further asked what tasks they completed on a computer, on the internet, on a mobile phone, and on a TV. Finally, participants were asked about the likelihood of their using a variety of assistive technologies as a reminder system for medication.

The following variables were derived:

Often – The frequency with which technologies such as computers, mobile phones and televisions are used. A higher score indicates a higher frequency. The possible range of scores was 5 – 20.

Difficulty – The difficulty with which participants rate using technology such as computers, mobile phones and televisions. A higher score indicates more difficulty. The possible range of scores was 6 – 24.

Confidence – The level of uncertainty the participant experiences when using technology. A higher score indicates more uncertainty; a lower score indicates more confidence. The possible range of scores was 6 – 24.

Participants were also asked about the number of tasks they completed on each of a computer, the internet, a mobile phone, and a TV. Thus four variables were created to describe the number of tasks completed with each.

2.4 Procedure

Questionnaires were distributed and returned via gatekeepers at the housing organisation and the IT club. Informed consent was granted from all participants. Data were entered into SPSS for analysis.

3 RESULTS

3.1 Overview

There are two parts to the data analysis. To begin with the relationships between the variables are

presented. This is followed by comparisons between the age groups surveyed.

3.2 Correlations

The frequency with which participants used technology was significantly correlated with the reported difficulty of use ($r = -.561$, $n = 18$, $p = .16$), and confidence in using technology ($r = -.533$, $n = 19$, $p = .19$). Those who use technology more frequently find it easier and have more confidence (see table 1). The easier using technology was perceived, the more confidence participants reported ($r = .947$, $n = 17$, $p < .001$).

Table 1: Correlations between frequency, ease and confidence of use with actual use of technology.

	Often	Difficulty
Difficulty (n)	-.561* (18)	
Confident (n)	-.533* (19)	.947** (17)

Key

* = $p < .05$

** = $p < .01$

3.3 Comparisons between Age Groups

Participants were grouped according to their age. Three groups were created; below 65, 65-74, and 75 plus. These groups were contrasted on their frequency of use, difficulty with technology, and confidence with technology, as well as the mean number of tasks completed online. See tables 2 and 3 for means and standard deviations.

Table 2: The mean (sd) frequency, difficulty and confidence with using technology according to age group.

	n	Often	n	Difficulty	n	Confident
Below 65	13	12.92 (2.33)	8	11.75 (4.53)	8	12.13 (4.85)
65-74	17	14 (1.80)	8	10.38 (2.07)	9	11.22 (2.44)
75 +	19	10.74 (2.70)	2	13 (0)	3	12.33 (1.53)
Total	49	12.45 (2.69)	18	11.28 (3.32)	20	11.75 (3.42)

Of the variables described in tables 2 and 3 significant differences were observed between the age groups regarding frequency of use ($F = [2, 46] 9.21$, $p < .001$), the number of functions used on a computer ($F = [2, 56] 4.18$, $p < .05$), and the number of functions used on a TV ($F = [2, 56] 3.35$, $p < .05$). Specifically, when compared to the 65-74 group, the over 75 group used technology less frequently and

fewer functions on a TV. They also used fewer functions on a computer than the under 65 group.

Table 3: The mean (sd) number of tasks completed with each technology according to age group[†].

	n	Computer	Internet	Mobile	TV
Below 65	13	2 (2.24)	1.08 (1.49)	1.85 (2.03)	2 (1.47)
65-74	25	1 (1.38)	.52 (1)	1.12 (1.48)	1.96 (1.27)
75 +	21	.48 (.98)	.38 (1.07)	.48 (1.36)	.95 (1.63)
Total	59	1.03 (1.58)	.59 (1.16)	1.05 (1.63)	1.61 (1.51)

[†] The possible range of scores were as follows; computer use = 0 – 12, internet use = 0 – 8, mobile phone use = 0 – 11, TV use = 0 – 8.

4 DISCUSSION

The present study indicated the frequency and ease of use are related to increased confidence. We therefore need to select a technology that is used frequently, and is easy to use. Of the technologies that were asked about, participants used the most functions on their TV. However, when considering the type of devices that could provide a suitable interface for older adults and mediate the assistive technology, the personal computer (PC) may be perceived as the ideal device. Currently, pre-installed operating environments (Windows, Mac OSX, and Linux) are too complex for untrained users to manage effectively, and as Oksanen-Sarela (Oksanen-Sarela, 2000) observes “(the) more complicated the technology, (the) less it gives space to different ways of using it as the user doesn’t have the skills or knowledge (to change its functionality)”. Furthermore, it is unlikely that the providers of AT systems will be able to provide the levels of training required by some users, as Dickinson et al (Dickinson et al., 2003) observe, “why should the user be re-educated and redesigned when it is the software that is inappropriate for their needs?”

It was found that that the technological artefacts that were used most often were the mobile phone and television. Given that frequency of use is associated with ease of use and confidence, these technologies have clear potential as applications for the AT. It has already been discussed (Millward & Nicholls, Submitted) that the telephone is not an acceptable medium to present the variety of data that can now be made available. For the proposed intervention, the television was chosen as the mode of delivery as it was the next most familiar. The group of participants with the least familiarity and

most difficulty with the technologies were those over the age of seventy-five. However, confidence levels were roughly equivalent between the three age groups which shows potential for training.

5 CONCLUSIONS

This paper has introduced a new programme of research that seeks to explore the use of assistive technologies (in particular telecare) in evaluating the psychosocial aspects of wellbeing. Barlow et al (2003) discuss that Telecare has been split into two areas – information provision and risk management. This split may be historical in nature, based on the evolution of Telecare systems from community alarm systems (Doughty & Williams, 2004), or it may be based on such systems originating from common requirements agreed between stakeholders from the health and care providers fields (Lines & Hone, 2004). Whatever the reason the result is that systems appear, at least to the end user, to be clinical in nature (Blythe, Monk, & Doughty, 2005). It is argued that it is time that the two facets of Telecare systems come together to form an EEAT system. The benefit of this type of system is that it may make Telecare systems more acceptable to older adults by being more attractive to older adults. This attribute could be stimulated by enhanced usability (this resonates with the definition of social dependability posited by (Blythe et al., 2005), (Dewsbury, Sommerville, Clarke, & Rouncefield, 2003). The specific aspirations of the client group will be assessed in stage 3 of the research plan. Further research will then be directed towards relating these ‘future requirements’ to the capabilities of the technologies currently available. The contribution of this research to the field is in identifying a potential new application of electronic assistive technology.

REFERENCES

- Barlow, J., Bayer, S., & Curry, R. (2003). Flexible homes, flexible care, inflexible attitudes? The role of telecare in supporting independence. *Paper presented at the Housing Studies Association Spring Conference*, University of York.
- Blythe, M. A., Monk, A. F., & Doughty, K. (2005). Socially dependable design: The challenge of ageing populations for HCI. *Interacting with Computers*, 17(6), 672-689.
- Curry, R., Trejo-Tinoco, M., & Wardle, D. (2002). *The use of information and communications technology to support independent living for older and disabled people*. London: Dept. of Health.
- Dewsbury, G., Sommerville, I., Clarke, K., & Rouncefield, M. (2003). A dependability model for domestic systems. In S. Anderson, M. Felici, & B. Littlewood (Eds.). *Computer Safety, Reliability, and Security: 22nd International Conference, Safecom 2003. Proceedings, Lecture notes in Computer Science*, 2788, pp 103-115. Heidelberg, Springer-Verlag.
- Dickinson, A., Eisma, R., & Gregor, P. (2003). *Challenging interfaces/redesigning users*. Vancouver, British Columbia, Canada: ACM Press.
- Doughty, K., & Williams, G. (2004). Practical solutions for the integration of community alarms, assistive technologies and telecare. *In s. c. o. s. a. technology* (Ed.):UK Government.
- Ellis, R. D., & Allaire, J. (1999). Modeling computer interest in older adults: The role of age, education, computer knowledge, and computer anxiety. *Human Factors*, 41(3), 345-355.
- Lines, L., & Hone, S. (2004). Eliciting user requirements with older adults: lessons from the design of an interactive domestic alarm system. *Universal Access in the Information Society*, 3(2), 141-148.
- Marquie, J. C., Jourdan-Boddaert, L., & Huet, N. (2002). Do older adults underestimate their actual computer knowledge? *Behaviour and Information Technology*, 21(4), 273-280.
- Millward, D. J., & Nicholls, W. (Submitted). The Invisible Observer: the challenges of human computer interaction with assistive technology systems for older adults. *Universal Access in the Information Society*
- Oksanen-Sarela, K. (2000). Invisible technologies, invisible boundaries? *Paper presented at the Home Informatics and Telematics*, Wolverhampton, England.
- Reed, J., & Monk, A. F. (2004). Using familiar technologies in unfamiliar ways and learning from the old about the new. *Universal Access in the Information Society*, 3(2), 114-121.
- Segrist, K. A. (2004). Attitudes of older adults towards a computer training program. *Educational Gerontology*, 30, 563-571.
- White, H., McConnell, E., Branch, L. G., Sloane, R., Pieper, C., & Box, T. L. (2002). A randomised controlled trial of the psychosocial impact of providing Internet training and access to older adults. *Ageing and Mental Health*, 6(3), 213-222.

TELEHEMATOLOGY

ICT Solution of a Shared Digital Image Repository

Daniel Schwarz

*Institute of Biostatistics and Analyses, Masaryk University, Kamenice 3, Brno, Czech Republic
schwarz@iba.muni.cz*

Miroslav Penka

*Department of Internal Hematooncology, Masaryk University, Brno, Czech Republic
m.penka@fnbrno.cz*

Ladislav Dusek, Petr Brabec

*Institute of Biostatistics and Analyses, Masaryk University, Brno, Czech Republic
dusek@iba.muni.cz, brabec@iba.muni.cz*

Keywords: Telehematology, Image Consultation, Web Based Learning.

Abstract: Telehematology is one of educational projects at the Faculty of Medicine of Masaryk University. The project has two various parts: the clinical part allows physicians to remotely consult their indefinite findings and the educational part brings new possibilities for contact tuition as well as for effective distant learning in a wide range of medical specialties. In this paper, the project is described mainly from the technical point of view. The design and implementation of the shared digital image repository are explained here.

1 INTRODUCTION

Hematology is the branch of medicine that is concerned with blood and its disorders. A subset of telemedicine services including digital transmission of visual information in hematology is usually referred to as telehematology (Beolchi, 2003).

Recent developments in digital photography together with the expanding use of the Internet have brought many new applications of image documentation, particularly in anatomical pathology (Leong and Leong, 2004) or in dermatopathology (Feit et al., 2005). First telehematology investigations were focused mainly on utilizing intrahospital networks and supporting clinical decision making (Mitsuhashi et al., 2000). Very specific hematologic issues are addressed in (Luethi et al., 2004), where the diagnostic accuracy in telehematology via email and in telehematology via real-time conference techniques was studied.

The goal of this paper is to describe the ICT solution of the Telehematology project at Masaryk University. The project has been initiated in cooperation with several hematooncology centres in the Czech Republic in 2006 with a long-term plan to

support medical students' access to digital microscopy images and to the documentation of diagnostics and treatment of serious diseases of the blood and the blood-forming organs.

2 ICT IN TELEHEMATOLOGY

Our Telehematology project consists of a clinical part and an educational part. It allows physicians to remotely consult their indefinite findings with the use of a shared digital image repository. The discussions as well as the images and their text descriptions are anonymous with a respect to patient information protection. In the educational part, it is available for contact tuition as well as for effective distant learning in a wide range of medical specialties.

2.1 Telehematology Workstations

Telehematology workstation is meant to be a site equipped with a laboratory microscope fitted with a digital camera which communicates through USB or Firewire interface with an ordinary personal compu-

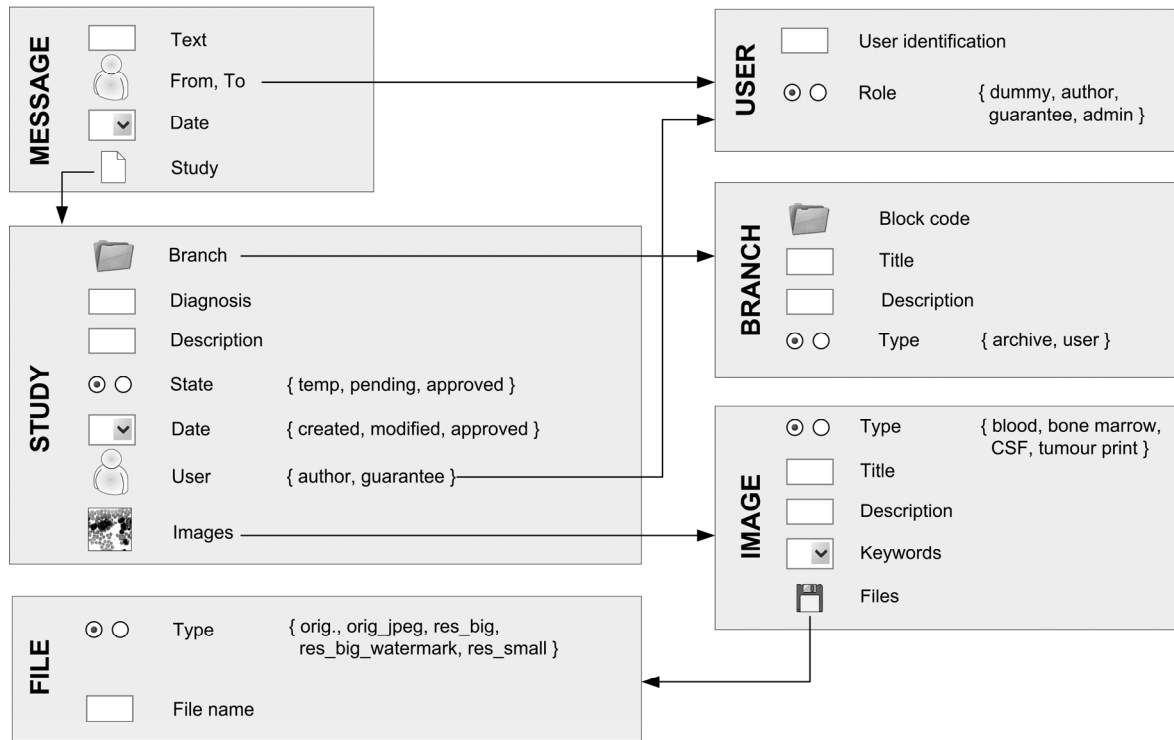


Figure 1: The simplified scheme of the shared image repository in the Telehematology project. Only the most important objects, their selected relations and properties are displayed.

ter (PC) connected to the Internet. The digital camera is usually attached to a trinocular head of the microscope with the use of a "c" mount video coupler. The communication between the camera and PC is controlled by a local software program which is usually supplied together with the camera. It is recommended to use also an additional piece of local software for advanced image editing. In order to send the captured images into the central repository and organize them in the image collection, a common internet browser is sufficient. A typical setup of a telehematology workstation involved in the Telehematology project includes a laboratory microscope Olympus BX41 (planachromat 10× and 20× objectives and a fluotar 100× objective), digital video camera Artcam 300 MI (3 Mpx, up to 46 fps, USB), optical coupler U-TV0.5X, PC control unit (Windows XP), software QuickPhotoMicro and freely distributed image editor GIMP.

2.2 Shared Image Repository

Despite the existence of various proprietary systems for digital image documentation and sharing, our original software has been developed, in order to

avoid over-reliance on technology by a single provider.

The shared digital image repository in our Telehematology project is implemented in the form of web-based application, what allows participation of any interested hematology site equipped with a telehematology workstation. The main components of the repository are shown in its simplified scheme in figure 1.

The image collection is organized in a tree-hierarchic structure. Particular tree branches represent either personal folders in the image consultation mode or diagnoses in the educational image atlas mode. In the latter case the WHO taxonomy of blood disorders is used. The tree-hierarchic structure is realized by the block code which is well-known in the field of communication techniques. It allows for arbitrarily deep branching.

The term “study” refers to a set of images captured from one case. Studies represent visual information about particular diagnoses of particular patients. Besides the images, each study contains a text description about the case. The number of images associated to a study is unrestricted. Each image is identified by its file name, text description and also by a single-choice parameter with a

following list of values: peripheral blood, bone marrow, cerebrospinal fluid, tumour print. In addition, a set of keywords selected from a predefined thesaurus may be added to each image object in order to allow advanced searching capabilities in the repository. Adding keywords from the thesaurus is implemented in the form of an interactive prompter which looks up a list of possible words compared to a searching phrase. The prompter is implemented with the use of AJAX technique to avoid extensive page reloading.

The images are stored in the repository in several versions: 1) the original file, 2) a copy with the original resolution converted to the JPEG format, 3) a converted copy resampled to VGA resolution (640×480), 4) a watermarked copy in VGA resolution and 5) a thumbnail resampled to 75 px along its longer side. All the image operations are done automatically without any user's intervention. Users may send images in arbitrary format. Format conversion, image resampling and watermarking are performed with the use of the freely distributed program ImageMagick. The computer on which the application is hosted is equipped with the Linux operation system, web server Apache and database server MySQL. Own programs of the shared image repository are written with the use of the scripting language PHP.

2.3 User Roles in Telehematology

The list of user roles in the image repository together with a description of user privileges follows. 1) Author: user may create new studies and edit existing own studies in the TEMP and PENDING state. The user may send messages to other authors and guarantees and associate the messages to particular studies. 2) Guarantee: user checks studies in the PENDING state, i.e. the studies which are completed by their authors. The user may change the state of studies from the PENDING to the APPROVED state. The user may send messages in the same manner as an author. 3) Dummy: the end-user of the shared image repository, who may only view the studies in the APPROVED state. User's view is restricted only to the resampled and converted images. Anonymous users may view only watermarked images. 4) Administrator: user with all privileges, who performs supervision over the repository. The user can extend the thesaurus of keywords or modify the hierarchic-tree structure of the image collection according to new requests from the guarantees.

All roles except the administrator role can be assigned not only in the global manner for the whole repository, but users may be granted also within a restricted subset of the tree-hierarchic structure.

3 PRACTICAL ASPECTS

During the initial phase of the Telehematology project there was a variety of difficulties which had to be cleared. The most important component of the project was always meant to be a stable and user-friendly shared image repository. Thus, the worst difficulties had been expected in this part of the project. However, the worst problems finally came out from the quality of captured images. Most of CCD as well as CMOS digital cameras which were tested required additional image editing in order to compensate for errors in white balancing. Images with an incorrect mixture of primary colors (red, green, blue) are not suitable for education purposes. In addition, such images with a color cast did not either motivate hematologists to extend the image collection nor to utilize the software resources of the project for image consultation. Having employed advanced image operations with the use of the software GIMP solved the problem.

During the first year of operation the authors have stored 336 images organized into 52 studies from which 31 have been still in the TEMP state and only 13 studies have been in the PENDING state. Only 7 studies have achieved the APPROVED state, what has showed us that the communication tools provided for authors and guarantees need to be improved, in order to speed up the process of revising studies into their final versions.

The importance of our Telehematology project lies in enabling students to access a well-organized and described collection of digital microscopy images for their undergraduate as well as postgraduate education in a wide range of medical specialties such as hematology, biology, histology, anatomical pathology, oncology etc. The web-oriented character of the project allows utilization of the gathered image material for the contact tuition as well as for self-study purposes.

Students at the Faculty of Medicine of Masaryk University usually find their multimedia and e-learning materials with the use of faculty's educational portal <http://portal.med.muni.cz>, which is described in (Dusek et al., 2006) in detail. The reference to the Telehematology project is included there among plenty of others.

4 CONCLUSIONS

In this paper, the Telehematology project at Masaryk University was presented, mainly from its ICT point of view. Key problems lying in the quality of captured images were mentioned.

The most original part of the developed image repository from the user's point of view is the thesaurus of keywords implemented in the form of the interactive prompter. Its usage avoids scrolling obscure long lists or time-consuming searching in tomes to get a right term.

The major part of the ICT solution for the shared image repository in the Telehematology project is used also in other web-based learning applications at Masaryk University, such as Image Atlas of Digital Mamography, Image Atlas of Dental Surgery, 3-D Models in Nerosurgery and many more, which are connected to medical image documentation.

REFERENCES

- Beolchi, L., 2003. *Telemedicine Glossary*, European Commision. Brusseles, 5th edition.
- Dusek, L., Schwarz, D., Racek, J., Brabec, P., Muzik, J., Regner, B., Zaloudik, J., 2006. Information and Communication Technologies in Education at the Faculty of Medicine of Masaryk University. In *Proc. of Information and Communication Technology in Education*. University of Ostrava.
- Feit J., Kempf, W., Jedlickova, H., Burg G., 2005. Hypertext atlas of dermatopathology with expert system for epithelial tumors of the skin. *Journal of Cutaneous Pathology*, 32, 433–437.
- Leong, F. J., Leong, A. S., 2004. Digital photography in anatomical pathology. *Journal of Postgraduate Medicine*, 50, 62–69.
- Luethi U., Risch L., Korte W., Bader M., Huber A.R., 2004. Telehematology: critical determinants for succesfull implementation. *Blood*, 103, 486–488.
- Mitsuhashi T., Kakai Y., Aral T., Shimizu N., Watanabe K., 2000. Telehematology Trials using the Sysmex LAFIA, Blood Cell Image Viking System. *Sysmex Journal International*, 10, 77–84.

A SURVEY OF INTEROPERABILITY IN E-HEALTH SYSTEMS

The European Approach

Marcos Da Silveira, Nicolas Guelfi

*Laboratory for Advanced Software Systems, University of Luxembourg, 6, rue Coudenhove-Kalergi, L-1359 Luxembourg
marcos.dasilveira@uni.lu, nicolas.guelfi@uni.lu*

Jerry-David Baldacchino, Pierre Plumer, Marc Seil, Anke Wienecke

*Centre de Recherche Public Henri Tudor 2A, rue Kalchesbrück, L-1852 Luxembourg
david.baldacchino@tudor.lu, pierre.plumer@tudor.lu, anke.wienecke@tudor.lu, marc.seil@tudor.lu*

Keywords: Interoperability, standards, e-health, information technology.

Abstract: The interoperability is often associated with the capacity of exchanging information that belongs to different workflows in a distributed environment. In e-health, the implementation of interoperable systems has a direct impact on the access of medical services, the costs and the quality of those services. This paper summarizes the efforts of standardization done by a selected group of Europeans researchers in the healthcare domain. The assessment of this standardization effort is made through a study of its impact in 20 FP5 and FP6 European projects that address the healthcare domain. The objectives are identify the new trends on interoperability technologies and to point out the importance of approximate industrials and research institutes.

1 INTRODUCTION

During the last century advances in medicine have significantly contributed to the reduction of the mortality rates. New medical practices, healthcare supporting tools (equipments, methods, medications, etc.), and the more participative role of the patients and their families in the treatments allow many patients to surpass physical and mental illnesses. But these practices are frequently isolated and information exchanges are done later when good results are obtained. The next challenge for healthcare domain is to implement collaborative works in a comprehensive environment where information exchanged between physicians are simple, fast and reliable. Consequently, this environment can contribute to reduce or to avoid misuse of new medicine practices, to improve the quality of healthcare services and to facility (and control) the medical information's access.

According to the lessons learned from Canada, Denmark and New Zeland (GAO, 2005), the standards' definition, the free access of national classification from all stakeholders, the education, the founding for implementation, and the proactive position of the government to solve privacy protection problems are important topics to be

considered during the definition of national e-health programs. It has been obtained by means of a common effort of industries and public institutions to define and to diffuse standards like HL7 (HL7, 2007), Cen/TC 251 (CEN/TC, 2007) and DICOM (DICOM, 2003) in order to improve the interoperability between healthcare systems. The aim of this paper is to verify if these in facto standards have been taking into account in researches funded by the European community.

The next section we present the impact that a set of standards has in a selected set of Framework Projects (FP5 and FP6).

2 STANDARDS IN RESEARCH PROJECTS

According to (IEEE, 1990), interoperability is *the ability of two or more systems or components to exchange information and to use the information that has been exchanged*. The scope of this paper covers only the interoperability between e-health IT systems. The importance of interoperability in the e-health market has increased with the rising number of healthcare professionals that use computers and

electronic devices to improve their services. Proprietary solutions have gradually been substituted by standard-based technologies. However, the implementation of these products has still facing a large number of barriers as well as unclear legislation, high implementation costs, physicians' reluctance (to adopt those technologies) and disorganization within healthcare structure. According to the Department of Health and Human Service of USA (HHS), healthcare is the largest sector of the economy that has not fully embraced information technology. The Medical Group Management Association and the Healthcare Information and Management Systems Society (GAO, 2005) reported that only 31% of physician group and only 19% of hospitals practices use fully operational Electronic Health Records (EHRs). This reality is often the result of adopting equipments and IT solutions that were not developed to interoperate with other systems. Changing these technologies is very expensive and takes time. The healthcare system is not ready to invest money without any guarantee of world interoperability and the market is not mature enough to give this guarantee.

This work contributes to analyze the European market trends based on the information took from a selected number of research projects in the e-health domain. We also point out if the interoperability challenge has really been attacked by scientist and industrials. During our researches, we evaluated two distinct levels of standards: One coming to integrate medical equipment (named physical level standards)

and the other used to integrate customers' applications (named application level). In the first level, the selected standards are: CAN, I2C, Bluetooth, Zigbee, USB, FireWire, RS232, IEEE 1284, Ethernet, GSM/GPRS, UMTS and IEEE 802.11x. In the second level the following standards are addressed: HL7, DICOM, ebXML and VITAL.

The analyze procedure starts by looking for references of this set of standards in the homepage of European projects (and in open source documents) available in the internet. In a second phase, we searched for partners of these projects that also participate on the specification of the selected group of standards. In the next section, the projects considered in our researches are presented.

2.1 Looking for Projects

In the previous sections we showed that the association of informatics with healthcare domains brings up many advantages and that the Healthcare industry is positioned to beneficiate of the advancements in technology and connectivity. High technological devices and software are available for healthcare services' providers and are adapted to the patient needs (patient-centered systems) (LAU et al., 2002). Customers of these technologies can expect to achieve greater performance, to reduce costs and to improve patient care. Consequently, they are expanding marketshare and pushing the transition to a digital era of e-Health. An example of the

Table 1: The selected list of European projects and the standards adopted.

#	Project	Area/objective	Supported Standards
1	C-CARE	EHR	XML, HL7, CEN, ISO
2	CHS	Home monitoring of Diabetes, heart failure, post trauma patients	UMTS, GPRS
3	HEALTHMATE	Telecare, Tele-consultation	XML, Bluetooth, GPRS, UMTS
4	HUMAN	Telemedicine, domotic	UMTS, GPRS
5	IDEAS	Multimedia architecture for e-health	XML, Bluetooth, GPRS, DICOM,
6	MOBIDEV	Secure access of medical database	Bluetooth, UMTS
7	MOBIHEALTH	Telemedicine, remote assistance	GPRS, UMTS
8	TOPCARE	Telecommunication support to Telecare	---
9	WIDENET	European EHRs interconnectivity	---
10	ARTEMIS	Semantic Web-Services	---
11	AUBADE	Neurology, psychology. Recognition of emotional state of the patient	Bluetooth, GPRS, UMTS
12	BIOPATTERN	Identification of European bioprofile	---
13	CLINICIP	Automatic injection of Insulin in ICU	---
14	COCOON	Healthcare risk management	---
15	DICOEMS	Integrated medical environment and database for critical situations	GPRS
16	INTREPID	Phobias' monitoring and treatment.	---
17	MYHEART	Intelligent clothes for heart failures prevention	Bluetooth, GPRS
18	NOESIS	Diagnosis supporting tools and Web-services	---
19	SEMANTICMINING	Data mining of medical information	---
20	PIPS	Generic medical database	---

Table 2: Synergy between standards and European projects FP5 & FP6.

PROJECT	STANDARD	STANDARD DEVELOPERS' MEMBER	INDUSTRIAL PARTNERS
C-CARE	XML, HL7, CEN, ISO	NO	HEALTH INFORMATION MANAGEMENT, STACKS CONSUTION E INGENIERIA EN SOFTWARE, OLIVETTI SANITA, MEDIGRIDGE, PROGEA, DATASOFT, INTRASOFT
CHS	UMTS, GPRS	YES (I2C, BLUETOOTH, ZIGBEE, FIREWIRE, USB, 802.11, HL7, DICOM)	PHILIPS, PAULADIS, AMERICAN MEDICAL DEVELOPMENT, CARD-GUARD
HEALTHMATE	XML, BLEUTOOTH, GPRS, UMTS	YES (CAN, I2C, BLUETOOTH, ZIGBEE, USB, 802.11, HL7, DICOM)	AIRTEL, BULL, SIEMENS, APLITEC, KNOSOS
HUMAN IDEAS	UMTS, GPRS	NO	SYNAPSIS, SEMA, EUTAELSAT, MBI
MOBIDEV	XML, BLEUTOOTH, GPRS, DICOM,	NO	APLITEC
MOBIHEALTH	BLEUTOOTH, UMTS	NO	AIRTEL, ARAKNE, INTERCEM, RELATIONAL TECHNOLOGY
MOBIHEALTH	GPRS, UMTS	YES (I2C, BLUETOOTH, ZIGBEE, FIREWIRE, USB, 802.11, HL7, DICOM)	PHILIPS, TELEFONICA, LOGICAM, TELIASONERA, HP, YUCAT, GESUNDHEITSCOUT
TOPCARE	---	NO	CALEA, GMD, DATAMED, DRAGER
WIDENET	---	NO	SADIEL, PROEC
ARTEMIS	---	NO	ALTEC, TEPE
AUBADE	BLEUTOOTH, GPRS, UMTS	YES (CAN, I2C, BLUETOOTH, ZIGBEE, USB, 802.11, HL7, DICOM)	ANCO, MARSERATI, SIEMENS
BIOPATTERN	---	NO	DAEDALUS, GAP, HELLENIC, HOEGSKOLAN, NEOVENTOR, SYNAPSIS
CLINICIP	---	NO	BRAUN, CARMEDA SISETRONIC, GAMBRO, SENSLAB
COCOON	---	YES (USB, FIREWIRE, 802.11, HL7)	AQUITAINE, CEFRIEL, ELYROS, ENPHASIS, MICROSOFT, ICSF, IDS, LOGICOM
DICOEMS	GPRS	NO	SYNERGIA 2000 , SSM
INTREPID	---	NO	AURELIA, ELYROS, INOCUS, PALLADION
MYHEART	BLEUTOOTH, GPRS	YES (I2C, BLUETOOTH, ZIGBEE, FIREWIRE, USB, 802.11, HL7, DICOM)	CSEM, HEIN, LINEAPIU, MANIFATURE, MEDGATE, MEDTRONIC, NOKIA, NYLSTAR, PHILLIPS
NOESIS	---	YES (CAN, I2C, BLUETOOTH, ZIGBEE, USB, 802.11, HL7, DICOM)	AIRAL, SIEMENS
SEMANTIC-MINING	---	NO	KITHAS, MERALL-ROSS
PIPS	---	YES (HL7)	ASTRAZENECA, ATENA, ATOS, GLAXOSMITHKLINE, MEDIC4ALL

increasing interest on e-health solutions are the quantity of research projects supported by the European Union, in particular in the *Framework Project 5 (FP5) and 6 (FP6)*.

The work presented on this section summarizes the main area/objectives of a selected list of e-health projects founded by FP5 and FP6. The objective is to verify how interoperability was considered in those works.

The selection procedure had started with a list of 80 projects and finished with 20 projects (after applying the filter). This filter consists on verify if the keywords e-services, standard, platform, telemonitoring, telemedicine, protocols, e-services, security and devices were present in the context of

the research projects. Table 1 shows the selected projects' name, their main area/objective and the standards that were explicitly indicated in their homepages. The evaluation methodology consists on the identification of industrial partners of each European project and cross it with the members' list of each emerging standard. The expected results are: Verification that European projects take into account emerging standards; Identification of new trends on the e-health domain; Point out the importance of industrial partners for those projects.

Table 2 presents the firsts results of these researches. The colon "*standards developers' members*" indicates that at least one of the partners belongs to

at least one of the standard members' list. The considered standards are: USB; CAN; ZIGBEE; Bluetooth; I2C (internet research with the keywords: company name + I2C); HL7; DICOM; FIREWIRE; IEEE 802.11; ebXML (the considered websites are in the references).

3 CONCLUSIONS

A list of interoperability standards used in healthcare systems was compared in this paper with a selected set of projects of framework programs 5 and 6. It is important to highlight that some projects have an incomplete/out-of-date website and the access to their technical reports were sometimes impossible. The available information led us to conclude that the international standards have not been taken into account by the majority of the considered FP projects. It identifies a gap between the industrial perspectives, the governmental efforts and the research trends. Lessons learned from New Zealand experiences (GAO, 2005) show that educating stakeholders about the value of developing health IT and using standards is very important to the success of a national plan. Diffusing the successes is also an important tool. This analysis allows us to highlight some weaknesses in the European research trends and can be useful to define new research strategies, in particular giving priority to projects that include international standards' promoters and participants (users, providers and developers). The participation of more research institutes in these committees can also contribute to the education of stakeholders.

REFERENCES

CEN/TC 251, European Standardization of Health Informatics. Retrieved June 10, 2007, from <http://www.cen251.org>

DICOM - Part 1 (2003). Introduction and Overview. Retrieved June 10, 2007, from http://medical.nema.org/dicom/2003/03_01PU.PDF

GAO - United States Government Accountability Office. Health information technology: HHS is taking steps to develop a national strategy. Report to the Chairman, Committee on the Budget, House of Representatives, May, 2005. Retrieved June 10, 2007, from <http://www.gao.gov/new.items/d05628.pdf>

HL7 - Health Level Seven. Retrieved June 10, 2007, from <http://www.hl7.org>

IEEE - Institute of Electrical and Electronics Engineers. IEEE Standard Computer Dictionary: A Compilation

of IEEE Standard Computer Glossaries. New York, NY: 1990.

Lau C., Churchill R. S., Kim J., Marsen F. A., Klim Y. (2001). Asynchronous Web-based patient-centered home telemedicine system. IEEE Transactions on biomedical engineering, Vol. 49, No. 12, December 2002, pgs 1452-1462.

The websites of standards members' lists are:
http://www.usb.org/members_landing/directory/complex_search_companies_form/process
<http://www.can-cia.org/cia/member-list/index.php?m=2&ss=&sa=Search&sl=100>
<http://www.zigbee.org/en/about/members.asp>
<https://programs.bluetooth.org/apps/directory/default.aspx>
<http://www.hl7.org/about/benefactors.htm>
<http://medical.nema.org/members.pdf>
<http://www.1394ta.org/About/Members/>
<http://www.ieee802.org/11/Voters/votingmembers.htm>
http://www.oasis-open.org/about/foundational_sponsors.php

INTEROPERABILITY IN SMART HOME MIDDLEWARE

The MPOWER Project

Sten Hanke and Thomas Fuxreiter

Austrian Research Centers GmbH – ARC, Biomedical Engineering, Viktor-Kaplan-Straße 2, Vienna, Austria
sten.hanke@arcsmed.at, thomas.fuxreiter@arcsmed.at

Keywords: Smart home, sensor applications, interoperability, middleware, MPOWER, interfaces, standards, service-oriented architecture, web services.

Abstract: The paper describes the use of interoperability standards and interoperability frameworks in a smart / sensor home project. During the EU funded project MPOWER an open middleware platform will be developed which should speed up the task of developing and deploying services for persons with cognitive disabilities and elderly. The developed middleware has different interfaces where interoperability standards are required. The paper points out these requirements and presents solutions in the different layers of the Service-Oriented Architecture approach. In future, standards and defined interfaces are more and more needed because of the need for a secure and easy data and messaging transfer. The middleware will be developed non-proprietary and is open for different applications and sensor integration.

1 INTRODUCTION

Interoperability in smart home and sensor systems is a rarely unvalued field in research and development so far. There is hardly any activity in using standards and defined interfaces although this area of research deals with a lot of different sensors and sensor data from the medical field as well as from the domestic domain and data have to be transferred to “outside” systems.

The operating experience has shown that a change of thinking is unavoidable for a faster and more cost efficient development and for a more safe data transfer.

For an efficient development in the future it will be important to overcome proprietary systems and solutions for every single smart home environment. With the experience from the medical field, where a lot of initiatives have been working on the development and promotion of standards and frameworks for years, an adaptation of the smart home context is desirable.

The following benefits can be mentioned: Interoperable and semantic systems are vendor independent and e.g. different kinds of sensors can easily be integrated. This opens the market and speeds up the development and inner data transfer. Furthermore the interfaces for any legacy system or nursing system are defined. Thus data can be

retrieved easily by any outer system, which fulfils the guidelines for the defined data format.

2 THE MPOWER PROJECT

MPOWER is a “Specific Target Research Project” (Contract number 034707), partially financed by the INFFSO DG of the European Commission.

The aim of the project is to define and implement an open platform to simplify and speed up the task of developing and deploying services for persons with cognitive disabilities and elderly.

With the start of the project an investigation of user needs was carried out. The aim was to gain knowledge about the needs of older people and people with dementia, in respect to technical IT solutions, which can support everyday living and to “aging in place”. Before developing new technological solutions, an investigation of the user groups’ needs and requirements was important.

When looking at the European society today, there are some striking trends: The growth of the older population, that is people over 65 years of age, is expected to rise in the decades to come. In general the health of older people is improving and life expectancy is rising in many countries.

This changing of the society in the next years reveals the need for a faster development of novel

and innovative technology solutions which are providing applications for an ambient assistive living.

The MPOWER middleware follows the IBM Service-Oriented Architecture (SOA) approach using web services. Applications built in the SOA are based on services. A service is an implementation of a well-defined business functionality, and such services can then be consumed by clients in different applications or business processes. SOA-based applications are distributed multi-tier applications that have a presentation, a business logic, and persistence layers. Services are the building blocks of SOA applications. While any functionality can be made into a service, the challenge is to define a service interface that is at the right level of abstraction (Mahmoud, 2005).

Web services are software parts designed to support machine-to-machine interaction over a network. This interoperability is gained through a set of XML-based open standards, such as WSDL, SOAP, and UDDI. These standards provide a common approach for defining, publishing, and using web services (Mahmoud, 2005).

The idea in the MPOWER project is to evaluate use cases and features for elderly people and people with dementia, based on user scenarios. The assigned features are leading to services, which will be implemented as web services. The application developer will be able to develop applications based on these services more efficient. For demonstrating the middleware there are two demonstration sites planned. One demonstration site will be based in Norway. The Norwegian pilot Point-of-Care (PoC) application will show the connection to a legacy system, storing information about medical treatment, social treatment and care planning. The other trial Point-of-Care application will be implemented in Poland. This demonstration site will show the integration of different domestic and medical sensors.

3 INTEROPERABILITY

An essential aspect of an architecture is the establishment of technical standards. In general, standards define common elements, such as user interfaces; system interfaces, representations of data, protocols for the exchange of data, and interfaces accessing data or system functions. Technical standards provide a number of advantages for the systems architect.

The partners will promote standardization through aligning their work with ongoing development of HL7, security and interoperability standards. Standards are important because they are accepted by multiple vendors, thereby increasing the likelihood that a collection of systems from diverse sources will be able to interoperate.

Definition of *interoperability* in ISO/IEC 23282-0: *Information Technology Vocabulary, Fundamental Terms*: “The capability to communicate, execute programs, or transfer data among various functional units in a manner that requires the user to have little or no knowledge of the unique characteristics of those units”

Interoperability can be achieved through:

- the implementation of standards,
- the usage of predefined sets of business procedures,
- the usage of accorded file formats and protocols for data transmissions.

4 INTEROPERABILITY IN MPOWER

Inside the MPOWER framework different interfaces for implementing interoperability standards and guidelines have been identified. Interoperability is important for every interface where data is transferred outside the closed system. Concerning the Service-Oriented Architecture (SOA) approach, messaging guidelines for data transfer between the different services are also a need of interoperability.

The following chapters will describe these three interfaces where interoperability needs in the MPOWER framework are identified and where and how standards are used.

4.1 Interoperability in Point of Care Systems

Homecare and Point-of-Care systems that are available today, usually provided by one single vendor, suffer under proprietary data and messaging transfer and lack very often in respect to data exchange and interoperability with other systems, especially third party systems.

A series of “health informatics, Point-of-Care” standards are being developed in a concerted approach between ISO, IEEE, IHE, and other major players in the field. These efforts are targeted towards implementing interoperable measurement

systems in healthcare, for example laboratory, intensive care and telemonitoring purposes as well as home care. The standards provide architectural components, information models and services (ISO/IEEE 11073-10201), and also methods for data exchange, both wire based, and wireless (ISO/IEEE 11073-30300).

Although the IHE started implementing the framework for health applications, the communication structure could also be used as well for homecare applications for elderly and persons with dementia. This could easily be done since there are already structures for Point-of-Care devices as well as for information sharing with legacy systems or care provider systems. The benefit would be the accordance to established standards and standard frameworks. It is clear that there is an advantage because of already established Point-of-Care communications from the medical field (blood pressure, temperature etc.). So the existing standards have to be extended for the homecare use cases and devices. Of course the whole IT structure and the cross-enterprise-document (XDS) sharing defined in the IHE framework could handle personal health data and information, which could be, depending on the use cases, important for medical help.

The ISO 11073 offers plug-and-play and a functional as well as a semantic interoperability between sensor systems and aggregation systems. In this standard all functions and use cases in patient oriented health care and of course in some aspects of smart homes for elderly are object orientated modelled already. That means a so-called domain information model is constructed where the device, the functionality, the measured data, settings, alarm functions, patient information and interfaces are defined. Furthermore there are codes for all information elements defined as “nomenclature” (ISO/IEEE 11073-10101) and “data dictionary” (ISO/IEEE 11073-10201).

The communication standard POCT-1A is implemented in the ISO 11073.9 and is specialized for patient near Point-of-Care. In principle the functionality of POCT1-A could be realized with HL7, but the functional range of the HL7 structure is in some cases (single sensors) too highly dimensioned. So it could depend on the special function that should be realized between the sensor (according sensors often have a restricted hardware) and the local controller. Because of a clear defined message communication the unique interpretation of the standard is guaranteed. POCT1-A is a flaring of HL7 not a competing standard.

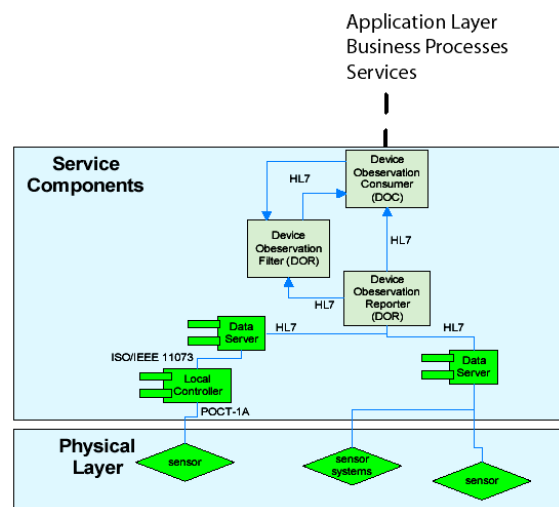


Figure 1: The Service Components Layer and the Physical Layer in the SOA approach.

Figure 1 shows the Service Components Layer and the Physical Layer in the SOA approach. The Service Components Layer contains the POC framework provided by the IHE with the Device Observation Filter, the Device Observation Reporter and the Device Observation Consumer.

The advantage of the CEN/ISO/IEEE 11073 is that it is the only comprehensive system of Point-of-Care medical device communication standards. The modality categories range from real-time-operating medical equipment to Point-of-Care test devices. Wired as well as wireless IR and RF network technologies are supported. If healthcare providers and management organizations want Point-of-Care to record transparency of information, then they must demand medical device interoperability. In addition to the core development bodies, the activity coordinates regularly with other health information activities (HL7, NCCLS, IHE and DICOM).

4.2 Interoperability to Legacy Systems

One of the key problems in healthcare and Point-of-Care is the lack of interoperability among different healthcare systems (service interfaces) at the one side and among diverse device examples and aggregation / computation systems (device interfaces) on the other side. Interoperability standards and frameworks would also advance the exchange of data and information between two system-interfaces respectively services as well.

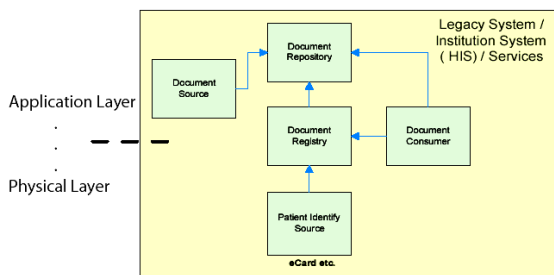


Figure 2: IT Infrastructure Technical Framework provided by the IHE.

Figure 2 shows the IHE Framework for Legacy Systems. The Document Source is connected to the middleware. The Documents Source could be any kind of health record as for example the Clinical Document Architecture, a XML schema developed by the HL7 group for store the HL7 messaging context in a document format. The Document consumer will be stakeholder who is interested in viewing the data or viewing parts of the patient related data, e.g. a doctor or a nurse. The identification is provided by the Patient Identity source; i.e. a registration as it could be managed by the eCard (www.chipkarte.at, 2007) in Austria for example.

4.3 Interoperability between Services

Between the services all medical data will be transferred in the well defined HL7 messaging. All contexts concerning non medical data e.g. the domestic sensor data will be transferred in a common XML schema with SOAP actions.

5 RESULTS

5.1 The Operational System Layer / Physical Layer

Based on the service oriented architecture model from IBM (Arsanjani, 2004) the interfaces for the interoperability have been integrated. For the development concerning the MPOWER middleware platform the layers of services and service or enterprise components are of interest.

The physical layer or operational layer consists of sensors as well as sensor systems, which provide a collected set of information. Through the operational layer there is of course also the possibility to have an access on existing custom built applications and legacy systems as well as for

instance other object-oriented implementation like older legacy which are providing miscellaneous data on the physical layer from bottom up.

The physical layer in the SOA approach represents the whole sensor layer, which provides the information of the smart house. This will be medical data like blood pressure as well as non-medical data like status sensors (door open, door closed, etc.).

5.2 The Enterprise Components Layer / Service Components Layer

The enterprise components or service components layer uses typically container-based technologies such as application servers to implement the components, workload management, high availability, and load balancing.

For future applications it must be guaranteed that different sensors or different sensor systems from different vendors could be connected to the system. In this layer device components will also be implemented. Based on the IHE framework the device components are defining the communication interface to the different services, which need detailed information of one or more sensors or sensor systems. The important interoperability interface for connecting different sensors and sensor systems to the system will be the data server in the service component layer. The data server will only accept data and information from registered sensors, which means that a sensor has to connect to the system. According to the ISO 11073 standards the sensor or the sensor system has to communicate its domain information model. The domain information model contains the following information: the device or sensor name, the functionality, the measured data the sensor will provide, settings, alarm functions and if possibly patient information (for instance when thinking of patient monitors or other complex measurement systems where such data are provided).

If a single sensor with a restricted possibility of processing power and memory capacity should be connected to the system the POCT-1A protocol standard is the recommended interface technology. The POCT-1A standard is optimized for the communication between sensors / devices and an observation reviewer or local controller. The protocol is XML, so it has the advantage for an easy interpretation to HL7. As mentioned before there are ambitions to integrate the POCT-1A into the HL7 v3.0 standard. The POCT-1A will be an add-on to the HL7 for e.g. memory restricted devices.

5.3 The Service Layer

In the service layer the different services for interoperability are present. The services could be used for the different business processes and interact with each other. The work of the interoperability group in the MPOWER project will be to provide services, which are the interface to the “outside” of the smart house. There will be services for document translation or extraction to provide medical relevant data for more comprehensive document sources like patient reports or health records. The export of the medical relevant data to a health record will be provided with ebXML and SOAP protocols and stored in a HL7 CDA adaptive format.

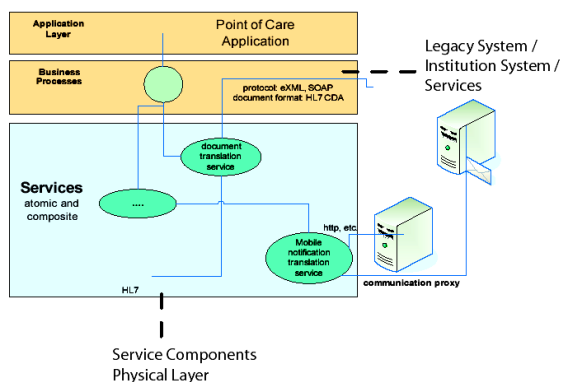


Figure 3: The Application Layer, the Business Processes Layer and the Service Layer in the SOA approach.

Figure 3 shows the Application Layer, the Business Processes Layer and the Service Layer and the connection to “outside” server proxies.

To have the option for providing information and data or a part of the smart house data to a healthcare delivery organization belonging to a clinical affinity domain (e.g. community of care) the cross-enterprise document sharing (XDS) defined by the IHE will be of interest. Within the XDS a federated document repository and a document registry create a longitudinal record of information about a patient within a clinical affinity domain. This profile is based upon ebXML registry standards, SOAP, HTTP and SMTP. By looking at these standards the MPOWER middleware platform could also provide interfaces and standards where a data exchange to such “outside” frameworks will be easily possible.

There will be services for mobile notification services like SMS and Email. The interfaces for these needs will be presented by services, which will make the translation of the different formless notification and alarming needs. It will be part of the Proof-of-Concept application, making the

demonstration, to set up the “outside” proxy, connected to the interface services provided by the interoperability working group, for the respective use case. The legacy system / institution is connected to the service layer of the MPOWER middleware because in this structure the legacy system is an external system where an interface for a possible data exchange should be provided. As mentioned above it must be clear that if the legacy system, which could be a hospital information system as well as any other proprietary system, the interface is in the operational layer.

5.4 The Business Processes Layer / Choreography Layer

In the business process composition or choreography layer different compositions and choreographies of services exposed in the services layer are defined. Services are bundled into a flow through orchestration or choreography, and thus act together as a single application. These applications support specific use cases and business processes (Arsanjani, 2004). According to the use cases a business process could access different services.

5.5 The Application Layer

The application developer will develop graphical interfaces for end users conjunct with invocations of service components as well as handling of business processes. Several use cases will be demonstrated during the MPOWER Proof-of-Concept applications.

6 CONCLUSIONS

Interoperability standards in smart home applications are not widely used and implemented. Different sensor systems and applications in the medical field have shown that standards and defined interfaces are needed for a satisfying message processing and data transfer. Proprietary systems are restrictive and there is no sharing of information. However the trend in Europe points to more non-central information admittance. More and more people are medically treated at home and telemonitoring is a growing IT field. Therefore novel sensor and medical data sharing applications have to be open for any data sharing. Integration of sensors from different vendors and information sharing to different legacy systems is required.

To meet these requirements it is not sufficient to have one fixed setup of individual components for the sensor system. Practical systems must be able to change within given limits, and still be able to fulfil their purposes. Different elements from the repository (sensors, processes, documents etc.) must be changeable by different vendors with similar properties. This goal can only be achieved through working together on integrating standards.

ACKNOWLEDGEMENTS

MPOWER is a “Specific Target Research Project” (Contract number 034707), partially financed by the INFFSO DG of the European Commission.

The authors would like to thank all persons and institutions in the four partner countries for participating in the investigation.

REFERENCES

- ISO/IEC JTC1 SC3, 2003. Information Technology for Learning, Education and Training. Proposed Draft Technical Report.
- Mahmoud, Q. H., 2005. Service-Oriented Architecture (SOA) and Web Services: The Road the Enterprise Application Integration (EAI).
- Arsanjani, A., 2004. IBM Service-oriented modelling and architecture.
- IHE Patient Care Device Technical Framework; Year 1: 2006-2007; Volume 1 Integration Profiles; Revision 1.1 Trail Implementation Version; Publication Date: 20060815
- IHE Patient Care Device Technical Framework; Year 1: 2006-2007; Volume 2 Transactions; Revision 1.1 Trail Implementation Version; Publication Date: 20060815
- IHE IT Infrastructure Technical Framework Volume 2 (ITI TF-1) Integration Profiles; Revision 2.0 – Final Text; August 15, 2005
- IHE IT Infrastructure Technical Framework; Volume 2 (ITI TF-2) Transactions; Revision 2.0 – Final Text; August 15, 2005
- eCard; www.chipkarte.at ; last seen: 06/07; Sozialversicherungs-Chipkarten Betriebs- und Errichtungsgesellschaft m.b.H. - SVC

WHO SHOULD ACCESS ELECTRONIC PATIENT RECORDS

A. Ferreira^{1,2}, L. Antunes², C. Pinho³, C. Sá³, E. Mendes³, E. Santos³, F. Silva³, F. Sousa³
F. Gomes³, F. Abreu³, F. Mota³, F. Aguiar³, F. Faria³, F. Macedo³
S. Martins³ and R. Cruz-Correia³

¹*Computing Laboratory, University of Kent, Canterbury, Kent, CT2 7NF, UK*

²*Department of Computer Science, Faculty of Science, University of Porto, Portugal, partially supported by KCrypt (POSC/EIA/60819/2004) and funds granted to LIACC through Programa de Financiamento Plurianual, FCT and Programa POSI. af84@kent.ac.uk, lfa@ncc.up.pt*

³*Faculty of Medicine, University of Porto and CINTESIS, Portugal
{med05167, med05169, med05184, med05185, med05186, med05189, med05190, med05187, med05188, med05191
med05193, med05237, med05096, rcorreia}@med.up.pt*

Keywords: Electronic Patient Record, access control, attitudes.

Abstract: Access control to Electronic Patient Records (EPR) may greatly depend on users' objectives and needs. The purpose of this study is to assess the opinions of medical doctors within a university hospital towards access control to an EPR. We selected a randomized sample of 58 doctors from a university hospital and 45 structured interviews were applied. 42 respondents (93%) agree with the existence of access control levels to patient information according to healthcare professionals' category and 31 (69%) think that more sensitive information (e.g. HIV) should be accessed only by doctors that treat those patients. As 24 doctors (53%) feel that there is no need for them to see all information about all the patients, 41 (91%) think that nurses should not be able to do it also. Further, 31 doctors (69%) believe that patients themselves should not access their full medical record. These results show that it is very hard to get to a consensual policy regarding access control to EPR by its regular users. There is therefore the need for a multidisciplinary agreement that can include healthcare professionals' experiences and needs in order to define the most appropriate and efficient way to perform access control to the EPR.

1 INTRODUCTION

Good communication between health providers is an essential component of high quality health care (Hassol et al., 2004). Paper-based medical record is still widely used in hospitals, where health professionals gather patient's clinical and administrative information. There is however some problems with this type of records and so computer-based medical records are being implemented and used in a more regular basis (Bakker et al., 2004).

The evolution of technology allows health providers to communicate electronically and to obtain information which includes patient's health story, examination findings, diagnosis and treatment over a period of time (Hassol et al., 2004) (Day, 2001).

This enabling technology that constitutes the informational basis for communication and cooperation in and between healthcare organizations is called Electronic Patient Records (EPR) (Ab et al., 2004).

However, this wide use of information systems and technologies shows the need for healthcare organizations to integrate and manage information from various sources, types and formats. This reflects the careful scrutiny that electronic access to medical information requires (Rogerson, 2000). Information security is then essential, moreover when people accessing the EPR can have varied objectives, different types of access and several processes to execute. Therefore, access control is essential to provide because it manages one of the first contacts between users of a system and its functionalities and features (Ferreira et al., 2005) (Ferreira et al., 2006).

According to a recent report, more than 1000 accidental deaths have been attributed to computer system failure (Gritzalis, 1997). Such occurrences must be present when considering the different interests and objectives that users want to achieve when using the EPR.

The Biostatistics and Medical Informatics Department of Porto's Faculty of Medicine

implemented a centralized EPR system (VEPR – Virtual EPR) between May 2003 and May 2004 in Hospital S. João (HSJ), Porto, Portugal. This hospital has more than 1300 beds and 5000 workers from 56 departments, where about 1000 are medical doctors, so any access to information needs to be properly defined, controlled and monitored. A generic but strong access control policy that reflects people's processes and interactions with the system, without incapacitating its use, is the basis for the VEPR success and, more importantly, acceptance, trust and use (Ferreira et al., 2005) (Ferreira et al., 2006). More than 900 doctors access this system on a daily basis, and this number is increasing, as healthcare professionals can feel the benefit from its use.

Even patient's access to their health records is now common in many places (Tracyl et al., 2004) (Pyper et al., 2004). How is access control going to be modelled in all these cases?

In this article we aim to get a small glimpse of what are the opinions of doctors working in HSJ towards who should access Electronic Patient Records, how should it be done and for whom this information should be (or not) restricted.

2 METHODS

2.1 Type of Study

This is an observational, descriptive, transversal study, in which the analysis unit is the individual.

2.2 Participants' Selection

Initially, we performed a bibliographic search of publications concerning access control to Electronic Patient Records. The next step was the selection of participants. Our target population was medical doctors. The available representative population was the medical doctors of the HSJ from a list available from the department of human resources at HSJ. From that list the medical doctors, department directors and pre-career doctors were selected. As a sampling method, from the filtered list, we selected a simple randomized sample of 92 elements.

2.3 Data Collection

The instrument used for data collection was a questionnaire with the characteristics of a structured interview, which was absolutely anonymous. The first steps in the questionnaire design were the research of questionnaires previously tested and the elaboration of a variable list.

The questionnaire was then pre-tested, in order to evaluate its validity and reproducibility. The pre-test's participant selection was made by a non-random accidental sampling process. The interviewer asked 10 HSJ doctors, who were at the hospital at that moment, to fill it in. Then, the final version of the questionnaire was elaborated with the pre-coded variables.

The questionnaire comprises 8 questions, some of them subdivided (see Appendix). The first 2 questions are global questions where doctors indicate the frequency they use the EPR and if there should be several access levels to records depending on the health professional's category (a Yes or No response). Question 3 refers to doctors' access control and question 4 refers to the access to more sensitive information about patients (like HIV tests). Question 5 demanded doctors' opinions about nurses' access to EPR. Questions 6, 7 and 8 describe other situations such as emergency situations, other uses of EPR and patient's access to their EPR.

The independent variables potentially relevant for the statistical analysis are: age, gender, professional category and department. This information was used to compare answers to the different questions (dependent variables) between these distinct groups in the statistical analysis.

The following step was the recruitment. Different departments were visited in order to find the doctors that were part of the sample. Those who did not work in HSJ anymore (29 people) or were already retired (5 doctors) were excluded, and the sample was reduced to 58 people. Then, the questionnaire was applied. If the doctors were not available at their department after three attempts, refused to answer the questionnaire or left it incomplete, they were eliminated from the study.

2.4 Statistical Analysis

In what concerns statistical analysis, we used SPSS to insert the collected data in a preformatted table.

We started to analyse our sample using absolute and relative frequency tables as well as pie graphs.

Chi-Square tests were also performed in order to evaluate the significance of the differences found between ages, genders, professional categories and departments, regarding the most relevant questions. As there are cases that do not respect the qui-square test's assumption (that require all expected values to be equal or superior to 5), some values are calculated using Fisher's exact test.

All the independent variables used in this study are categorical variables, except the age. In order to facilitate the data analysis, we transformed this numerical variable in a categorical one.

Furthermore, some independent variables were attached in categories so that we could perform a chi-square test. The variable age was separated in two categories: under 35 and over 35. We chose 35 as the dividing age because most doctors become specialists at that age. Professional categories were also divided in two categories: pre-career doctors and medical doctors. Departments were categorized in medical departments or medical – surgical. The significance level used in this study was 0.05.

3 RESULTS

Of the 58 applied questionnaires, 45 were fully answered, so the response rate was 78%. 10 doctors were not available in the department for three consecutive times and 3 refused to answer.

Participants’ characteristics are listed in Table 1. Most doctors were over 35 years old and there were more female doctors than male doctors.

Table 1: Respondents’ demographics (N=45).

Age	<35	22%
	>35	78%
Gender	Male	58%
	Female	42%
Department	Chirurgic	38%
	Medical	62%
Professional Category	General Intern	7%
	Specialist Intern	9%
	Specialist	38%
	Graduated Specialist	36%
	Service Director	11%

All doctors confirmed that they have already used EPR. Most of them said that they use this kind of records daily and that they agree with the existence of different access levels of information depending on the healthcare professional’s category. 93% (3) of the doctors said they agree and 7% (42) answered they do not agree or have no opinion regarding this issue.

In what concerns doctors’ access to information, the answers are summarized in Figure 1. More than a half of the respondents thought that doctors should not have full access to patients’ information. While some thought that doctors should only have access to the information of the patients they treat, others considered that they should have access to all the information of their department.

Further, 31 (69%) respondents thought that sensitive information such as HIV tests, venereal or cancer diseases should only be accessed by doctors who treat those patients.

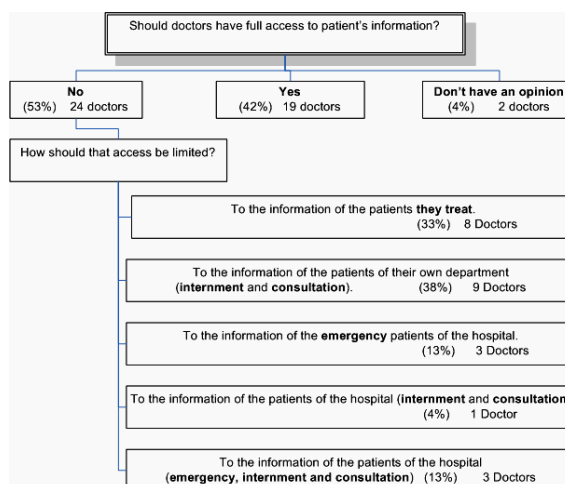


Figure 1: Answers for doctors’ access to a full EPR.

In what concerns nurses (Figure 2), a vast majority of doctors (41 - 91%) thought that they should not have full access to patients’ information. The majority believe that nurses should only have access to the information of the patients they treat.

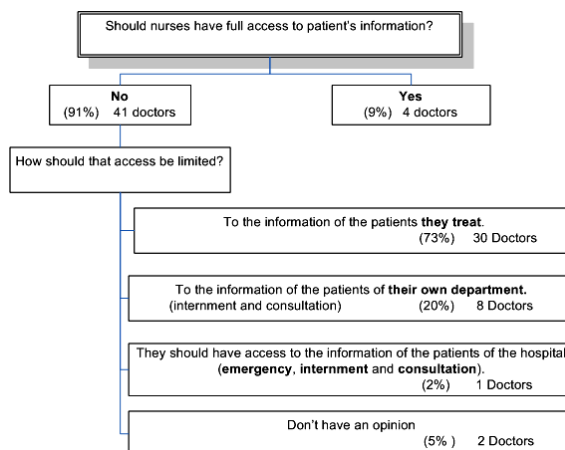


Figure 2: Answers for nurses’ access to a full EPR.

Most doctors also agreed that, in emergency situations, non authorized doctors and nurses must have access to patients’ information, but that access must be registered and controlled (Ferreira et al., 2006). The majority of respondents found pertinent to use the patients’ records to other purposes such as clinical or epidemiologic investigation.

Regarding now patients, most doctors thought that patients should not have full access to their clinical information, 69% (31) thought that they should not be able to access it while 31% (14) said that they should.

4 CONCLUSIONS

From these results we can see that EPR are intensely used by doctors. We can also discuss that doctors are mostly concerned with situations regarding sensitive information (e.g. HIV tests), and patients' access to these type of records. This is why they see access control as an essential part of the EPR.

Also, doctors do not agree with the fact that patients should be able to access the whole of their healthcare record, thinking probably some of the notes they make should be for their own use only. This opinion is also demonstrated in another study where they seem to be worried about the information accessed by their patients.

Further, doctors were reluctant in what concerns nurses' access to patients' information. They think they should only access the information of the patients they treat. This can be problematic as nurses spend more time dealing and treating patients than the doctors themselves and may need all the information about the patient relating to other types of treatment they can had been undergoing. It should be noted that all doctors had an opinion regarding this matter.

Our study also shows a tendency between some variables. It is interesting to note that, within the 4 doctors who think that nurses should have total access to information, 3 were male doctors and 3 were specialists.

Finally, doctors' attitudes towards the use of information for other purposes such as research were mostly positive. They also vastly agreed with the existence of different levels of access to EPR.

In conclusion, these results show that it is very hard to get to a consensual policy regarding access control to EPR by its regular users.

There is therefore the need for a multidisciplinary agreement that can include healthcare professionals' experiences and needs in order to define the most appropriate and efficient way to perform access control to the EPR. Several issues concerning the type of information, location, type of user and other situations (e.g. emergency or other unanticipated) may influence the way access control should be made.

We believe that this is a very important issue to be pursued and further studied. There is the need to evaluate more healthcare professionals and patients' attitudes and needs in order to define a better way to perform access control to EPR (Ferreira, Cruz-Correia et al., 2006).

ACKNOWLEDGEMENTS

We would like to thank class 8 of the 1st year medical students from the 2005/2006 academic year at the Biostatistics and Medical Informatics Department of the Faculty of Medicine of Porto for their work and enthusiasm in the development of this project.

REFERENCES

- Ab, B., Addabit, B.V., 2004. Access to EHR and access control at a moment in the past: a discussion of the need and an exploration of the consequences. *International Journal of Medical Informatics*, 73, 267-270.
- Blobel, B., 2004. Authorisation and access control for electronic health record systems. *International Journal of Medical Informatics*, 73(3): 251-257.
- Day, J., 2001. Privacy and Personal Health Data in Cyberspace: the Role and Responsibility of Healthcare Professionals. *The journal of contemporary Dental Practice*, 2(1).
- Ferreira, A., Cruz-Correia, R., Antunes, L., Chadwick, D.W., 2007. Access Control: how can it improve patients' healthcare? *Studies in Health Technology and Informatics*. IOS Press, 127:65-76.
- Ferreira, A., Correia, R., Antunes, L., Palhares, E., Farinha, P., Costa-Pereira, A., 2005. How to start modeling Access Control in a Healthcare Organization. *Proceedings of the 10th International Symposium on Health Information Management Research*.
- Ferreira, A., Cruz-Correia, R., Antunes, L., Farinha, P., Oliveira-Palhares, E., Chadwick, D. W., Costa-Pereira, A., 2006. How to break access control in a controlled manner? *Proceedings of the 19th IEEE Symposium on Computer-Based Medical Systems*, 847-851.
- Gritzalis, D., 1997. A Baseline Security Policy for Distributed Healthcare Information Systems. *Computers and security*, 16(8):709-719.
- Hassol, A., Walker, J., Kidder, D., Rokita, K., Young, D., Pierdon, S., Deitz, D., Kuck, S., Ortiz, E., 2004. Patient Experiences and Attitudes About Access to Patient Electronic Healthcare Record and Linked Web Messaging. *J Am Med Inform Assoc*, 11, 505-513
- Pyper, C., Amery, J., Watson, M., Crook, C., 2004. Access to electronic health records in primary care- a survey of patients' views. *Med Sci Monit*, 10(11):SR17-22.
- Rogerson, S., 2000. Electronic Patient Records. *IMIS*, 10(5).
- Tracyl, S., Dantas, C., Upshur, R., 2004. Feasibility of a patient decision aid regarding disclosure of personal health information: qualitative evaluation of the Health Care. *BMC Medical Informatics and Decision Making*.

CLASSIFICATION OF PREMATURE VENTRICULAR BEAT USING BAYESIAN NETWORKS

Lorena S. C. de Oliveira, Rodrigo V. Andreão and Mário Sarcinelli-Filho
Federal University of Espírito Santo, Graduate Program on Electrical Engineering
Av. Fernando Ferrari, 51, Goiabeiras, CEP 29075-910, Vitória, ES, Brazil
Tel.: 00 55 27 4009 20 77, Fax: 00 55 27 4009 2737
lorena@ele.ufes.br, rodrigo@ele.ufes.br, mario.sarcinelli@ele.ufes.br

Keywords: Artificial Intelligence, Medical Informatics, Bayesian Networks, Decision-Support Systems, PVC detection.

Abstract: This paper presents a system based on Bayesian networks (BN) to support medical decision-making. The proposed approach is able to learn from available data, and provides an intuitive graphical interpretation of the problem, which can be easily configured by a physician. This approach is evaluated for the first time in the problem of premature ventricular contraction (PVC) detection, using a representative set of records of the MIT-BIH database. The results obtained emphasize the capability of the Bayesian network to make decisions even when the information about some symptoms or events is not complete. Moreover, the good performance obtained opens many perspectives for the use of BN to deal with beat classification.

1 INTRODUCTION

In the last two decades a great effort has been invested to develop systems to automatically interpret long-term electrocardiogram (ECG) records. Two important reasons are the great demand for those exams and the long time specialists spend to make a complete diagnosis based on such records. The particular interest for ECG is due to its efficiency in the diagnosis of arrhythmia and the great incidence of cardiac diseases in industrialized countries (Kadish et al., 2001).

Most of the automatic analyses of ECG are made through rule-based systems conceived by experts in the fields of artificial intelligence and pattern recognition. In general, a rule-based system consists in acquiring the knowledge about a given process or problem through a set of examples or facts already happened in order to apply it to new situations related to the same problem. Several works in this field employ heuristic rules, neural networks and statistical approaches to build such systems.

Heuristic approaches model the human reasoning through a set of deterministic rules, which are very dependent on how the individual deals with a particular problem. The cause-effect relations among

the rules can be graphically represented, allowing the individual to follow the decision logic and criticize the results. However, this kind of approach does not necessarily consider the uncertainty, a key feature when regarding decision-making systems.

An example of a heuristic approach for PVC classification is presented in (Andreão, Dorizzi and Boudy, 2006), where the limitation of the heuristic rule is dealt with through using regions of certainty related to the possible values of a certain variable.

On the other hand, statistical approaches are built after a learning phase based on a set of selected examples. Thus, the classification capability of this class of approaches is highly dependent on the information learned a priori. However, they embed a certain potential of evolution (through using a new set of examples in the learning phase).

A method which is in evidence nowadays to deal with arrhythmia classification is the use of neural networks (Farrugia, Yee and Nickolls, 1991; Kuppuraj, 1993). However, such approaches perform as black boxes, making too hard to an expert in cardiology to interpret and configure the classifier. Another weak point associated to the use of neural networks is their inflexibility to adapt themselves to new examples, since the learning procedure demands a large set of examples. Moreover, the uncertainty related to the classification problem is not really treated. That is

why many systems employing neural networks should be improved before aiming at being considered as a tool for the decision-making associated to a diagnosis (Crawford et al., 1999).

In order to overcome the limitation of the NN to deal with uncertainty, Gao et al. (2005) used a probability measure in the output layer of an NN for each beat class. However, the other limitations of the NN above mentioned continue unsolved.

Finally, other statistical approaches, such as those employing Bayesian networks, have shown to be promising to address most of the limitations of the NN. Indeed, Bayesian networks are quite suitable to deal with uncertainty (Pearl, 1988), they are flexible enough to learn from new examples and they provide a very intuitive graphical representation that could easily be configured by a cardiologist.

Following this reasoning, a rule-based system is here proposed to assist the cardiologist in the diagnosis of premature ventricular beats, a quite common arrhythmia present in long-term electrocardiograms. Such a system is based on the Bayesian network framework, which is quite suitable to take into account the uncertainty associated to human interpretation. To the extent of the author's knowledge, this is the first time such framework is used for the particular problem of beat-classification. The major characteristics of such approach are hereinafter stressed, and experiments validating it are presented as well.

2 METHODOLOGY

2.1 Bayesian Networks

Probabilistic methods are a well-known topic in the field of artificial intelligence (AI), and are quite useful to model uncertainty.

To know probability is a need when one should deal with the idea of a random experiment, which generates events having an assigned uncertainty of occurrence (Clarke and Disney, 1970). In the theory of probability one finds the Theorem of Bayes, which is employed to compute $P(A|B)$ the conditional probability of the occurrence of an event A given the evidence B , i.e., the event B was observed. Such a theorem states that

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}, \quad (1)$$

where $P(A|B)$ is the posterior probability of A given B , $P(B|A)$ is the prior probability of B given A , $P(A)$ is the prior probability of A , and $P(B)$ is the prior probability of B .

The Bayesian networks are based on the Theorem of Bayes, working on the causal relations among random variables. They are direct acyclic graphs, whose nodes represent random variables with assigned uncertainty whose arcs represent the direct causal relation between the connected nodes. These causal relations are quantified through conditional probability distributions (Pearl, 1988).

The network does not necessarily have nodes corresponding to all causes of a given event, since the influence of irrelevant factors is modelled by the probability. Thus, using just a few variables it is possible to deal with a large number of causes.

The Bayesian networks can be composed of discrete random variables (multinomial distribution), continuous random variables (Gaussian and exponential distributions (Shachter, 1986; Buntine, 1991)), or a mixture of both.

2.2 ECG Analysis

The automatic analysis of the ECG signal has been a topic of research over the three last decades. The ECG is a record of the heart electrical activity in which changes in the elementary waveforms of the signal (P, QRS complex and T waves) characterizes an abnormal beat (see Figure 1).

In particular, the premature ventricular contraction (PVC) is a heart beat which is generated by an electrical impulse which does not follow the normal electrical conduction path through the heart (sinus node, atrioventricular node, and ventricles). Instead, it starts at the ventricles earlier than expected. A PVC beat is characterized by a heart beat without the P wave (atrial contraction), a premature and large QRS complex and a compensatory pause just after, as shown in Figure 2.




Figure 1: Heartbeat observed on an ECG with elementary waveforms and intervals identified (Andr e, Dorizzi and Boudy, 2006).




Figure 2: Electrocardiogram containing PVC beat.

Because of its characteristics and the availability of ambulatory ECG signal databases rich on this arrhythmia, most work in the field of ECG analysis evaluate the performance classifying systems through the detection of PVC beats.

Taking the characteristics of the PVC beat, however, it is very difficult to precise how much premature it is. Actually, a beat is premature when the R-R interval, which means the time interval between its peak and the peak of the previous beat, is shorter than it should be. The problem here is that the decision of how long should this interval be to characterize a premature beat depends on the point of view of the cardiologist, which characterizes an uncertainty that should be taken into account by the system during classification.

2.3 Bayesian Network for PVC Classification

The first step to be followed when building a Bayesian network is to identify, from the problem given, the random variables and their causal relations. Figure 3 illustrates the graphical representation of the Bayesian network implemented here. The nodes, represented as rectangles, are discrete random variables, while the nodes, represented as circles, are continuous random variables. RR is the random variable modelling the time interval between two consecutive QRS-complexes (or heart beats), whose probability density function (pdf) is a Gaussian function. The node LL is also a random variable with a Gaussian pdf, now representing the measure of the duration of the QRS-complex, in terms of likelihood (Andreão, Dorizzi and Boudy, 2006).

Every time a PVC episode occurs (PVC node is true) the related beat is premature (*Premature Beat*




Figure 3: Graphical representation of the Bayesian Network for PVC beat classification.

node is true), and if its QRS-complex is larger than the normal one the *Ventricular Node* is also true. Since the nodes PVC , *Premature Beat* and *Ventricular Beat* are discrete, their conditional probabilities are represented by a table in which the binary possibilities true (T) and false (F) have their related probabilities.

After identifying all random variables, it is necessary to estimate their respective probabilities. This procedure can be accomplished by a specialist having the prior knowledge about each variable of the system. In our case, the knowledge of the specialist about the distributions (pdfs) of the RR and LL random variables was obtained through a labelled database, where each heart beat of the ECG record has a *Normal* or a *PVC* label (see Section 3). However, the RR and LL values are not provided by the database. In this paper, it was used the HMM-ECG system developed by Andreão, Dorizzi and Boudy (2006), which automatically labels the signal, returning as a result the RR and LL values of each detected heartbeat. The values are normalized according to the value of the last normal beat detected.

The mean and variance of the LL and RR pdfs were estimated over a training set extracted from the MIT-BIH (1997) database, and the results are shown in Figures 4 and 5, respectively. There, one can observe that the normal values are all normalized as one and the abnormal ones are spread to the right for the LL values and to the left for the RR values.

The other variables of the network are modelled through tables of conditional probability, which are shown in Tables 1 and 2.




Figure 4: Histogram of the likelihood of the Normal and Ventricular QRS complexes. Two Gaussian functions are used to approximate both histograms.

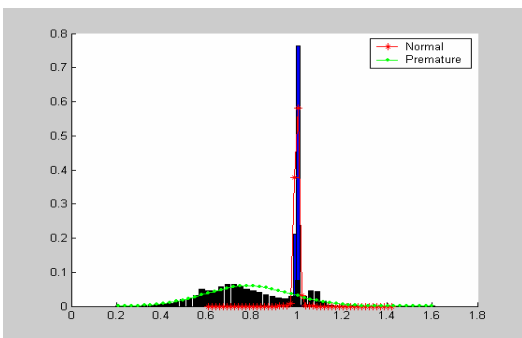


Figure 5: Histogram of the RR interval of the Normal and Premature beats. Two Gaussians functions are used to approximate both histograms.

The estimation of the variables individually does not model the relations of dependency among variables suitably. This is why we have implemented a learning step where all network parameters are adjusted from a training set of examples. Such a learning step is performed in two different ways: 1) it is firstly considered that the five nodes are all observable. This means that from the labels of the database we can identify the right value for each node, given the observations of *RR* and *LL*; 2) the second learning strategy considers that some nodes are non-observable (hidden) nodes. In this paper, the nodes *RR*, *LL* and *PVC* are observable, while the nodes *Premature Beat* and *Ventricular Beat* are non-observable. The necessary information is provided by the HMM-ECG system (Andreão, Dorizzi and Boudy, 2006) and the cardiologist labels as well. Since there are non-observable nodes in the second strategy, the expectation-maximization learning algorithm has been used to train the network based on a training data set. For the first strategy the classical junction-tree method (Pearl, 1988) was adopted, which also maximizes the probability of the observations given the model.

Table 1: Probability of PVC after a process of heuristic learning.

PVC	Probability
T	60%
F	40%

Table 2: Probability of Premature Beat and Ventricular Beat after a process of heuristic learning.

PVC	Prem. Beat	%	PVC	Ventr. Beat	%
F	F	95%	F	F	95%
T	F	5%	T	F	5%
F	T	5%	F	T	5%
T	T	95%	T	T	95%

3 RESULTS

All experiments have employed the MIT-BIH database, which possesses forty eight sequences of heart beats. However, only forty three of such sequences were used here, because recorded sequences containing pace beats or too much signal amplitude distortion were removed.

The forty three sequences used were split in a training set, which was used to adjust the parameters of the Bayesian network, and a test set, necessary to evaluate the performance of the trained network in terms of PVC classification. The total number of beats (normal and PVC) correspond to 95,257 beats, where 64,074 beats were used for training and 31,183 were used for testing. It is important to remark that the MIT-BIH database contains other types of beats, which were considered as normal in our experiments.

The Bayesian network was built using a MATLAB toolbox called BNT (2002). The Expectation Maximization and Junction Tree learning algorithms were used to train the network.

The performance of the network is assessed in terms of: 1) confusion matrix; 2) sensibility, here understood as the capacity of the system to correctly identify normal beats (true positives); 3) specificity, here understood as the probability of classifying the PVC beats (true negatives); 4) positive predictive, which is the probability that an event detected as normal effectively belongs to this class of beats; 5) negative predictive, which is the probability that an event detected as PVC effectively belongs to this class of beats.

The first experiment evaluates the performance of the network after manually estimating the probabilities of each random variable based on the training data set (a try-and-error methodology), as

described in the previous section. The results obtained for the test data set are in Table 3.

The second experiment considers the effect of the learning-from-data step, which results in a better modelling of the relationship among variables, for which just observable nodes are considered. The network parameters were adjusted based on the same training data set, and the results for the same test data set are shown in Table 4.

The third experiment also performs parameter estimation through a learning strategy. However, in this case, observable and non-observable nodes were considered, for the same training data set used in the previous experiments. The results for the same test data set are shown in Table 5.

One can observe from Table 3 that the negative predictive value is very low, showing that the PVC detection provided by the system is not trustful enough yet, since these type of beat has a probability of just 32,34% of being the correct one. On the other hand, the adoption of a learning-from-data step improves significantly the system performance. This means that the strategy to be adopted for estimating the contribution of each variable requires some knowledge about the other variables, which is fulfilled by the training method adopted, as well as means that the estimation of the parameters for each variable independently is a quite poor strategy.

Table 3: Classification results for the Bayesian network without the learning-from-data step.

Confusion Matrix		
	Classification N	Classification V
Label N	26.427	3.007
Label V	312	1.437
Sensibility		98,83%
Specificity		82,16%
Positive Predictive		89,78%
Negative Predictive		32,34%

Table 4: Classification results for the Bayesian network after the learning from data step and considering only observable nodes.

Confusion Matrix		
	Classification N	Classification V
Label N	29.244	190
Label V	388	1.361
Sensibility		98,69%
Specificity		77,82%
Positive Predictive		99,35%
Negative Predictive		87,75%

Table 5: Classification results for the Bayesian network after the learning from data step and considering observable and non-observable nodes.

Confusion Matrix		
	Classification N	Classification V
Label N	29.344	90
Label V	358	1.391
Sensibility		98,79%
Specificity		79,53%
Positive Predictive		99,69%
Negative Predictive		93,92%

When comparing the two training strategies, one can observe that the one with hidden nodes is significantly better in terms of negative predictive (a smaller number of false positive PVC beats has been identified). The main reason for this improvement is the inclusion of a certainty zone in the values of *LL* and *RR* generated by the HMM-ECG system (Andrião, Dorizzi and Boudy, 2006). Thus, the hidden nodes have been left free to be estimated by the learning method and, hence, a more appropriate value is computed regarding the observed events.

Finally, the good results (see Table 5) obtained confirmed that the Bayesian network is a powerful tool to acquire knowledge from an available data set labelled by an expert. Moreover, when used as a tool for diagnostic aid it can indicate the degree of certainty associated to each result, through using the concept of probability.

In Christov et al (2006), the PVC wave is classified using two approaches, knowing Morphological Descriptors (MD) and Matching Pursuits for extracting time-frequency beat descriptors (TFD), and the result found for MD and TFD were, respectively, 96,27% and 94,77% for Sensibility; 99,13% and 99,08% for specificity; 89,87% and 89,19% for positive predictive and finally 99,70% and 99,58% for negative predictive. As a comparison, the Bayesian approach presents a better result in terms of sensibility and predictive negative. However, it is necessary to stress that in Christov et al (2006) more than one channel is used, while in this work it was used just one channel (regarding the recorded ECG).

In the approach proposed by Andrião, Dorizzi and Boudy (2006), the result for just one channel is 64.36% for Sensibility and 66.14% for Positive Predictive. When using two channels, however, the results are significantly improved: 87,20% for Sensibility and 85,64 % for Positive Predictive

The performance of the Bayesian network for only one channel has confirmed that the method is

suitable for the proposed application. Although the performance is not yet as good as those of the best systems, some improvements can be carried out in the model, through the use of channel fusion.

4 CONCLUSIONS

For the first time this framework is employed in this particular problem. The Bayesian network presents a more comprehensive graphical representation, deals with uncertainty through its probabilistic representation, and can work with incomplete data through its inference engine.

The capability of learning the network parameters from a training data set was verified using two training strategies, and the result is that when working with non-observable nodes the training method based on the EM algorithm produces a better modelling of the uncertainty related to the observed data and the labels defined by a cardiologist.

Our future work will focus on evaluating the performance of this system using a fusion strategy in order to explore information obtained from multiple channels. On the other hand, this network will be extended to classify more arrhythmias, as ischemic episodes. We hope that this system can be further developed and then implemented to assist an expert in the analysis of such events in ECG signals.

ACKNOWLEDGEMENTS

This research has received financial support from CAPES (a foundation of the Brazilian Ministry of Education), CNPq (an agency of the Brazilian Ministry of Science and Technology) and FAPES (a foundation of the Secretary of Science and Technology of the Government of the State of Espirito Santo, Brazil).

REFERENCES

- Andreão, R. V., Dorizzi B. and Boudy, J. (2006) 'ECG Signal Analysis Through Hidden Markov Models', *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 8, August, pp. 1541-1549.
- BNT, How to use the Bayes Net Toolbox (2002) viewed 06/07/2007 <http://boole.cs.iastate.edu/book/1-Science/1-ComputerScience/3-Paper/1-AI/Bayesian/How%20to%20use%20the%20Bayes%20Net%20Toolbox.htm>.
- Buntine, W. L. (1991) 'Theory Refinement on Bayesian Networks', *Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence*, 13-15 July, 1991, Morgan Kaufmann, Los Angeles, California, pp. 52-60.
- Clarke, A. B. and Disney, R. L. (1970) *Probability and Random Processes for Engineers and Scientists*, New York: John Wiley & Sons, Inc.
- Crawford M. H. et al. (1999) 'ACC/AHA Guidelines for Ambulatory Electrocardiography', *Journal of the American College of Cardiology*, vol. 34, no. 3, pp. 912-948.
- Christov I. et al (2006), 'Comparative study of morphological and time-frequency ECG descriptors for heartbeat classification', *Journal of Biomedical Engineering, Medical Engineering & Physics*, vol 28, no.7, pp. 876-887.
- Farrugia S., Yee, H. and Nickolls, P. (1991) 'Neural Network Classification of Intracardiac ECGs', *Proceedings of the IEEE and INNS Int. Joint Conf. on Neural Networks*, 18-21 November, 1991, Singapore, pp. 1278-1283.
- Gao, D. et al. (2005) 'Bayesian ANN Classifier for ECG Arrhythmia Diagnostic System: A Comparison Study', *Proceedings of International Joint Conference on Neural Networks*, July 31 - August 4, 2005, Montreal, Canada, vol. 4, pp. 2383-2388.
- Kadish A. et al. (2001) 'ACC/AHA Clinical Competence Statement on Electrocardiography and Ambulatory Electrocardiography', *Journal of the American College of Cardiology*, vol. 38, no. 7, pp. 2091-2100.
- Kuppuraj, R. N. (1993) 'A Neural Network System to Classify Simulated ECG Rhythms', *Proceedings of the IEEE Biomedical Engineering Conference*, New Orleans, Louisiana, USA.
- Lauritzen, S. L. and Spiegelhalter, D. J. (1988) 'Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems (with discussion)', *Journal of the Royal Statistical Society Series B (Methodological)*, vol. 50, no. 2, pp. 157-224.
- MIT-BIH Harvard-MIT Division of Health Sciences and Technology Biomedical Engineering Center (1997) 'Arrhythmia DataBase Directory', viewed 06/07/2007, <http://www.physionet.org/physiobank/database/html/mitdbdir/mitdbdir.htm>.
- Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, 2nd printing, San Francisco: Morgan Kaufmann.
- Shachter, R. D. (1986) 'Evaluating Influence Diagrams', *Operations Research*, vol. 34, no. 6, December, pp. 871-882.

USING A COMPLEX NUMERICAL AND MULTIMEDIA DATABASE IN INTERNAL MEDICINE

Pre-release Evaluations

Liana Stanescu, Dumitru Burdescu, Cosmin Stoica-Spahiu, Anca Ion

University of Craiova, Faculty of Automation, Computers and Electronics, Romania

stanescu@software.ucv.ro, burdescu@software.ucv.ro, stoica_cosmin@software.ucv.ro, anca.ion@software.ucv.ro

Dorin Stanescu

S. C. EDA SOFT, Craiova, Romania

dorin@eda-soft.ro

Keywords: Medical application, medical images, color and texture features, content-based image query.

Abstract: The paper presents a software tool implemented using Firebird and Delphi technologies, dedicated for managing and querying medical multimedia databases. The database contains images related to the internal medicine area. This on-line application allows creation of complex medical files of patients that can be viewed and updated both by internist and general practitioner. The main functions of the application are: managing patients contact information, examinations, imagery and personal folders; simple text based query; content based query using color characteristic for images provided by medical devices. It can be used in individual offices, laboratories or in the hospital clinics and departments. The application provides security and confidentiality for patient's data.

1 INTRODUCTION

Internal medicine is an important component of the general medicine that can be regarded as basics for many specialties: cardiology, pneumology, gastroenterology, nephrology, haematology, rheumatology, etc. It needs large amounts of paraclinical exploration performed by different devices that generates visual and numerical data used for diagnosis and for follow-up of treatment or evolution. Numerical information is given by functional and biological assays while various imagistic devices provide image data. In his reports, the doctor is making qualitative descriptions of abnormalities he found and elaborates reports that will be forwarded to the patients or to his colleagues from other departments. That is why internal medicine departments usually accumulate huge quantities of medical data, including thousands of image files, millions of numerical values and thousands of written reports.

Because most of the patients in these services are chronically ill patients, with frequent visits and frequent use ambulatory services, the rapid access to

an archive containing both numerical and imagistic data would be an advantage. More than that, because access to high tech medical services is quite limited in some countries, many patients are frequently investigated in geographically dispersed medical centres. By using a dedicated application, the local doctor will be able to check simultaneously all the investigations and to give a diagnosis. Thus, he can integrate better all diagnostic data, manage the visits or the therapy taking into account all the concomitant pathology, or send the patient to another department for other investigations. As consequence, other specialists are able to see information from local or regional database and can query the archives.

These are some reasons for creating a complex application for managing and querying a database containing information and images from medical domain. Our database is implemented in Firebird (Interbase) that is a free and modern database management systems (Interbase, 2006).

The application is implemented using Delphi programming language that has important facilities for database management (Kermann, 2001). It can

be used on-line, so the users have access from distance.

Besides managing patient information, including their consultations, the application has the possibility for managing and viewing the images provided by medical devices.

An element of originality is content-based visual query using color characteristic. It permits selection of a query image and finding all similar images from the database (Del Bimbo, 2001; Smith, 1997). This option could be very helpful for establishing the diagnosis.

2 THE DATABASE STRUCTURE

In this section, database structure used by the application including tables and logical connections between them is presented in detail. The database contains a number of tables populated when installing the application. The tables contain a series of codes that makes easier the work of updating patient information and the investigations. These tables are:

- Medical units with a specific code and name: the medical units can be an individual office, a laboratory or hospital clinics and departments.
- Users groups: this table stores information about user groups, administration rights for each group (administrators, doctors, nurses etc.). The table also contains a group identifier and a group name that can be specified for each group. Each unit has specific management rights.
- Users: is a table that contains user identification, name, password, and corresponding group.

The data confidentiality is ensured by user name and password that are provided separately for each unit. In order to increase the data security, the password is encrypted. Each doctor has access to information regarding his own patients, but he can share some data that can be seen by other specialists. He can also access both statistical and scientific data regarding all patients in database but in this case the identification data about patients is hidden and the ID number, name and address are blinded. At information management level, anonymizing all information concerning clinical data ensures confidentiality; diagnosis, paraclinical information, treatments are also blinded and only statistical data can be viewed. On the other hand, the office secretary can see personal information about patients

and referring doctors, but no diagnosis and treatment elements are accessible.

- Diagnosis table is used for storing the diagnosis code and name.
- Analyses table codifies paraclinical and biological data. It has the following structure: code, description, and minimal and maximal value.
- Clinical examination table includes the following elements: code, description, and value.
- Patient groups include code and description. The patients can be grouped by the category of disease (digestive, cardiology, renal etc.), by the medical insurance category or by participation to different programs (national health programs or clinical studies).

The following tables are the most important in the database because they store information about patients, examinations, investigations and results:

- Patients table is used for storing information about patient's visits: personal ID number, name, doctor, county, city, address, phone/fax number, email and program – if any.
- A patient might have several examinations, for each of them storing in the Consulting table, the diagnosis, date and treatment. Each examination might contain one or several clinical examinations (it is stored code, description, a series of analyses identified by code and obtained value).
- Images table is storing information about still or moving images, obtained from a patient during his whole disease history. These data are: path and name of the image file, type of image and color information automatically extracted for later content-based image query on color feature.

3 THE FUNCTIONS OF THE APPLICATIONS

3.1 Set-up

This function permits updating auxiliary tables in database as for example tables containing diagnosis codifications, clinical examinations, analyses codifications, departments, user groups and users.

3.2 Patients Information Management

This function is one of the main functions in the application, and the information about patients has the following organization:

a) **Contact information** (personal ID number, name, address, phone, fax, email, category and National License Number of the examining doctor).

b) Examinations

The management of this information is implemented in a window that contains several secondary windows, as seen in figure 1. The first secondary window contains a record for each patient examination with the following information: examination date, diagnosis and results of the visit (solved/unsolved). This window is associated with four secondary ones, having the following functions:

1. Collecting data from clinical examination. For each analyzed system or segment the user can specify if he found normal or abnormal relations. A short description of the found abnormalities can also be added.
2. Collecting numerical data from laboratories and adding other information about important data.
3. Storing, as descriptive text, the results of various investigations: radiological, echography, endoscopy descriptions.
4. Storing treatment recommendations and prescriptions resulting from diagnosis.

Data from each secondary window can be easily updated using coding tables that were created in the set-up phase. The solution with secondary windows was chosen because the doctor should have an instant and easy access to the whole evaluation of the patient from one examination to another.

c) Imagery

This option gives access to all the functions of the application referring to the imagistic data concerning a patient, images provided by different devices (echograph, endoscope, MRI, CT, etc). These images can be loaded from saved files or can be imported directly from medical devices using a real time acquisition system. The system can be launched directly from this window. Imported images will be saved directly in the patient folder. It is possible to see the images directly, or to select one as a query image and to execute a content-based image query for the whole database to search similar images.

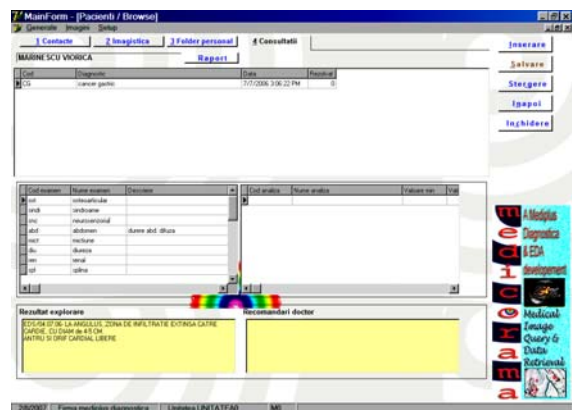


Figure 1: Window for managing patient examination data.

3.3 Database Query

It is one of the most important functions of the application. There are two types of queries: text-based query and content-based image query on color feature.

In the first case there are several search criteria that might be composed using “and”. These criteria are: patient name, personal numerical number, address, clinical exam and analyses. For the first three criteria it is used the “LIKE pattern” operator. Using this operator the Select command permits searching a string in all the values existing in database. For the last criteria, where information is codified in the database, the user can display a list of options to select one for search. Some useful queries that can be used are: list all the patients with a specified diagnosis, or list of all the patients that undergone a certain investigation or examination.

Visual data needs more evolved access methods. Such a method is content-based retrieval, which takes into consideration attributes or characteristics extracted from multimedia information. If we take into consideration the images, the technique is called content-based visual retrieval (Del Bimbo, 2001; Smith, 1997). This type of query implies selecting an image as query image, and finding all the images in database that are similar with it.

The medical areas where content-based visual queries methods can bring advantages are well known (Müller et al, 2004; Lehmann et al, 2004; Shyu et al, 1999):

- Diagnostic aid
- Medical teaching
- Medical research
- Electronic patient records

The reasons presented above generate the need to implement in the application the content-based visual query methods using color characteristic.

The color is the visual feature that is immediately perceived on an image. The color space used for representing color information in an image has a great importance in content-based image query, so this direction of research was intensely studied (Del Bimbo, 2001).

There is no color system that it is universal used, because the notion of color can be modelled and interpreted in different ways (Gevers, 2004).

It was proved that the HSV color system has the following properties: it is close to the human perception of colors; it is intuitive; it is invariant to illumination intensity and camera direction;

The operation of color system quantization is needed in order to reduce the number of colors used in content-based visual query: from millions to tens. The quantization of the HSV color space to 166 colors, solution proposed by J.R. Smith, is the idea used in this application (Smith, 1997).

The intersection of the histograms is used for computing the similitude between the query image Q and the target image T for color feature (Smith, 1997):

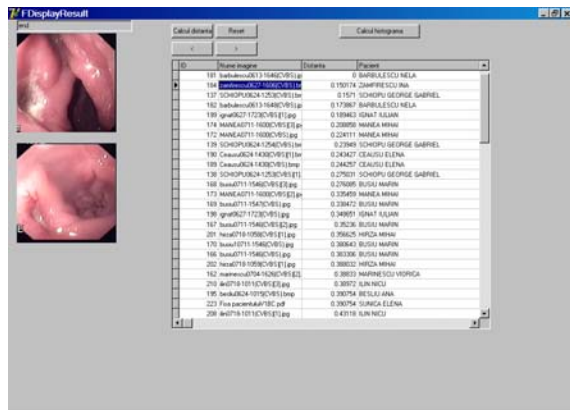


Figure 2: Window for content-based visual query.

The results of the experiments performed on a database with 960 images from the field of the digestive area are summarized in table 1.

The values in the table represent the number of relevant images in the first 5 retrieved images. For these five types of images (each type representing one diagnosis) the results are encouraging, but the experiments must be performed on a larger imagistic database and taking into account more types of color images. Only this way the conclusions can be reliable.

Table 1: Content-based image query experimental results.

Query	Nr. of retrieved images
Polyps	3
Colitis	4
Ulcer	3
Ulcerous Tumor	3
Esophagitis	4

4 CONCLUSIONS AND FUTURE WORK

The paper presents the design of a Firebird database, containing medical alphanumerical and imagistic information. The application is implemented in Delphi and can be accessed on-line. The main functions of the application are:

- Managing information about patients: contact information, visits, imagistic recordings, laboratory results, treatment
- Simple text based query that might combine several criteria
- Content-based visual query using color characteristic.
- Generates complete or synthetic reports

REFERENCES

Interbase Documentation, <http://info.borland.com/techpubs/interbase/>, 2006.

Kermann, M., 2001. *Programming and Problem Solving with Delphi*, Addison-Wesley.

Del Bimbo, A., 2001. *Visual Information Retrieval*, Morgan Kaufmann Publishers. San Francisco USA.

Smith, J.R., 1997. *Integrated Spatial and Feature Image Systems: Retrieval, Compression and Analysis*, Ph.D. thesis, Graduate School of Arts and Sciences, Columbia University.

Müller, H. et al., 2004. A review of content-based image retrieval systems in medicine-clinical benefits and future directions. *International Journal of Medical Informatics* 73, 1-23.

Lehmann, T.M. et al., 2004. IRMA - Content-based Image Retrieval in Medical Applications. *Proceedings of the 14th World Congress on Medical Informatics*, 842-848.

Shyu, C. et al., 1999. ASSERT, A physician-in-the-loop content-based image retrieval system for HRCT image databases. *Computer Vision and Image Understanding* 75, 111-132.

Gevers, T., 2004. *Image Search Engines: An Overview. Emerging Topics in Computer Vision*, Prentice Hall.

A LIFE SUPPORT NETWORK FOR ELDERLY PEOPLE LIVING IN A RURAL AREAS

Bayme Abaydulla, Jun Sasaki, Michiru Tanaka, Keizo Yamada and Yutaka Funyu
Faculty of Software and Information Science, Iwate Prefectural University, 152-52 Sugo, Takizawa, Iwate, Japan
g231e029@edu.soft.iwate-pu.ac.jp, {jsasaki, m-tanaka, k-yamada, funyu}@iwate-pu.ac.jp

Keywords: Independent Senior, Remote Healthcare, Life Support.

Abstract: This paper proposes a new concept, a Life Support Network (LSN), for elderly people living in rural areas. The network is an intranet that incorporates a safety confirmation system, a remote healthcare system and an emergency information system. We developed an experimental LSN system and carried out a field experiment in a typical rural town “Shiwa” in Iwate Prefecture of Japan. We demonstrated the experimental LSN, called “Yui Net,” performs well in the field.

1 INTRODUCTION

Recently, the number of elderly people living in the rural areas of the world, especially in developed countries, has been increasing. In Japan, the number of families which include people over 65 years old is 17,273,000 (37.7% of all homes), the percentage of senior citizens who are in a couple is 28.7%, and that of elderly people who live alone is 19.7%. These rates are still increasing, and at an unexpectedly high speed (The Cabinet Office, 2005).

Support systems are necessary to address problems associated with advanced depopulation and aging in rural regions. These systems might, for example facilitate senior citizens’ being able to live alone. In addition, various research approaches are necessary in order to help prevent the solitary death of elderly people. Some researchers have asserted that information technology (IT) can be effective in improving the welfare situation of the elderly.

In section two of this paper, we present the concept of a Life Support Network. Then in Section 3 we introduce an experimental system we developed, called “Yui Net,” and the results of the field experiments performed on the system are discussed in Section 4. We present our conclusions and suggestions for future work in Section 5.

2 THE LIFE SUPPORT NETWORK

We propose a type of intranet for a closed rural area to support elderly life. We call this intranet a “Life Support Network” (LSN). The LSN creates a constant connection among volunteers, home-visiting nurses, the regional hospital, drugstores and other life and health related organizations. The LSN is a high-security and high-speed network, which is like a Local Area Network (LAN) carried over optical-fiber cables. With current technology, it is possible to construct such a network within an appropriate budget (Komine et al., 2001 in Japan).

With the LSN, a senior citizen living alone can send a daily safety confirmation message to his/her remote contact person by way of an easy-to-use terminal. The remote contact person could be, for example, a helper or a care provider in a social welfare council, a doctor or nurse in a healthcare center or hospital, a friend, a neighbor or a family member.

An elderly person can use the healthcare terminal in his/her home every day. Each person measures his/her vital data and confirms it through a public facility such as a healthcare center connected to the LSN (Sasaki et al., 2006). In this case, nurse home-visits, home medical services, residential care, emergency calls, and other necessary services remain available. In case of an emergency, necessary personnel such as a doctor or helper can quickly

engage in correspondence with the elderly client (Shinagawa et al., 2006).

This plan was designed to achieve all stages or services in the LSN through an information system. The system ensures a safer living environment for residents and senior citizens living alone in rural areas, not only in Japan but throughout Asia (The Cabinet Office, 2004).

3 EXPERIMENTAL SYSTEM

3.1 Structure of “Yui Net”

As the concept of the LSN is quite broad, we decided to focus on some specific important functions: safety confirmation for elderly people living alone, and the remote healthcare service. We call the set of the systems “Yui Net,” where “Yui” means “helpful relationship” in Japanese.

The development project team consists of the NEC Group for hardware development, Iwate Prefectural University for software development, and the Shiwa town in the Iwate Prefecture of Japan as the feasibility test field. We developed each system as an open source Web application, which is available over a Virtual Private Network (VPN) in the Shiwa town.

The Local Authorities Satellite Communications Organization (LASCOM) of the Japanese Support Organization supported our project by funding the development. We would like to thank LASCOM for their contribution.

3.2 Safety Confirmation System

The most commonly used technology to ensure the safety of elderly people is sensor-type systems. One Japanese company produces an interesting sensor-type system. In this system, when the elderly client uses the electronically-equipped kettle that has telecommunication functions, the system sends a text message to a remote family member’s cellular phone (Zojirushi Corporation, 2005).

Many people dislike using such sensor-type systems because of privacy concerns. Furthermore, there are occasional occurrences of false alarm messages being sent. Alternative safety systems with human behaviour sensibility and high reliability are required to augment sensor type systems (Shinagawa et al., 2006).

We have reported on the development and operation of a “Mimamori” (meaning “watching over someone”) network system for the elderly in

the Village of Kawai in the Iwate Prefecture of Japan. In developing the Mimamori system, an L-mode terminal produced by NTT (Nippon Telephone and Telegram Ltd.) was chosen. An L-mode terminal is a telephone with a touch panel display and internet connection function. When the elderly client touches the display, the terminal sends a message to the social welfare council and remote family members (Yoneda et al., 2006). Unfortunately, NTT has decided to stop producing the L-mode terminal in 2010.

We therefore proposed a new method to confirm elderly people's safety by using TV terminals instead of L-mode terminals. We chose TV terminals because most elderly people in Japan watch TV routinely. According to a report by the NHK (Japan Public Broadcasting Corporation), elderly people, particularly those 70 years or older, watch TV for five hours or more a day on average.

Figure 1 shows the structure of the safety confirmation system. The television with STB (Set Top Box) and the Intranet environment are set up in the home of an elderly person. The Web server with the custom developed support system is located at Iwate Prefectural University.

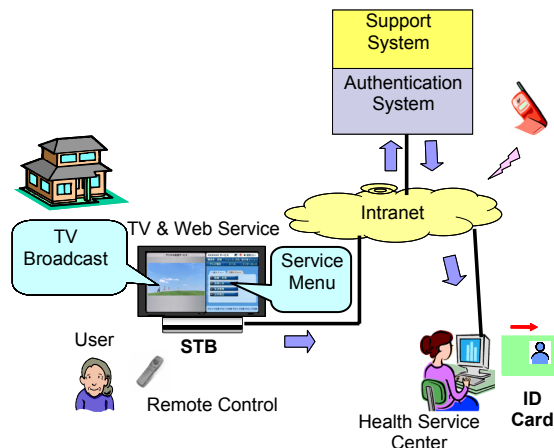


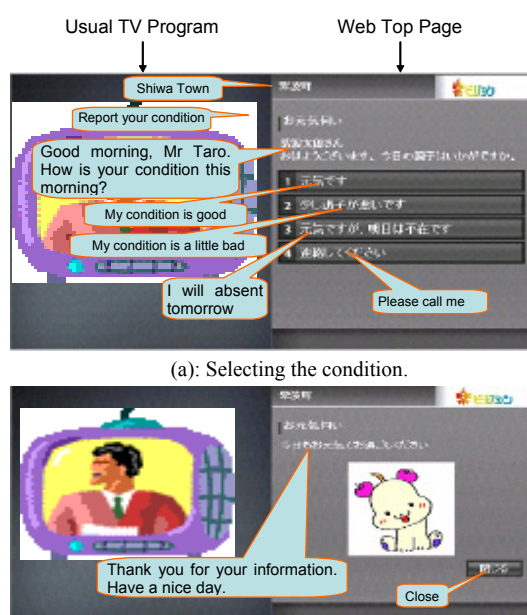
Figure 1: Safety confirmation system.

We can see the display is divided into two sections. The left-hand section represents the current TV program and the right-hand section represents the first page of the Web-based safety confirmation system. This page shows a brief message, below which there are four response buttons to select from:

- (1): My condition is good
- (2): My condition is a little bad
- (3): I will be absent tomorrow
- (4): Please call me

The user can easily select a button using the STB's remote control device. After the user selects his/her current condition, a Web page displays a confirmation message as shown in Figure 2(b). After the user presses the close button, the TV screen reverts to display the usual program broadcast. The information selected by the user is transmitted to the healthcare center and/or the remote family member over the intranet for display on a PC as well as a text message for display on a cellular phone.

If the user does not use the system to report their condition in the morning, the related life supporter (home helper, care provider, family member, neighbour, etc.) uses a telephone to confirm the user's condition.



(a): Selecting the condition.
(b): After selecting the condition.
Figure 2: Samples of the TV display.

3.3 Remote Healthcare System

Figure 3 shows the structure of the remote healthcare system.

The user can measure his/her vital data such as blood pressure, heart rate, electrocardiogram, body fat ratio and internal organ fats every day using ahealth checker'; a stand alone device. The items to be measured are determined according to user's specific health condition and interest.

The measured data stored in the health checker is automatically transmitted to the PC via infra-red. The PC, in turn transmits this data to the server located at Iwate Prefectural University through the intranet.

If a user has measures his/her vital data at a facility, the system transmits the data to the user's file in the database. All the data for a user can be input and accessed from anywhere.

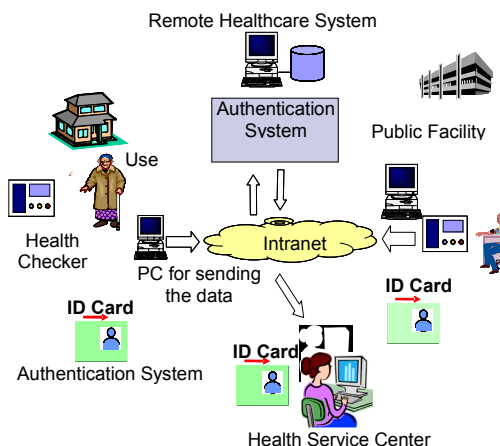


Figure 3: Remote healthcare system.

4 EXPERIMENT

4.1 Experimental Method

We carried out the experiment from December 18, 2006 to February 9, 2007 to confirm the effectiveness of the system.

We selected and requested the participation of three persons (a male aged 80 and two females aged 82 and 76) who live alone to use the safety confirmation system. The users sent their condition information through the system every day using the STB's remote control.

We selected and requested the participation of three different persons (three males aged 73, 59 and 82) to use the remote healthcare system. We set up the remote healthcare system in the office of three public facilities located in Shiwa town. There, any person with an ID card can use the remote healthcare system.

Table 1 shows the categories and numbers of system users in the experiment.

Table 1: System users in the experiment.

Name of system	Category of users	Number of users
Safety confirmation system	Home user living alone	3
	Family user living apart from the home user	2
	System manager	6
Remote healthcare system	Home user	3
	User in each office of three public facilities	36

Following the experimental period we obtained the users' opinions on the systems by way of discussion and a survey questionnaire.

4.2 Experimental Results and Discussion

4.2.1 Safety Confirmation System

Regarding the operability when using the safety confirmation system, two home users answered "comprehensible" and one female home user answered "incomprehensible" because of her low IT literacy. Two users said the TV screen display is easy to understand and they hope to continue using it. The system managers responded that they were able to operate the system satisfactorily. Family users living apart from the home users reported they felt a "feeling of safety" regarding the home users.

We confirmed that users could use the system with little difficulty and the life supporters and their families felt assured by checking the information sent from the user every day by cellular text message and web page.

The assimilation of the system into the responsible organizations and optimal management in cases of a lack of safety information are issues which remain to be resolved.

4.2.2 Remote Healthcare System

Sixty percent of all users said the remote healthcare system was "comprehensible" and easy to use, including the health checker for measuring health data and the PC for sending the data on.

We confirmed that the users could operate the system comparatively easily and they could take interest in their own health condition by using the remote healthcare system.

In future we plan to propose systems to effectively use the vital data captured for users in the town.

5 CONCLUSIONS

This paper proposes a new concept: a Life Support Network (LSN) for elderly people living in rural areas. We developed two experimental systems as part of the LSN and carried out a field experiment. The experimental system, which has a safety confirmation system and a remote healthcare system, obtained a satisfactory evaluation from users and shows good feasibility with satisfactory results in the field.

The experiment is the first step in constructing the LSN concept. It is important to prepare several life-support services, which the users and their relatives can select according to their particular health condition and social environments.

REFERENCES

- Cabinet Office, Government of Japan, 2005. The 6th life and international comparison investigation result summary concerning consideration of senior citizen. In *online*, http://www8.cao.go.jp/kourei/ishiki/h17_kiso/index.html.
- Komine, Y., Nakazono, M., Ikaruga, S., Iwamoto, S., Sato, T., 2001. Residential network in rural area neighboring city. In *Journal of City Planning Review*, Vol.; No. 230, pp: 77-82.
- Nakazawa, J., Takahashi, K., 2006. Policy and planning for ensuring information and communication networks/services in disasters. In *The Journal of the Institute of Electronics, Information, and Communication Engineers*, Vol.89, No.9, pp: 782-786.
- Sasaki, J., Wano, K., Yoneda, T., Kamata, H., Mizunuma, Y., Funyu, Y., 2006. A medical consultation support system to reduce time and costs for life-style related disease. In *Japan Journal of Medical Informatics*, Vol. 26, No. 1, pp.13-22.
- Shinagawa, Y., Kishimoto, T., Ohta, S., 2006. Development of an automatic alarm system for elderly living alone. In *Japan Journal of Medical Informatics*, Vol. 26, No. 1, pp.1-12.
- Shinagawa, Y., Kishimoto, T., Ohta, S., 2006. Development of an algorithm for automatic emergency calls using non-response intervals of infrared sensors. In *Kawasaki Medical Welfare Journal of Japan*, Vol. 15, No. 2, pp: 553-563.
- Yoneda, T., Ogawa, A., Sasaki, J., Funyu, Y., 2006. Development and operation of mimamori network system for elders in Kawai Village, Iwate Prefecture. In *Personal computer users' application technology association*, Vol. 16, No. 3, pp:31-38
- Zojirushi Corporation, Mimamori Hotline, 2005. In *online*, <http://www.mimamori.net>.

PERFORMANCE ANALYSIS OF WIRELESS SYSTEMS IN TELEMEDICINE

Hybrid Network for Telemedicine with Satellite and Terrestrial Wireless Links

M. Luglio and F. Zampognaro

*University of Rome Tor Vergata, Via del Politecnico 1, 00156 Rome, Italy
luglio@uniroma2.it, zampognaro@ing.uniroma2.it*

Keywords: Telemedicine, hybrid networks, DVB-RCS, wireless communications.

Abstract: Telemedicine services represent a valuable opportunity to provide medical assistance ensuring high flexibility and prompt set up and to significantly reduce costs. The use of hybrid networks based on satellites and terrestrial wireless systems can be extremely advantageous in terms of flexibility, capillarity and integration with modern medical equipment, in particular representing a suitable solution in case of disasters. In the paper such an architecture is described and key performance for some reference applications, evaluated through simulation, are shown and discussed.

1 INTRODUCTION

Nowadays telemedicine is gaining increasing interest, in particular in remote and disaster-struck areas where telecommunication infrastructure can be respectively missing or compromised.

The aim of this paper is to show feasibility and effectiveness of an hybrid network architecture for telemedicine, through performance evaluation carried out by simulations. The proposed hybrid network, selected also for the Telesal project (Arenaccio, Aversa and Luglio, 2006), is composed of a satellite core network, interconnected with terrestrial wireless tails using commonly available and consolidated wireless technologies.

Satellite systems are extremely suitable to represent the core infrastructure of the network for their capability to provide data access ubiquitously and in mobility over very large areas, including remote or impervious locations where typically terrestrial telecommunication infrastructures are not present or economically not viable. To complement such characteristics, terrestrial wireless systems can be fruitfully utilized realizing the terrestrial tails to ensure capillarity and to improve efficiency and flexibility (Luglio and Vatalaro, 2002).

Performances of some reference telemedicine applications will be evaluated using NS2 as simulation tool (Fall and Varadhan, 2007).

In section 2 an overview of the involved technologies is introduced, in section 3 the network architecture is described along with some telemedicine applications. In section 4 the simulations setup is presented offering some the simulation outputs summaries. Finally in section 5 conclusions are drawn.

2 TECHNOLOGY REVIEW

Different kinds of satellite configuration (geostationary and low orbit) can be utilized in the proposed hybrid architecture. As concerns terrestrial systems in particular WiMax, WiFi and Bluetooth looks the most suitable technologies for our scope.

2.1 Satellite Networks

Two kinds of satellite systems are considered:

- a) wide band VSAT systems using a geostationary (GEO) satellite and
- b) narrow band global communication system using a low orbit (LEO) satellite constellation.

VSAT systems (Very Small Antenna aperture Terminal) are characterized by the use of directional fixed or steerable (to allow mobility) dish antennas with a size of around 80-120 cm. They usually adopt star or mesh topology, using GEO satellites, which

are suitable for multi-nation coverage area (usually at continental level). They can offer uplinks of up to 2 Mbit/s, with downlink up to 40 Mbit/s. Communications suffer of a physical delay of about 550 ms round trip, since the GEO satellite is placed in a 36000 km orbit. Star systems need a double hop to allow two terminals to directly communicate, from the terminal to the star-centre (called HUB) and from it to the other terminal, resulting in physical delay of four times 125-130 ms. Mesh systems, instead, allow direct communication between two terminals without crossing the HUB (physical delay of two times 125-130 ms).

For the specific scenario that have been simulated, DVB-RCS standard (ETSI EN 301 790, 2003), developed for VSAT systems, has been selected. The architecture referred to the standard usually applies to star topology, with a central HUB called NCC (Network Control Centre). Downlink channel is broadcasted to all users using DVB-S standards (ETSI EN 300 421, 1997), while return channel is shared with a MF-TDMA technique. DVB-RCS allows each terminal to negotiate capacity requests on demand for transmission on the shared return link according to pre-defined service level agreement:

- volume based dynamic capacity (VBDC), to issue bandwidth requests based on the actual volume of traffic needed;
- rate based dynamic capacity (RBDC), to issue bandwidth request based on the estimation of transmission rate;
- constant rate assignment (CRA), to obtain guaranteed bandwidth assignment.

Such an assignment scheme is called DAMA (Demand Assignment Multiple Access) and it is used to share the same upload channel dynamically and efficiently among several terminals. According to the request policy, different cost may be charged by the satellite operator.

On the other hand LEO constellations are composed of several satellites at low orbit (between 700 and 1500 km), which are in continuous movement with respect to terrestrial Earth surface. The system is designed to maximize the probability of user-satellite line of sight even at high latitudes and handover functionalities must be implemented in order to keep connection when changing serving satellite. LEO terminals use omnidirectional antennas and offer limited bit rate, usually dimensioned for voice communications (similar to GSM). Latency is limited to a few ms, but variable in time and affected to big spikes due to the handover execution.

Globalstar has been selected for the simulation campaign (<http://www.globalstar.com/en/>), due to its common availability in Europe.

2.2 Terrestrial Wireless Networks

To realize the terrestrial component PANs (Personal area networks), LANs (local area networks) or WANs (wide area networks) concepts can be adopted. The first two are usually associated to license free bands (IMS), with data throughput ranging from 1 up to tens of Mbit/s (with a coverage from few meters to some tens of meters). In particular Bluetooth (IEEE Std 802.15.1) and Wi-Fi (IEEE Std. 802.11) are representative technologies of PAN and LAN, respectively. A WAN is instead capable of long range coverage with higher throughput and it usually works either in licensed or free bands. WiMAX (IEEE Std. 802.16e-2005) is an example of WAN with allocation of commercial frequency bands around 3.5 GHz. HIPERLAN (ETSI EN 300 652, 1998) represents another example of WAN system working in the unlicensed band of 5.4 GHz.

For the purpose of our hybrid network proposal, only the license-free LAN and PAN technologies will be included, leaving to a future study a more comprehensive integration of WAN, LAN and PAN networks together with the satellite segment.

Wi-Fi is a widespread wireless technology that provides wired-LAN-like connection service to mobile devices in the range of around 100 m. Maximal bandwidth available on Wi-Fi variants ranges from 11 Mbit/s (standard 802.11b) to 54 Mbit/s (standard 802.11a or 802.11g). So far the infrastructure mode, with a central base station (called Access Point), has been widely deployed in most cases, although Wi-Fi foresees an ad-hoc direct connectivity. A set of base stations can serve up to 128 user terminal each, guaranteeing local mobility. Newer standard 802.11i and 802.11e are defining respectively stronger algorithms for security (WPA2) and QoS at MAC layer.

Bluetooth is a PAN ad-hoc wireless system which allows terminals to flexibly and autonomously configure themselves to communicate without a pre-existing infrastructure in a peer-to-peer fashion. Bluetooth Standard version 1.1 is the actual reference implemented in commercial products such as headsets, GPS devices, etc. It is designed to offer a total 1 Mbit/s data rate with a coverage of 10 meters maximum. When Bluetooth terminals get close enough, they can cluster into a piconet and temporarily designate one master unit to

coordinate transmissions with up to seven slave units. The time needed to join a piconet and start service is in the order of some seconds.

Bluetooth is based on packet transmission and frequency hopping (FH) technologies to provide channelization among different piconets within the same area, to form the so called scatternets. Each Bluetooth service has a pre-defined QoS profile to announce during setup, and it is accepted in the piconet only if there are enough resources.

3 TELEMEDICINE NETWORK

The set up of a telecommunications network as support to telemedicine can be extremely difficult in remote or in disaster-struck areas. For instance, the installation of a single dedicated point to point radio link to restore or deliver GSM communication channels can take several hours. In this context, the use of satellite terminals can shorten this time to a few minutes, thanks to the intrinsic broad coverage of a satellite service. The core satellite network can be complemented by terrestrial wireless tails composed of Bluetooth piconets and Wi-Fi links.

3.1 Architecture

The proposed architecture is shown in Figure 1. Connection between satellite terminal and HUB is assumed to be alternatively realized with either DVB-RCS or Globalstar.

An application client on the disaster-struck area is assumed to be reached directly by the satellite terminal, or being part of a Bluetooth piconet. In both cases the connection with the Satellite Terminal can be wired or realized via a Wi-Fi wireless link (dashed line). Nodes of different segments are connected with Ethernet cables or with internal bus if integrated in the same hardware. In all cases connectivity is implemented at IP layer to leverage on IP built-in routing functionalities and address resolution.

Security and QoS must also be carefully considered in hybrid networks offering telemedicine services. QoS is a key issue for real time services, and must be offered end-to-end along the whole data path. This means that each segment must be coordinated centrally for its specific QoS management setup. Security and encryption, usually available for each technology independently, must be ensured also end-to-end, due to the sensitivity of data transmitted. End-to-end QoS and security could be handled at IP layer, since it can be considered too

complex an adaptation of different QoS and security procedures offered by the different technologies at layer 2. Solutions like DiffServ and secure tunneling (VPNs) could be adopted.

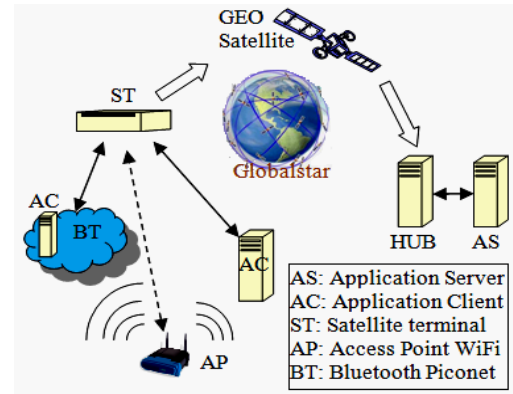


Figure 1: Network Architecture.

3.2 Reference Applications

Two kinds of applications, both real time and non real time, will be tested over the proposed network. Table 1 shows a list of these applications with the most significant characteristics. The last two rows show dimensions of representative files which can be transferred by non real time telemedicine applications.

Table 1: Telemedicine applications.

Real time			
Application	Protocol	Codec	Bitrate
Voice call	RTP	G729	8÷12 kbit/s
Video call	RTP	MPEG4	>384 kbit/s
Non Real time			
Application Data	Protocol	Size Raw	Size compressed
Radiography	FTP	5.7 Mbytes	380 kbytes
ECG trace	FTP	90 kbytes	45 kbytes

4 SIMULATIONS

Simulations have been performed using NS2 platform. The architecture introduced in section 3.1 has been set up and verified with the help of NAM (NS2's visual output), as shown in Figure 2. Application clients 4-6 have Bluetooth connectivity to the Home Gateway. The Home Gateway and the Application clients 7-8 have a Wi-Fi and wired Ethernet connectivity to the satellite terminal, respectively.




Figure 2: NAM output of NS2 simulation.

Traffic sources originated from Application Clients are compatible with the applications listed in Table 1. All the traffic is delivered from the application clients to the satellite terminal via a wired or wireless link, through the satellite (in alternative DVB-RCS or Globalstar) up to the application server at the other end of the network. The application server is representing the operative centre for emergency handling. TCP New Reno has been used as transport protocol for file transfer and standard UDP for real time traffic.

For simulations using DVB-RCS as satellite technology, a return link capacity of 512 kbit/s with the correct physical delay has been considered. For DAMA capacity requests, RBDC has been simulated according to (Roseti and Kristiansen, 2006) while CRA consists in a granted capacity, similar to a SCPC service (Single Channel Per Carrier, unshared uplink channel). For all the other technologies involved (Bluetooth, WiFi, Globalstar), common operational values as seen in section 2 have been used in NS2 simulated links.

Non real time application performances are summarized in Table 2. For this class of applications the performance index considered is the average time needed by the application on a Bluetooth node to send one data file (including reception and acknowledgement) to the application server at the other end of the network.

Table 2: Non Real Time Average Data Transfer Time.

Data	Globalstar	DVB-RCS w/ RBDC requests	DVB-RCS w/ CRA requests
Radiography RAW	Not performed	140s	115.6s
Radiography Compressed	333.5s	16s	12.5s
ECG trace RAW	79.6s	6.9s	4.9s
ECG trace Compressed	45.9s	4.8s	3.4s

Figures clearly show the differences in performances between GEO and LEO satellite system, and also between the two different request policies for DVB-RCS. The use of satellite to transfer medical data of limited size is acceptable under all conditions, while bigger size data transfer is not practical for narrow band Globalstar satellite system.

For real time applications, two different setup have been considered for the two alternative GEO and LEO satellite systems:

- When using DVB-RCS, a video call with five simultaneous voice calls with higher quality profile (12 kbit/s each) have been initiated from the Bluetooth nodes 4-6. Other two voice calls have been initiated from Application clients 7 and 8
- Over the Globalstar system, only one voice call originated by a Bluetooth node could be performed at the minimum codec rate (8 kbit/s)

In both cases the packet error rate has been verified to remain under 1%.

The one way average delay of RTP packets delivery from source to destination, measured at the two ends of the network, has been used as performance index for the real time applications. The averaged delay values are shown in Table 3.

Table 3: Real Time average packets delay (one way).

Data	Globalstar	DVB-RCS w/ RBDC requests	DVB-RCS w/ CRA requests
Voice Call	0.27s	0.48s	0.36s
Video Call	n.a.	0.61s	0.35s

Globalstar and DVB-RCS with CRA profiles show similar performances for the voice call average delay. In fact, in both cases there is no need for explicit bandwidth requests which introduce an additional access delay. RBDC request policy has a bigger average delay compared to the other two cases, because periodic requests based on estimated rate must be issued by the terminal to the NCC, thus increasing the perceived delay by an additional factor. RBDC is usually associated to a cheaper contract with satellite operator in comparison with CRA.

Jitter has resulted limited when simulation adopted the Globalstar network and the DVB-RCS with CRA profile. In particular some significant variations are present during LEO satellite handover, which was not simulated in our NS set up and usually occurs every 20 minutes in average.

Instead, when DVB-RCS system adopts RBDC access scheme, significant jitter variations are observed. Please note that RBDC allocation mechanism is not standardized by DVB and the reported effect might vary depending on the selected implementation.

In particular the jitter is due to voice calls packets pattern, a small packet each 10 ms, according to the codec standards. Small packets can trigger the request of a bigger capacity of the simulated DVB-RCS system which remains assigned to the terminal for a longer time (in this simulation 100 ms), resulting in temporary extra capacity. As consequence, the time needed for the delivery of packets decreases, resulting in a lower perceived rate needed. This vicious loop makes RBDC requests oscillating, together with system capacity assigned. This affects packet delivery delay too, which is shown in Figure 3 for a voice call.




Figure 3: Packet delay oscillations on RBDC requests.

A proper de-jitter buffer (Ferrari and Verma, 1991) must be designed at receiver side in order to compensate the packet arrival latency variations for this case.

5 CONCLUSIONS

The paper describes a hybrid network for telemedicine applications adopting satellite networks and wireless technologies of different kinds. Both sample real time and non real time applications has been run in a simulated scenario including all the described network links, to assess performances.

The results obtained has proven the feasibility of such an hybrid network including satellite links. In particular positive results has been obtained with both GEO and LEO systems, taking into account limitations of LEO narrow band capacity. The main differences between different systems and the use of

different request policies for DVB-RCS has been discussed.

REFERENCES

- S. Arenaccio, F. Aversa and M. Luglio, 2006. A satellite based wireless network for telemedicine: the Telesal Project. In *12th Ka and Broadband Communications Conference*, Naples, Italy, pp.749-754.
- M. Luglio, F. Vatalaro, 2002. The Use of Wireless Technology for Telemedicine. In *Symposium on Telemedicine In Care Delivery, Technology and Application*, pp.353-359
- K. Fall and K. Varadhan, 2007. *The NS Manual*, <http://www.isi.edu/nsnam/ns/doc/index.html>
- ETSI EN 301 790, 2003, Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution System, V1.3.1.
- ETSI EN 300 421, 1997, Digital Video Broadcasting (DVB); Framing structure, channel coding and modulation for 11/12 GHz satellite service, V1.1.2.
- ETSI EN 300 652, 1998, Broadband Radio Access Networks (BRAN) - High PErformance Radio Local Area Network (HIPERLAN) Type 1, v1.2.1
- IEEE Std 802.15.1, 2002, Wireless MAC and PHY Specifications for Wireless Personal Area Networks (WPANs™)
- IEEE Std 802.11, 1999, Wireless LAN Medium access control (MAC) and Physical Layer (PHY) Specifications.
- IEEE Std 802.16e-2005, Air Interface for Fixed Broadband Wireless Access Systems- Physical and MAC Layers for Combined Fixed and Mobile Operation in Licensed Bands.
- C. Roseti and E. Kristiansen, 2006, TCP behaviour in a DVB-RCS environment. In *Proceedings of 24th AIAA International Communications Satellite Systems Conference (ICSSC)*, San Diego.
- D. Ferrari and D. Verma, 1991. Buffer Space Allocation for Real-Time Channels in a Packet Switching Network. *Technical report*, Information Computer Science Institute, Berkeley, California.

DEALING WITH THE COMPLEXITIES WHEN IMPLEMENTING INFORMATION SECURITY PRACTICES IN HEALTHCARE ORGANIZATIONS

Heitor Gottberg and Ivan Torres Pisa
Federal University of São Paulo (UNIFESP), Brazil
heitor@cisco.com

Keywords: Information Security Guidelines, Information Storage and Retrieval, Computer Security, Information Management.

Abstract: With the increasing use of electronic healthcare records and other medical systems, private and confidential information are electronically stored on different databases in several computers. A new set of processes and controls are necessary to assure the information system security and personal privacy. One of the approaches to meet these demands is to establish information security practices based on international standards. Due to the complexity of healthcare operations, managers must be aware that there are additional complexities on implementing those practices. This article depicts these additional efforts, highlighting four of the extra controls that shall be implemented: disposal of media, clock synchronization, backup, and network services – as well as threats as repudiation, theft, and terrorism that must be taken into consideration by healthcare CIOs in order to become compliant to the information security standards and, therefore, fostering the use of IT on medical practice.

1 INTRODUCTION

Currently, information security issues are increasingly calling the attention of system administrators, IT managers, software and IT services vendors. Organizations and public infrastructures depend on information systems and services that store and retrieve information in a digital way and might be vulnerable to information security threats (Cavalli, 2004).

There is also an increasing pressure for organizations to fulfil the requirements determined by the regulatory agents and institutions (e.g. Health Information Portability and Accountability Act – HIPAA) in order to assure trust and recognitions from customers and business partners. The corporate security policies must then be compliant on renowned standards, guidelines and best practices, e.g. ISO17799 (ISO, 2005), so that the organization certification of the information security controls can be accepted by other entities. Regulations are designed to prevent disastrous incidents and reduce security threats. By the nature of information security – that must prevent information leaks to happen – the corporation is impelled to act and

implement a proactive information security enforcement strategy (Yip, 2006).

The wider use of clinical and administrative health information systems imposes a careful management of all aspects of security within the IT environment, including assets as hardware, software, data entry and storage, networking, and also processes and controls of information security. The essential objectives are to guarantee: (1) data integrity: i.e. the data must be protected from accidental or intentional alterations or losses; (2) data availability whenever needed by authorized users; and (3) data confidentiality: i.e. the data must only be accessible to authorized users - people or programs (Ravera, 2004). Due the reasons presented above, companies start focusing on implementing information security controls within an Information Security Management System (ISMS).

This article – based on the existing works of the International Standards Organization (ISO) – depicts part of the extra effort that is necessary when implementing information security controls and practices in healthcare organizations due the special business characteristics and demands of provisioning healthcare services.

2 INTERNATIONAL INFORMATION SECURITY STANDARDS

From the extent content of material produced by the ISO, including published standards, standard's drafts, and technical committees' debates; two documents will support this article to show some of the existing content around information security controls. These documents are the ISO 17799 and the ISO/DIS 27799 which are further explained below.

ISO/IEC 17799:2005 – ISO/IEC 17799 version 2005 is an internationally recognized standard that establishes guidelines and general principles for implementing information security controls as part of a corporation's ISMS. Within the document, the objectives are outlined to provide a general guidance and best practices for information security management that are commonly accepted. Control objectives and implementation guidelines are stated to drive the development of the organizations processes. Organizations seeking for ISO 17799's certification will typically develop and modify their organization security policy and processes to match the controls specified in the ISO document. In this task, those organizations will make use of assessment toolkits consisting of the standards itself and templates that ease the work of designing a control in according to the standard. While internationally, over 80,000 firms are said to be ISO 17799 compliant, apparently the companies use the ISO 17799 as a guideline, selecting some of the controls applicable to their environment. They usually do not seek the certification of the entire standard but to portions of the standard relevant to their operations (ISO, 2005). In this direction, healthcare organizations will also make a selection of the controls to be applied.

To support this adaptation task of healthcare organizations – to ISO 17799 and other themes – and provide additional and internationally accepted standards for healthcare informatics, the International Standards Organization established the Technical Committee “TC 215 - Health Informatics” with the main aim of contributing to the improvement and maintenance of health by producing those ISO standards, which the international community regards as necessary to enable the successful utilization of health Information and Communication Technology (ISO/TC215, 2000).

ISO/DIS 27799 - ISO/DIS 27799(ISO/TC215, 2006) is a standard under development by the ISO/TC 215, currently available in a “Draft of International Standard” status. It explores the specificities of implementing security management of the health sector and its unique operating environments.

According to the ISO 27799 all of the security control objectives described in ISO/IEC 17799 are relevant to health informatics but some controls require additional explanation in regard to how they can best be used to protect the confidentiality, integrity, and availability of health information. There are also additional health sector specific threats that must be analyzed when designing the institution's ISMS. The goal of this international standard is to provide additional guidance and controls, adjusting the format and terminologies. Making use of this standard, people responsible for health information security can understand and adopt in a faster (and possibly cheaper) way. Besides giving extra details to the ISO 17799 controls, ISO 27799 also states some threats to information security specifically to healthcare environments.

It is important to be aware that this standard is under development but still it can be used as reference to make tangible the extra complexity of implementing information security controls in health services provision. To do that we will choose some controls and some of the threats stated in it.

3 INFORMATION SECURITY CONTROLS

Information security controls are part of an ISMS, which implementation phases involves several steps as: ISMS design, risk identification and assessment, control selection, documenting security measurements, implementation, evaluation of measurements versus documentation, auditing, and improvement (Posthumus, 2004). These phases can be related to a Plan, Do, Check, and Act (PDCA) model, which guide the establishment and constant development of the ISMS (ISO/TC215, 2006). Figure 1 relates these steps to the PDCA cycle.

Based on Figure 1, we see that each organization must define the controls that will be part of its ISMS.

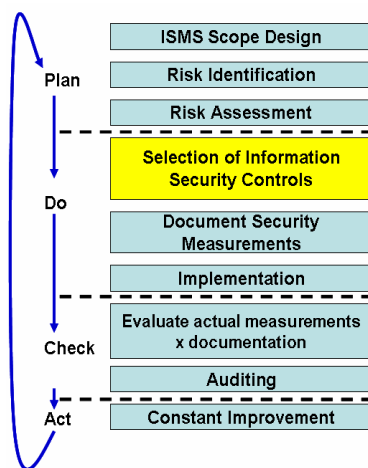


Figure 1: ISMS implementation phases.

The ISO 17799 itself states that the investments to put controls in place must be balanced to the equivalent risk that it prevents and, at the Controls Selection chapter, reinforces the idea that not all the controls of the standard must be actually implemented. In healthcare environment, this selection is equally critical and, according to ISO 27799, enhanced by the need of additional controls.

Below, four of these controls and 3 additional threats are explained, which have been selected by our comprehension that they can best exemplify the additional complexities when implementing information security practices in healthcare organizations:

Disposal of media: When disposing an information storage media or medical devices that record and report data, all personal health information must be securely overwritten or deleted, evaluating also the possibility of destroying the media. Not to pay the correct attention to this disposal continues to be a source of serious breaches of patient confidentiality. It is important to note that this control should be applied not only on disposal but also prior to the repair of any associated equipment (ISO/TC215, 2006).

An example of how critical this issue can be in healthcare is a report that focused on discarded hard drives and disk sanitization practices. It disclosed that in August 2002, the United States Veterans Administration Medical Center in Indianapolis sold or donated 139 of its old computers without removing confidential information contained on their hard drives, including the names of veterans who had AIDS and mental illnesses (Hoffman, 2006).

Clock synchronization: In healthcare operations, being able of tracing the exact timeline of

happenings and activities is a must for health information systems supporting time critical care activities. The institution's IT infrastructure shall provide time synchronization services to support tracing and reconstitution of activity timelines. This information must be certified in a way that can be accepted by outside and legal entities, once the timing of events as electronically recorded in personal health information and in audit records can be relevant in processes such as coroners' inquests, investigations into medical malpractice, and other judicial proceedings (ISO/TC215, 2006).

Health information backup: In healthcare organizations, the backup of information serves not only to free storage space but attend also the regulations of future availability. If we merge this demand to the necessity of confidentiality of the stored information, we are driven to the need of a control that assure that back up of all personal health information is stored it in a physically secure environment and backed up in an encrypted format (ISO/TC215, 2006).

Security of network services: In addition the guidelines of ISO/IEC 17799, institutions dealing with personal health information should be aware of the higher impact of loosing the network service availability, e.g. impacts upon clinical practice (ISO/TC215, 2006).

These additional controls remind us that once we migrate the patient record to a digital format in an information system, the clinical practice will have this system as the fundamental tool to perform the day to day operational work. Unlike other segments, in healthcare, breaks in processes workflow due system or network crashes, are usually very near of costing lives.

4 INFORMATION SECURITY THREATS

Figure 1 above shows that to plan the ISMS it is necessary to identify and assess the risks involved at the company operations, related to digital information. Those risks relate to the threats existing on the corporate environment and processes setup.

In healthcare, there are also some additional threats that must be considered when planning an ISMS as:

Repudiation: This threat refers to users denying that they have sent or received a message. Nowadays, informatics provides us with tools to assure the preventions of this threat, such as digital signatures on prescriptions and receive/read receipts messages to

emails. This assurance can become critical on flows of personal health information from one provider to the other and/or on investigations of medical practice (ISO/TC215, 2006).

Theft: Theft of data and equipment is a serious problem in hospitals. Theft may cause breaches of confidentiality, either because confidential data resides on a server or laptop computer that is subsequently stolen or, else because the data itself is the target of the theft. The threat of theft personal health information increases with the fame or notoriety of the data subject (e.g. a celebrity or head of state) and decreases with the potential severity of punitive consequences - e.g. the loss by a physician of his or her license to practice (ISO/TC215, 2006).

Terrorism: Even having no notice of wide terrorist acts to healthcare institutions, once the healthcare infrastructure is usually part of the national or regional community sustainability infrastructure, once large scale health information systems are planned, the terrorist threat must be assessed due the possible effects on increased effectiveness of bioterrorist and other attacks that cause a health-related crisis (ISO/TC215, 2006).

5 CONCLUSIONS

From the exposed above, we explain that the establishment of an information security management system, compliant to international standards, gains complexity and scope extent when we are in a healthcare organization. This statement is reinforced by the existence of a technical committee within ISO to study the specificities of the use of informatics by healthcare service providers – the ISO TC 215 – and within this group, a subgroup focusing the information security needs.

On the other hand, this additional complexity is a price to pay for the benefit of converting the patient information to an electronic form and so have the possibility of storing, retrieving, and distributing this information in an easier, faster, and cheaper way.

We must remember that provisioning healthcare services itself is one of the most complex duties in terms of managing the needs and legal regulations of integrity, confidentiality, and availability of patient information. Therefore, the additional tasks that come with the informatization of these data are a natural consequence of its nature.

This work will now continue in two steps. First we will try to identify the existing and used tools to implement the information security controls (e.g.

standards toolkits and risks assessment tools). Second we will move on trying to map and score how compliant are healthcare organizations in our region to the international standards that are suggested by the Brazilian National Council of Medicine to allow the migration of hospitals to full electronic healthcare records.

ACKNOWLEDGEMENTS

We thank Mrs. Beatriz Leão (SBIS - Brazilian Association for Health Informatics) for supporting the development of this paper.

We thank also CISCO do Brasil Ltda that is supporting the participation of the main author at HealthINF 2008.

REFERENCES

- Cavalli, E, et al, 2004; *Information security concepts and practices: the case of a provincial multi-specialty hospital*; International Journal of Medical Informatics (2004) 73, 297-303.
- ISO, 2005; *ISO/IEC 17799:2005 Information technology - Security techniques - Code of practice for information security management*; International Standard Organization.
- Yip, F, et al, 2006; *Enforcing business rules and information security policies through compliance audits*; IEEE 1-4244-0176 - 03/2006.
- Ravera, L, et al, 2004; *Security and privacy at the private multispecialty hospital instituto clinico humanitas: strategy and reality*; International Journal of Medical Informatics (2004) 73, 321—324.
- ISO TC 215, 2000; *ISO/TC 215 business plan template*; International Standard Organization – Technical Committee 215 – Health Informatics (extracted from: http://isotc.iso.org/livelink/livelink/fetch/2000/2122/687806/ISO_TC_215_Health_informatics.pdf?nodeid=1001750&vernum=0; in Feb, 24th, 2007).
- ISO TC 215, 2006; *ISO/Draft of international standard 27799: health informatics — security management in health using ISO/IEC 17799, ICS 35.240.80*; 2006; International Standard Organization – Technical Committee 215 – Health Informatics.
- Posthumus, L., 2004; *Use of the ISO/IEC 17799 framework in healthcare information security management*; Stud Health Technol Inform. 2004, 103:447-52 (PMID: 15747954).
- Hoffman, S.; Podgurski, A., 2006, *In sickness, health, and cyberspace: protecting the security of electronic private health information*; Social Science Research Network Electronic Paper Collection (<http://ssrn.com/abstract=931069>).

HEALTH CARE PROCESS MODELLING AND IMPROVEMENT

Nadja Damij

*Faculty of Economics, University of Ljubljana, Kardeljeva ploscad 17, 1000 Ljubljana, Slovenia
nadja.damij@gmail.com*

Janez Grad

*Faculty of Administration, University of Ljubljana, Gosarjeva 5, 1000 Ljubljana, Slovenia
janez.grad@fu.uni-lj.si*

Keywords: Process Modelling, activity table technique, process improvement.

Abstract: The paper discusses the problem of process modelling and aims to introduce a new technique called the activity table to find a better solution for the problem mentioned. The activity table is a technique for process modelling and improvement. Business process modelling is done by identifying the business processes and is continued by choosing a process, defining its work processes and activities. Process improvement is achieved by simulation of the activity table, suggesting changes and improvements, and giving solutions for existing problems. To do this, we concentrate our work on understanding and analyzing the activity table. A complete understanding of the activity table is an essential precondition to moving forward with the simulation, which enables us to make improvements of the process modelled. The problem of conducting a surgery is used as an example to test the technique.

1 INTRODUCTION

Business process modelling is a complex and difficult problem. A process model, which truly represents the business process discussed, is essential for carrying out business process improvement and information system development successfully. Business process improvement became a very important way of ensuring changes in an organisation's structure and functioning in order to create better and more efficient processes which consequently lead to a competitive and successful organization.

There are many methods and techniques which cover the field of business process modelling. The aim of this work was to introduce a new technique called the activity table to develop a process model which truly represents the original business process.

In Section 2, the problem of business process modelling is discussed, different techniques and approaches which deal with this field are mentioned, and the studies of a number of researchers are addressed. In this section, the activity table technique is introduced. The activity table enables us to develop a process model by linking each of the process activities to its resource (performer). In Section 3, business process improvement is stressed. This is

done by analysis and simulation of the process model developed. The last section contains some useful remarks and conclusions. The process is applied to a specific problem to illustrate the implementation of the technique.

2 BUSINESS PROCESS MODELLING

The recent literature offers various definitions of and the extent of a process or process modelling. Throughout the last decades, the fields of business process modelling and consequently business process renovation have been gaining recognition and acceptance. The reasons for such evolution are found in the literature, academic publications and research studies that deal with the theme, as well as in the increasing involvement of consultancy and software development companies. A comparative study that closely examined 25 methodologies, 72 techniques and 102 tools was conducted (Kettinger et al., 1997). Furthermore, business process modelling is one of the requirements of the ISO 9000 international standard for quality management and assur-

ance (Ould, 1995). Both business process modelling and business process renovation are based on the fact that a business process is the key element in the analysis of the organisation.

A process is defined as a structured, measured sets of activities designed to produce a specified output for a particular customer or market (Davenport, 1993). Hence, a process converts inputs by summing their value through various activities into outputs. A business process is a collection of activities that takes one or more kinds of input and creates an output that is of value to the customer (Hammer, 1990). However, others stressed that a business process is related to the enterprise, as it defines the way in which the goals of the enterprise are achieved (Aguilar-Saven, 2003).

Successful business process modelling depends on the appropriate selection of available modelling methods, techniques or process flow analyses. There are many techniques or analyses used in this field, such as general process charts, process activity charts, flowcharts, dataflow diagrams, quality function deployment, the integrated definition of function modelling, coloured Petri-nets, object-oriented methods, seven management and planning tools and so forth.

In this paper, we introduced a technique which could be used to produce a process model that represents a true reflection of the reality of the process discussed.

2.1 Activity Table

Business process modelling is a complex and difficult task. We are looking for a technique, which could produce a process model that represents a true reflection of reality.

The activity table is a technique for process modelling, analysis and improvement. This is achieved by identifying the business processes and is continued by defining the work processes and activities of the process discussed. To do that, we have to conduct interviews with the management at different levels. The purpose of these interviews is to identify the organization's business processes, the work processes related to each business process, and the activities related to every work process identified.

The activity table uses the term "entity" to define a user, group of users or other system of importance in the organization's functioning. An entity is any source of information that is part of the system or is connected with the system by some interaction. Therefore, an entity may be internal or external. An internal entity is inside the system and takes part in

the system's operation. An external entity is not part of the system, but it has one or more interactions with the system (Damij, 2000).

A work process is the lowest-level group activity within the organisation (Watson, 1994). A work process is a collection of activities followed in a determined order in carrying out distinguishable work to produce a certain output.

The activity table is organised as follows: the first column represents business process, the second column shows work processes, the activities are listed in the rows of the third column, and the entities are introduced in the remaining columns of the table grouped by the departments to which they belong. Such organisation of the activity table enables us to create a clear and visible picture of every business process and its work processes, and also of each work process and its activities (see Table 1).

Each activity occupies one row of the table. A non-empty square(i,j) links the activity defined in row i with its source, this is an entity defined in column j . Developing the activity table is a result of interviews organised with the internal entities defined in the columns of the table. In the rows of the activity table we first register each activity identified during an interview and then link this activity with the entities in the columns, which cooperate in carrying it out. To make the activity table represent the real world, we link the activities horizontally and vertically. The purpose of defining horizontal and vertical connections is to define their similarity to the real world in which they occur.

Horizontal linkage means that each activity must be connected with those entities in the columns which are involved in it. To indicate this, symbols \square , \diamond , \rightarrow and \leftarrow are used. Symbol \square or \diamond in square(i,j) indicate that entity(j) is a resource of activity(i), where j ranges from 1 to the number of internal entities and i ranges from 1 to the number of activities. An arrow drawn from square(i,j) to square(i,k) indicates an input enters activity(i) from another activity performed by entity(j), where i ranges from 1 to the number of activities, j and k range from 1 to the number of entities, and $j \neq k$.

Vertical linkage is used to define the order in which the activities are performed. Vertical linkage is used only in connection with the internal entities. This is achieved by using arrows \uparrow or \downarrow to connect the activities.

An arrow \uparrow or \downarrow from square(i,j) to square(m,j) means that activity(i) is a predecessor to activity(m). Two activities, which are not indicated in the same column, may be connected by horizontal and vertical arrows. For example, to connect square(i,j) to

square(i,k), we use two arrows. A horizontal arrow to connect square(i,j) to square(i,k) and then a vertical arrow ↓ to continue from square(i,k) to square(m,k). This means that activity(i) is a predecessor to activity(m), which is performed by entity(k). Of course, these two horizontal and vertical arrows can be replaced by a diagonal arrow from square(i,j) to square(m,k).

The main difference between the introduced technique and others is that this technique requires linking each activity defined in the rows of the activity table with its resource (an internal entity) defined in the columns.

Linking the activities with their resources in a visual manner helps a great deal in identifying the activities and tracing their order, which leads to discover the process as it occurs in reality, and enables us to develop a model that is a true reflection of the original process.

Surgery: The management of a clinic wished to improve the “Surgery” process by making it more efficient and less time consuming.

The process Surgery leads the patient, who needs to have surgery, through a number of activities in different departments of the hospital such as Reception Office, Clinic, Laboratory, X-Ray, Anaesthesia and Surgery Block.

The process “Surgery” was modelled using the activity table technique, see Table 1. This table shows that process Surgery consists of 4 work processes, which contain 36 activities.

2.2 Property Table

As we develop the activity table we simultaneously develop another table, the property table, which is very important in describing activities in detail. So, for each activity inserted in the activity table, we open a new row in the property table, which shows detailed information about this activity.

The property table is organized as follows: the activities are represented in the rows of the table and the characteristics of the activities are defined in the columns. Description: this is used to write a short description of the activity defined in the current row of the table.

Resource: this is used to determine the entity, which performs of the activity defined in the current row of the table.

Time: this is used to denote that the activity discussed needs a determined time to be accomplished. Time may become a very useful parameter should we wish to use it to improve business processes.

Rule: this is used to define when performance of the activity requires that one or more rules must be fulfilled. Rule is a precise statement that defines a constraint, which must be satisfied in order for a certain activity to be executed.

Input/Output: this is used to indicate which inputs or outputs are connected with the activity described.

Cost: this is the sum of the costs of the resources needed to accomplish an activity. This parameter is used to calculate the cost of work and business processes and therefore is important in improving business processes.

Developing the activity and property tables is an iterative process. Some of the interviews have to be repeated to arrive at a precise understanding of the user’s work. If anything is misunderstood, then we have to organize new interviews with the responsible users until everything is clear.

Surgery: Because of space limitations, only ten activities defined in the activity table are described in detail in the property table, Table 2. The values shown in the column Time in Table 2 are approximate values obtained from the medical staff. Unfortunately, we could not get any information concerning the costs of the listed activities from the management of the hospital.

3 PROCESS IMPROVEMENT

The aim of process improvement is to improve the organisation’s processes in achieving greater customer satisfaction by developing, reforming and advancing their quality, effectiveness, availability, and in lowering the cost of business processes within the organisation. Also, process improvement helps in increasing the efficiency of the process, improving customer service, sharing data and information, effectively deploying information technology, and reducing duplicate processes.

According to Harrington et al. (1997) a great deal of effort is being focused on continuous improvement of subprocesses, activities, and tasks. If the management of the organisation stops the evolution of the process once process improvement has been completed, the organisation will lose the value gained. Consequently, continuous improvement tasks need to be performed and as a result, as stated in Harrington et al. (1997), this should result in a 10 – 15 % yearly ongoing improvement in the process.

Table 1: Activity table of the process "Surgery".

Business Process	Work Process	Department	Reception Office		Clinic			Lab	X-Ray	Anaesthesia	Surgery Block		
		Entity	1.	2.	3.	4.	5.	6.	7.	8.	9.	10.	11.
			Nurse	Doctor	Nurse-In	Nurse-Cl	Surgeon	Technician	Doctor	Doctor	Anaesthetist	Surgeon	Patient
		Activity											
Surgery	Registration	1. Register patient	Entity										
		2. Forward patient	Entity										
		3. Examine patient	Entity										
		4. Send blood	Entity										
		5. Test blood	Entity										
		6. Forward blood findings	Entity										
		7. Decide type of treatment	Entity										
		8. Issue a release report	Entity										
		9. Order hospitalization	Entity										
	Hospitalization	10. Accept hospitalization order	Entity										
		11. Prepare examination order	Entity										
		12. Make x-ray examination	Entity										
		13. Create anaesthetic report	Entity										
		14. Forward medical findings	Entity										
		15. Analyze findings	Entity										
		16. Decide on surgery	Entity										
		17. Explain surgery	Entity										
		18. Schedule surgery	Entity										
		19. Get information for anaesthesia	Entity										
		20. Sign documents	Entity										
	21. Wait for surgery	Entity											
	Carrying out Surgery	22. Prepare patient	Entity										
		23. Carry out anaesthesia	Entity										
		24. Carry out surgery	Entity										
		25. Wake up patient	Entity										
		26. Post-surgery recovery	Entity										

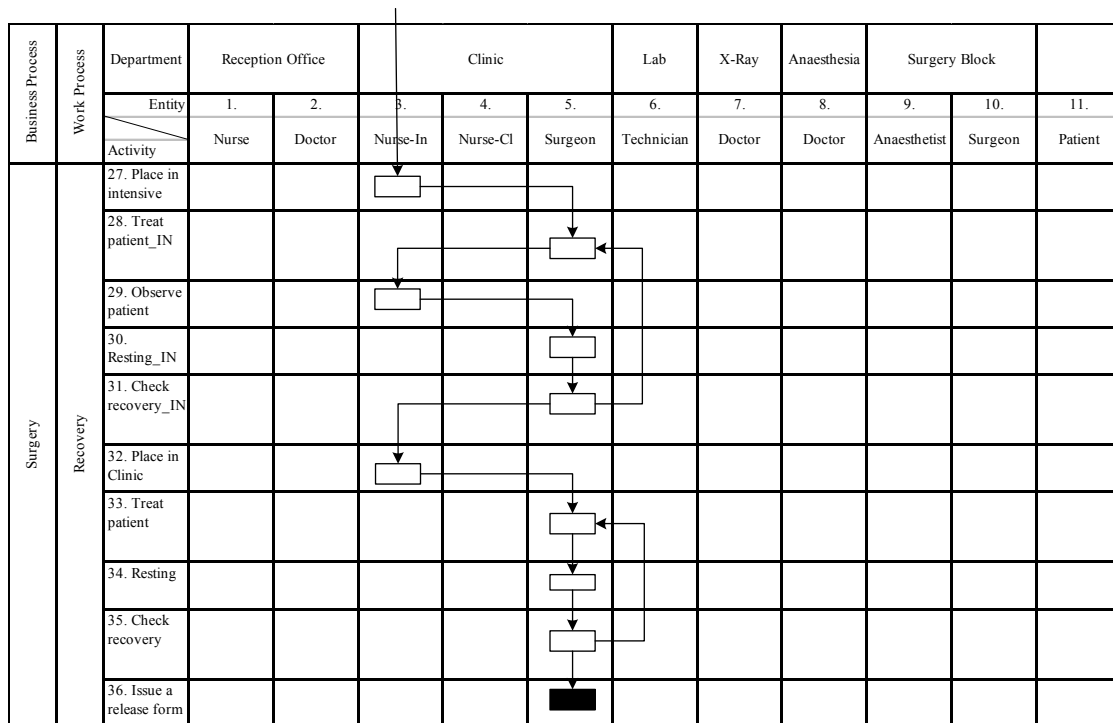


Table 2: Property table of the process "Surgery".

Characteristic Activity	Description	Resource	Time	Rule	Input/Output	Cost
1. Register patient	Nurse in Reception Office accepts patient's medical card, Doctor's order, registers her/him	Nurse	10 min	Check medical card validity	Doctor's order, Medical card	
2. Forward patient	Forward the patient and patient's documents to the doctor	Nurse	5 min		Medical card	
3. Examine patient	Doctor in Reception Office examines the Patient	Doctor	10-20 min	Check patient medical record	Medical record	
4. Send blood	Nurse in Reception Office takes patient's blood sample and send it to Laboratory	Nurse	10 min	Indicate needed blood examination order	Blood examination order	
5. Test blood	Technician in Laboratory tests blood example and sends results back to reception office	Technician	30 min	Check blood examination order	Blood exam. order, blood findings	
6. Forward blood findings	Nurse in Reception Office prints patient's blood findings and gives it to Doctor	Nurse	5 min		Blood findings	
7. Decide type of treatment	Doctor in Reception Office decides for a conservative treatment or for surgery after analyzing blood findings	Doctor	10 min	Check blood findings	Medical record, Blood findings	
8. Issue a release report	Doctor in Reception Office issues a release report and prescribes needed medications	Doctor	20-40 min	Prescribe medications	Medical report	
9. Order hospitalization	Doctor in Reception Office asks Nurse to prepare hospitalization order	Doctor	30 min		Hospitalization order	
10. Accept hospitalization order	Nurse in Clinic accepts hospitalization order from Nurse in Reception office to hospitalize the patient	Nurse	30-60 min	Check hospit orders & register the patient	Hospitalization order	

The relationship between the essence of process modelling and overall business effectiveness and the efficiency of the organization depends on the consumer's satisfaction with the desired output. If the latter is everything the consumer required and aimed for, business processes are well-designed, efficient, as well as effective and will in time result in successful organizations (Aguilar-Saven, 2003). On the other hand, if the consumer lacks appropriate satisfaction or the organization's growth and profit are decreasing, it is crucial to understand that improvement of business processes has to be planned and carefully carried out.

The goal of process improvement is achieved by simulation of the process presented in activity and property tables in order to suggest changes and improvements, and giving solutions for existing problems.

3.1 Simulation

Business processes are modelled with the aim of analyzing their current states within the organization, as well as improving them through the execution of potential "what-if" simulation scenarios.

Simulation modeling according to Pidd (1998) is based on very simple principles: the analyst builds a model of the system of interest, writes a computer program which embodies the model and uses a computer to initiate the system's behaviour when subject to a variety of operating policies.

Simulation is the imitation of the operation of a real-world process or system over time (Banks et al., 2001). A simulation model enables the analyst to observe and study the system's behaviour as it advances through time.

Surgery: We ran the simulation of the process "Surgery" shown in Table 1, taking into consideration a Clinic for abdominal surgery with a capacity of 30 beds. We simulated the process "Surgery" with 20 patients, who were already in the Clinic in different phases of the process, and with 30 patients who were scheduled for different operations. In addition to this, we postulated that 3 patients, from the planned 30 patients were hospitalized every day.

4 CONCLUSIONS

The aim of this work was to study the possibility of developing an effective technique for carrying out business process modeling and improvement. The technique has to enable the analyst to develop a visi-

ble and comprehensible model, which represents a true reflection of the real business process. This fact is essential in making the task of process analysis and identification of the necessary changes possible, so as to carry out a successful business process improvement.

We are certain that including resources (entities) in the process model is a new and important additional modelling dimension, which makes the modelling process easier and more precise.

To continue with the improvement of the business process "Surgery", Table 1 was transformed into a diagram of iGrafx software to run a simulation of the process. The results of the process simulation are very encouraging and show that the process "Surgery" is well planned and does not have major problems. Nevertheless, the process could be improved by shortening the time of 2.94 days spent before surgery. Some of the medical examinations could be done before hospitalization and also the time of waiting for surgery could be shortened. We are aware that these suggestions cannot be generalized for all patients, but they are good points for the medical staff to rethink.

REFERENCES

- Aguilar-Saven, R.: Business Process Modeling. Review and Framework. *International Journal of Production Economics*, Vol. 90, No. 2, (2003) 129-149
- Damij, T.: An Object-Oriented Methodology for Information Systems Development and Business Process Re-engineering. *Journal of Object-Oriented Programming*, Vol. 13, No. 4, (2000) 23-34
- Davenport, T.H., Short, J.: The New Industrial Engineering: Information Technology and Business Process Redesign. *Sloan Management Review*, Vol. 31 No. 4 (1990) 11-27
- Hammer, M.: Reengineering work: Don't automate. Obliterate. *Harvard Business Review*, Vol. 68 No. 4 (1990) 104-112
- Kettinger, W.J., Teng, J.T.C., Guha, S.: Business Process Change: A Study of Methodologies, Techniques and Tools. *MIS Quarterly*, Vol. 21, (1997) 55-80
- Ould, M.A.: Business Processes: Modelling and Analysis for Re-engineering and Improvement. John Wiley & Sons, Chichester (1995)
- Watson H.G.: Business Systems Engineering. Managing Breakthrough Changes for Productivity and Profit. John Wiley & Sons, New York (1994)
- Harrington H.J., Esseling E. and van Nimwegen H.: Business Process Improvement Workbook. Documentation, Analysis, Design, and Management of Business Process Improvement. The McGraw-Hill Companies, Inc, (1997)

A FIRST APPROACH FOR A REGIONAL WIDE VEPR

Pedro Vieira-Marques

*Informatics Department, Faculty of Medicine, University of Porto
CINTESIS, Faculty of Medicine, University of Porto, Portugal
pmarques@med.up.pt*

Arthur Cunha, Luís Antunes

*LIACC, Faculty of Science, University of Porto, Portugal
abcunha@fc.up.pt, lfa@ncc.up.pt*

Ricardo Cruz-Correia, Altamiro Costa-Pereira

*Department of Biostatistics and Medical Informatics, Faculty of Medicine University of Porto
CINTESIS, Faculty of Medicine, University of Porto, Portugal
rcorreia@med.up.pt, altamiro@med.up.pt*

Keywords: Electronic Health Records; Agents and Cooperative Systems; Integration and Interoperability.

Abstract: Patients visit multiple health institutions and leave a trail of information scattered around hospitals, healthcare centres and laboratories. Information availability is of major importance in healthcare delivery. Most of the Electronic Patient Record systems are unarticulated and usually address only the specificities of a single medical specialty. Virtual Electronic Patient Records such as MAID (Multi Agent system for the Integration of Data) system provide for the necessary means for intra-institutions departmental information integration. In this paper is presented a mobile agent based extension to the agent based MAID system in order to enable inter-institution patient data integration. This system was designed as a MAID extension with additional patient data integration features. In order to accomplish this, modules for external data discovery and collection were developed using mobile agents. Data collection activities are triggered by scheduled clinical events. The system is intended to enhance an existing institutional system taking it beyond the institutional barrier providing health professionals with a more complete patient clinical history.

1 INTRODUCTION

Through the years Electronic Patient Record (EPR) systems have been developed in order to provide physicians with structured and helpful information. However most of these systems are unarticulated and usually address only the specificities of a single medical specialty. Integration of healthcare Information Systems (IS) is essential to support shared care in hospitals, to provide proper care to mobile individuals and to make regional healthcare systems more efficient.

Healthcare is recognised as one of the most important areas for applications and services integration.

However, to integrate clinical IS in a way that will improve communication and data use for

healthcare delivery, research and management, many different issues must be addressed (Berg 2001; Littlejohns, Wyatt et al. 2003). Consistently combining data from heterogeneous sources takes a great deal of effort because the individual feeder systems usually differ in several aspects, such as functionality, presentation, terminology, data representation and semantics (Lenz and Kuhn 2002).

Successful integration seeks ubiquity – data stored at one place should be available elsewhere; consistency and integrity – different data items stored within the system should not be contradictory and integrity constraints should be maintained; synchronization – in particular, concurrent data access should not result in inconsistent data; single system image – the user should be able to interact with the system via a single standard user interface; and finally transparency – the user should not be aware of heterogeneous system components.

Work partial granted by FCT (POCTI/0753/2004) through CINTESIS n° 753/2004 FCT, by KCrypt (POSC/EIA/60819/2004) and LIACC through the Programa de Financiamento Plurianual, FCT and Programa POSI.

Multiagent systems can successfully solve problems where the required knowledge is physically distributed in several places (for example, they can gather patient data from different medical institutions or discover distributed e-health services) (Moreno 2006). Previously, we have developed MAID (Multiagent System for Integration of Data) (Cruz-Correia, Vieira-Marques et al. 2005), an agent-based system that provides access to patient data that's scattered between different units of major Portuguese hospital. It has been running for two years, having integrated (searched, verified, collected and stored) over 2 millions clinical records.

In (Vieira-Marques, Robles et al. 2006) we proposed a model for a medical-information-gathering system that addresses issues such as inter-institutional patient health data integration, the retrieval of momentarily unavailable online remote data, and secure data access and transport. In this model, clinical events scheduling trigger agents to gather information for a particular patient clinical history, making that history as complete as possible when those events occur. Integration efforts are directed at clinical documents and not at the data themselves.

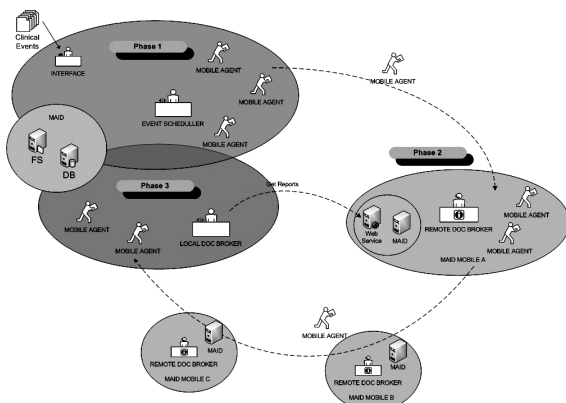


Figure 1: MAID Mobile model description.

In this paper, we describe a pilot implementation of the model which is focused in gathering information from multiple institutions, aiming at creating a first approach to a national wide Virtual Electronic Patient Record (VEPR).

2 METHODS

The system design was divided in three phases. To each phase corresponds one functional moment of the system (Fig. 1).

2.1 Events Management Phase

This is the first phase and corresponds to the management and scheduling of new events like consultations or surgeries. These events trigger actions for information search and retrieval. The Scheduler Agent (*SA*) manages the events and a set of information search and retrieval Mobile Agents (*MA*). To each patient a set of remote places where it is known to exist clinical information is stored. This list will grow as new places with information are discovered.

2.2 Information Search Phase

The second phase consists on a search performed by *MA* that will move between remote systems looking for information (exams, lab reports, etc) of a designated patient. If any documents are discovered references are collected and stored.

Besides clinical information, *MA* will ask for additional platforms where it is known for the patient data to exist. If other platforms are provided then they are added to the agent itinerary.

2.3 Document Retrieval Phase

At this stage, the process of searching for information references has ended and the collected references have been stored locally for retrieval. This retrieval will make use of basic remote transfer services using ftp or http urls or trough interfacing web services. After collection actions all the information is made available to the end user.

3 RESULTS

A set of agents was implemented using JADE platform and are described in the following sections along with their interactions while pursuing the designated objectives for each stage.

3.1 Scheduler Agent

This agent (*SA*) is in charge of scheduled events and *MA* management.

When some event is scheduled the *SA* constructs two itineraries, one with the remote systems that are known to have patient information and other with the complementary MAID systems retrieved from a directory services agent. These itineraries are delivered to two independent *MA*. Upon *MA*'s return this agent is notified that the process is finished and

the *MA* have returned with the collected information. This agent will also receive a message from the Local Broker Agent (*LBA*) informing about referenced documents retrieval success in order to eliminate the event from its list.

3.2 Mobile Agent

The *MA* actions are crucial to the purpose of this system as they are the ones in charge of discovering and carrying new information about a given patient.

These agents' actions are implemented using a Finite State Machine (*FSM*) which describes its roles and roles transitions during the different stages of moving, remote systems interaction, etc.

As soon as the agent receives the message from the *SA*, it saves the necessary data (itinerary and patient id) and moves to the first place. While in the remote platform, messages are exchanged with the remote broker agent (*RBA*). This message exchange has the purpose of obtaining the new local interface role, allowing the *MA* to retrieve the document references. During the initial exchange of messages a simple authentication is performed.

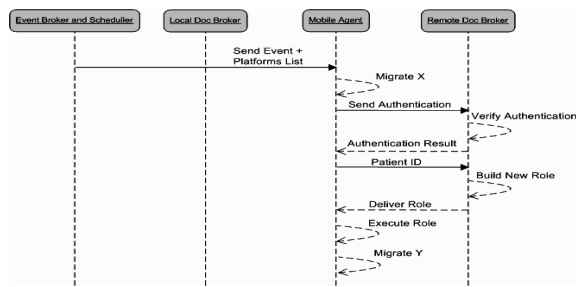


Figure 2: Agent interactions taking place during Information Search Phase.

As soon as the *MA* receives the role, it adds it to its *FSM* defining the necessary additional state transitions. From this moment on the agent is ready to get existing data through the execution of the new role. The retrieved data include a document reference list and a list of platforms that the visited system knows to have documents related to the patient the *MA* represents. By doing so, it will discover new sources of information. When the *MA* finishes its itinerary it will go back to the original platform. Upon arrival he notifies the *SA* of its return and delivers the gathered information to the *LBA*.

3.3 Remote Broker Agent

This agent stays in the local system and acts as an interface agent to incoming agents. It has to be implemented in accordance to each platform specificity, as all remote system at this point are

MAID systems they are implemented the same way in all nodes. It is in charge of authenticating and providing the interface behavior to the incoming agent.

3.4 Local Broker Agent

This agent is in charge of receiving the result of the references retrieval process. After receiving new references from the *MA*, it stores them into the database and starts the document retrieval process through the external services provided by each platform. After the documents being retrieved they are made available to the end user (Doctor, Nurse, etc) through existing user interfaces. The documents are saved in the file system digitally signed, guaranteeing the documents integrity.

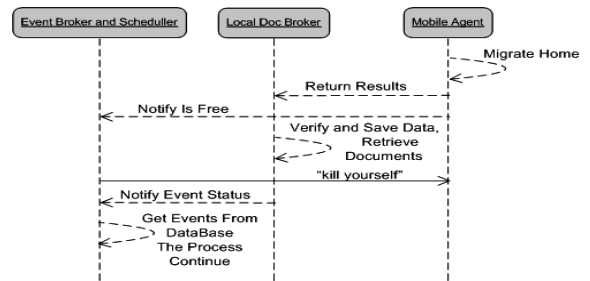


Figure 3: Agent interactions taking place during Document Retrieval Phase.

3.5 Agent Interactions

The UML activity diagrams describe the interactions between the different agents during information references search (Fig. 2) and document retrieval (Fig. 3).

4 DISCUSSION

This prototype extends the MAID systems by implementing a set of interfacing agents enabling the retrieval of remote documents. It addresses a scenario where multiple MAID systems coexist. We feel that multi-agent technologies can help implement integration between heterogeneous healthcare Information Systems in a satisfactory manner.

4.1 Integration Achieved

Regarding **ubiquity**, our implementation enables the exchange of information between each integrated system in a similar way. All information is

equidistant to each system. To face the problem of latency of communications, or even momentary lack of connectivity between different institutions when the healthcare professional needs the patient information, the system tries to collect useful patient data a priori, i.e., before the user request. Scheduled appointments trigger data discovery and report collection. Agents facilitate ubiquity as they allow MAID to operate in an asynchronous fashion, which is more powerful than other technology solutions like web services that rely on synchronous communication.

Regarding **synchronization**, our proposal takes two different approaches. When integrating patient data in an institution, MAID collects all updated reports from the departmental IS so that its local repository stays updated. In this scenario on the other hand, when integration occurs between different MAID systems our method is more like **lazy synchronization**, i.e., the information is updated only when a healthcare professional access to the system is scheduled. Although, this approach may raise problems in unscheduled accesses, we feel that full synchronization stresses network resources too much.

Although **single system image** is usually regarded as being a user interface issue, MAID delivers the same functionality of single image not to human users but to requesting systems that may have user interfaces themselves.

Consistency and integrity were not a major concern in this implementation. Nevertheless it should be noticed that MAID has already some data quality checking in place, namely to detect wrong patient identifications (Cruz-Correia, Vieira-Marques et al. 2006). For full consistency checking, it is essential to have documents introspection which stresses the use of informatics standards like XML to describe documents, or even healthcare related standards like HL7-CDA or OpenEHR. Unfortunately, the Portuguese reality is still far from having healthcare IS that give access to their data through this standards.

Transparency is very difficult to obtain when dealing with documents generated in heterogeneous IS, because they lack normalization regarding presentation. To achieve a high level of transparency it is essential to use structured documents (e.g. XML).

4.2 Implementation Issues

By providing interfacing behaviours to incoming mobile agent there is no need for complex

interactions reducing interface agent congestion which would arise from multiple agents requesting database queries.

REFERENCES

- Berg, M. (2001). "Implementing information systems in health care organizations: myths and challenges." *Int J Med Inf* 64(2-3): 143-156.
- Cruz-Correia, R., P. Vieira-Marques, et al. (2005). "Integration of hospital data using agent technologies - a case study." *AI Communications* 18: 191-200.
- Cruz-Correia, R., P. Vieira-Marques, et al. (2006). "Monitoring the integration of hospital information systems: how it may ensure and improve the quality of data." *Stud Health Technol Inform* 121: 176-82.
- Lenz, R. and K. A. Kuhn (2002). "Integration of Heterogeneous and Autonomous Systems in Hospitals." *Business Briefing: Data management & Storage Technology*.
- Littlejohns, P., J. C. Wyatt, et al. (2003). "Evaluating computerised health information systems: hard lessons still to be learnt." *BMJ* 326: 860-3.
- Moreno, A. (2006). "On the evolution of applying agent technology to healthcare." *IEEE Intelligent Systems* 21(6): 8-10.
- Vieira-Marques, P., S. Robles, et al. (2006). "Secure Integration of Distributed Medical Data using Mobile Agents." *IEEE Intelligent Systems Journal* 21(6): 47-54.

FOR HOW LONG IS DATA FROM PREVIOUS ADMISSIONS ACCESSED BY HOSPITAL DOCTORS?

Ricardo João Cruz-Correia

*Biostatistics and Medical Informatics Department, Faculty of Medicine, University of Porto
CINTESIS, Faculty of Medicine, University of Porto, Portugal
rcorreia@med.up.pt*

Altamiro Costa-Pereira

*Biostatistics and Medical Informatics Department, Faculty of Medicine, University of Porto
CINTESIS, Faculty of Medicine, University of Porto, Portugal
altamiro@med.up.pt*

Keywords: Electronic patient records, Data integration, Information use.

Abstract: Distinguishing relevant information enables for better user interfaces, as well as better storage management. However, it is hard to distinguish between information really important to clinical care and only occasionally desirable. We aim to answer for how long are clinical documents useful for health professionals in a hospital environment considering its' content and the context of information request. We have studied the databases of a Virtual Electronic Patient Record that included (1) patient identification and the list of clinical documents integrated, (2) the visualization logs; and (3) a hospital encounters database that includes the list of encounters since 1993. Our results show that some clinical reports are still used after one year regardless of the context in which they were created, although significant differences exist in reports created in distinct encounter types. The half-life of reports by encounter type is 1.7 days for emergency, 3.9 days for inpatient and 27.7 for outpatient encounters. We conclude that the usage of patients past information (data from previous hospital encounters), varied significantly according to the setting of healthcare and content.

1 INTRODUCTION

Patient records exist to memorize and communicate the data existing on a particular individual, to help deliver care to him or her. Records are not only an information system but also a communication system that enables communication between different health professionals and between the 'past and present' (Dick & Steen, 1997; Nygren, Wyatt, & Wright, 1998).

Currently there are great quantities of stored data regarding patients. Although great advances have been made over the years (Cruz-Correia et al., 2007), on-demand access to clinical information is still inadequate in many settings, contributing to duplication of effort, excess costs, adverse events, and reduced efficiency (Feied et al., 2004). While it is widely accepted that full access to integrated electronic patient records and instant access to up-to-date medical knowledge significantly reduces faulty

decision making resulting from lack of information (Dick & Steen, 1997; Miller & Sim, 2004; Overhage et al., 2002), there is still very little evidence that life-long Electronic Health Records (EHR) improve patient care (Clamp & Keen, 2007).

Distinguishing between relevant and useless information enables for better user interfaces by highlighting most relevant information, as well as better storage management by choosing storage devices with better performance for relevant data. However, it is hard to understand what information is really important to clinical care, and what is simply occasionally desirable (Coiera, 1997).

Data age is usually one of the factors used to assess importance, making new information more relevant to the current search. But different data ages differently according to its type, i.e., some clinical reports describe situations less ephemeral than other and so are found useful longer than others. Also, the context of healthcare (e.g.: hospital environment,

primary care, oncology) probably influences the way information maintains its relevance.

We aim to study for how long are clinical documents useful for health professionals in a hospital environment.

2 BACKGROUND

In May 2003, the Department of Biostatistics and Medical Informatics implemented a Virtual Electronic Patient Record (HSJ-VEPR) (Cruz-Correia et al., 2005) for the Hospital S. João (HSJ), a university hospital with over 1 350 beds. The system integrates clinical data from 12 legacy departmental IS and the Diagnosis Related Groups and Hospital Administrative databases, aiming to deliver the maximum information possible to health professionals. Over 700 medical doctors use the system on a daily basis and the HSJ-VEPR retrieves an average of 3000 new reports each day (in PDF or HTML formats) (Cruz-Correia et al., 2005; Cruz-Correia et al., 2006), adding up to 2 million reports collected so far.

Each health professionals may access clinical information either by reading the paper patient record, using the HSJ-VEPR or using other IS available on the hospital.

3 METHODS

3.1 Participants

This study is concentrated in the report visualizations occurred in a two years period (2005 and 2006). In this period the hospital had 978 553 outpatients visits, 464 683 emergency visits and 82 444 inpatient visits. Reports' half-life by feeder system is analysis is based on the 3rd quarter of 2006 view results.

3.2 Data Preparation

The data considered in this study existed in three different Oracle schemas: (1) the HSJ-VEPR patient database, which included patient identification and the list of clinical documents integrated and; (2) the visualization logs including sessions, health professionals' identification and category and document views; (3) a hospital encounters database that includes patient identification, the list of encounters since 1993. These schemas use slightly different patient identification numbers, so transformation of these values was necessary to create relations between the tables.

HSJ-VEPR system does not know in what context (inpatient, outpatient or emergency) is the user accessing each report. The context was induced by confronting the date of view and the dates of the different patient encounters. When the date of view matches an encounter, that encounter is associated with the visualization. When no match is made no assumption is made regarding the encounter.

3.3 Clinical Report Half-life

Clinical reports' percentile is calculated by grouping all report views by type of encounter, ordering all visualizations by date, and calculating its relative position (current visualization position / number of visualizations). This technique allows us to compare the different encounter type groups by standardizing the position of each view. Reports half-life refers to the age of the report in percentile fifty.

4 RESULTS

Table 1 shows the number of visualizations taking in consideration the context of report creation and the context of report visualization. It should be noted that more than a half of the reports seen in 2005/2006).

Table 1: Number and *percentage* of visualizations grouped by context of visualization and report creation in 2005 and 2006.

Year	Report viewed in	Concomitant encounter		Previous encounter			Total			
		Emergency	Inpatient	Emergency	Inpatient	Outpatient				
2005	Emergency	861	40	334	16	511	24	447	21	2 153
	Inpatient	18 929	62	4 794	16	3 337	11	3 352	11	30 412
	Outpatient	154	1	1 158	4	5 150	18	22 043	77	28 505
	Total	19 944	33	6 286	10	8 998	15	25 842	42	61 070
2006	Emergency	2 973	49	743	12	1 129	19	1 202	20	6 047
	Inpatient	43 328	65	9 618	14	6 453	10	7 432	11	66 831
	Outpatient	290	0	2 543	4	10 804	18	46 874	77	60 511
	Total	46 591	35	12 904	10	18 386	14	55 508	42	133 389




Figure 1: Reports half-life grouped by episode type according to views in 2005 and 2006.

In inpatient encounters more than 35% (38/35%) report views regard previous encounters, without any clear distinction of which are found more relevant (previous emergency 16/14%, inpatient 11/10% and outpatient 11/11%).

In outpatient encounters almost all of the report visualizations were of reports created in previous encounters (99/100%). Most of them created in a previous outpatient encounter (77/77%).

Figure 1 illustrates reports' half-life by the type of encounter when report was generated. It shows that some clinical reports are still used after one year regardless of the context in which they were created. Nevertheless, outpatient reports are in average more durable than inpatient reports and emergency reports. The half-life of reports (percentile 50) by encounter type is 1.7 days for emergency, 3.9 days for inpatient encounters and 27.7 days for outpatient encounters.

Table 2 describes the reports' half-life (median of report age when viewed) group by department of feeder system in the 3rd quarter of 2006. It should be noticed the great difference in reports' half-life regarding feeding system (e.g. half-life of the pathology lab is 10 times greater than the clinical pathology lab).

Table 2: Reports half-life (median of report age when viewed) by department of feeder system in the 3rd quarter of 2006.

Feeder system	Views (n)	Half-life (days)
Clinical Pathology	18 261	4.4
Imuno-hemotherapy	23 691	4.6
Obstetrics	241	8
Pneumology	457	15
Intensive Care	141	26
Gastroenterology	1 773	38
Gynaecology	100	44
Pathology	16 567	47

5 DISCUSSION

Our results show that many report visualizations refer to previous encounters. Although the Hospital has not a unique patient record (in paper or electronic form), it is obvious that doctors which to access to previous encounter reports. It is also relevant that even older reports (more than one year) are still found useful by doctors.

As more and more patient information is stored, it is very important to efficiently select which one is more useful and promote it in a scenario where the scarceness of resources (screen space, disk space, bandwidth and doctors' time) is very real.

We intend to take in consideration reports' half-life in the next version of our system replacing the first patient record screen, reports collected in the last 24 hours, by a table in which the time interval is different for each type of report. Outpatient reports will be maintained in the list of last reports longer than inpatient and emergency reports.

This study rises new questions regarding what type of characteristics help maintain a report useful over the years.

6 CONCLUSIONS

We conclude that the usage of patients past information (data from previous hospital encounters), varied significantly according to the setting of healthcare and content.

ACKNOWLEDGEMENTS

This study was supported by POCTI/0753/2004 - Unit I&D CINTESIS n° 753/2004 from Fundação para a Ciência e a Tecnologia.

REFERENCES

- Clamp, S., & Keen, J. (2007). Electronic health records: Is the evidence base any use? *Medical Informatics and the Internet in Medicine*, 32(1), 5-10.
- Coiera, E. (1997). *Guide to medical informatics, the internet and telemedicine*: London: Chapman & Hall.
- Cruz-Correia, R. J., Vieira-Marques, P., Costa, P., Ferreira, A., Oliveira-Palhães, E., Araújo, F., et al. (2005). Integration of hospital data using agent technologies - a case study. *AI Communications*, 18, 191-200.
- Cruz-Correia, R. J., Vieira-Marques, P., Ferreira, A., Almeida, F., Wyatt, J., & Costa-Pereira, A. (2007). Reviewing the integration of patient data: how systems are evolving in practice to meet patient needs. *BMC Medical Informatics and Decision Making*, 7(1), 14.
- Cruz-Correia, R. J., Vieira-Marques, P., Ferreira, A., Oliveira-Palhães, E., Costa, P., & Costa-Pereira, A. (2006). Monitoring the integration of hospital information systems: how it may ensure and improve the quality of data. *Stud Health Technol Inform*, 121, 176-182.
- Dick, R., & Steen, E. (1997). *The Computer-based Patient Record: An Essential Technology for HealthCare*.
- Feied, C. F., Handler, J. A., Smith, M. S., Gillam, M., Kanhouwa, M., Rothenhaus, T., et al. (2004). Clinical Information Systems: Instant Ubiquitous Clinical Data for Error Reduction and Improved Clinical Outcomes. *Acad Emerg Med*, 11(11), 1162.
- Miller, R. H., & Sim, I. (2004). Physicians' use of electronic medical records: barriers and solutions. *Health Aff (Millwood)*, 23(2), 116-126.
- Nygren, E., Wyatt, J. C., & Wright, P. (1998). Helping clinicians to find data and avoid delays. *Lancet*, 352, 1462-1466.
- Overhage, J. M., Dexter, P. R., Perkins, S. M., Cordell, W. H., McGoff, J., McGrath, R., et al. (2002). A Randomized, Controlled Trial of Clinical Information Shared From Another Institution. *Ann Emerg Me*, 39(1), 14-23.

USING EXPLANATION FACILITIES IN HEALTHCARE EXPERT SYSTEMS

Keith Darlington

*London South Bank University, BCIM, 103, Borough Road, London, SE1 0AA, UK
keithd@lsbu.ac.uk*

Keywords: Explanation, Expert Systems, NHS Direct, Prodigy.

Abstract: A great deal has been written about healthcare expert systems in recent years. This paper examines a particular feature of expert systems: namely explanation facilities. A limited explanation capability is an integral part of a rule based expert system. The role of explanation in expert systems has been largely ignored in healthcare literature, since the MYCIN system and its derivatives were developed in the mid 1980s. However, empirical research has shown that users are more likely to adhere to recommendations made by expert systems when explanation facilities are available. Furthermore, explanation provision have been shown to improve performance and aid the user with a better understanding of the subject domain as well as result in more positive user perceptions of an expert system. This paper looks at the evolution of explanation facilities in healthcare expert systems, and investigates user requirements for explanation facilities in the healthcare domain.

1 INTRODUCTION

The medical expert system MYCIN (Shortliffe 1981) was amongst the first of a number of decision support diagnostic systems developed in the late 1970s. Since this time, the use of expert systems as an IT decision making aid in healthcare has grown rapidly (Schank, Doney and Seizyk, 1988), (Thornett 2001). For example, NHS Direct Hotline uses an expert system in basic patient diagnostics. NHS Direct has been at the forefront of 24-hour health care in the UK - delivering telephone and e-health information services direct to the public, and is accessed by over two million people every month. PRODIGY (Thornett 2001) is another example of an expert system that is used in primary care in the UK. Prodigy provides decision support to general practitioners within consultations regarding drug prescribing.

But many researchers have been sceptical about expert system usage in healthcare. For example, Delaney et al, (1999) believe that computerised decision support systems have great potential but have largely failed to live up to their promise. However, Walton et al, (1997) presents evidence to suggest that advice from a computer will be more convincing if presented simultaneously with an explanation for that advice. In their evaluation of

CAPSULE, an expert system giving advice to general practitioners about prescribing drugs, they say that "Finding the most effective way of presenting the explanation is an important goal for future studies of computer support for prescribing drugs". Yet, an explanation component is a standard feature of expert systems in that the systems problem solving behaviour can be observed during a consultation.

The inclusion of explanation facilities can enhance performance of the decision making and lead to greater adherence to the recommendations of the expert system (Gregor and Benbasat, 1999), (Arnold et al, 2006). Indeed, many studies have demonstrated the importance of a system being able to explain its own reasoning. For example, in a study of physician's expectations and demands for computer based consultation systems it was found that explanation was the single most important requirement for advice giving systems in medicine (Buchanan and Shortliffe, 1984). Also, according to (Berry et al., 1995), explanation is seen as a vital feature of expert systems – particularly in high risk domains, such as medicine, where users need to be convinced that a system's recommendations are based on sound and appropriate reasoning. However, despite the importance attached to explanations, few expert systems provide acceptable

explanation. Surprisingly, research in explanations has been largely ignored in healthcare expert systems since the development of MYCIN. Mao and Benbasat (2000) cite reasons why explanations in expert systems had failed to appeal to the user: that they were difficult to understand, and that they ignored the needs of different users. The following sections examine ways in which some of these shortcomings have been overcome in the healthcare domain, beginning with a look in some detail at one of the first ever expert systems to incorporate explanation – the medical expert system MYCIN.

2 THE MYCIN EXPLANATION FACILITY

The explanation facilities in MYCIN (Shortliffe, 1981) were presented in a natural language that was translated from the rules making the explanations easier to follow for the user. Explanations were also supplemented with certainty factors that numerically expressed the degrees of certainty attached to conclusions or outcomes. This meant that users of MYCIN could get an understanding of the likelihood of the advice given. However, there were many shortcomings identified with MYCIN's explanations. The following sections describe these shortcomings and how they were overcome.

2.1 Rule Trace

MYCIN is a rule based expert system – which means that knowledge is stored in the form of rules (Darlington, 2000). The explanation facilities provided in MYCIN and the other first wave of rule-based expert systems would have been a rule trace. This is, essentially, a record of the system's run-time rule invocation history during a consultation.

2.1.1 Feedforward and Feedback Explanations

A feedforward explanation provides the user with a means to find out why a question is being asked during a consultation (i.e., during the data input stage). The feedforward explanation will retain the manner in which input information is to be obtained for use in finding a solution.

A feedback explanation provides the user with a record of problem solving action during a consultation: i.e., how a conclusion was reached when the data has been completely input. The feedback explanations will present a trace of the

rules that were invoked during the consultation and display intermediate inferences in getting to a particular conclusion.

2.2 Strategic Explanations

Rule trace methods formed the basis of explanations in MYCIN (Shortliffe, 1981), but Clancey (1983) tried to adapt MYCIN from its diagnostic role to that of tutorial role in a system called GUIDON. The purpose of GUIDON was to provide a training system for junior consultants. What was thought to be a simple task turned out to be very difficult because MYCIN did not contain knowledge which explicitly contained strategic knowledge. This is knowledge about how to approach a problem by choosing an ordering for finding subgoals to minimise effort in the search for a solution. For example, the rule of thumb that alcoholics are likely to have an unusual aetiology can lead the expert to focus on less common causes of infection first – thereby pruning the search space to find a solution. The strategic knowledge in MYCIN was implicitly incorporated in the problem solving rules. However, Clancey realised that this knowledge needed to be explicitly represented, so that it could become transparent to students training to use the system. The lessons of GUIDON led Clancey to develop a follow up system called NEOMYCIN (Clancey & Letsinger, 1981): this was a consultation system whose medical knowledge base contained the tutorial strategic knowledge.

2.3 Justification Knowledge

An expert system can only reason with the knowledge contained in its knowledge base. Thus, the designer of a diagnostic medical expert system would ensure the knowledge base contains enough problem solving knowledge to ensure the system can arrive at the correct conclusions. However, the rule trace can only reconstruct a trace from what knowledge is contained in the expert system knowledge base. If the builder has not included the knowledge to justify the knowledge in the rule-base, then the system will not be able to justify the existence of the knowledge? Clancey (1983) realized the importance of this justification knowledge when attempting to extend the MYCIN system to support the training of junior physicians. Again, he found that MYCIN failed to do this because it did not contain justification knowledge. Justification knowledge is often unavailable because the rules which model the domain do not capture all the forms

of knowledge used by experts in their reasoning. This is because builders of expert systems capture “rules of thumb” shallow knowledge that only enable the system to solve diagnostic problems. Empirical research has consistently shown that user acceptance of expert systems increases for non-expert users when this justification knowledge is present, and that justification is the most effective type of explanation to bring about positive changes in user attitudes toward the advice giving system (Ye & Johnson 1995).

2.3.1 Capturing Deep Knowledge to Represent Justification Knowledge

Expert physicians would, of course, use rules of thumb themselves in solving problems, but they would also – as a result of their training and experience – possess a deep theoretical understanding of their subject domain. They may, for example, use “rules of thumb” or heuristic knowledge when performing a diagnosis. These “rules of thumb” may be sufficient to enable the physician to carry out a diagnosis, and therefore, this is the knowledge that is captured in rule-based expert systems because it is clearly much easier to obtain and code, and is sufficient for problem solving. However, this justification knowledge would have to be explicitly captured by the system designer if it were required for explanation.

3 USER REQUIREMENTS FOR EXPLANATION

In the healthcare domain, user requirements would vary according to employment categories which may include physician (including junior physician), nurses, administrators and also, as we will see later in some application examples, patients.

Considering the expert physician vs. non expert divide, research has shown that expert physicians do make use of explanation facilities but their requirements are very different to that of other users. Experts tend to have a preference for feedback rule trace explanations and are more likely – than non-experts – to use explanations for resolving anomalies (disagreement), verification and exploring alternative diagnoses (Arnold et al 2006), (Mao and Benbasat, 2000). Non-experts such as trainee physicians, on the other hand, are more likely to use explanations for short and long term learning. Moreover, (Arnold et al 2006), have

shown that non-experts tend to use both feedback and feedforward justification explanations, as well as terminological feedforward explanations, i.e., definitions of domain terms, etc., to assist during the data input stage. Patients – or other non-experts – using the NHS Direct system may benefit from this type of explanation because these explanations would facilitate learning of the subject domain. Some examples of these user-centred applications of explanation are briefly described in the next section.

3.1 Some User-Centred Explanation Prototypes in the Healthcare Domain

A number of healthcare projects involving the use of user modelling for explanation have been prominent in recent years. OPADE (Carolis et al, 1996) is a European Community Project funded expert system for generating beneficiary centred explanations about drug prescriptions that take into account the user requirements. The main objective of OPADE is to improve the quality of drug treatment by supporting the physician in their prescription process and by increasing compliance with the therapy (Berry et al, 1995). OPADE supports two types of user: those who directly interact with the system such as general practitioners and nurses, and those who receive a report of results – i.e., the patients. The explanations that are generated are dynamic (unlike static canned text explanations) in that a “user model” is maintained containing the characteristics of the user. A “text planner” component plans the discourse during a consultation. The text planner will build a tree containing the discourse plan which will depend on the objectives that are to be met by the user model. The explanation is then delivered in natural language by taking as input the tree generated by the text planner and transforming it using text phrases into the appropriate format.

HEALTHDOC (Marco et al, 1995) is another user-centred expert system whose main purpose is to generate reports for patients and materials for health education. This system enables the production of health-information and patient-education documents that are tailored to the individual personal and medical characteristics of the patients. Health-education documents can be much more effective in achieving patient compliance if they are customized for individual readers. The documents are presented textually and the text can be adapted to different patients, because the system contains a database

containing information about the clinical data of every patient – such as their personal and medical characteristics.

4 CONCLUSIONS

This paper recognises the role of expert systems in healthcare. However, one of their main features - explanation facilities – has been largely ignored in healthcare systems to date. Yet, empirical research has consistently shown, in recent years, that users are more likely to adhere to expert system recommendations when explanation facilities are available. Furthermore, explanation provision have been shown to improve performance and aid the user with a better understanding of the subject domain as well as result in more positive user perceptions of an expert system.

However, users will not use an explanation unless it addresses their basic information needs. This means that system designers must involve users in the evaluation of explanation facilities to ensure that they serve the needs of specific user groups. As this paper has shown, providing designers submit the effort, explanations can be tailored to the needs of different users. Perhaps the time has come for healthcare expert system designers, and users of such systems to re-evaluate the potential of explanation facilities.

REFERENCES

- Arnold, V., Clark, N., Collier, P.A., Leech, S. A., Sutton, S.G., 2006. The Differential Use and Effect of Knowledge-Based System Explanations in Novice and Expert Judgment Decisions. *MIS Quarterly*, Vol 30 No 1.
- Berry, D., Gillie, D. T., and Banbury. S., 1995. What Do Patients Want to Know: an Empirical Approach To Explanation Generation and Validation. *Expert Systems with Applications*, 8(4):419-428.
- Buchanan, B., G., Shortliffe, E., H., 1984. *Rule-based Expert Systems: The MYCIN Experiments of the Stanford heuristic programming project*. Addison Wesley, Reading, MA.
- Carolis, B.D., Rosis, F.D., Grasso, F., Rossiello, A., Berry, D., and Gillie, T., 1996. Generating recipient-centered explanations about drug prescription. *Artificial Intelligence in Medicine* 8, 123-145.
- Clancey, W. J., 1983. Epistemology of a rule-based Expert System: A framework for explanation. *AI magazine*, 20(3) pp 215-251.
- Clancey, W. J., Letsinger, R., 1981. NEOMYCIN; reconfiguring a rule-based expert system for application to teaching. Pages 361–381 in W.J. Clancey and E. H. Shortliffe (eds.), *Readings in medical artificial intelligence: The first decade*, Addison-Wesley, Reading, Massachusetts
- Darlington, K., 2000. *The Essence of Expert Systems*. Prentice Hall, London. Pearson Education.
- Delaney, B.C., Fitzmaurice, D. A., Riaz, D., and Hobbs, F. D., 1999. “Can computerised decision support systems deliver improved quality in primary care?” *Journal of BMJ*. November 13; 319(7220): 1281.
- Gregor, S., & Benbasat, I., 1999. Explanations from intelligent systems: theoretical foundations and implications for practice. *Management Information Systems Quarterly*, 23, 497-530.
- Mao, J., Benbasat, I., “The Use of Explanations in Knowledge-Based Systems: Cognitive Perspectives and a Process-Tracing Analysis”, *Journal of Management Information Systems*, Vol. 17(2), pp153-180, 2000.
- Marco, C.D., Hirst, G., Wanner, L., and Wilkinson, J., 1995. *HealthDoc: Customizing patient information and health education by medical condition and personal characteristics*. In *First International Workshop on Artificial Intelligence in Patient Education*, Glasgow, UK.
- Schank, M. J., Doney, L. D., Seizyk, J., 1988. The potential of expert systems in nursing. *J Nurse Administration*. June; 18(6):26-31. PMID: 3286842 [Pub Med - indexed for MEDLINE].
- Shortliffe, E. H., 1981. *Computer Based medical consultations - MYCIN* Published by Elsevier, USA.
- Thornett, A., 2001. Computer decision support systems in general practice. *International Journal of Information Management*; 21:39-47.
- Ye, L.R., Johnson, P.E., 1995. “The impact of explanation facilities on user acceptance of expert systems advice”, *MIS Quarterly*, Vol. 19(2), pp157-172.
- Walton, R., 1997. An evaluation of CAPSULE, a computer system giving advice to general practitioners about prescribing drugs. *J INF Primary Care*.

MDFLUXO: OPHTHALMOLOGY EDUCATION WITH A PDA EFFICACY AND USABILITY EVALUATION

Vladimir Camelo Pinto, Thiago Martini da Costa, Marcelo Contardo Moscoso Naveira

Daniel Sigulem, Paulo Schor and Ivan Torres Pisa

Medical Informatics Department, Federal University of São Paulo (UNIFESP), Rua Botucatu 862, São Paulo, SP, Brazil

vladimir-pg@dis.epm.br, thiago-pg@dis.epm.br, navsmed41@hotmail.com

sigulem@unifesp.br, pschor@pobox.com, ivan.pisa@unifesp.br

Keywords: PDA, Students, Health professionals, Usability.

Abstract: In the last decade, handheld computers, also known as Personal Digital Assistants (PDAs), have become popular among physicians, residents and medical students. We have developed a PDA-software, called MDFluxo, to assist ophthalmologic teaching as a guide. We used a user-centered design to try to diminish usability problems. It's interesting that the literature concerning PDA use on health area doesn't emphasize usability attributes and it's evaluation, which is important to widespread PDA use among health professionals. MDFluxo efficacy was evaluated comparing it to traditional book guide and a non-guide self-learning. Usability inspection methods, such as Think Aloud Method, Nielsen Heuristic Evaluation, QUIS 5.0 adaptation, and Cognitive Walkthrough are being used to identify usability problems. The usability inspection methods help us to identify usability problems and correct them. We evidenced that the students who used MDFluxo improved the resolution of the clinical case when compared to participants which did not use any material as support, however there was no significant difference when comparing the MDFluxo students to students using a book, thus, as a support to ophthalmology learning, MDFluxo is an efficacy tool.

1 INTRODUCTION

In the last decade, handheld computers, also known as Personal Digital Assistants (PDAs), have become popular among physicians, residents and medical students. It is estimated that near 60-70% of the medical students and residents on the USA use PDA for some purpose related to health (Kho *et al.*, 2006). Some equipment characteristics such as: size, mobility, objectivity at offering relevant clinical information can be related to the observed spread (Kho *et al.*, 2006). The main applications used among medical students and residents are: research tools (Sutton, 2004), medical calculators (Honeybourn *et al.*, 2006), bibliographical references (Rudkin *et al.*, 2006), and drugs databases (Stroud *et al.*, 2005). We didn't find papers related to health professionals and medical students PDA usage on Brazil, although we are managing an electronic research to get this data (<http://telemedicina.unifesp.br/palm>).

Widespread use and growth in the use of PDA can be due to it usability. On a recent review (<http://telemedicina.unifesp.br/mdfluxo>) that we have performed, we have evidenced the increasing use of PDA in the health environment, as well as

PDA-software efficiency and effectiveness in the aid of teaching and medical practice. There were few studies regarding PDA-software development centering on health professional opinions. This could generate problem in understanding results, since Nielsen (1993) asserts "is virtually impossible to design a user interface that has no usability problems from the start.

The aim of this paper is to present MDFluxo, for learning ophthalmology based on handheld computers, and its evaluation among medical students. There will also be introduced preliminaries results concerning usability inspection methods we are applying.

2 MATERIALS AND METHODS

The MDFluxo software is based on the book "Federal University of São Paulo (UNIFESP) Guide for Ambulatories Medicine - Ophthalmology" written by Schor *et al.* (2004). In the opening pages Schor *et al.* (2004) mention that the book is both a clinical and surgery manual and that it was created to help medical practice, describing the most common le-

sions in each human eye segment, followed by anatomical location. Each chapter possesses a flowchart that allows grouping manifestations by their common characteristics. The book offers a total of 55 flowcharts regarding each disease. Each flowchart summarizes the medical knowledge discussed in the current chapter and offers mnemonics to compact the text with objectiveness.

The flowcharts extracted from the book were, first revised by specialists, and then transcribed to eXtensible Markup Language (XML) due to its flexibility, enhanced integration and SOA-oriented services for different environments (Pinto *et al.*, 2006).

On this study, we used Palm OS Simulator[®] version 5.2.1 to simulate a Palm environment on desktop PC. We also used palmOne[™] Tungsten E handheld devices with 320x320 screens. On the desktop PC, we had installed Microsoft Visual Basic[®] 6.0 with the plug-in MobileVB[®] 4.0.0 of the AppForge[®], and database Microsoft Access[®] 2003.

According to general Human-Computer Interface (HCI) literature and our expertise on developing health-applications, a close relationship between health-professionals is vital to avoid failures. Thus, we had chosen a user-centered design process (Norman, 1983) in which we included ophthalmologists, physicians and residents on prototyped development process.

2.1 Usability Inspection

Trying to evaluate the five usability quality components: learnability, efficiency, memorability and satisfaction (Nielsen, 2003), we planned to use the following methods: Think Aloud Method, Nielsen's Heuristic Evaluation, Questionnaire for User Interface Satisfaction (QUIS 5.0) and Cognitive Walkthrough.

A Think Aloud Method (Someren *et al.*, 1994) was applied to ophthalmology specialists, on a laboratorial environment, to evaluate learnability and efficiency in a simulated case-solving.

Nielsen's heuristic evaluation (Nielsen, 1994) was applied by 5 biomedical informatics students to identify problems on the five usability components.

There are two other usability inspection methods that will be applied. We translated to Brazilian Portuguese language and adapted the QUIS 5.0 (Chin *et al.*, 1988). We planed to apply it on several voluntry physicians to assess their satisfaction using MDFluxo. A Cognitive Walkthrough Evaluation Process (Lewis *et al.*, 1990) is being planned to quantify the efficiency, the learnability, and memorability with medical student in a PBL evaluation.

2.2 MDFluxo Comparing

We didn't intend to evaluate MDFluxo as a substitute to the traditional learning. On this paper we compared if MDFluxo is a valid tool to complement the ophthalmologic learning.

One of the authors (Paulo Schor) developed three different, equivalent clinical cases to assess MDFluxo efficacy. Patient's main and secondary complains, personal info, clinical history, and related issues were included on the clinical cases.

From UNIFESP and Faculdade de Ciências Médicas de Santos – Centro Universitário Lusiada (FCMS-UNILUS), fourth and fifth year medical students were divided in three groups. The first group used MDFluxo as a didactic support, the second group used the Guide of Ophthalmology book (Schor *et al.*, 2004) and the last one group didn't use any didactic material as support to solve the clinical case.

Each participant was required to answer three questions: the diagnostic hypothesis, etiology for them, and conduct. To each right answer, half right and wrong it will be attributed values 1, 0.5 and 0 respectively and the accomplishment of the trial was timed.

The students' response was corrected and the data tabulated in an electronic spreadsheet. For each question of the test (diagnostic hypothesis, etiology, conduct) the grades and time spent to answer were evaluated. We used the Friedman's test (Siegel, 1988) to verify the variance among the groups.

3 RESULTS

The MDFluxo project resulted in 22 prototype versions, each one built considering the physicians opinion, regarding flowcharts content, usability attributes, including layout design. We evidenced great differences between initial and final prototypes, as shown in Figure 1 and Figure 2 respectively.

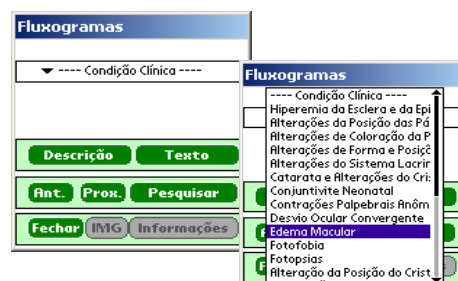


Figure 1: Two screenshots from 5th prototype, with several buttons and few information.

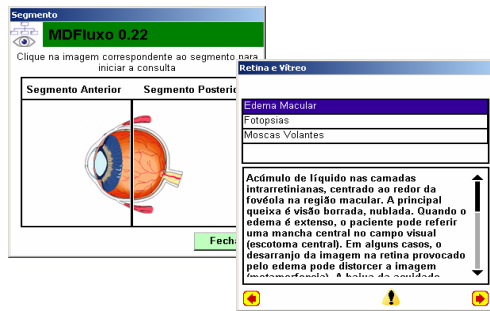


Figure 2: Two screenshots from 22th prototype, with more information and improved navigability.

3.1 Usability Inspection

Some preliminaries results of usability inspection that we are conducting are listed below:

1. We diminished the amount of options present on preliminary releases to clear the screen.
2. There were no icons on the system. Icons increases the mapping with the real world.
3. The exclamation point (“!”) used for “View info on the Author” and the button “Close” have not a consistent use on the system.
4. The red color is often used in the system for normal instruction of pure text, however, this color should be used only for critical alerts.
5. It is not obvious that the “Eye Segment” is selectable.
6. There were grammar and writing mistakes on titles and menus and generated texts.
7. For the “gender” and “age” we used Select Box, which does not follow the WIMP standard.
8. The “Instructions” option doesn’t offer a step-by-step manual, but a short explanation on each button function. The names “Instructions” and “Project Information” are not appropriated to their meaning.
9. In some screens the user could not click on the text and change it, however, this feature was wrongly enabled.
10. Error reports were shown several times, such as “Form was unable to load” and “subscript out of range” during the system loading.
11. The “Help” function is only presented on the starting screen. During navigation there is no other “Help” button or “Return to Home/Main”.

3.2 MDFluxo Comparing

The descriptive statistic, which describes the study results, is present on Table 1. On this table we present average results for the groups that used MDFluxo on PDA, book and the group that didn’t use anything for each question and time.

Table 1: Descriptive statistics containing average values for each studied feature.

	Diagnostic hypothesis	Etiology	Conduct	Time	General Result
PDA	0.600	0.350	0.450	16.980	1.400
Book	0.650	0.325	0.425	14.420	1.400
Without nothing	0.475	0.050	0.125	12.220	0.650

Applying Friedman’s test to verify statistical differences in the values presented on Table 1, we evidenced that using PDA leads to a better etiology resolution and conducts when compared to group that doesn’t use any material.

Beyond what has been stated, there was no significant result for resolution of diagnostic hypothesis in comparison to the other two groups. In the analysis among the medium time required for the accomplishment of the clinical case among three groups, the use PDA as reference material was better than the book and without the aid of any material. However it should be taken into consideration that the students did not receive previous training for the use of the software.

On Table 2 we present the statistical difference in diagnostic hypothesis, etiology, conduct, time and general result for MDFluxo comparing to book and students with nothing.

Table 2: Analysis of variance according to Friedman’s test. (*) p-value > 0.050 (without significance).

Diagnostic hypothesis	Etiology	Conduct	Time	General Result
$x^2_{calc} = 1.720$ ($p < 0.423$)	$x^2_{calc} = 9.140$ ($p < 0.010$)	$x^2_{calc} = 6.140$ ($p < 0.050$)	$x^2_{calc} = 8.400$ ($p < 0.015$)	$x^2_{calc} = 4.630$ ($p < 0.099$)
	PDA & Book > Without nothing	PDA > Without nothing	Book > PDA & Without nothing	*

4 CONCLUSIONS

There are many boundaries on PDA-software development, especially when we consider the health environment and its applications. For MDFluxo project, these restraints only reinforced the role of the participating user during the development process. A weak mnemonic gap could have been created because, due to the PDA screening limitation, flowcharts are showed in parts, instead of the whole, as in the book. However the group that used PDA-software MDfluxo achieved better result than the group that did not use any material for the resolution clinical case.

The participants that used PDA-software MDFluxo did not have any difficulties to navigate through the different. However the amount of flow-charts presented by the software did not correspond to the users' expectations.

The ophthalmologic clinicians and students found MDFluxo as a portable, quick, and intuitive guideline and a method of learning on managing clinical cases. Improvements, such as considering the usability on the development process, could turn the PDA into a more useful resource to physicians, enhancing the assistance of patients.

A research was accomplished in medical database with objective of finding in the medical literature representations of ophthalmologic information similar to the representations of the PDA-software MDFluxo that were extracted of the book of Schor et al., (2004). However, it was not found similar studies it to allow accomplishing comparisons with the results of this study.

ACKNOWLEDGEMENTS

We thank the support provided by the Ophthalmology Department, Medical Informatics Department, and Biostatistics and Technology Information Division of UNIFESP, and the volunteers. This project was accomplished with CAPES master's scholarship.

REFERENCES

- Chin, J.P., Diehl, V.A., Norman, K.L., 1988. Development of an instrument measuring user satisfaction of the human-computer interface. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. Washington, D.C., United States. ACM Press, New York, pp.213-218.
- Honeybourne, C., Sutton, S., Ward, L., 2006. Knowledge in the Palm of your hands: PDAs in the clinical setting. *Health info libr J.* 23(1), pp.51-9.
- Kho, A., Henderson, L.E., Dressler, D.D., Kripalani, S., 2006. Use of handheld computers in medical education. *J Gen Intern Med.* 21(5), pp.531-537.
- Lewis, C., Polson, P.G., Wharton, C., Rieman, J., 1990. Testing a walkthrough methodology for theory-based design of walk-up-and-use interfaces. In Proceedings of the SIGCHI. Seattle, Washington, D.C., J. C. Chew and J. Whiteside, Eds. CHI '90. ACM Press, New York, pp.235-242.
- Nielsen, J., 1993. Iterative user-interfaces design. *IEEE Computer.* 26(11), pp. 32-41.
- Nielsen, J., 1994. Heuristic evaluation. In: Nielsen J, Mack RL (Eds.). *Usability Inspection Methods*, John Wiley & Sons, New York, NY.
- Nielsen, J., 2003. Usability 101: Introduction to Usability. (online). Available from: <http://www.useit.com/alertbox/20030825.html>. (cited 11 July 2007)
- Norman, D.A., 1983. Design principles for human-computer interfaces. In Proceedings of the SIGCHI. Boston, Massachusetts, United States, A. Janda, Ed. CHI '83. ACM Press, New York, NY, pp.1-10.
- Pinto, V.C., Neira, R.A.Q., Bottós, J.M., Bottós, K.M., Macedo, A.A., Schor, P., Pisa, I.T., 2006. MDFluxo: Ferramenta de auxílio ao ensino em oftalmologia baseada em computador de mão. XCBIS'2006. Florianópolis, SC.
- Rudkin, S.E., Langdorf, M.I., Macias, D., Oman, J.A., Kazzi, A.A., 2006. Personal digital assistants change management more often than paper texts and foster patient confidence. *Eur J Emerg Med.* 13(2), pp.92-6.
- Schor, P., Chamon, W., Belfort, R., 2004. Guia de Oftalmologia - Série Guia de Medicina Ambulatorial e Hospitalar da UNIFESP/EPM. 1 ed. Manole, São Paulo.
- Siegel, S., Castellan, Jr. N.J., 1988. *Nonparametric Statistics.* 2 ed. McGraw-Hill Int. Ed, New York, pp-399.
- Sutton, J., Stockton, L., McCord, G., Gilchrist, V.J., Fedyna, D., 2004. Handheld computer use in a family medicine clerkship. *Acad Med.* 79(11), pp.1114-9.
- Stroud, S.D., Erkel, E.A., Smith, C.A., 2005. The use of personal digital assistants by nurse practitioner students and faculty. *J Am Acad Nurse Pract.* 17(2), pp.67-75.
- van Someren, M.W., Barnard, Y.F., Barnard, J.A.C., 1994. *The Think Aloud Method: A Practical Guide to Modelling Cognitive Processes.* ISBN-0-12-714270-3. Academic Press, Inc., 525 B Street, Suite 1900, San Diego.

TELEMEDICINE TEST-BED

A Tool for Determination of Accuracy in Asynchronous Collaborative Method

Josceli Tenório, Cristina Lucia Feijó Ortolani, Paulo Roberto de Lima Lopes and Ivan Torres Pisa

Federal University of São Paulo (UNIFESP), Brazil

josceli-pg@dis.epm.br, ortocris@pop.com.br, plopes@unifesp.br, ivan.pisa@unifesp.br

Keywords: Teledermatology, telemedicine.

Abstract: Store-and-forward method has contributed for telemedicine as an efficient method to the clinical diagnosis acquisition. Studies in teledermatology demonstrated important outcomes, such as the use of computer systems based in asynchronous method for long-distance diagnosis, categorization, triage, or clinical outcomes agreements. However, for some issues there is no consensus between studies, such as the image parameters or the data that must be submitted to analyze. The aim of this study is propose the development of a web system that establishing patterns to construct web asynchronous system. A prototype was constructed to test a structure (clinical data and images) to obtain clinical diagnosis concordance in teledermatology. The preliminary outcomes showed that rate clinical diagnosis is nearly when we compare face-to-face method (76.6%) and store-and-forward (70.3%). The most effectiveness occurs to some diseases, who present easy clinical diagnosis in traditional consult. We expect that this system could introduce a numerical method, established in statistical data, to construct store-and-forward system.

1 INTRODUCTION

The telemedicine was developed and based under the concept of the access to health-care using information technology in situations where the long distance is a critical factor (Craig and Patterson, 2005). The Internet extended this concept under the paradigm of the widely collaboration between health professionals.

Due to traditional face-to-face method consultation, using visual observation of the patient and method to obtain date to realize an initial diagnosis hypothesis, specialties such as dermatology and ophthalmology (Finamor *et al.*, 2005) are further adjusted to these requirements, widely documented (Ferguson, 2006).

Several studies based in store-and-forward teledermatology have been performed and the outcomes are:

- higher concordance rates in comparison to outcomes histological diagnosis (Du Moulin *et al.*, 2003; Moreno-Ramirez *et al.*, 2005; Chao *et al.*, 2003);
- categorization of outcomes (Piccolo *et al.*, 2004);
- clinical case triage, according to urgency level (Oakley *et al.*, 2006);

- fast reply (80% of submitted consultations in one day) to requisitions submitted to experts using a non-commercial network based in use of easy platform to teleconsultations services (Massone, 2006);
- number of patients (51%) reduction referral for the consultation with the expert, when teleconsultations (data and image) are submitted from general practioner by email, to expert. This value is similar to the real time cases (Knol *et al.*, 2006);
- similar clinical outcomes in comparison to virtual teleconsultation (Pak *et al.*, 2007).

However, there is no agreement in several issues as the influence and minimum level of images quality (resolution and compression) for a compose efficient clinical diagnosis, minimum data set to be submitted, the composition of the patients history clinical and reply, the adequate specialties to asynchronous method, the restrictions to use, the level of complexity of the clinical cases that must be diagnostics efficiently, and the influence of expert's experience.

It's necessary to implement a structure and vocabulary to the establishment of a computation system to realize experiments that allows the

evaluation of the necessary and enough parameters for efficient clinical diagnosis.

The aim of this study is to propose the basis to development of system web multi-specialty and to show the outcomes of a prototype web system, applied to teledermatology, which allows to establishment of a virtual collaborative research environment, to realize tests for the evaluation of the telemedicine accuracy, an important subject for the reliability of the asynchronous method.

2 MATERIALS AND METHODS

The model is the reproduction of a real experimental environment, where a clinical data set and images, of diagnostic performed by face-to-face consultations can be available and evaluated by specialists, whom opinion could be validated according to gold-standard.

The physical architecture will contemplate the client-server model. According to established aims, the use of open source software is necessary.

To assessment of platform will realize experiments: a control group composed for clinical cases recorded in clinical archives with data and images obtained in a face-to-face consultation, selected by the manager health professional for test control, with previously known clinical diagnosis, what it will compose the standard-gold of this tests. The initial focus of study will be dermatology and ophthalmology.

The proposed web platform will provide the manager health professional a computational system to deliberate data set and diagnosis options, including false options (tricks). Besides, it will be able to change images size, resolution, and compression rate. To further reliability, the manager health professional will be able available to randomize for choosing the clinical cases.

Potential data sets to be tried are composed of minimum personal data, clinical history, and clinical data from face-to-face consultation, image quality parameters (minimum resolution, and compression). The combination of these factors will compose other experimental data sets.

2.1 Application

A prototype was performed to application in teledermatology. The aim was the determination of rate diagnosis concordance among initial diagnosis hypothesis face-to-face consultations and store-and-forward in comparison of both cases with gold-

standard. The literature (Bowns *et al.*, 2006) reveals that the diagnosis concordance rate will must be more than 60 %.

To conduct this study, it was defined two standard users: evaluator (a dermatologist responsible for a test control and access to statistical outcomes) and specialist (two dermatologists analyzed the 64 cases together, according to real environment).

To realize the trials, a control group was performed for 64 randomized clinical cases registered in paper, composed by personal and clinical data, dermatoscopic images, a initial diagnoses hypothesis and a result of clinical analysis (gold-standard), selected by evaluator.

Figure 1: Web environment to dermatologist.

This prototype allowed trial a minimum data set to teleconsultation in dermatology. Personal data (gender and age), clinical (body's part and time of occurrence of lesion) and dermatological images (2-6 images, mean score 2.6). There were two options to resolution: standard (640 x 480) or the captured image. The figure 1 shows the web system environment available to specialist.

The specialist task was choice the option of a diagnoses hypothesis, image's quality, difficulty degree for diagnosis and available data in a list. Commentaries can be performed if necessary.

3 RESULTS

The diagnostic concordance rate in face-to-face consultation (initial diagnoses hypothesis) was 76.6% (49/64) and using web system (asynchronous method) was 70.3% (45/64). This outcome is according to literature (Bowns, 2006) to web system.

The trial allowed to obtain the diagnosis concordance among consultation face-to-face and using web system according to diagnoses hypothesis, showed in table 1. The most significant results occurs in absolutely success in basal cell carcinoma diagnosis (face-to-face: 11 match/11 choices; web: 11 match/11 choices) and atypical nevus failure (face-to-face: 2 match/10 choices; web: 0 match/6 choices).

Table 1: Diagnoses concordance rates in comparison to gold-standard.

Diagnoses Hypothesis	Face-to-face (%)	Web (%)
Basal Cell Carcinoma	100.0	100.0
Haematoma	100.0	100.0
Malignant Lentigo Melanoma or Malignant Lentigo	77.8	75.0
Simple Lentigo	0.0	0.0
Solar Lentigo or Solar Melanosys	100.0	100.0
Malignant Melanoma	70.0	75.0
Atypical Nevus	20.0	0.0
Blue Nevus	100.0	100.0
Congenital Nevus	100.0	71.4
Reed's Nevus	0.0	0.0
Melanocytic Nevus	88.9	80.0
Actinic Keratosis	0.0	0.0
Seborrheic Keratosis	100.0	83.3

The specialist's opinion about the available data using the web system was showed in table 2.

This trial allowed us to measure of a specialist opinion about the diagnosis difficulty. It was possible compares the face-to-face and web analysis. Table 3 shows this comparison.

The table 4 shows comparison among the specialist opinion about the diagnosis difficulty and the success of diagnoses hypothesis. This is according our forecast, because the rate for easy diagnosis is major.

Table 2: Specialists opinion rates about data available in web system.

Data was enough to realize a diagnostic?	Rate (%)
No, it's very important includes more data.	1.6
Yes, certain.	70.3
Yes, but I had doubt.	28.1
Image's Quality	
Very Nice	47.1
Nice	6.5
Bad	5.3
No important	0.0

Table 3: Diagnosis difficulty degree rates.

Diagnosis Difficulty Degree	Face-to-face (%)	Web (%)
Easy	46.9	43.8
Medium	21.9	28.1
Difficult	31.3	28.1
Sum	100.0	100.0

Table 4: Success diagnosis hypothesis in relation with the diagnosis difficulty degree.

Diagnosis Difficulty Degree	Face-to-face (%)	Web (%)
Easy	96.7	96.4
Medium	64.3	55.6
Difficult	55.0	44.4

These trials conduct us to basic structure to web system to remote diagnosis web system in dermatology:

- Personal data minimum: gender and age.
- Clinic data minimum: body' part and time.
- Image's quantity: 2.6 (mean-score).

This structure presents more efficiency (around 70%) to diseases with easy diagnosis when traditional consultation.

4 CONCLUSIONS

Studies based on the asynchronous method has not been enough yet to establish it as an greatest choice for the resolution of difficulty access health-care to populations in isolated areas geographically, triage of clinical cases for reduction of specialist access delay, available of the second opinion or discussions

of difficult clinical cases. There is only a potential indication for the use of asynchronous method and the success factors are being established yet, but only for specific specialties (Finamor *et al.*, 2005 and Ferguson, 2006). The studies of the diagnosis agreement performed still not presented a consensus.

This prototype showed us a basic structure to have success in remote telediagnosis in dermatology. We expect that a web system based in these characteristics must have rate success nearly face-to-face consult. This is the aim of test-bed: trial a structure, to obtain data to analyse the outcomes and construct the web systems based in tested values.

In the experiments, found in literature, that use email to submit data and attach images files, it there is not a pattern that it indicates that set data must be submitted. The text is based only on professional common-sense and medical practice. Bergus *et al.* (2006), demonstrated that the quality of information is very important to referral reduce the need for clinic consultation. In this prototype data standardization was used to facility the data input and establish a vocabulary.

In the next trials, we expected that test-bed shows efficiently to obtain a list of another standard indicators that will answer previously to the raised issues: what size, minimum image resolution and compression will be able to realize a diagnosis, minimum personal data, what data will must be compose recent clinical history, the minimum amount of images to diagnosis, that images are essential, the efficiency rate if images and data will be analyzed separately. Indications to the use for specific clinical cases also could be known.

We expect that this study will contribute for the establishment of the asynchronous flexible collaboration as efficient, reliable, comfortable, and economic method for health professionals in specific medical specialty.

ACKNOWLEDGEMENTS

Sérgio Hirata, Sérgio Yamada and Priscila Yshioka, Department of Dermatology - Federal University of São Paulo (UNIFESP), Brazil.

REFERENCES

- Bowns, I.R., Collins, K., Walters, S.J., McDonagh, A.J.G. *Telemedicine in dermatology: a randomised controlled trial*. Health Technology Assessment 2006; 10; 43:
- Bergus, George R., Myra E., Reed, David A. e Attaluri, A. *Email teleconsultations: well formulated clinical referrals reduce the need for clinic consultation*. Journal of Telemedicine and Telecare 2006; 12: 33–38.
- Chao, L. W., Cestari, T. F., Bakos, L. et al. *Evaluation of an Internet-based teledermatology system*. Journal of Telemedicine and Telecare 2003; 9: Suppl 1:S9-12.
- Craig, J., Patterson, V. *Introduction to the practice of telemedicine*. Journal of Telemedicine and Telecare 2005; 11: 3-9.
- Du Moulin, M. F. M. T., Bullens-Goessens, Y. I. J. M., Henquet, C. J. M. et al. *The reliability of diagnosis using store-and-forward teledermatology*. Journal of Telemedicine and Telecare 2003; 9: 249-252.
- Ferguson, J. *How to do a telemedical consultation*. Journal of Telemedicine and Telecare 2006; 12: 220-227.
- Finamor, L. P. dos S, Martins, M C, Muccioli, C, Sigulem, D. Lopes, P R L, Belfort, R *Teleoftalmologia como auxílio diagnóstico nas doenças infecciosas e inflamatórias oculares*. Revista da Associação Médica Brasileira, 2005; 51(5):279-84. <http://www.scielo.br/scielo.php?pid=S010442302005000500020&script=sci_arttext>. Accessed: 03/20/2006.
- Knol, A., Van der Akker, T.W., Damstra, R.J., Haan, J de. *Teledermatology reduces the number of patient referrals to a dermatologist*. Journal of Telemedicine and Telecare 2006; 12; 75-78.
- Massone, C., Soyer, H.P., Hofmann-Wellenhof, R. *et al. Two years' experience with web-based teleconsulting in dermatology*. Journal of Telemedicine and Telecare 2006; 12; 83-87
- Moreno-Ramirez, D., Ferrandiz, L., Galdeano, R., Camacho, F. M. *Teledermatoscopy as a triage system for pigmented lesions: a pilot study*. Clinical and Experimental Dermatology 2005; 31: 13-18.
- Oakley, A.M.M, Reeves, F., Bennet, J., Holmes, S.H., Wickham, H. *Diagnostic value of written referral and/or images for skin lesions*. Journal of Telemedicine and Telecare 2006; 12: 151-158.
- Pak, H., Crystal, A.T., Lindquist, J.H., Grambow, S.C., Whited, J.D. *Store-and-forward teledermatology results in similar clinical outcomes to convencional clinic-based care*. Journal of Telemedicine and Telecare 2007; 13; 26-30.
- Piccolo, D., Soyer, H. P., Chimenti, S. *et al. Diagnosis and categorization of acral melanocytic lesions using teledermatoscopy*. Journal of Telemedicine and Telecare 2004; 10): 346-350.

BRAZILIAN TELEMEDICINE NETWORKS

Ricardo Alfredo Quintano Neira, Leandro Galassi Zavitoski, Paulo Roberto de Lima Lopes
Henrique Manuel Lederman, Paulo Schor, Ivan Torres Pisa and Daniel Sigulem

Medical Informatics Department, Federal University of São Paulo (UNIFESP), Rua Botucatu 862, São Paulo, Brazil
rquintano-pg@dis.epm.br, lgalassi@saude.sp.gov.br, plopes@dis.epm.br
lederman@unifesp.br, pschor@pobox.com, ivanpisa@dis.epm.br, sigulem@dis.epm.br

Keywords: Telemedicine, Networks, Public Health.

Abstract: One of the outcomes of the telemedicine evolution was the emergence of telemedicine networks which can be characterized as the interconnection of centres or healthcare professionals to provide telemedicine services. This paper aims to present a study of the characteristics of the Brazilian telemedicine networks. For the implementation of this study and the selection of the projects, the authors proceeded a search and revision of articles and sites obtained in the Google and PubMed repositories. As conclusion, the existence of two types of networks was verified: infra-structure and telemedicine service; all of the networks are based on asynchronous telemedicine, most of them acting in inter-cities, nevertheless there are not telemedicine cases inside the own municipal district to eliminate socioeconomic barriers; there were not find inside these networks projects or initiatives to regulate the remuneration of the professionals involved in the processes as well as, the payment for the services that were used.

1 INTRODUCTION

Telemedicine is defined as the offer of services related to health care, as well as the changing of information in cases where the distance is a critical factor (Neira, 2006). These services are provided by health professionals, using information and communication technologies for exchanging valid information for diagnoses, and for the health care providers continuous education, as well as, for research and evaluations purposes.

Actions in telemedicine are being done in Brazil since the 90's decade, and although the efforts demonstrate exciting potential results and the great medical centers possess plenty of advanced hospital units, it persists a growing demand for medical care in remote places, executed by skilled professionals and, in this way, initiatives and projects appeared in telemedicine (Rede Nacional de Ensino e Pesquisa, 2006).

Among these projects we can observe the telemedicine networks implementation which can be characterized as the interconnection of health professionals, universities, diagnoses centers, libraries, hospitals, as well as government entities. These networks intend to eliminate the physical and

social barriers to the healthcare services. Their objectives are (Anderson, 2006): 1) improve access to the specialists' services; 2) allow the patients' attendance minimizing displacements; 3) improve teaching and updating in medicine and health.

The objective of this article is to present a study and analysis of the characteristics of the Brazilian telemedicine networks.

2 MATERIALS AND METHODS

To start this work, first of all we collected information from Annals of Brazilian Health Informatics Congresses from 2006. As well as, we have collected more information by contacting local authorities, national administration and other informal sources of information, like websites, newspaper and magazines.

For this formal study we have selected telemedicine network projects considering papers by systematic review from PubMed (PubMed Central, 2007) and systematic search in Google (Google, 2007). With the objective to use the specific descriptors for the searches above, first it was done a review on DeCS - Health Sciences Descriptors (Descritores em Ciências da Saúde, 2007) – a Latin

American version of MeSH – to identify keywords for the search. Only the term *telemedicina* (telemedicine) was found. Although we didn't find a representing term for telemedicine networks, the following terms were adopted for this research: *rede telemedicina* (telemedicine network) and *rede telemedicina Brasil* (telemedicine network Brazil). Then, on February 2007, we searched on Google (Google, 2007) - with the web option selected - and on PubMed (PubMed Central, 2007) repositories using the identified terms described above.

2.1 Search and Inclusion Criteria

The first 20 summaries results of selected the repositories were analyzed. When the summary presented information about a Brazilian telemedicine network, the article was considered for this study. Just the articles that had objectives and goals about Brazilian telemedicine networks were added to this review. In those cases where article made reference to a new telemedicine network that wasn't found previously it was made a new specific search. We search that new network on same repositories proceeding to verify the 10 first results following the same criteria of the main search.

Due to the lack of telemedicine networks information standardization, without any condition to determine similar points of comparison, these criteria were established.

3 RESULTS

The telemedicine networks found based on the established criteria were eight. Following the results:

3.1 Academic Network of Telemedicine – *Rede Universitária de Telemedicina (RUTE)*

The RUTE objectives include to support the improvement of existent telemedicine projects in university hospitals and to motivate the arising of futures inter-institutions works (Rede Universitária de Telemedicina, 2007).

The initiative intends to create the use of advanced network services to promote the arising of new applications and tools that explore innovative mechanisms for the education in health, for distance cooperation in pre-diagnosis and for the remote evaluation of medical service data. (Santos, 2007).

3.2 Catarinense's Telemedicine Network Project - *Rede Catarinense de Telemedicina (RCTM)*

It consists of an infrastructure for distributing digital diagnosis services. Its objectives are to be used as sample for the formulation of a new technological model and standards that can be used by institutions of public health in other states of Brazil.

The project includes all of the functionalities of the assistance telemedicine, integrating on a single on-line platform the acquisition of data or images of exams, the dispatch of results, the request of patient internment and the decisions to be taken. (The Cyclops Group, Rede Catarinense de Telemedicina, 2007) (Sociedade Brasileira de Informática em Saúde, 2007).

3.3 Medical Digital Station - *Estação Digital Médica (EDM)*

The project links hospitals, universities, and other institutions through the telecommunications, informatics and other technological solutions.

The EDM objectives include fomenting the practice of medical care and education in distance at Brazil, optimizing the health system. Its actions, among others, include programs of professional improvement, technical training, information and prevention, second specialized medical opinion, clinical discussions (Estação Digital Médica, 2007).

3.4 ONCONET

The ONCONET includes a network of universities, and research medical institutions. It consists of a system developed in a client-server model using open software tools.

The ONCONET objective consists in supporting the medical practice in pediatric oncology. This project presents the electronic patient records for child cancer, distance education, and the data mining and statistics quantifications (Yuuji, 2005). It began its operations in 2004. Today it has 30 hospitals affiliated to the Brazilian Society of Pediatric Oncology (SOBOPE - *Sociedade Brasileira de Oncologia Pediátrica*).

3.5 HEALTHNET

HealthNet is a telemedicine application that interconnects institutions between São Paulo and Recife. More specifically, it gives support to the remote diagnosis and to the second medical opinion.

It is a project from NUTES (NUTES - Núcleo de Telessaúde, 2007).

Its objectives includes to improve the health services for distant and lacking areas, as well as to allow implanting a medical cooperation process among great specialist centers. (Rede Nacional de Ensino e Pesquisa, 2007).

3.6 Telemedicine Brazilian System - *Sistema Brasileiro de Telemedicina (SBTM)*

The SBTM objectives are to include all the functionalities of the assistance telemedicine in a unique on-line platform. It involves the creation of a national model of telemedicine with the definition of a Brazilian standardization for the section and the construction and refining a group of software technologies model for the implantation of a public network of telemedicine (The Cyclops Group, Sistema Brasileiro de Telemedicina – SBTM, 2007).

3.7 T@LEMED

This network intends to demonstrate the supply of telemedicine services in critical areas of Colombia and Brazil. The implementation of this e-health model is based on modern technologies of remote health as well as on medicine based on evidences. (Sachpazidis, 2006; Santos, 2007).

3.8 BH-Telessaúde (BHT)

Its goal consists in developing a low cost project that solves two problems on the Unique System of Health (SUS - *Sistema Único de Saúde*): physicians and other professionals precarious formation and the visible choke of the secondary attention. The network is implanted in 76 basic health units (UBS – *Unidade Básica de Saúde*), and focuses in the primary attention, by offering assistance support through second opinion, consultancy and discussion of clinical cases, besides activities of permanent education (Santos, 2007).

In order to facilitate the observation and the analysis of the results of this study, Table 1 was generated. In this way, it can be seen which networks practice distance education, which practice the asynchronous telemedicine, among other characteristics. The identified items with an 'X' indicate that the network has such characteristic; those identified with '?' mean that there was not found any information about the characteristic; the

lack of any identification type means that the network does not have such characteristic.

Terms used in Table 1: *Distance Education* - the telemedicine network provides the remote education for the health professionals involved. *Implanted and Operating* - the telemedicine network is firmly established. *Asynchronous Telemedicine* – type of telemedicine that not occurs in real time. That means that the response for a request of health care occurs later without synchronicity (Wootton, 2006). *Synchronous Telemedicine* – type of telemedicine that occurs in real time. That means that the entities involved in the process need to be online at the same time, with synchronicity (Wootton, 2006). *Inter-cities* – the telemedicine process occurs between different cities. *Telemedicine Service* - systems and processes that promote telemedicine services. *Infrastructure Network* - physical structures (equipments and communication resources) that allow and facilitate the use of telemedicine.

4 CONCLUSIONS

Analyzing the data presented on Table 1, it is verified the existence of two types of networks: infra-structure and telemedicine service. The first one can be characterized just by network infra-structure that allow and facilitate the use of telemedicine. This infra-structure is represented as LAN (Local Area Network) or WAN (Wide Area Network) equipments and communication resources. The second one can be characterized by systems and processes that promote telemedicine services and may come in computational systems format or as an agreement between institutions. Normally, the service networks runs over the network infra-structure. It's possible to verify that most of the Brazilian networks do the service type.

It can be observed that all of the networks use the asynchronous telemedicine and only four of them use the synchronous form. This difference may be justified due to the low cost and the easiness of use that the asynchronous telemedicine proportionate (Wootton, 2006). For example, there is no need to schedule its use and its participants do not need to be online at the same time.

The majority of the networks act in inter-cities. However there are few cases of telemedicine inside the own municipal district to eliminate socioeconomic barriers and/or to provide safety to the health professionals, for example, in the

Table 1: Main characteristics of the Brazilian telemedicine networks.

Characteristic	RUTE	RCTM	EDM	ONCONET	HEALTHNET	SBTM	T@LEMED	BHT
Distance Education	X		X	X	X	?	-	X
Implanted and Operating		X		X	X	-	X	X
Asynchronous Telemedicine	X	X	X	X	X	X	X	X
Synchronous Telemedicine	X	X	X	-	X	?	X	X
Inter-cities	X	X	X	X	X	X	X	X
Telemedicine Service	X	X	X	X	X	X	X	X
Infra-structure Network	X	X	-	-	-	-	-	-

penitentiary health care. The majority part of Brazilians perceives a 2440 dollars/year salary and don't have sufficient financial resources to pay the transportation to specialized medical centers.

Finally, there were not find inside these networks any project or initiative to regulate the remuneration of the professionals involved in the processes as well as, the payment for the services that were used. Those points are important because of professionals financial resources needs and it always exists costs of maintenance of equipments and systems. The outcome could be disinterest and discontinuity of services that aim to improve the people's health and of their communities. The cost-effectiveness of any of the eight telemedicine solutions was not found.

REFERENCES

- Neira, R. A. Q., Lopes, P. R. L., Jesus, A., et al., 2006. Construção e Avaliação de Ferramenta de Teleconsulta Assíncrona Multi-especialidades. In IX Congresso Brasileiro de Informática em Saúde. Infantsi AFC.
- Rede Nacional de Ensino e Pesquisa. (n.d.). Retrieved December 20, 2006, from <http://www.rnp.br/imprensa/2006/rel-60420.html>.
- Anderson, D., Ashworth, V., Battcock, A., et al., 2000. CA Telehealth Report. In The 5th international conference on the medical aspects of telemedicine. Canadian Society of Telehealth.
- PubMed Central. (n.d.). Retrieved June 7, 2007, from <http://www.pubmedcentral.nih.gov>.
- Google. (n.d.). Retrieved June 7, 2007, from <http://www.google.com>.
- Biblioteca virtual em saúde, DeCS - Descritores em Ciências da Saúde. Retrieved June 7, 2007, from <http://decs.bvs.br>.
- Rede Universitária de Telemedicina. (n.d.). Retrieved June 7, 2007, from <http://rute.rnp.br/sobre/rute/index.php>.
- Santos, A. F., Souza, C., Alves, H. J., Santos, S.F., et al., 2007. Telessaúde - Um Instrumento de Suporte Assistencial e Educação Permanente, UFMG. Belo Horizonte - Brazil.
- The Cyclops Group, Rede Catarinense de Telemedicina. (n.d.). Retrieved June 7, 2007, from <http://cyclops.telemedicina.ufsc.br/index.php?lang=en&page=projects&subpage=rctm>.
- Sociedade Brasileira de Informática em Saúde, Fornecimento digital de diagnósticos médicos amplia acesso ao serviço público de saúde. (n.d.). Retrieved June 7, 2007, from <http://www.sbis.org.br/site/site.dll/noticia?pagina=2&item=29>.
- Estação Digital Médica. (n.d.). Retrieved June 7, 2007, from <http://www.estacaodigitalmedica.org.br>.
- Yuuji, A. H., Tognoli, T. L., Nebel, A. M., Odone, F. V., Knörich, M. Z., Deus, R.L., 2005. Establishment of the Brazilian telehealth network for paediatric oncology. In Journal of telemedicine and telecare.
- NUTES - Núcleo de Telessaúde. (n.d.). Retrieved June 7, 2007, from <http://www.nutes.ufpe.br>.
- Rede Nacional de Ensino e Pesquisa, HealthNet: um sistema integrado de telediagnóstico e segunda opinião médica. (n.d.). Retrieved June 7, 2007, from <http://www.rnp.br/newsgen/0107/healthnet.html#ng-1>.
- The Cyclops Group, Sistema Brasileiro de Telemedicina - SBTM. (n.d.). Retrieved June 7, 2007, from <http://cyclops.telemedicina.ufsc.br/index.php?lang=en&page=sbtm&subpage=sbtm>.
- Sachpazidis, I., Kiefer, S., Selby, P., Ohl, R., Sakas, G., 2006. A medical network for teleconsultations in Brazil and Colombia. In Proceedings of the Second IASTED International Conference of TELEHEALTH.
- Wootton, R., Craig, J., 2006 Introduction to Telemedicine, Royal Society of Medicine Press. 2nd ed.

MEDICAL DIAGNOSIS ASSISTANT BASED ON CATEGORY RANKING

Fernando Ruiz-Rico, Jose Luis Vicedo and María-Consuelo Rubio-Sánchez

University of Alicante, Spain

frr@alu.ua.es, vicedo@dlsi.ua.es, mcrrs7@alu.ua.es

Keywords: Medical diagnosis, category ranking.

Abstract: This paper presents a real-world application for assisting medical diagnosis which relies on the exclusive use of machine learning techniques. We have automatically processed an extensive biomedical literature to train a categorization algorithm in order to provide it with the capability of matching symptoms to MeSH diseases descriptors. To interact with the classifier, we have developed a web interface so that professionals in medicine can easily get some help in their diagnostical decisions. We also demonstrate the effectiveness of this approach with a test set containing several hundreds of real clinical histories. A full operative version can be accessed on-line through the following site: www.dlsi.ua.es/omda/index.php.

1 INTRODUCTION

Text categorization consists of automatically assigning documents to pre-defined classes. It has been extensively applied to many fields such as search engines, spam filtering, etc. and in particular, some efforts have been focused on MEDLINE abstracts classification (Ibushi et al., 1999). However, as far as we are concerned, it has never been used to assist medical diagnosing by using the textual information provided by biomedical literature together with patient histories.

Every year, thousands of documents are added to the *National Library of Medicine* and the *National Institutes of Health* databases¹. Most of them have been manually indexed by assigning each document to one or several entries in a controlled vocabulary called MeSH² (Medical Subject Headings). The MeSH tree is a hierarchical structure of medical terms which are used to define the main subjects that a medical article or report is about. Due to the wide use of this terminology, we can find translations into several languages such as Portuguese and Spanish (i.e. DeCS³ - Health Science Descriptors). The diseases sub-tree not only defines on its own more than 4,000 pathological states, but also offers the chance to search for documented case reports related to each of them.

¹<http://www.pubmed.gov>

²<http://www.nlm.nih.gov/mesh>

³<http://decs.bvs.br/I/homepagei.htm>

Our proposal tries to estimate a ranked list of diagnoses from a patient history. To tackle this problem, we have selected an existing categorization algorithm, and we have trained it using the textual information provided by lots of previously reported cases. This way, a detailed symptomatic description is sufficient to obtain a list of possible diseases, along with an estimation of probabilities.

We have not used binary decisions from binary categorization methods, since they might leave some interesting MeSH entries out, which should probably be taken into consideration. Instead, we have chosen a category ranking algorithm to obtain an ordered list of all possible diagnoses so that the user can finally decide which of them better suits the clinical history.

In this paper, first of all, we will explain the way we have developed our experiments, including a full description of the sources and methods used to get both training and test data. Secondly, we will provide an example of a patient history and both the expected and provided diagnoses. We will finish by showing and commenting several evaluation results on.

2 PROCEDURES

We have extracted the training data from the *PubMed* database¹ by selecting every case reports on diseases written in English including abstract and related to humans beings. These documents were retrieved by using the “diseases category[MAJR]” query, where

[MAJR] stands for “MeSH Major Topic”, asking the system for retrieving only documents whose subject is mainly a disease. The query provided us with 483,726 documents⁴ that we downloaded by sending them to a file in MEDLINE format. We automatically processed that file to obtain the titles and abstracts with their corresponding MeSH topics. This led us to 4,024 classes with at least one training sample each.

With respect to the test set, we have used 400 medical histories from the School of Medicine of the University of Pittsburgh (Department of Pathology⁵). Although, so far the web page contains 500 histories⁴, not all of them are suitable for our purposes. There are some which do not provide a concrete diagnosis but only a discussion about the case, and some others do not have a direct matching to the MeSH tree. We downloaded the HTML cases and afterwards we converted them to text format by using from each document the title and all the clinical history, including radiological findings, gross and microscopic descriptions, etc. To get the expected output, we extracted the top level MeSH diseases categories corresponding to the diagnoses given on the titles of the “final diagnosis” files (dx.html).

To select a proper ranking algorithm, we have looked up the most suitable one through several decades of literature about text classification and category ranking. We have chosen the Sum of Weights (SOW) approach (Ruiz-Rico et al., 2006), that is more suitable than the rest for its simplicity, efficiency, accuracy and incremental training capacity. Since medical databases are frequently updated and they also grow continuously, we have preferred using a fast and unattended approach that lets us perform updates easily with no substantial performance degradation after incrementing the number of categories or training samples. The restrictive complexity of other classifiers such as SVM could deviate to an intractable problem, as stated by (Ruch, 2005).

To evaluate how worth our suggestion is, we have measured accuracy through three common ranking performance measures (Ruiz-Rico et al., 2006): Precision at recall = 0 ($P_{r=0}$), mean average precision (AvgP) and Precision/Recall break even point (BEP). Sometimes, only one diagnosis is valid for a particular patient. In these cases, $P_{r=0}$ let us quantify the mistaken answers, since it indicates the proportion of correct topics given at the top ranked position. To know about the quality of the full ranking list, we use the AvgP, since it goes down the arranged list averaging precision until all possible answers are covered. BEP is the value where precision equals recall, that

⁴Data obtained on February 14th 2007

⁵<http://path.upmc.edu/cases>

Order	Diseases	Probability	Expand
1	Nervous System Diseases	100	+
2	Neoplasms	93	+
3	Pathological Conditions, Signs and Symptoms	61	+
4	Congenital, Hereditary, and Neonatal Diseases and Abnormalities	25	+
5	Endocrine System Diseases	23	+
...

Figure 1: Example of the first level of a hierarchical diagnosis.

Nervous System Diseases			
Nervous System Neoplasms			
Central Nervous System Neoplasms			
Order	Diseases	Probability	Expand
1	Brain Neoplasms	100	+
2	Central Nervous System Cysts	44	+
3	Meningeal Neoplasms	13	+
4	Spinal Cord Neoplasms	3	+

Order	Diseases	Probability
1	Craniopharyngioma	100
2	Cysts	63
3	Pituitary Neoplasms	59
4	Central Nervous System Cysts	48
5	Headache	46
6	Xanthomatosis	44
...

Figure 2: Output example after manual expansion of high ranked topics (up) and by selecting the flat diagnosis mode (down).

is, when we consider the maximum number of relevant topics as a threshold. To follow the same procedure as (Joachims, 1998), the performance evaluation has been computed over the top diseases level.

2.1 Availability and Requirements

No special hardware nor software is necessary to interact with the assistant. Just an Internet connection and a standard browser are enough to access on-line through the following site: www.dlsi.ua.es/omda/index.php.

By using a web interface and by presenting results in text format, we allow users to access from many types of portable devices (laptops, PDA's, etc.). Moreover, they will always have available the latest version, with no need of installing specific applications nor software updates.

3 AN EXAMPLE

One of the 400 histories included in the test set looks as follows:

Case 177 – Headaches, Lethargy and a Sellar/Suprasellar Mass

A 16 year old female presented with two months of progressively worsening headaches, lethargy and visual disturbances. Her past medical history included developmental delay, shunted hydrocephalus, and tethered cord release ...

The final diagnosis expected for this clinical history is: “Rathke’s Cleft Cyst”, which is a synonym of the preferred term “Central Nervous System Cysts”. Translating this into one or several of the 23 top MeSH diseases categories we are lead to the following entries:

- Neoplasms
- Nervous System Diseases
- Congenital, Hereditary, and Neonatal Diseases and Abnormalities.

In hierarchical mode, our approach provides automatically a first categorization level with expanding possibilities as shown in figure 1. We provide navigation capabilities to allow the user to go down the tree by selecting different branches, depending on the given probabilities and his/her own criteria. Moreover, a flat diagnosis mode can be activated to directly obtain a ranked list of all possible diseases, as shown in figure 2.

After an individual evaluation of this case, we have obtained the following values: $P_{r=0} = 1$, $AvgP = 0.92$, and $BEP = 0.67$, since the right topics in figure 1 are given at positions 1, 2 and 4.

4 RESULTS

Last row in table 1 shows the performance measures calculated for each medical history and its diagnosis, averaged afterwards across all the 400 decisions. $P_{r=0}$ indicates that we get 69% of the histories correctly diagnosed with the top ranked MeSH entry. $AvgP$ value means that the rest of the list also contains quite valid topics, since it reaches a value of 73%.

First row in table 1 provides a comparison between SVM (Joachims, 1998) and sum of weights (Ruiz-Rico et al., 2006) algorithms using the well known OHSUMED evaluation benchmark. Even using a training and test set containing different document types, BEP indicates that the performance is not far away from that achieved in text classification tasks, meaning that category ranking can also be effectively applied to our scenario.

Table 1: Averaged performance for both text categorization and diagnosis.

Corpus	Algor.	$P_{r=0}$	AvgP	BEP
OHSUMED	SVM	-	-	0.66
	SOW	-	-	0.71
Case reports and patient histories	SOW	0.69	0.73	0.62

5 CONCLUSIONS

We believe that category ranking algorithms may provide a useful tool to help in medical diagnoses from clinical histories. Although the output of the categorization process should not be directly taken to diagnose a disease without a previous review, the accuracy achieved could be good enough to assist human experts. Moreover, our implementation demonstrates that both training and classification processes are very fast, leading to an accessible and easy upgradable system.

REFERENCES

- Ibushi, K., Collier, N., and Tsujii, J. (1999). Classification of medline abstracts. *Genome Informatics, volume 10*, pages 290–291.
- Joachims, T. (1998). Text categorization with support vector machines: learning with many relevant features. In *Proceedings of ECML-98, 10th European Conference on Machine Learning*, pages 137–142.
- Ruch, P. (2005). Automatic assignment of biomedical categories: toward a generic approach. *Bioinformatics, volume 22 no. 6 2006*, pages 658–664.
- Ruiz-Rico, F., Vicedo, J. L., and Rubio-Sánchez, M.-C. (2006). Newpar: an automatic feature selection and weighting schema for category ranking. In *Proceedings of DocEng-06, 6th ACM symposium on Document engineering*, pages 128–137.

TOWARD MOBILE HEALTHCARE SERVICES BY USING EVERYDAY MOBILE PHONES

Akio Sashima, Yutaka Inoue, Takeshi Ikeda

Tomohisa Yamashita, Masayuki Ohta and Koichi Kurumatani

National Institute of Advanced Industrial Science and Technology (AIST)/CREST, JST

2-41-6, Aomi, Koto-ku, Tokyo 135-0064, Japan

{sashima-akio, yutaka.inoue, ikeda-takeshi, tomohisa.yamashita, m-ohta, k.kurumatani}@aist.go.jp

Keywords: Healthcare, mobile phone, wireless biological sensor, electrocardiograph, skin thermometer, and 3-axis accelerometer.

Abstract: More than two billions people use mobile phones in the world of today. The mobile phones are not just potable telephones but portable computers which have WWW browsers with multi-task OS. In this paper, we specifically examine the possibility of mobile healthcare services by using everyday mobile phones. We describe a prototype system of the mobile healthcare services. It consists of the following components which cooperatively work with mobile phones: wireless biological sensors, mobile sensor routers, and sensor middleware. The service of the system aims to maintain and improve user's condition by monitoring one's biological sensing data, such as ECG, skin temperature, and body movement.

1 INTRODUCTION

Aged people have rapidly increased in most of advanced nations. In such matured societies, mobile and ubiquitous computing technologies which support healthcare services are important. In addition, more than two billions people use mobile phones in the world of today. As popularity of the mobile phones have increased, healthcare services using the mobile phones have drawn much attention from researches (Leijdekkers and Gay, 2006)(Oliver and Flores-Mangas, 2006). Although mobile services have been studied in mobile computing for many years, researches using everyday mobile phones with sensor technologies are very few. Recently, several researches which intended to detect user's contexts by using sensor devices embedded in a mobile phones are proposed in the field of ubiquitous computing (Lester et al., 2006)(Kawahara et al., 2007).

This research specifically examines the possibility of ubiquitous healthcare service by using everyday mobile phones. Although current mobile phones are regarded as portable computers which have various computational facilities, they have limited computational powers. Smartphones which have sufficient computational powers are proposed for several years. However, they have got little popularity for ordinary people.

In this paper, we describe a prototype for the mobile healthcare system by using everyday mobile phones and wireless biological sensors. To compensate the limited abilities of the everyday mobile phones, the prototype system consists of the following components which cooperatively work with mobile phones: *wireless biological sensors*, *mobile sensor router*, and *sensor middleware*. The system provides the healthcare service to maintain and improve user's conditions by monitoring one's biological information, such as heartbeat, posture, and movement. The system uses wearable wireless sensors, e.g., electrocardiograph, skin thermometer, and 3-axis accelerometer.

2 HEALTHCARE SERVICES USING MOBILE PHONES

To develop mobile healthcare systems, we have considered the following two healthcare scenarios.

Self monitoring services for physical exercises

When people are getting physical exercises, such as doing aerobics, it is important to monitor their current physical conditions for keeping appropriate strengths of the exercises. If their everyday mobile phones can show the conditions

without special self monitoring devices, they can easily control the strength of the exercises. Thus, self monitoring for physical exercises using their mobile phones is an important mobile healthcare service.

Remote and self healthcare services for aged people

Aged people and patients of heart diseases require preparing for unexpected health troubles. Their physical conditions should be remotely monitored by their doctors and family members. Thus, monitoring their current physical conditions, such as abnormality of ECG, is an important healthcare service using their mobile phones.

Under the vision of the above scenarios, we have developed a prototype system of the mobile healthcare services. Figure 1 shows an outline of the mobile healthcare system.

The system monitors user's physical conditions using an everyday mobile phone¹ and a wireless biological sensor². To sense electrocardiograph correctly, it requires to be attached to user's chest by sticking electrodes of the sensor on tight with a peel-off sticker. Once it is attached to user's chest, it can detect the inclination and movement of the upper half of user's body by the 3-axis accelerometer.

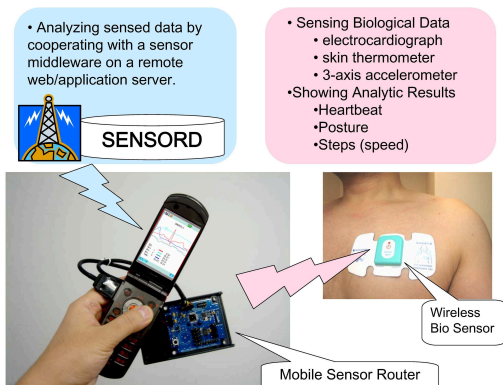


Figure 1: Healthcare services by using everyday mobile phones.

The system consists of *wireless biological sensors*, *mobile sensor router* and *sensor middleware*, and provides following facilities: 1) communicating to a wireless bio sensor via a mobile sensor router attached to user's mobile phone, 2) analyzing the sensed

¹NTT-DoCoMo FOMA N903i: <http://www.nttdocomo.co.jp/english/product/foma/903i/n903i/index.html>

²MESI RF-ECG: http://www.natureinterface.com/index.files/EK_panf.pdf (Japanese)

data by cooperating with sensor middleware on a remote server to capture one's conditions, and 3) providing personalized information for the user using GUI on the mobile phone.

Most of mobile phones popularly used in Japan can execute Java programs downloaded from Web sites. Thus, service processes in mobile phones are implemented by multiple threads running on the Java runtime environment. The processes includes collecting surrounding sensor data via a mobile sensor router, sending sensed data to a remote sensor middleware for analyzing them, and graphically showing analytic results of the sensed data.

The wireless biological sensor for sensing user's conditions is a small wearable sensor which integrates three kinds of built-in biological sensors: electrocardiograph, skin thermometer, and 3-axis accelerometer (see Figure 2).



Figure 2: Wireless bio sensor.

Figure 3 shows a mobile sensor router attached to a serial port (UART) of the mobile phones. The router has three main functions: *communicating with sensor networks*, *reducing the number of sensed data*, and *communicating with a mobile phone*.




Figure 3: Mobile Sensor Router.

Although the router is able to receive full rates of sensor data stream from the wireless bio sensor, the

mobile phone is not able to receive such amount of data stream. Thus, the router sends not raw data received from the sensors but reduced data of them at the router. It repeats the following processes to reduce the data:

1. clearing a sample buffer;
2. storing obtained sensing data for a certain period of time in the sample buffer;
3. calculating a representative value (e.g., average, maximum, minimum, latest) of the stored data in the sample buffer;

The size of the sample buffer is configurable. If the size is N , the process reduces by one N th the number of sensed data.

To analyze sensed data for the healthcare service, sensor middleware, called SENSORD (Sashima et al., 2006), performs several signal processing and classification processes based on machine learning techniques, such as Discrete Fourier Transform (DFT), support vector machine (SVM), nearest neighbor learning. Using such algorithms, the sensed data is statistically analyzed or classified to some qualitative categories. The results of the analysis are used for the service to create suitable contents, such as HTML documents to be shown by the mobile phones.

3 USER INTERFACE OF THE SERVICE

User interface of the prototype system has the following three modes: *configuration*, *graph*, and *monitor*. In the rest of this section, we describe outlines of the service at the each mode.

3.1 Configuration Mode

In “configuration mode”, users can configure several parameters of the services, such as transmission rates and sensor device id. Considering limited computational power of a mobile phone, we have decided a best transmission rate and a type to reduce the data for each sensor. We have also decided best transmission rates of the router for each sensor. Table 1 shows a default configuration of the sensors and the router for the healthcare service. Notice that all transmission rates of the router are smaller than the rates of sensors. This means that the data sent to the mobile phone is not raw sensor data but reduced sensor data.

In our implementation, the default configuration shows best results about processing the data. For example, when the transmission rates are more than the

configuration values, the response time of the system is slow down, and showing classification results becomes delayed for several dozen seconds.

Table 1: Default configuration of the data transmission for the mobile healthcare service.

Sensing Data	Transmission Rate		Reduction Type
	From Sensor to Router	From Router to Phone	
Skin temp.	204 Hz	2 Hz	latest
ECG.	204 Hz	8 Hz	maximum
3-axis acc.	204 Hz	8 Hz	average

3.2 Graph Mode

In “graph mode”, the system graphically shows user’s biological statuses, such as ECG, in real time (see left side of Figure 4). The user can know their current physical conditions graphically, and control the strength of the exercise. In this mode, the received data from the router is directly shown on the display without data processing by the sensor middleware.

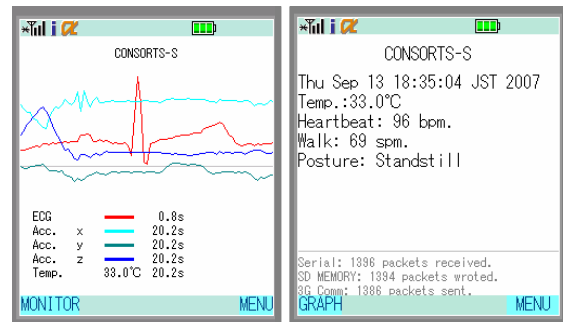


Figure 4: GUI of the healthcare service: graph mode (left side) and monitor mode (right side).

3.3 Monitor Mode

In “monitor mode”, the system shows analytical results of user’s sensing data. The results include numerical values representing the user’s conditions, such as heartbeat. It also categorizes user’s conditions into some qualitative statuses: running, walking, standing-still, etc. (see left side of Figure 4). The following analytical results are shown in this mode.

3.3.1 Posture Recognition

This process analyzes the sensed data to categorize user's postures into three categories: *standing-still*, *facing downward*, and *facing upward*. The postures are determined by calculating inclination of the upper half of user's body based on the values of the 3-axis accelerometer.

3.3.2 Movement Recognition

This process analyzes the sensed data to monitor user's movements, such as step speeds. The steps are recognized by calculating Discrete Fourier Transform (DFT) of finite length time series of sensed data of the y-axis accelerometer. In the current implementation, the length of the time series is 8 seconds (64 samples). When the middleware is asked to calculate the current steps, it retrieves the latest 64 samples from the data storage of the middleware, and calculate it. Then, it classifies user's statuses into three categories: *staying*, *walking*, *running* based on the results of DFT. When user's status is "walking" or "running," it also shows an average speed of user's movements as Steps Per Minute (SPM). The user can control his/her movements based on the information.

3.3.3 Heartbeat Monitoring

This process analyzes the sensed data to monitor the user's current status of heartbeat as beat per minute (BPM). The heartbeats are calculated by DFT of finite length time series of sensed data of the electrocardiograph. In the current implementation, the length of the time series is about 16 seconds (128 samples). When the sensor middleware is asked to calculate the heart beat, it retrieves the latest 128 samples from the data storage of the middleware, and calculates it.

As a prototype implementation of the remote healthcare scenario, emergency messages (email) are sent to a doctor and family members when the abnormality of the heartbeat and the posture are recognized.

3.3.4 Skin Temperature Monitoring

The biological sensor is able to monitor user's skin temperature. If the temperature is lower than 31 °C, the system recognizes that the sensor is not attached to a human body.

ing not smart-phones but mobile phones with attachments, namely mobile sensor routers. Smartphones are still special devices for limited types of persons, such as businesspersons. By using the existing mobile phones, we have aimed to enable ordinary people to always collect and analyze their health information derived from wireless biological sensors. We have confirmed that cooperations between the current mobile phones, the mobile sensor routers, and the sensor middleware is able to provide the mobile healthcare service for the ordinary people. In future work, we plan to examine possibility of a general-purpose service platform for various mobile healthcare services by using everyday mobile phones.

REFERENCES

- Kawahara, Y., Kurasawa, H., and Morikawa, H. (2007). Recognizing user context using mobile handsets with acceleration sensors. In *IEEE International Conference on Portable Information Devices. PORTABLE07*, pages 1–5.
- Leijdekkers, P. and Gay, V. (2006). Personal heart monitoring and rehabilitation system using smart phones. In *ICMB '06: Proceedings of the International Conference on Mobile Business*, page 29, Washington, DC, USA. IEEE Computer Society.
- Lester, J., Choudhury, T., and Borriello, G. (2006). A practical approach to recognizing physical activities. In *Proceedings of The Fourth International Conference on Pervasive Computing (PERVASIVE 2006)*, pages 1–16.
- Oliver, N. and Flores-Mangas, F. (2006). Healthgear: A real-time wearable system for monitoring and analyzing physiological signals. In *Proceedings of the International Workshop on Wearable and Implantable Body Sensor Networks (BSN'06)*, pages 61–64, Washington, DC, USA. IEEE Computer Society.
- Sashima, A., Inoue, Y., and Kurumatani, K. (2006). Spatio-temporal sensor data management for context-aware services: designing sensor-event driven service coordination middleware. In *ADPUC '06: Proceedings of the 1st international workshop on Advanced data processing in ubiquitous computing (ADPUC 2006)*, page 4, New York, NY, USA. ACM Press.

4 CONCLUSIONS

In this paper, we have focused on implementing a prototype system of the mobile healthcare service by us-

ARE BETTER FEATURE SELECTION METHODS ACTUALLY BETTER?

Discussion, Reasoning and Examples

Petr Somol, Jana Novovičová

*Inst. of Information Theory and Automation, Pattern Recognition Dept., Pod vodárenskou věží 4, Prague 8, 18208, Czech Republic
somol@utia.cas.cz, novovic@utia.cas.cz*

Pavel Pudil

*Prague University of Economics, Faculty of Management, Jarošovská 1117/II, Jindřichův Hradec, 37701, Czech Republic
pudil@fm.vse.cz*

Keywords: Feature Selection, Subset Search, Search Methods, Performance Estimation, Classification Accuracy.

Abstract: One of the hot topics discussed recently in relation to pattern recognition techniques is the question of actual performance of modern feature selection methods. Feature selection has been a highly active area of research in recent years due to its potential to improve both the performance and economy of automatic decision systems in various applicational fields, with medical diagnosis being among the most prominent. Feature selection may also improve the performance of classifiers learned from limited data, or contribute to model interpretability. The number of available methods and methodologies has grown rapidly while promising important improvements. Yet recently many authors put this development in question, claiming that simpler older tools are actually better than complex modern ones – which, despite promises, are claimed to actually fail in real-world applications. We investigate this question, show several illustrative examples and draw several conclusions and recommendations regarding feature selection methods' expectable performance.

1 INTRODUCTION

Dimensionality reduction (DR) concerns with the task of finding low dimensional representation for high dimensional data. DR is an important step in data preprocessing in pattern recognition applications. It is sometimes the case that such tasks as classification of the data represented by so called feature vectors, can be carried out in the reduced space more accurately than in the original space. There are two main ways of doing DR depending on the resulting features: DR by *feature selection* (FS) and DR by *feature extraction* (FE). The FS approach does not attempt to generate new features, but tries to select the “best” ones from the original set of features. The FE approach defines a new feature vector space in which each new feature is obtained by transformations of the original features. FS leads to savings in measurement cost and the selected features retain their original physical interpretation, important e.g., in medical applications. On the other hand, transformed features generated by FE may provide a better discriminative ability than the best subset of given features, but these new features may not have a clear physical mean-

ing. A typical feature selection process consists of four basic steps: *feature subset selection*, *feature subset evaluation*, *stopping criterion*, and *result validation*. Based on the *selection criterion* choice, feature selection methods may roughly be divided into three types: the *filter* (Yu and Liu, 2003; Dash et al., 2002), the *wrapper* (Kohavi and John, 1997) and the *hybrid* (Das, 2001; Sebban and Nock, 2002; Somol et al., 2006). The filter model relies on general characteristics of the data to evaluate and select feature subsets without involving any mining algorithm. The wrapper model requires one predetermined mining algorithm and uses its performance as the evaluation criterion. It attempts to find features better suited to the mining algorithm aiming to improve mining performance. This approach tends to be more computationally expensive than the filter approach. The hybrid model attempts to take advantage of the two approaches by exploiting their different evaluation criteria in different search stages. The hybrid approach is recently proposed to handle large datasets.

In recent years FS seems to have become a topic attracting an increasing number of researchers. Among the possible reasons the main one is certainly

the importance of FS (or FE) as an inherent part of classification or modelling system design. Another reason, however, may be the relatively easy accessibility of the topic to the general research community. Apparently, many papers have been published in which any substantial advance is difficult to identify. One is tempted to say that the more papers on FS that are published, the fewer important contributions actually appear.

Certainly many key questions remain unanswered and key problems remain unsolved to satisfaction. For example, not enough is known about error bounds of many popular feature selection criteria, especially about their relation to classifier generalization performance. Despite the huge number of methods in existence, it is still a very hard problem to perform FS satisfactorily, e.g., in the context of gene expression data, with enormous dimensionality and very few samples. Similarly, in text categorization the standard way of FS is to completely omit context information and to resort to much more limited FS based on individual feature evaluation. In medicine these problems tend to become emphasized, as the available datasets are often incomplete (missing feature values in sample vectors), continuous and categorical data is to be treated at once, and the notion of feature itself may be difficult to interpret.

Among many criticisms of the current FS development there is one targeted specifically at the effort of finding more effective search methods, capable of yielding results closer to optimum with respect to some chosen criterion. The key argument against such methods is their alleged tendency to “over-select” features, or to find feature subsets fitted too tightly to training data, what degrades generalization. In other words, more search-effective methods are supposed to cause a similar unwanted effect as classifier over-training. Indeed, this is a serious problem that requires attention.

In recent literature the problem of “over-effective” FS has been addressed many times (Reunanen, 2003; Raudys, 2006). Yet, the effort to point out the problem (which seems to have been ignored, or at least insufficiently addressed before) now seems to have led to the other extreme notion of claiming that most of FS method development is actually contra-productive. This is, that older methods are actually superior to newer methods, mainly due to better over-fitting resistance.

The purpose of this paper is to discuss the issue of comparing actual FS methods’ performance and to show experimentally what impact of the more effective search in newer methods can be expected.

1.1 FS Methods Overview

Before giving overview of the main methods to be discussed further we should note that it is not generally agreed in literature what the term “FS method” does actually describe. The term “FS method” is equally often used to refer to a) the complete framework that includes everything needed to select features, or b) the combination of search procedure and criterion or c) just the bare search procedure. In the following we will focus mainly on comparing the standard search procedures, which are not criterion- or classifier dependent. The widely known representatives of such “FS methods” are:

- Best Individual Features (IB) (Jain et al., 2000),
- Sequential Forward Selection (SFS), Sequential Backward Selection (SBS), (Devijver and Kittler, 1982),
- “Plus l -take away r ” Selection (+L-R) (Devijver and Kittler, 1982),
- Sequential Forward Floating Selection (SFFS), Sequential Backward Floating Selection (SFBS) (Pudil et al., 1994),
- Oscillating Search (OS) (Somol and Pudil, 2000).

Many other methods exist (in all senses of the term “FS Method”), among others generalized versions of the ones listed above, various randomized methods, methods related to use of specific tools (FS for Support Vector Machines, FS for Neural Networks) etc. For overview see, e.g., (Jain et al., 2000; Liu and Yu, 2005). The selection of methods we are going to investigate is motivated by their interchangeability – any one of them can be used with the same given criterion, data and classifier. This makes experimental comparison easier.

2 PERFORMANCE ESTIMATION PROBLEM

FS methods comparison seems to be understood ambiguously as well. It is very different whether we compare concrete method properties or the final classifier performance determined by use of particular methods under particular settings. Certainly, final classifier performance is the ultimate quality measure. However, misleading conclusions about FS may be easily drawn when evaluating nothing else, as classifier performance depends on many more different aspects than just the actual FS method used. Nevertheless, in the following we will adapt classifier accuracy as the main means of FS method assessment.

There seems to be a general agreement in the literature that wrapper-based FS enables creation of more accurate classifiers than filter-based FS. This claim is nevertheless to be taken with caution, while using actual classifier accuracy as the FS criterion in wrapper-based FS may lead to the very negative effects mentioned above (overtraining). At the same time the weaker relation of filter-based FS criterion functions to particular classifier accuracy may help better generalization. But these effects can be hardly judged before the building of classification system has actually been accomplished.

In the following we will focus only on wrapper-based FS. Wrapper-based FS can be accomplished (and accordingly its effect can be evaluated) using one of the following methods:

- Re-substitution – In each step of the FS algorithm all data is used both for classifier training and testing. This has been shown to produce strongly optimistically biased results.
- Data split – In each step of the FS algorithm the same part of the data is used for classifier training and the other part for testing. This is the correct way of classifier performance estimation, yet it is often not feasible due to insufficient size of available data or due to inability to prevent bias caused by unevenly distributed data in the dataset (e.g., it may be difficult to ensure that with two-modal data distribution the training set won't by coincidence represent one mode and the testing set the other mode)
- 1-Tier Cross-Validation (CV) – Data is split to several parts. Then in each FS step a series of tests is performed, with all but one data part used for classifier training and the remaining part used for testing. The average classifier performance is then considered to be the result of FS criterion evaluation. Because in each test a different part of data is used for testing, all data is eventually utilized, without actually testing the classifier on the same data on which it had been trained. This is significantly better than re-substitution, yet it still produces optimistically biased results because all data is actually used to govern the FS process.
- Leave-one-out – can be considered a special case of 1-Tier CV with the finest data split granularity, thus the number of tests in one FS step is equal to the number of samples while in each test all but one sample are used for training with the one sample used for testing. This is computationally more expensive, better utilizes the data, but suffers the same problem of optimistic bias.
- 2-Tier CV – Defined to enable less biased esti-

mation of final classifier performance than it is possible with 1-Tier CV. The data is split to several parts, FS is then performed repeatedly in 1-Tier CV manner on all but one part, which is eventually used for classifier accuracy estimation. This process yields a sequence of possibly different feature subsets, thus it can be used only for assessment of FS method effectivity and not for actual determination of the best subset. The average classifier performance on independent test data parts is then considered to be the measure of FS method quality. This is computationally demanding.

In our experiments we accept 2-Tier CV as satisfactory for the purpose of FS methods performance evaluation and comparison. Due to the fact that 2-Tier CV yields a series of possibly different feature subsets, we define an additional measure to be called *consistency*, that expresses the stability, or robustness of FS method with respect to various data splits.

Definition: Let $Y = \{f_1, f_2, \dots, f_{|Y|}\}$ be the set of all features and let $\mathcal{S} = \{S_1, S_2, \dots, S_n\}$ be a system of $n > 1$ feature subsets $S_j = \{f_i | i = 1, \dots, d_j, f_i \in Y, d_j \in \langle 1, |Y| \rangle\}$, $j = 1, \dots, n$. Denote \mathcal{F}_f the system of subsets in \mathcal{S} containing feature f , i.e.,

$$\mathcal{F}_f = \{S | S \in \mathcal{S}, f \in S\}. \quad (1)$$

Let F_f be the number of subsets in \mathcal{F}_f and X the subset of Y representing all features that appear anywhere in system \mathcal{S} , i.e.,

$$X = \{f | f \in Y, F_f > 0\}. \quad (2)$$

Then the *consistency* $C(\mathcal{S})$ of feature subsets in system \mathcal{S} is defined as:

$$C(\mathcal{S}) = \frac{1}{|X|} \sum_{f \in X} \frac{F_f - 1}{n - 1}. \quad (3)$$

Properties of $C(\mathcal{S})$:

1. $0 \leq C(\mathcal{S}) \leq 1$.
2. $C(\mathcal{S}) = 0$ if and only if all subsets in \mathcal{S} are disjoint from each other.
3. $C(\mathcal{S}) = 1$ if and only if all subsets in \mathcal{S} are identical.

The higher the value, the more similar are the subsets in system to each other. For $C(\mathcal{S}) \approx 0.5$ on average each feature present in \mathcal{S} appears in about half of all subsets. When comparing FS methods, higher *consistency* of subsets produced during 2-Tier CV is clearly advantageous. However, it should be considered a complementary measure only as it does not have any straight relation to the key measure of classifier generalization ability.

Remark: In experiments, if the best performing FS method also produces feature subsets with high *consistency*, its superiority can be assumed well founded.

3 EXPERIMENTS

To illustrate the differences between simpler and more complex FS methods we have collected experimental results under various settings: for two different classifiers, three FS search algorithms and eight datasets with dimensionalities ranging from 13 to 65 and number of classes ranging from 2 to 6. We used 3 different mammogram datasets as well as wine and wave datasets from UCI Repository (Asuncion and Newman, 2007), satellite image dataset from ELENA database (ftp.dice.ucl.ac.be), speech data from British Telecom and sonar data (Gorman and Sejnowski, 1988). For details see Tables 1 to 8.

Note that the choice of classifier and/or FS setup may not be optimal for each dataset, thus the reported results may be inferior to results reported in the literature; the purpose of our experiments is mutual comparison of FS methods only. All experiments have been done with 10-fold Cross-Validation used to split the data into training and testing parts (to be denoted “Outer CV” in the following), while the training parts have been further split by means of another 10-fold CV into actual training and validation parts for the purpose of feature selection and classifier training (to be denoted “Inner CV”).

The application of SFS and SFFS was straightforward. The OS algorithm as the most flexible procedure has been used in two set-ups: slower randomized version and faster deterministic version. In both cases the *cycle depth* set to 1 [see (Somol and Pudil, 2000) for details]. The randomized version, denoted in the following as OS(1,r3), is called repeatedly with random initialization as long as no improvement has been found in last 3 runs. The deterministic version, denoted as OS(1,IB) in the following, is initialized by means of Individually Best (IB) feature selection.

The problem of determining optimal feature subset size was solved in all experiments by brute force. All algorithms were applied repeatedly for all possible feature sizes whenever needed. The final result has been determined as that with the highest classification accuracy (and lowest subset size in case of ties).

3.1 Notes on Obtained Results

All tables clearly show that more modern methods are capable of finding criterion values closer to optimum – see column Inner-CV in each table.

The effect pointed out by Reunanen (Reunanen, 2003) of the simple SFS outperforming all more complex procedures (regarding the ability to generalize) takes place in Table 4, column Outer-CV, with Gaus-

sian classifier. Note the low *consistency* in this case. Conversely, Table 2 shows no less outstanding performance of OS with 3-Nearest Neighbor classifier (3-NN) with better *consistency* and smallest subsets found, while Table 3 shows top performance of SFFS with both Gaussian and 3-NN classifiers. Although it is impossible to draw decisive conclusions from the limited set of experiments, it should be of interest to extract some statistics (all on independent test data – results in the column Outer-CV):

- Best result among FS methods for each given classifier: SFS 11×, SFFS 17×, OS 11×.
- Best achieved overall classification accuracy for each dataset: SFS 1×, SFFS 5×, OS 2×.

Average classifier accuracies:

- Gaussian: SFS 0.652, SFFS 0.672, OS 0.663.
- 1-NN: SFS 0.361, SFFS 0.361, OS 0.349.
- 3-NN: SFS 0.762, SFFS 0.774, OS 0.765.

4 DISCUSSION AND CONCLUSIONS

With respect to FS we can distinguish the following entities which all affect the resulting classification performance: search algorithms, stopping criteria, feature subset evaluation criteria, data and classifier. The impact of the FS process on the final classifier performance (with our interest targeted naturally at its generalization performance, i.e., its ability to classify previously unknown data) depends on all of these entities.

When comparing pure search algorithms as such, then there is enough ground (both theoretical and experimental) to claim that newer, often more complex methods, have better potential of finding better solutions. This often follows directly from the method definition, as newer methods are often defined to improve some particular weakness of older ones. (Unlike IB, SFS takes into account inter-feature dependencies. Unlike SFS, +L-R does not suffer the nesting problem. Unlike +L-R, Floating Search does not depend on pre-specified user parameters. Unlike Floating Search, OS may avoid local extremes by means of randomized initialization etc.) Better solution, however, means in this context merely being closer to optimum with respect to the adopted criterion. This may not tell much about final classifier quality, while criterion choice has proved to be a considerable problem in itself. Vast majority of practically used criteria have only insufficient relation to correct classification rate,

Table 1: Classification performance as result of wrapper-based Feature Selection on wine data.

<i>Wine data: 13 features, 3 classes containing 59, 71 and 48 samples, UCI Repository</i>									
Classifier	FS Method	Inner 10-f. CV		Outer 10-f. CV		Subset Size		Consistency	Run Time h:m:s.ss
		Mean	St.Dv.	Mean	St.Dv.	Mean	St.Dv.		
Gaussian	SFS	0.599	0.017	0.513	0.086	3.1	1.221	0.272	00:00:00.54
	SFFS	0.634	0.029	0.607	0.099	3.9	1.136	0.370	00:00:02.99
	OS(1,r3)	0.651	0.024	0.643	0.093	3.1	0.539	0.463	00:00:34.30
1-NN scaled	SFS	0.355	0.071	0.350	0.064	1	0	1	00:00:00.98
	SFFS	0.358	0.073	0.350	0.064	1	0	1	00:00:02.27
	OS(1,r3)	0.285	0.048	0.269	0.014	1.1	0.3	0.5	00:00:15.61
3-NN scaled	SFS	0.983	0.005	0.960	0.037	6.5	1.118	0.545	00:00:01.10
	SFFS	0.986	0.005	0.965	0.039	6.6	0.917	0.5	00:00:03.75
	OS(1,r3)	0.986	0.004	0.955	0.035	6.1	0.7	0.505	00:00:45.68

Table 2: Classification performance as result of wrapper-based Feature Selection on mammogram data.

<i>Mammogram data, 65 features, 2 classes containing 57 (benign) and 29 (malignant) samples, UCI Rep.</i>									
Classifier	FS Method	Inner 10-f. CV		Outer 10-f. CV		Subset Size		Consistency	Run Time h:m:s.ss
		Mean	St.Dv.	Mean	St.Dv.	Mean	St.Dv.		
Gaussian	SFS	0.792	0.028	0.609	0.101	9.6	3.382	0.156	00:12:07.74
	SFFS	0.842	0.030	0.658	0.143	12.8	2.227	0.179	00:46:59.06
	OS(1,IB)	0.795	0.017	0.584	0.106	7.2	2.638	0.139	01:29:10.24
1-NN scaled	SFS	0.335	0.002	0.337	0.024	1	0	1	00:00:30.05
	SFFS	0.335	0.002	0.337	0.024	1	0	1	00:00:59.72
	OS(1,IB)	0.335	0.002	0.337	0.024	1	0	1	00:01:45.63
3-NN scaled	SFS	0.907	0.032	0.856	0.165	15.3	6.001	0.361	00:00:31.10
	SFFS	0.937	0.017	0.896	0.143	7.7	3.770	0.206	00:03:03.16
	OS(1,IB)	0.935	0.014	0.907	0.119	5.3	0.781	0.543	00:04:18.10

Table 3: Classification performance as result of wrapper-based Feature Selection on sonar data.

<i>Sonar data, 60 features, 2 classes containing 103 (mine) and 105 (rock) samples, Gorman & Sejnowski</i>									
Classifier	FS Method	Inner 10-f. CV		Outer 10-f. CV		Subset Size		Consistency	Run Time h:m:s.ss
		Mean	St.Dv.	Mean	St.Dv.	Mean	St.Dv.		
Gaussian	SFS	0.806	0.019	0.628	0.151	20.2	12.156	0.283	00:08:41.83
	SFFS	0.853	0.016	0.656	0.131	22.8	8.738	0.326	01:51:46.31
	OS(1,IB)	0.838	0.018	0.649	0.066	21.5	10.366	0.315	03:36:04.92
1-NN scaled	SFS	0.511	0.004	0.505	0.010	1	0	1	00:01:51.78
	SFFS	0.511	0.004	0.505	0.010	1	0	1	00:03:10.47
	OS(1,IB)	0.505	0.001	0.505	0.010	1	0	1	00:08:06.63
3-NN scaled	SFS	0.844	0.025	0.618	0.165	15.2	7.139	0.273	00:02:15.84
	SFFS	0.870	0.016	0.660	0.160	18.9	7.120	0.293	00:12:26.01
	OS(1,IB)	0.864	0.016	0.622	0.151	15.8	5.474	0.247	00:25:55.39

while their relation to classifier generalization performance can be put into even greater doubt.

When comparing feature selection methods as a whole (under specific criterion-classifier-data settings) the advantages of more modern search algorithms may diminish considerably. Reunanen (Reunanen, 2003) points out, and our experiments confirm, that a simple method like SFS may lead to better classifier generalization. The problem we see with the ongoing discussion is that this is often claimed to be

the general case. But this is not true, as confirmed by our experiments.

According to our experiments the “better” methods (being more effective in optimizing criteria) also tend to be “better” with respect to final classifier generalization ability, although this tendency is by no means universal and often the difference is negligible. No clear qualitative hierarchy can be recognized among standard methods, perhaps with the exception of mostly inferior performance of IB (not shown

Table 4: Classification performance as result of wrapper-based Feature Selection on mammogram data.

WPBC data, 31 features, 2 classes containing 151 (nonrecur) and 47 (recur) samples, UCI Repository									
Classifier	FS Method	Inner 10-f. CV		Outer 10-f. CV		Subset Size		Consistency	Run Time h:m:s.ss
		Mean	St.Dv.	Mean	St.Dv.	Mean	St.Dv.		
Gaussian	SFS	0.807	0.011	0.756	0.088	9.2	4.534	0.241	00:00:21.24
	SFFS	0.818	0.012	0.698	0.097	15.4	5.731	0.441	00:04:07.81
	OS(1,r3)	0.826	0.010	0.682	0.062	12.6	5.219	0.356	00:34:07.20
1-NN scaled	SFS	0.251	0.020	0.237	0.018	1	0	1	00:00:14.93
	SFFS	0.251	0.020	0.237	0.018	1	0	1	00:00:39.71
	OS(1,r3)	0.332	0.021	0.237	0.018	7.3	4.776	0.169	00:03:19.70
3-NN scaled	SFS	0.793	0.013	0.712	0.064	9.4	5.869	0.226	00:00:15.56
	SFFS	0.819	0.008	0.722	0.086	11.7	4.797	0.322	00:01:48.94
	OS(1,r3)	0.826	0.007	0.687	0.083	11	3.550	0.325	00:14:44.24

Table 5: Classification performance as result of wrapper-based Feature Selection on mammogram data.

WDBC data, 30 features, 2 classes containing 357 (benign) and 212 (malignant) samples, UCI Rep.									
Classifier	FS Method	Inner 10-f. CV		Outer 10-f. CV		Subset Size		Consistency	Run Time h:m:s.ss
		Mean	St.Dv.	Mean	St.Dv.	Mean	St.Dv.		
Gaussian	SFS	0.962	0.007	0.933	0.039	10.8	6.539	0.303	00:00:22.21
	SFFS	0.972	0.005	0.942	0.042	10.6	2.653	0.36	00:03:24.90
	OS(1,r3)	0.973	0.004	0.943	0.039	10.3	2.147	0.366	00:36:36.49
1-NN scaled	SFS	0.373	0.000	0.373	0.004	1	0	1	00:01:33.07
	SFFS	0.421	0.022	0.373	0.004	1	0	1	00:03:26.00
	OS(1,r3)	0.435	0.001	0.373	0.004	7.6	2.871	0.202	00:25:31.84
3-NN scaled	SFS	0.981	0.002	0.967	0.020	15.3	4.451	0.456	00:01:32.19
	SFFS	0.983	0.001	0.970	0.019	13.7	4.220	0.414	00:08:16.72
	OS(1,r3)	0.985	0.002	0.959	0.025	13.4	3.072	0.421	01:41:02.62

Table 6: Classification performance as result of wrapper-based Feature Selection on speech data.

Speech data, 15 features, 2 classes containing 682 (yes) and 736 (no) samples, British Telecom									
Classifier	FS Method	Inner 10-f. CV		Outer 10-f. CV		Subset Size		Consistency	Run Time h:m:s.ss
		Mean	St.Dv.	Mean	St.Dv.	Mean	St.Dv.		
Gaussian	SFS	0.773	0.008	0.770	0.052	9.6	0.917	0.709	00:00:03.28
	SFFS	0.799	0.008	0.795	0.042	9.3	0.458	0.684	00:00:20.51
	OS(1,r3)	0.801	0.008	0.793	0.041	9.5	0.5	0.642	00:02:46.16
1-NN scaled	SFS	0.522	0.001	0.519	0.002	1	0	1	00:01:27.25
	SFFS	0.521	0.001	0.519	0.002	1	0	1	00:03:07.95
	OS(1,r3)	0.556	0.011	0.519	0.002	8.6	2.577	0.526	00:22:55.49
3-NN scaled	SFS	0.946	0.003	0.935	0.030	7	1.483	0.487	00:01:33.55
	SFFS	0.948	0.003	0.939	0.030	6.7	1.1	0.509	00:05:54.57
	OS(1,r3)	0.949	0.003	0.937	0.029	7	1.095	0.537	01:08:39.20

here). It has been shown that different methods become the best performing tools in different contexts, with no reasonable way of predicting the winner in advance (note, e.g., OS in Table 1 – gives best result with Gaussian classifier but worst result with k-NN).

Our concluding recommendation can be stated as follows: only in the case of strongly limited time should one resort to the simplest methods. Whenever possible try variety of methods ranging from SFS to more complex ones. If one method only has to be cho-

sen, than we would stay with Floating Search as the best general compromise between performance, generalization ability and search speed.

4.1 Quality of Criteria

The performance question of more complex FS methods is directly linked to another question: How well do the available criteria describe the quality of evaluated subsets ? The contradicting experimental results

Table 7: Classification performance as result of wrapper-based Feature Selection on satellite land image data.

<i>Satimage</i> data, 36 features, 6 classes with 1072, 479, 961, 415, 470 and 1038 samples, ELENA database									
Classifier	FS Method	Inner 10-f. CV		Outer 10-f. CV		Subset Size		Consistency	Run Time h:m:s.ss
		Mean	St.Dev.	Mean	St.Dev.	Mean	St.Dev.		
Gaussian	SFS	0.509	0.016	0.516	0.044	19	7	0.643	00:05:21.77
	SFFS	0.525	0.011	0.528	0.034	13.7	3.743	0.474	00:41:25.60
	OS(1,IB)	0.527	0.010	0.517	0.055	12.2	3.311	0.410	01:57:06.71
1-NN scaled	SFS	0.234	0.000	0.234	0.001	1.6	1.2	0.244	03:05:20.63
	SFFS	0.234	0.000	0.234	0.001	1	0	0.444	08:00:19.17
	OS(1,IB)	0.234	0.000	0.217	0.001	1.2	0.6	0.222	19:32:09.52
3-NN scaled	SFS	0.234	0.000	0.234	0.001	1	0	1	03:16:08.09
	SFFS	0.234	0.000	0.234	0.001	1	0	1	07:51:08.98
	OS(1,IB)	0.234	0.000	0.234	0.001	1.1	0.3	0.296	19:09:44.29

Table 8: Classification performance as result of wrapper-based Feature Selection on wave data.

<i>Waveform</i> data, 40 features, 3 classes containing 1692, 1653 and 1655 samples, UCI Repository									
Classifier	FS Method	Inner 10-f. CV		Outer 10-f. CV		Subset Size		Consistency	Run Time h:m:s.ss
		Mean	St.Dev.	Mean	St.Dev.	Mean	St.Dev.		
Gaussian	SFS	0.505	0.002	0.493	0.015	2.1	0.3	0.222	00:08:38.86
	SFFS	0.506	0.003	0.492	0.016	2.4	0.663	0.185	00:42:36.39
	OS(1,IB)	0.506	0.002	0.489	0.015	2.7	1.005	0.222	01:57:58.04
1-NN scaled	SFS	0.356	0.009	0.331	0.000	1	0	1	07:29:40.76
	SFFS	0.356	0.009	0.331	0.000	1	0	1	16:09:52.71
	OS(1,IB)	0.331	0.000	0.331	0.000	1	0	1	35:55:50.53
3-NN scaled	SFS	0.826	0.002	0.810	0.024	17.4	2.332	0.411	08:08:17.25
	SFFS	0.829	0.003	0.808	0.020	17.4	1.020	0.475	38:46:26.60
	OS(1,IB)	0.830	0.002	0.816	0.016	17.1	2.022	0.593	95:12:19.24

seem to suggest, that the criterion used (classifier accuracy on testing data in this case) does not relate well enough to classifier generalization performance. Although we do not present any filter-based FS results here, the situation with filters seems similar. Thus, uneven performance of more complex FS methods may be viewed as a direct consequence of insufficient criteria. In this view it is difficult to claim that more complex FS methods are problematic per se.

4.2 Does It Make Sense to Develop New FS Methods?

Our answer is undoubtedly yes. Our current experience shows that no clear and unambiguous qualitative hierarchy can be established within the existing framework of methods, i.e., although some methods perform better than others more often, this is not the case always and any method can prove to be the best tool for some particular problem. Adding to this pool of methods may thus bring improvement, although it is more and more difficult to come up with new ideas that have not been utilized before. Regarding the performance of search algorithms as such, developing

methods that yield results closer to optimum with respect to any given criterion may bring considerably more advantage in future, when better criteria may have been found to better express the relation between feature subsets and classifier generalization ability.

ACKNOWLEDGEMENTS

The work has been supported by EC project No. FP6-507752, the Grant Agency of the Academy of Sciences of the Czech Republic project A2075302, grant AV0Z10750506, and Czech Republic MŠMT grants 2C06019 and 1M0572 DAR.

REFERENCES

- Asuncion, A. and Newman, D. (2007). UCI machine learning repository, <http://www.ics.uci.edu/~mlrepository.html>.
- Das, S. (2001). Filters, wrappers and a boosting-based hybrid for feature selection. In *Proceedings of the 18th International Conference on Machine Learning*, pages 74–81.

- Dash, M., Choi, K., P., S., and Liu, H. (2002). Feature selection for clustering - a filter solution. In *Proceedings of the Second International Conference on Data Mining*, pages 115–122.
- Devijver, P. A. and Kittler, J. (1982). *Pattern Recognition: A Statistical Approach*. Prentice-Hall International, London.
- Gorman, R. P. and Sejnowski, T. J. (1988). Analysis of hidden units in a layered network trained to classify sonar targets. *Neural Networks*, 1:75–89.
- Jain, A. K., Duin, R. P. W., and Mao, J. (2000). Statistical pattern recognition: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22:4–37.
- Kohavi, R. and John, G. (1997). Wrappers for feature subset selection. *Artificial Intelligence*, 97:273–324.
- Liu, H. and Yu, L. (2005). Toward integrating feature selection algorithms for classification and clustering. *IEEE Transactions on Knowledge and Data Engineering*, 17(3):491–502.
- Pudil, P., Novovičová, J., and Kittler, J. (1994). Floating search methods in feature selection. *Pattern Recognition Letters*, 15:1119–1125.
- Raudys, S. (2006). Feature over-selection. *Lecture Notes in Computer Science*, (4109):622–631.
- Reunanen, J. (2003). Overfitting in making comparisons between variable selection methods. *Journal of Machine Learning Research*, 3(3):1371–1382.
- Sebban, M. and Nock, R. (2002). A hybrid filter/wrapper approach of feature selection using information theory. *Pattern Recognition*, 35:835–846.
- Somol, P., Novovičová, J., and Pudil, P. (2006). Flexible-hybrid sequential floating search in statistical feature selection. *Lecture Notes in Computer Science*, (4109):632–639.
- Somol, P. and Pudil, P. (2000). Oscillating search algorithms for feature selection. In *Proceedings of the 15th IAPR Int. Conference on Pattern Recognition, Conference B: Pattern Recognition and Neural Networks*, pages 406–409.
- Yu, L. and Liu, H. (2003). Feature selection for high-dimensional data: A fast correlation-based filter solution. In *Proceedings of the 20th International Conference on Machine Learning*, pages 56–63.

IMPLEMENTATION OF EPS_T2DM

Implementation of Early Prediction System for Type 2 Diabetes Mellitus

Dan Bi Kim, Ju-young Lee, Eun-jung Jang, Ji-young Lee and Taek Kim
Division of Bio-Medical Informatics, Center for Genome Science
Korea Centers for Disease Control and Prevention, Korea National Institute of Health
194, Tongil-ro, EunPyong-Gu, Seoul, Korea
danbi_kim@hotmail.com, jylee_monte@naver.com

Keywords: Type 2 Diabetes Mellitus, Early Prediction System, KoGES.

Abstract: This paper describes the implementation of an early prediction system for Type 2 diabetes mellitus. Type 2 diabetes mellitus is a multifactorial disease. It is not only associated with an unhealthy lifestyle but also has a strong genetic component. Accordingly, in order to decrease an incidence rate of T2DM, it is important to predict T2DM risk with using multifactors which are supposed to affect T2DM. We have implemented a prediction system for T2DM, and it employs several statistical prediction models. These models are produced by statistical analysis about cohort data of Korean Genome and Epidemiology Study (KoGES), and include risk factors which are adequate for preventing T2DM in Korean populations. The prediction system is written in JSF and Java, and developed into web application which is designed through object oriented modeling. Web application of this system offers user interfaces in order to input data which is needed for predicting risk group, select predefined prediction models, and so on. The system provides the results which are predicted by selected models using inputted information.

1 INTRODUCTION

Type 2 diabetes mellitus (T2DM) is a multifactorial disease in which environmental triggers interact with genetic variants in the predisposition to the disease. In the last thirty years, due to rapid industrialization and the lifestyle change, the number of Korean patients with T2DM is on an increasing trend. In addition, more serious problem is that some significant number of T2DM patients do not aware of their disease before onset it. Therefore, it is not only important to treat T2DM, but also necessary to predict its risk and prevent it in advance.

To prevent and control a certain disease for specified population effectively, it is necessary to understand epidemiological features of patients. So, we have analyzed the data collected by cohort study from Korea National Institute of Health (KNIH). The primary prevention method of a disease is prediction of its risk group or patients unaware of their disease, and decreasing the incidence of the disease or complications by reduction of risk factors for that group. For that reason, we have implemented the early prediction system for T2DM which is named EPS_T2DM.

2 MATERIALS AND METHODS

At first, we have analyzed KoGES data to find risk factors for T2DM, and built prediction models by statistical analysis. Then, for constructing database schema, input values including risk factors are defined. Finally, the entire system is implemented as input modules, a select model module, model processing modules and displaying module for results.

2.1 Building Statistical Prediction Models

This study has been supported by community based cohort study Ansong and Ansan, Korean Genome and Epidemiology Study (KoGES) from KNIH. The data for this study consists of epidemiological information and genetic information (SNP values) for 212 T2DM public and 472 general (non-T2DM) public (Table 1).

Table 1: Basic characteristics. (Means±SD).

Characteristic	Diabetes	Non Diabetes
Numbers	212	472
Sex(F/M)	133/79	264 / 208
Age(Years)	64.50±2.82	64.03±2.86
BMI(kg/m²)	25.07±3.44	23.31±3.11
Triglyceride (mg/dl)	196.60±131.69	149.52±70.97
Fasting glucose(mg/dl)	111.20±28.96	74.54±3.55
Total cholesterol(mg/dl)	202.78±44.79	180.66±31.64
HDL cholesterol(mg/dl)	44.56±10.59	44.34±9.91

The Ansung and Ansan cohort data have been processed using the Statistical Analysis System for Windows (Ver. 9.1, SAS institute Inc., Cary, NC, U.S.A), and mined by several algorithms such as QUEST, C4.5, logistic regression, SVM, and KNN algorithm. Among data mining results, the results of QUEST, a binary-split decision tree algorithm for classification and data mining, is applied to EPS_T2DM. From this result, risk factors for T2DM are defined and prediction models are produced.

2.2 Design of Database for Input Values

According to data mining results, risk factors for T2DM are as follows:

- Clinical information – Height and weight for BMI, Waist circumference for abdominal fatness, Blood pressure, Total cholesterol, High-density cholesterol, Triglyceride, Fasting glucose, and etc;
- History of diseases information – Diabetes Mellitus, Cerebrovascular disease and other vascular diseases, Hypertension;
- Family history of diseases information – Diabetes Mellitus;
- Genetic Information – Selected single nucleotide polymorphisms (SNPs) in 15 genes in Insulin pathway, 8 genes in fatty acid binding/translocation, and 13 genes in GLUT4 translocation and 51 more genes related to T2DM.

Each category is represented in a table, and each factor is corresponds to each field (Figure 1).

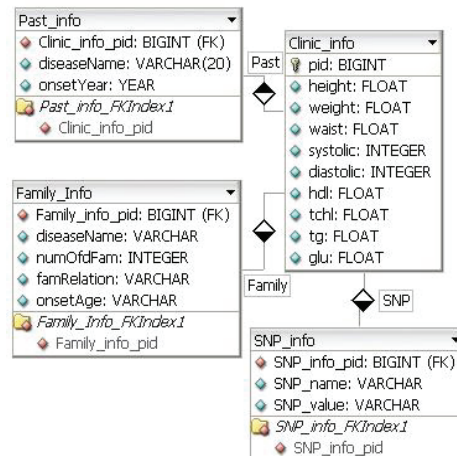


Figure 1: Database schema for input values.

2.3 Implementing EPS_T2DM

EPS_T2DM is based on Object Oriented Modeling. It is developed on Fedora core 6, written in HTML, JSF, JavaScript, Java, and etc. MySQL is used as a database to store and access data. The entire system is implemented with Spring Framework which is a layered Java/J2EE application framework, and then Model-View-Controller (MVC) design pattern is applied. MVC is an architectural pattern which encapsulates some data together with its processing (the model) and isolates it from the manipulation (the controller) and presentation (the view) part. Figure 2 illustrates the system architecture which is composed of five layers with UI layer, 3 layers above, and persistence layer.




Figure 2: System architecture of EPS_T2DM.

The function of EPS_T2DM consists of 4 parts: (1) user interfaces to get input data and user’s selection, (2) management modules of inputted data, (3) processing and management modules of statistical prediction models, and (4) interfaces to offer prediction results. JSF is used to implement user interfaces, and XML is employed to manage input data, SNP information, and prediction models.

3 RESULTS

EPS_T2DM is implemented as a web application. It offers interfaces in order to input SNP values and epidemiology data which include clinical information, history of diseases information and family history of diseases information. After getting input data, the system shows prediction models to users, and users can select them. The system applies selected models to user's input values, and displays prediction results that are represented whether or not a person corresponding to the inputted data belongs to risk group of T2DM. Figure 3 gives a flow chart of this system.

The web interfaces consist of epidemiology information input page with tabbed pane (Figure 4), SNP values input page, selecting models page, displaying results page, and so on.

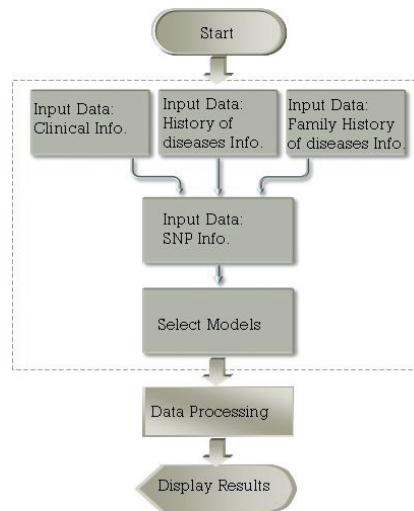


Figure 3: Web application's flow chart of EPS_T2DM.

Disease Name	Existence Of Disease	Diagnosis Year
Hypertension*	<input type="radio"/> No <input type="radio"/> Yes	<input type="text"/>
Diabetes Mellitus*	<input type="radio"/> No <input type="radio"/> Yes	<input type="text"/>
Hyperlipidaemia*	<input type="radio"/> No <input type="radio"/> Yes	<input type="text"/>
Cerebrovascular Disease*	<input type="radio"/> No <input type="radio"/> Yes	<input type="text"/>
Other Vascular Diseases*	<input type="radio"/> No	<input type="text"/>

Disease Name	The Number of Diseased Family	Family Relations & Incidence Age
Hypertension*	2persons	Select! <input type="text"/> years Select! <input type="text"/> years
Diabetes Mellitus*	Select!	
Hyperlipidaemia*	Select!	
Cerebrovascular Disease*	Select!	
Other Vascular Diseases*	Select!	

Health Condition*	Select!
Height*	0.0 <input type="text"/> cm
Weight*	0.0 <input type="text"/> kg
Waist girth*	0.0 <input type="text"/> cm
Blood Pressure *	<input type="text"/> / <input type="text"/> mmHg
High-density Cholesterol*	0.0 <input type="text"/> mg/dl
Total Cholesterol*	0.0 <input type="text"/> mg/dl
Triglyceride*	0.0 <input type="text"/> mg/dl
Fasting Blood Sugar *	0.0 <input type="text"/> mg/dl

Figure 4: The snapshots of EPS_T2DM's web application. The interface for inputting epidemiology data which is defined as risk factors of T2DM and other additive information such as sex, birth, etc.

4 DISCUSSION

In many cohort studies, various T2DM predicting models have been developed to guide intervention and inform health policy. Most of these studies tested models that only used personal information and clinical variables, not including genetic information. However, since development of T2DM is influenced by a complex interaction between genetic and environmental factors, genetic information is also needed to develop the prediction models for T2DM.

In this study, we implemented the system to predict the T2DM risk group with applying T2DM risk factors and statistical prediction models that include clinical and genetic information. The risk factors and prediction models are based on Ansong and Ansan, KoGES data and decision tree learning method, and these can be updated or changed by based data or analysis methods. EPS_T2DM is developed as object-oriented program, so it is easy to extend and enhance the system.

Our next step will be to expand and improve the system. This includes followings:

- The betterment of statistical models. This means improving accuracy or building new prediction models by analyzing other data or applying other data mining methods
- The enhancement of web application for uploading or processing massive data sets.

REFERENCES

- Zimme P, Alberti KG, Shaw J. (2001). Global and societal implications of the diabetes epidemic. In *Nature*, 414, 782-787.
- Ju-young Lee, Eun-jung Jang, Ji-young Lee, Kyung-soo Oh, Hyun-woo Han (2007). PTPN1 Gene: The haplotype-based interaction with Type 2 Diabetes. In *ISMB 2007, Poster 182*.
- Peter W.F. Wilson, James B. Meigs, Lisa Sullivan, Caroline S. Fox, David M. Nathan, Ralph B. D'Agostino. (2007). Prediction of Incident Diabetes Mellitus in Middle-aged Adults. In *Archives of Internal Medicine*, 167, 1068-1074.
- The Diabetes Prevention Program Research Group. (2005). Strategies to Identify Adults at High Risk for Type 2 Diabetes. In *Diabetes Care*, 28, 138-144.
- Alok Bhargava. (2003). A longitudinal analysis of the risk factors for diabetes and coronary heart disease in the Framingham Offspring Study. In *Population Health Metrics*, 1:3.
- Enzo Bonora, Stefan Kiechl, Johann Willeit, Friedrich Oberhollenzer, Georg Egger, James B. Meigs, Riccardo C. Bonadonna, Michele Muggeo. (2004).

Population-Based Incidence Rates and Risk Factors for Type 2 Diabetes in White Individuals. In *Diabetes*, 53, 1782-1789.

Kyong Soo Park. (2004). Prevention of Type 2 diabetes mellitus from the viewpoint of genetics. In *Diabetes Research and Clinical Practice*, 66S, S33-S35.

PET-CT IMAGING AND DIAGNOSIS SYSTEM FOLLOWING DOCTOR'S METHOD

Hiroshi Arisawa, Takako Sato and Sinya Hata

*Graduate School of Yokohama Nat. Univ., 79-7, Tokiwadai, Hodogayaku, Yokohama, 240-8501, Japan
arisawa@ynu.ac.jp, sugar@arislalab.ynu.ac.jp, sinya@arislalab.ynu.ac.jp*

Keywords: Medical Imaging, PET-CT, Expert Systems.

Abstract: Computer Assisted Diagnosis (CAD) is one of the promising technologies for the future Medical Image Processing Systems. Among them, whole-body PET (Positron Emission Tomography) and X-ray CT (Computer Tomography) image based cancer detection has been playing an essential role in the modern medical world. Using PET-CT images the Radiologist can find a very small cancer or a malignant tumor. On the other hand, this diagnosis process is very stressful work, because such area is too small and localized but may appear at any place of patient bodies. This paper presents an automated diagnosis system in order to improve the above difficulties. The system consists of three parts, Diagnosing Algorithm, Algorithm Interpreter(Engine) and Image Viewer. The algorithm and the engine can reproduce doctor's methods faithfully as the rule-based inference system. Using this system, we made an retrospective studies for the actual group of patients and the results shows the usefulness of this approach.

1 INTRODUCTION

Computer assisted diagnosis systems are an important issue in the world of modern medicine. Computer-based interpretations of 2-D images (such as CT, MRI, etc) are among this group of systems which have been utilized by physicians with promising results. Because of this, a great deal of research on computer-assisted-diagnosis support systems (CAD; Computer Aided Diagnosis) have been proposed (Jiang et al, 2000) (Toriwaki et al, 2000) (Tsai et al, 2001) (Cheng et al, 1998) (Ukai et al, 2000), especially for diagnostic imaging tools.

The diagnostic method for cancer detection using the PET (Positron Emission Tomography) and X-ray CT (Computed Tomography) images is a core technology, which attracts the interests of many medical scientists (Murakami, 2003). During cancer inspection by a PET scan, drugs called FDG (Fluorodeoxyglucose: a glucose-mimicking radioactive element) are administered to the patient, and the gamma rays emitted from the patient are photographed by the nuclear imaging system. The PET images show the various level of absorption (SUV: Standard Uptake Value) of the FDG through out the body. As a result, we can observe the FDG concentration absorbed by the tissues and organs. This is useful because malignant cancer cells have

an increased glucose metabolism, so much more FDG is taken into a cancer cell, and so, a SUV value will become much higher than a normal cell. This is referred to as an "abnormal accumulation." However, more FDG will also be taken into areas where inflammation has occurred or organs (such as kidneys, urinary bladder, liver, etc.) which take in more glucose naturally even without the effects of cancer. This latter is called a "physiological accumulation." The purpose of an automated diagnosis system is finding out which areas have signs of an abnormal accumulation based on the images of the whole body PET scan.

During a PET scan of the whole body, tomography is performed by rotating the camera around the axis of the body at intervals of about 3mm. From the results of a whole body tomography (from a femoral region to the parietal region), a physician receives about 3000 slice images per patient. In order to manually analyze these images a physician needs to have much knowledge and experience of PET scans as well as time and effort to interpret these images. This creates a very large burden for the physician, there by increasing fatigue and decreased concentration which may lead to a misdiagnosis. Moreover, the number of physicians who can perform an interpretation of PET image is insufficient.

On the other hand, in general during a medical check up for cancer by PET diagnostic imaging, 90 percent or more of the samples are normal images. These images lack any indication of cancer. Even still searching for cancer cell is stressful work for the radiologists.

Based on this information, an automated diagnostic support system is a very effective tool for analyzing and pinpointing potentially cancerous areas of the human body.

For the above purpose, existing works are mostly focusing on the specific organs such as the lung (Takeo et al, 2005). Also most researches are based on CT images and the recognition of area are on the combination of filters. On the other hand, our method are based on PET images. However, as for several critical area such as bone and lung, 3D shape and scale are extracted from CT slices. We confirmed that a similar process has been executed in the actual radiologist's diagnosing.

In this paper, design principle of the diagnosis system is introduced in the section 2, and the experimental results are presented and evaluated in the section 3.

2 A SYSTEM WHICH IMITATES DIAGNOSIS OF PHYSICIAN

2.1 Basic Principles for the System's Configuration

The basic principles of the automated diagnosis system which we created are summarized by the following two points.

- In order to use this diagnosis system for mass screenings, the fundamental function is to 'classify' cases into two categories:
 - "The possibility of abnormality followed by a careful examination"
 - "Normal with no need for further inspection"

Then the next step is to reduce the chance of a misdiagnosis by preventing false-negative results from occurring.

- In order to get the trust of physicians and satisfy the two principles above, the system must reproduce the physician's diagnostic process as faithfully as possible. Our strategy for building this system is to interview many physicians and radiologists to have a greater understanding as well as the feasibility of

creating a diagnostic support system. We are establishing a replica of the physician's diagnostic process with an automated computer system. To meet this end we have developed a description language to reproduce the process of a PET scan diagnosis through computer algorithms.

2.2 Feature Analysis of a PET Image

A PET image, unlike CT or MRI, does not express morphological information such as the shape of an organ, but expresses functional values (for example the differences in the amount of glucose metabolized) in the undefined areas. The images are often low resolution and very coarse. Therefore, it is difficult to determine the specific cancerous region within tissue with just PET scan image.

While interpreting the PET images, a physician utilizes his/her knowledge of previous PET scan test cases in order to make a proper diagnosis. Also, the outline image of organs from a CT scan is very useful to confirm the location of the accumulation of cancer. With this in mind, we decided to create a software package that interprets a physician's methodology during PET scan diagnosis and uses the proper PET terminology to describe the program's actions so that it would mimic the performance of the physician. This system is designed so that a physician can monitor the process of diagnosis from the local to global (whole body) level, to evaluate diagnosis's validity and to recommend improvement.

2.3 Architecture of PET Automated Diagnosis System

In order for a computer to faithfully reproduce physician's judgment, this based on experience and knowledge of physician as well as a strict usage of the grammar and language of PET scan diagnosis. The person who translates a physician's methodology into a description language will be called a knowledge engineer. Figure 1 shows the complete concept of the automated diagnosis system.

The knowledge engineer interviews the physician about knowledge and techniques used for the interpretation of PET images. When the diagnostic method is understood, the knowledge engineer translates this information into a diagnostic algorithm. This automated procedures preformed by the computer must be understood and be able to be examined by the physician directly. In order to do

this, the system imitates the diagnosis procedure of a physician step by step. Furthermore, image processing functions are needed to visualize the automated diagnosis. So we divide the diagnosis into several parts and provide a proper description of the diagnostic process at each step as well as for each data-processing function that is called.

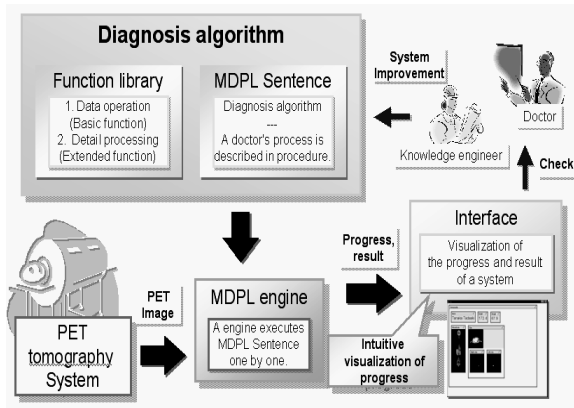


Figure 1: Architecture of PET automated diagnosis system.

And by allowing a physician to view each portion of the diagnostic process, the physician can examine and validate each part of the diagnosis by checking the functions of that particular step of the process and viewing intermediate results, without having to understand the underlying computer program language (Endo et al, 2004). Herewith, a physician can understand and check each individual diagnostic process which the system performs.

In order to express intermediate result data, we developed the logical data structure NEW (Nested Entity Window) (Hasegawa et al, 2005), and we regard the diagnostic process as data manipulation to NEW. Then we also propose programming in language MDPL (Multimedia Data Processing Language) (Hasegawa et al 200) in order to describe the diagnostic process. With MDPL, we can express a particular section of the process in NEW, which has a complex structure.

3 AN EXPERIMENT ON THE COMPUTER-BASED DIAGNOSTIC METHOD

We experimented about the described diagnostic method.

In our laboratory, the MDPL interpreter, NEW manager, and Database System are currently being developed as an independent project. Therefore, in

this experiment, we implemented the basic data-processing functions required for building the diagnostic method with suitable parameter. Moreover, we recorded various states of NEW as it was processed by a sequence of functions. Then we could reproduce the flow of the structural change of NEW at the time a particular MDPL statement is called.


3.1 Basic Experimental Data

The data used for the experiment are as follows.

- 12 normal example
- 18 cases of cancer
- Total number of accumulations which have possibility of cancer is 22

3.2 Experimental Result

After applying the computer-based diagnostic method to all examples, the system pointed out 110 'abnormal' accumulations (having possibility of cancer). On the other hand, the professional radiologist pointed out 22 abnormal accumulations. Fig.2 shows the above results as a Venn diagram.



U: total accumulation
 A: accumulations pointed out by automatic diagnostic system
 D: accumulations pointed out by professional radiologists

$D - A$: false negative

$A - D$: false positive

$$\frac{|A \cap D|}{|D|} \text{ recall}$$

$$\frac{|A \cap D|}{|A|} \text{ precision}$$

Figure 2: Experimental result based on a diagnostic method.




Figure 3: Result of Automated Diagnostic System (thyroid gland cancer).

In the diagram, for example, D-A indicates a set for ‘false negative’ accumulations. That is, each instance of D-A indicates the accumulation that were pointed out as abnormal one pointed by professional radiologist but not pointed out by the automated diagnostic system. Such instance must be kept from occurring in the Computer Assisted Diagnosis.

In our preliminary experiment, the results are as follows:

False negative	0 %
False positive	80.0 %
Precision ratio	20.0 %
Recall ratio	100 %

4 CONCLUSIONS

During this research, we developed and configured a cancer automated diagnosis system and tested its capabilities. It is able to imitate the “real” diagnosis of a physician. The physicians can evaluate the validity of the result by themselves. Also, through the constant feedback and discussions with physicians, we could acquire more information, and, as a result, make improvements to the diagnosis system. Currently we are working to ensure that this system will be able to properly diagnosis and provide a detailed description about any abnormal spots any where in the human. We continue to improve the system in order for the successful utilization in the medical field.

ACKNOWLEDGEMENTS

We express our appreciation to Dr. Akiko Suzuki who cooperated with Professor Tomio Inoue of the Yokohama City University graduate school of medical studies.

REFERENCES

Jiang, H., Masutoh, N., Nishimura, O., Okumura, T., Yamamoto, S., Iisaku, S., Matsumoto, T., Tateno, Y., Inuma T., Matsumoto, M., 2000. Computer-Aided Diagnosis System of Lung Cancer Screening by X-Ray CT, *IEICE TRANS.(D-II)* , *J83-D-II (1)*, 333-341.

Toriwaki, J., 2000. Trends and Future Problems in Computer Aided Diagnosis of X-ray Images, *IEICE TRANS.(D-II)* , *J83-D-II (1)*, 3-26.

Tsai, D., Lee, Y., 2001. Fuzzy Reasoning Based Computer-Aided Diagnosis for Automated Discrimination of Myocardial Heart Disease from Ultrasonic Images, *IEICE TRANS.(A)* , *j84-A(12)* , 1431-1438.

Cheng, X., Akiyama, I., Wang, Y., Itoh, K., 1998. Breast Tumor Diagnosis System Using Three Dimensional Ultrasonic Echography , *IEICE TRANS.(D-II)* , *J81-D-II (1)*, 1402-1410.

Ukai, Y., Niki, N., Satoh, H., Watanabe, S., 2000. Computer Assisted Diagnosis System for Coronary Calcifications Based on Helical CT Images , *IEICE TRANS.(D-II)* , *J83-D-II (1)*, 342-350.

Murakami, K., 2003. Base and Clinical Application of PET for Lung Cancer, *Japanese Journal of Lung Cancer*, *43(7)*, 811-817.

Takeo H., Shimura K., Imamura T., Shimizu A., Kobatake H., 2005, Detection System of Clustered Microcalcifications on CR Mammogram (Biological Engineering), *IEICE transactions on information and systems*, *E88-D(11)*, 2591-2602

Endo, T., Hashimoto, R., Arai, J., Arisawa, H., Suzuki, A., Inoue, T., 2004. Construction of a Medical Database for Automatic Diagnosis System of Cancer Using PET Images , *DEWS2004 papers 1-6-03 (IEICE technical research report)*.

Hasegawa, T., Yaguchi, Y., Endo, T., Arai, J., Arisawa, H., 2005. Decision Support System with the Description of Thinking and Estimation Process of Experts , *DEWS2005 papers 4A-i7 (IEICE technical research report)*.

**SPECIAL SESSION ON
KNOWLEDGE DISCOVERY AND DECISION
SUPPORT SYSTEMS IN HEALTH
INFORMATION SYSTEMS**

**CHAIRS:
DOMENICO CONFORTI
MANOLIS TSIKNAKIS
JOERG HABETHA
JOËL BACQUET**

INITIAL RESULTS ON KNOWLEDGE DISCOVERY AND DECISION SUPPORT FOR INTRACRANIAL ANEURYSMS

Christoph M. Friedrich, Martin Hofmann-Apitius

*Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Department of Bioinformatics
Schloss Birlinghoven 53754 Sankt Augustin Germany
friedrich@scai.fraunhofer.de, hofmann-apitius@scai.fraunhofer.de*

Robert Dunlop

*Infermed Ltd, 25 Bedford Square, London, UK
robert.dunlop@infermed.com*

Ioannis Chronakis

*Department of Engineering, Oxford University, UK, Department of Academic Oncology, UCL, UK
ioannis.chronakis@eng.ox.ac.uk*

Miriam C. J. M. Sturkenboom, Roelof Risselada

*Erasmus MC, Medical Informatics, 3000 CA Rotterdam, Netherlands
m.sturkenboom@erasmusmc.nl, r.risselada@erasmusmc.nl*

Baldo Oliva, Ferran Sanz

*Research Unit on Biomedical Informatics (GRIB) IMIM/UPF, C/Dr. Aiguader, 88, 08003 - Barcelona, Spain
baldo@imim.es, fsanz@imim.es*

Keywords: Knowledge Discovery, Decision Support, Intracranial Aneurysm.

Abstract: Intracranial Aneurysms are bulbous expansions of the intracranial vessels, that may rupture and lead to subarachnoid haemorrhage, which can result in severe disability or death of the affected person. The prediction of the individual rupture risk of a patient based on information from images, haemodynamic simulations, clinical parameters and genetic markers is one of the aims of the European Integrated Project @neurIST. The predicted rupture risk is meant to support decision making on clinical treatment. We will present initial results on Knowledge Discovery through a combination of text-mining, data integration from public bioinformatics data sources, and database mining. Additionally, we provide first results for decision support through knowledge based clinical guidelines and Bayesian networks.

1 INTRODUCTION

The advent of improved medical imaging facilities and their routine use in clinical practice increases the number of accidentally detected asymptomatic Intracranial Aneurysms (IA). Intracranial Aneurysms are bulbous expansions of the intracranial vessels, that may rupture and lead to intracranial bleeding (subarachnoid haemorrhage), which can result in severe disability or death of the affected person. In (Rinkel et al., 1998) the prevalence of the disease for adults without risk factors for subarachnoid haemorrhage is reported with approximately 2 % and the annual risk rate for a rupture with 0.7 %. This relatively high prevalence with low incidence of the dangerous event

leads to controversial discussions on treatment decisions. In general there are three treatment options:

1. do not treat the asymptomatic aneurysm with low risk
2. conduct neurosurgical clipping
3. deploy a platinum coil via endovascular intervention

One of the targets of the European Integrated Project @neurIST¹ is to support decision making on IA treatment options by building a distributed environment for healthcare. This environment will allow access to patient related information from images,

¹<http://www.aneurist.org>

haemodynamic simulations, clinical parameters, genetic markers and epidemiological data.

Additionally, a set of application suites will be developed, that are based on this infrastructure and directly support the goal of improving clinical decision making. A draft architecture of the distributed system is described in (Arbona et al., 2007). Considering the target of this paper, two application suites are of interest, @neuRisk, a decision support system based on clinical guidelines and @neuLink, a research oriented application targeted at linking genetics to disease. The @neuLink suite supports Knowledge Discovery for the detection of genetic risk factors.

2 INITIAL RESULTS

In the following we will give examples and preliminary results of Decision Support and Knowledge Discovery that have been developed during the first year of the project.

2.1 Decision Support based on the Proforma Language and the REACT Application

In this section we will describe the work carried out in the development of the first @neuRisk prototype. This version of the prototype employs a mixed quantitative and qualitative approach to provide risk assessment and decision support in the treatment of cerebral aneurysms. The knowledge base for the approach is derived from two trials: The International Study on Unruptured Intracranial Aneurysms (ISUIA) (Wiebers et al., 2003) and International Subarachnoid Aneurysm Trial (ISAT) (Molyneux et al., 2005). For demonstration purposes, some additional test data were also included to show how future research results could be incorporated to the final @neuRisk suite.

Qualitative decision support in @neuRisk has been implemented using the PROforma method and tools (Sutton and Fox, 2003) while quantitative decision support was implemented by adapting an existing treatment planning application called REACT (Risk, Events, Actions and their Consequences over Time) (Glasspool et al., 2006).

PROforma is a well established technology, first published in 1996 (Fox et al., 1996) and described in detail in 2000 (Fox and Das, 2000; Sutton and Fox, 2003), is a well established clinical decision support technology, that has been tested in a number of trials with promising results (Fox et al., 2006; Hurt et al.,

2003; Patkar et al., 2006). There are two major implementations available, the Tallis implementation from Cancer Research UK² and the Arezzo implementation by InferMed Ltd in London³. REACT technology is based on PROforma concepts and has been tested in one trial in the area of genetic counselling with encouraging results (Glasspool et al., 2006).

In the qualitative part of the prototype, we used the PROforma language (Sutton and Fox, 2003) to model both the workflow involved in patient management and a set of treatment decisions. The resulted computer-executable guideline application is then enacted by the Tallis PROforma engine. It guides the user through the workflow, provides a set of data capture services to collect data from the various disciplines (clinician, radiologist, geneticist, etc) and finally, it offers support for the treatment decisions. It does this by offering a set of logical arguments (rules) to the clinicians which either support or oppose each of the available treatment actions. The system suggests the most appropriate action but the final decision is taken by the clinician (Fox et al., 2006).

In the quantitative part of the prototype, we used an adapted version of the REACT tool (Glasspool et al., 2006). This tool provides support for planning the treatment of the patient based on the effect that each treatment action has on the risk. The REACT user interface is divided into 4 major parts:

1. The treatment plan,
2. The graph area,
3. The argumentation area and
4. The notification area.

The treatment plan area provides the clinician with a set of available treatment options and a timeline where she/he can schedule them, similar to a Gantt chart. As the user adds or removes events from the treatment plan, the graph area plots the expected consequences in real time (in the @neuRisk prototype these are life expectancy and years gained or lost if an aneurysm were left untreated). This allows the user to explore the space of available options with immediate feedback of the various interactions and consequences. Information directly relevant to each option

²Tallis is an implementation of the PROforma engine written in Java which was developed at Cancer Research UK by the Advanced Computation Laboratory (<http://acl.icnet.uk/>) that was led by Prof John Fox. The engine is supported by a suite of tools including an engine, an authoring tool, a tester tool and a web based enacting application (<http://www.cossac.org/tallis.html>).

³Arezzo is an implementation of the PROforma engine created by InferMed (<http://www.infermed.com>) and it has been used in a number of commercial products.




Figure 1: Example of the @neuRisk prototype, comparing treatment versus non treatment of a patient. The output shows the expected Life years.

is summarized in the form of arguments for or against each action in the argumentation area while the notification area shows warnings when any planned actions may pose danger to the patient (Glasspool et al., 2006). Figure 1 shows a screenshot of the user interface.

In the future we intend to expand the @neuRisk decision support functionality in various ways. We will extend the workflow model to mirror more accurately the management of an @neurIST patient. We will also integrate @neuRisk with other services under development by the project (such as @neuInfo (Arbona et al., 2007)). We will enrich both the qualitative and the quantitative decision support models with new arguments based on risk factors that will be established from literature review, epidemiology analysis and clinical studies. Finally we will receive feedback from the end users aiming to improve @neuRisk usability.

2.2 Decision Support with Bayesian Networks

Complementary to the rule-based Decision Support approach as described in section 2.1, we implemented Bayesian Network models. These are probabilistic graphical models, also referred to as Belief Networks (Russel and Norvig, 2003). Due to the lack of sam-

pled data at this state of the project, we had to generate artificial training and testing data. Similar to the PROforma model we made use of the data distributions and associations reported in the ISUIA and ISAT study. For the development of the Bayesian Networks, we used the GeNIe⁴ (Druzdzel, 1999) system and implemented two models, estimating the risk of rupture and predicting the treatment outcome. The chosen risk factors have been described in the mentioned studies e.g. size and location of the aneurysm and state of prior haemorrhages. In figure 2 an example of such a model is given. Due to the lack of independent data, only an evaluation of the model performance based on bootstrapping (Efron and Tibshirani, 1993) could be conducted, which results in an accuracy of approximately 98 % in a 50-fold bootstrap for the rupture risk model (Han et al., 2006).

In (Ghinea and van Gelder, 2004) a related work has been published, which also lacks data and the authors filled in hypothetical data from their clinical practice.

We look forward to validate our models on real data, that will be gathered during the project and will add other published risk factors, that have been reported in the Cochrane Review on risk factors (Clarke, 2007).

⁴<http://genie.sis.pitt.edu>

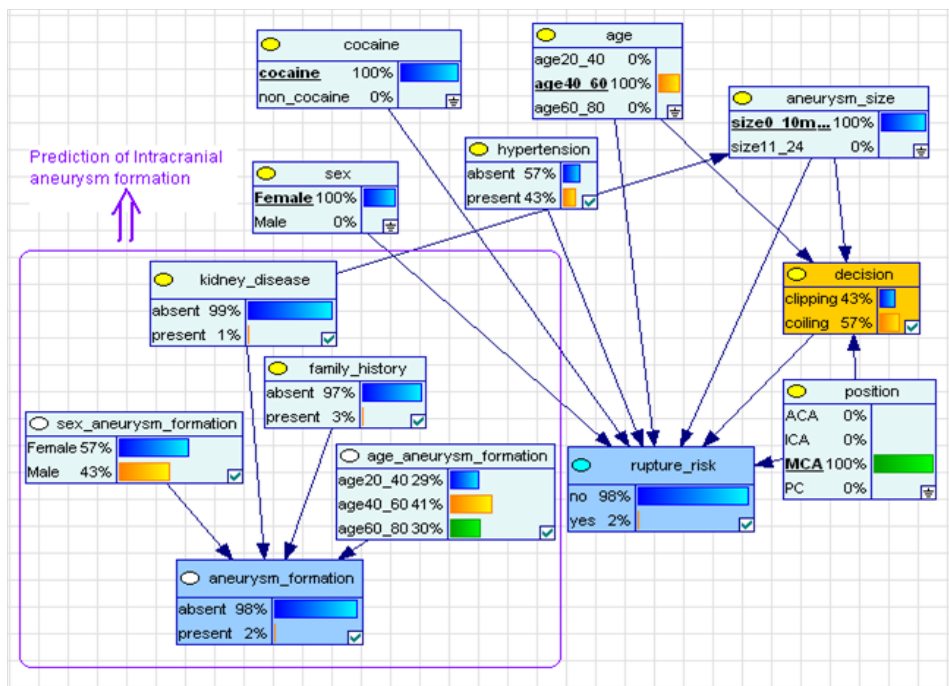


Figure 2: Example of an interconnected Bayesian network to predict formation and rupture risk as well as treatment outcome for Intracranial Aneurysms based on data from ISUIA and ISAT studies.

2.3 Database Mining in Epidemiological Databases

As another source for risk factor assessment, we mined the Integrated Primary Care Information⁵ (IPCI) database (van der Lei et al., 1993). IPCI is a longitudinal observational database, which contains data from computer-based medical records of a group of approximately 400 General Practitioners (GP) (800,000 patients) throughout the Netherlands who voluntarily agreed to supply data to the database. These data contain extensive and detailed information on GP visits, hospital discharges, medication subscriptions and biometrical measurements (e.g. blood pressure, cholesterol levels, glucose levels). Using this information, (Risselada et al., 2007) were able to confirm the relative risks of the established risk factors smoking and hypertension on occurrence of an aneurysmatic subarachnoid haemorrhage. In a study population of women under 50 years of age, the relative risks were 5.0 (95 % confidence interval [CI] 2.2-11.4) for hypertension and 2.2 (95 % CI 1.0-4.9) for smoking. These figures are similar to those found in literature .

Besides these established risk factors (Risselada et al., 2007) have been able to estimate the associa-

⁵<http://www.ipci.nl>

tion between the use of oral contraceptives and the risk of rupture. In a study population of women under 50 years of age, In the aforementioned study population, the risk of an aneurysmatic subarachnoid haemorrhage was doubled in recent users of hormonal contraceptives.

Subdivision of the hormonal products according to their progestin content revealed relative risks of 2.2 (1.3–18) for recent levonorgestrel users and 4.8 (1.3–18) for lynestrenol users, compared to non-users of progestins. These database mined relative risks, and additional ones currently under investigation will be used in the future versions of the @neuRisk suite described in section 2.1.

2.4 Finding Candidate Genes and Variation Mentions in Biomedical Text

Similar to other complex diseases, it is known that Intracranial Aneurysm formation and risk of rupture is associated with environmental, lifestyle and genetic factors. To find the susceptibility genes, many studies have been conducted. In (Krischek and Inoue, 2006) and (Nahed et al., 2007) reviews on the results are given. The problems of these expert generated reviews amongst others, are timeliness, amount of ex-

pert work, specialisation on one topic and possible selection bias. As an alternative to this manual extraction, we consider text-mining (Jensen et al., 2006). This helps to get an overview on genes possibly involved in a disease and to find potential new genes from publications. We implemented a Find Candidate Genes module in the @neuLink application suite. This part of the application suite is based on two text-mining systems, ProMiner (Hanisch et al., 2005) and OSIRIS (Bonis et al., 2006). ProMiner finds entities (Genes/Proteins, Drugnames, Chromosomal Locations ...) and links them to unique database identifiers, e.g. EntrezGene (Maglott et al., 2005). OSIRIS finds and disambiguates mentions of genetic variations in text to dbSNP (Smigielski et al., 2000) identifiers with a query-expansion approach.

To support the focussed view of the user on relevant information in the disease context, we developed a ranking mechanism based on Relative Entropy (Kullback and Leibler, 1951), also known as Kullback/Leibler divergence. In this ranking mechanism, we use the complete MedLine as reference corpus and contrast it with the specific corpus derived from full-text search.

Finding and disambiguating variation mentions in text with the OSIRIS system, needs a high-quality gene-mention machinery. We therefore combined our text-mining tools and complemented them with a machine learning variation finding engine based on Conditional Random Fields (CRF) (Lafferty et al., 2001). The improved results of this approach have been described in (Klinger et al., 2007).

One of the crucial questions of all Discovery methods is their validation. For the finding and disambiguation of gene mentions in text, we have been able to get an independent assessment of the performance of our approach by participating in the BioCreative II assessment (Morgan and Hirschmann, 2007). Our ProMiner system assessed as described in (Fluck et al., 2007), has been ranked 3rd of 21 submissions.

In a second evaluation, we tested whether our system, given the keyword search “intracranial AND aneurysm*”, was able to detect the same related susceptibility genes, that have been found by human experts. The review on genetics (Krischek and Inoue, 2006) mentions 18 associated genes in the context of Intracranial Aneurysms. In our evaluation (as of 2007-10-01) (Gattermayer, 2007), we find 16,548 documents in PubMed related to the keyword and 596 documents, that mention 316 different genes/proteins. We find and could disambiguate all 18 genes in publications and rank them to the first 238 hits with 7 hits among the top 16 candidates. See figure 3 for a screenshot of the interface. Among the high-ranked

false positives we find frequently used therapeutic proteins like the plasminogen activator (PLAT), but also new true positives like the JAG1 gene, that have not been mentioned in the genetic reviews.

2.5 Generating Protein-Protein Interaction Networks

We used “Protein interactions and network analysis” (PIANA) (Aragues et al., 2006) to combine data from the Database of Interacting Proteins (DIP) (Salwinski et al., 2004), the MIPS database of interactions (Pagel et al., 2005), the Molecular INteractions database (MINT) (Chatr-aryamontri et al., 2007), IntAct (Kerrien et al., 2007), the Biomolecular Interactions Database (BIND) (Alfarano et al., 2005), the BioGrid (Stark et al., 2006) and Human Protein Reference Database (HPRD) (Peri et al., 2003) and the human interactions from two recent high-throughput experiments (Rual et al., 2005), (Stelzl et al., 2005). We also provide the interactions obtained from STRING (von Mering et al., 2005) and methods of protein-protein interaction prediction based on sequence/structure patterns (Espadaler et al., 2005), (Cockell et al., 2007). The integration of many different sources of interactions into a single repository allowed us to work with an extensive set of 363,571 interactions between 42,040 different protein sequences.

PIANA represents protein interaction data as a network where the nodes are proteins and the edges interactions between them. In such a network, a set of proteins linked to protein p_j (i.e. physically interacting with p_j) is named “partners of p_j ”. PIANA builds the network by retrieving partners for a initial set of seed proteins (i.e. the relevant proteins, here referred as “seed proteins”) that were obtained from the Find Candidate Genes module in section 2.4. A network is generated for the set of proteins that contains them and their partners. In this network, a protein that is connected to more than one “seed” is referred as a linker- N , with N being the number of seed proteins to which it is connected. Finally, proteins only connected to one seed protein are named leafs. This allowed us to enlarge the interaction network and detect new putatively relevant proteins for the biological pathway.

3 CONCLUSIONS

We have presented initial results on Decision Support, Database Mining and Knowledge Discovery for Intracranial Aneurysms. Due to the lack of patient data,

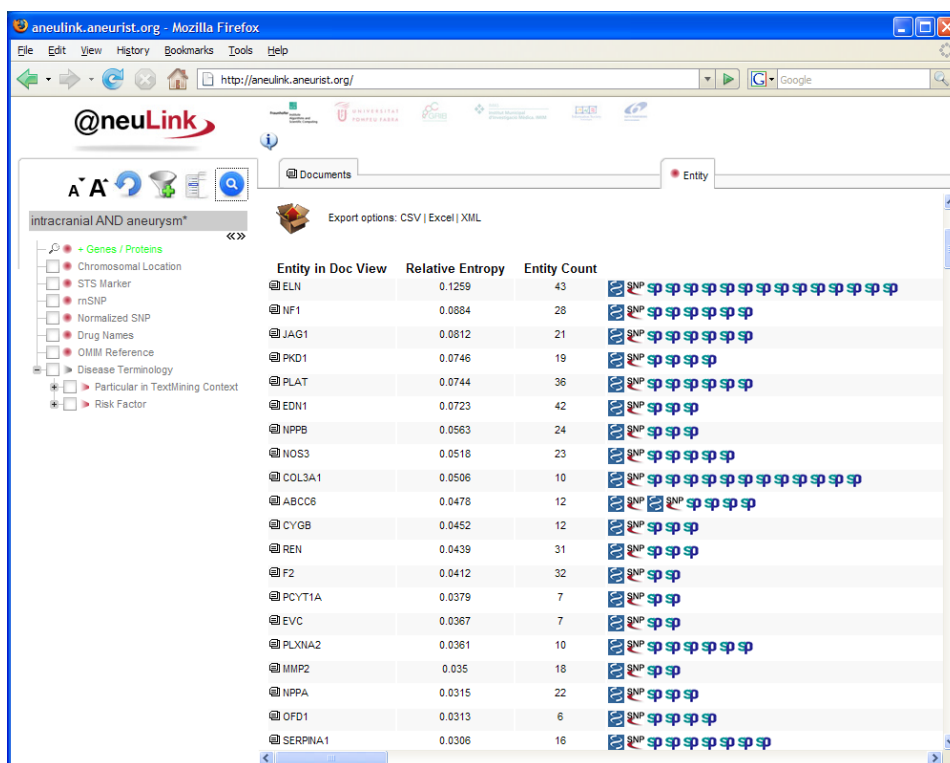


Figure 3: Example screen of the @neuLink Find Candidate Genes module.

at this phase of the project, we could only validate our findings on literature data.

Knowledge Discovery is defined in (Fayyad et al., 1996) as:

... the non-trivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data.

We do not claim to have fulfilled this definition with our initial solutions completely, but aim at the fulfillment in later stages of the project.

ACKNOWLEDGEMENTS

This work has been funded in the framework of the European integrated project @neurIST, which is co-financed by the European Commission through the contract no. IST-027703 (see <http://www.aneurist.org>). Additionally we like to acknowledge the excellent work of our co-workers: BeiBei Han, Roman Klinger, Holger Dach, Laura I. Furlong, Tobias Gattermayer and Oriol Fornes.

REFERENCES

Alfarano, C., Andrade, C. E., Anthony, K., Bahroos, N., Bajec, M., Bantoft, K., Betel, D., Bobechko, B., Boutilier, K., Burgess, E., Buzadzija, K., Cavero, R., D'Abreo, C., Donaldson, I., Dorairajoo, D., Dumontier, M. J., Dumontier, M. R., Earles, V., Farrall, R., Feldman, H., Garderman, E., Gong, Y., Gonzaga, R., Grytsan, V., Gryz, E., Gu, V., Haldorsen, E., Halupa, A., Haw, R., Hrvojic, A., Hurrell, L., Isserlin, R., Jack, F., Juma, F., Khan, A., Kon, T., Konopinsky, S., Le, V., Lee, E., Ling, S., Magidin, M., Moniakis, J., Montojo, J., Moore, S., Muskat, B., Ng, I., Paraiso, J. P., Parker, B., Pintilie, G., Pirone, R., Salama, J. J., Sgro, S., Shan, T., Shu, Y., Siew, J., Skinner, D., Snyder, K., Stasiuk, R., Strumpf, D., Tuekam, B., Tao, S., Wang, Z., White, M., Willis, R., Wolting, C., Wong, S., Wrong, A., Xin, C., Yao, R., Yates, B., Zhang, S., Zheng, K., Pawson, T., Ouellette, B. F. F., and Hogue, C. W. V. (2005). The biomolecular interaction network database and related tools 2005 update. *Nucleic Acids Res*, 33(Database issue):D418–D424.

Aragues, R., Jaeggi, D., and Oliva, B. (2006). Piana: protein interactions and network analysis. *Bioinformatics*, 22(8):1015–1017.

Arbona, A., Benkner, S., Engelbrecht, G., Fingberg, J., Hofmann, M., Kumpf, K., Lonsdale, G., and Woehrer, A. (2007). A service-oriented grid infrastructure for

- biomedical data and compute services. *IEEE Trans Nanobioscience*, 6(2):136–141.
- Bonis, J., Furlong, L. I., and Sanz, F. (2006). OSIRIS: a tool for retrieving literature about sequence variants. *Bioinformatics*, pages 2567–2569.
- Chatr-aryamontri, A., Ceol, A., Palazzi, L. M., Nardelli, G., Schneider, M. V., Castagnoli, L., and Cesareni, G. (2007). Mint: the molecular interaction database. *Nucleic Acids Res*, 35(Database issue):D572–D574.
- Clarke, M. (2007). Cochrane literature review snapshot v2.0 - aneurysm analysis rupture risk review. Deliverable of the @neurIST project D18, University of Oxford.
- Cockell, S. J., Oliva, B., and Jackson, R. M. (2007). Structure-based evaluation of in silico predictions of protein-protein interactions using comparative docking. *Bioinformatics*, 23(5):573–581.
- Druzdel, M. J. (1999). GeNIe: A development environment for graphical decision-analytic models. In *Proceedings of the 1999 Annual Symposium of the American Medical Informatics Association (AMIA-1999)*, page 1206.
- Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman & Hall/CRC.
- Espadaler, J., Romero-Isart, O., Jackson, R. M., and Oliva, B. (2005). Prediction of protein-protein interactions using distant conservation of sequence patterns and structure relationships. *Bioinformatics*, 21(16):3360–3368.
- Fayyad, U. M., Piatetsky-Shapiro, G., and Smyth, P. (1996). Knowledge Discovery and Data Mining: Towards a Unifying Framework. In *Knowledge Discovery and Data Mining*, pages 82–88.
- Fluck, J., Mevissen, H. T., Dach, H., Oster, M., and Hofmann-Apitius, M. (2007). ProMiner: recognition of human gene and protein names using regularly updated dictionaries. In (Hirschmann et al., 2007), pages 149–151.
- Fox, J. and Das, S. (2000). Safe and sound: Artificial intelligence in hazardous applications. In *Proceedings of AAAI 2000*.
- Fox, J., Johns, N., Rahmzadeh, A., and Thomsen, R. (1996). PROforma: A method and language for specifying clinical guidelines and protocols. In Brender, J., Christensen, J., Scherrer, J.-R., and McNair, P., editors, *Proceedings of the Medical Informatics Europe 1996*, pages 516–520. IOS Press.
- Fox, J., Patkar, V., and Thomson, R. (2006). Decision support for health care: the proforma evidence base. *Inform Prim Care*, 14(1):49–54.
- Gattermayer, T. (2007). SCAIView: annotation and visualization system for knowledge discovery. Master's thesis, Life Science Informatics at Bonn-Aachen International Center for Information Technology (B-IT); Germany.
- Ghinea, N. and van Gelder, J. M. (2004). A probabilistic and interactive decision-analysis system for unruptured intracranial aneurysms. *Neurosurg Focus*, 17(5):E9.
- Glasspool, D. W., Fox, J., Oettinger, A., and Smith-Spark, J. H. (2006). Argumentation in decision support for medical care planning for patients and clinicians. In *Proceedings of the AAAI Spring Symposium Series 2006*.
- Han, B., Friedrich, C. M., Manthey, R., and Hofmann-Apitius, M. (2006). Bayesian network modeling for the prediction of treatment outcomes of intracranial aneurysms. In Huson, D., Kohlbacher, O., Lupas, A., Nieselt, K., and Zell, A., editors, *Short Papers and Poster Abstracts of the German Conference on Bioinformatics (GCB2006)*, page 122.
- Hanisch, D., Fundel, K., Mevissen, H.-T., Zimmer, R., and Fluck, J. (2005). ProMiner: rule-based protein and gene entity recognition. *BMC Bioinformatics*, 6 (Suppl 1)(S14).
- Hirschmann, L., Krallinger, M., and Valencia, A., editors (2007). *Proceedings of the Second BioCreative Challenge Evaluation Workshop*. Centro Nacional de Investigaciones Oncológicas, CNIO.
- Hurt, C., Fox, J., Bury, J., and Saha, V. (2003). Computerised advice on drug dosage decisions in childhood leukaemia: a method and a safety strategy. In Dojat, M., Keravnou, E., and Barahona, P., editors, *Proceedings of the 9th Conference on Artificial Intelligence in Medicine in Europe (AIME'03)*, LNAI 2780, pages 158–163. Springer Verlag.
- Jensen, L. J., Saric, J., and Bork, P. (2006). Literature mining for the biologist: from information retrieval to biological discovery. *Nature Reviews; Genetics*, 7:119–129.
- Kerrien, S., Alam-Faruque, Y., Aranda, B., Bancarz, I., Bridge, A., Derow, C., Dimmer, E., Feuermann, M., Friedrichsen, A., Huntley, R., Kohler, C., Khadake, J., Leroy, C., Liban, A., Lieftink, C., Montecchi-Palazzi, L., Orchard, S., Risse, J., Robbe, K., Roechert, B., Thorneycroft, D., Zhang, Y., Apweiler, R., and Hermjakob, H. (2007). Intact–open source resource for molecular interaction data. *Nucleic Acids Res*, 35(Database issue):D561–D565.
- Klinger, R., Furlong, L. I., Friedrich, C. M., Mevissen, H. T., Fluck, J., Sanz, F., and Hofmann-Apitius, M. (2007). Identifying gene specific variants in biomedical text. *Journal of Bioinformatics and Computational Biology*, 5(6):in print.
- Krischek, B. and Inoue, I. (2006). The genetics of intracranial aneurysms. *J Hum Genet*, 51(7):587–594.
- Kullback, S. and Leibler, R. (1951). On information and sufficiency. *Ann. Math. Stat.*, 22:79–86.
- Lafferty, J., McCallum, A., and Pereira, F. C. N. (2001). Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*, pages 282–289. Morgan Kaufmann Publishers.
- Maglott, D., Ostell, J., Pruitt, K. D., and Tatusova, T. (2005). Entrez Gene: gene-centered information at NCBI. *Nucleic Acids Research*, 0:D1–D6.

- Molyneux, A. J., Kerr, R. S. C., Yu, L., Clarke, M., Sneade, M., Yarnold, J. A., and Sandercock, P. (2005). International subarachnoid aneurysm trial (ISAT) of neurosurgical clipping versus endovascular coiling in 2143 patients with ruptured intracranial aneurysms: a randomised comparison of effects on survival, dependency, seizures, rebleeding, subgroups, and aneurysm occlusion. *Lancet*, 366:809–817.
- Morgan, A. A. and Hirschmann, L. (2007). Overview of BioCreative II gene normalization. In (Hirschmann et al., 2007), pages 17–27.
- Nahed, B. V., Bydon, M., Ozturk, A. K., Bilguvar, K., Bayrakli, F., and Gunel, M. (2007). Genetics of intracranial aneurysms. *Neurosurgery*, 60(2):213–25; discussion 225–6.
- Pagel, P., Kovac, S., Oesterheld, M., Brauner, B., Dunger-Kaltenbach, I., Frishman, G., Montrone, C., Mark, P., Stümpflen, V., Mewes, H.-W., Ruepp, A., and Frishman, D. (2005). The mips mammalian protein-protein interaction database. *Bioinformatics*, 21(6):832–834.
- Patkar, V., Hurt, C., Steele, R., Love, S., Purushotham, A., Williams, M., Thomson, R., and Fox, J. (2006). Evidence-based guidelines and decision support services: A discussion and evaluation in triple assessment of suspected breast cancer. *Br J Cancer*, 95(11):1490–1496.
- Peri, S., Navarro, J. D., Amanchy, R., Kristiansen, T. Z., Jonnalagadda, C. K., Surendranath, V., Niranjan, V., Muthusamy, B., Gandhi, T. K. B., Gronborg, M., Ibarrola, N., Deshpande, N., Shanker, K., Shivashankar, H. N., Rashmi, B. P., Ramya, M. A., Zhao, Z., Chandrika, K. N., Padma, N., Harsha, H. C., Yatish, A. J., Kavitha, M. P., Menezes, M., Choudhury, D. R., Suresh, S., Ghosh, N., Saravana, R., Chandran, S., Krishna, S., Joy, M., Anand, S. K., Madavan, V., Joseph, A., Wong, G. W., Schiemann, W. P., Constantinescu, S. N., Huang, L., Khosravi-Far, R., Steen, H., Tewari, M., Ghaffari, S., Blobel, G. C., Dang, C. V., Garcia, J. G. N., Pevsner, J., Jensen, O. N., Roepstorff, P., Deshpande, K. S., Chinnaiyan, A. M., Hamosh, A., Chakravarti, A., and Pandey, A. (2003). Development of human protein reference database as an initial platform for approaching systems biology in humans. *Genome Res*, 13(10):2363–2371.
- Rinkel, G. J., Djibuti, M., Algra, A., and van Gijn, J. (1998). Prevalence and risk of rupture of intracranial aneurysms. *Stroke*, 29:251–256.
- Risselada, R., van der Lugt, A., Niessen, W. J., and Sturkenboom, M. C. J. M. (2007). Hormonal contraceptives and risk of aneurysmatic subarachnoid haemorrhage. *Basic & Clinical Pharmacology & Toxicology, Special Issue on 8th Congress of the European Association for Clinical Pharmacology and Therapeutics*, 101 (Suppl.):40.
- Rual, J.-F., Venkatesan, K., Hao, T., Hirozane-Kishikawa, T., Dricot, A., Li, N., Berriz, G. F., Gibbons, F. D., Dreze, M., Ayivi-Guedehoussou, N., Klitgord, N., Simon, C., Boxem, M., Milstein, S., Rosenberg, J., Goldberg, D. S., Zhang, L. V., Wong, S. L., Franklin, G., Li, S., Albala, J. S., Lim, J., Fraughton, C., Llamas, E., Cevik, S., Bex, C., Lamesch, P., Sikorski, R. S., Vandenhaute, J., Zoghbi, H. Y., Smolyar, A., Bosak, S., Sequerra, R., Doucette-Stamm, L., Cusick, M. E., Hill, D. E., Roth, F. P., and Vidal, M. (2005). Towards a proteome-scale map of the human protein-protein interaction network. *Nature*, 437(7062):1173–1178.
- Russel, S. and Norvig, P. (2003). *Artificial Intelligence a modern Approach*. Prentice Hall, 2nd edition.
- Salwinski, L., Miller, C. S., Smith, A. J., Pettit, F. K., Bowie, J. U., and Eisenberg, D. (2004). The database of interacting proteins: 2004 update. *Nucleic Acids Res*, 32(Database issue):D449–D451.
- Smigielski, E. M., Sirotkin, K., Ward, M., and Sherry, S. T. (2000). dbSNP: a database of single nucleotide polymorphisms. *Nucleic Acids Research*, 28(1):352–355.
- Stark, C., Breitkreutz, B.-J., Reguly, T., Boucher, L., Breitkreutz, A., and Tyers, M. (2006). Biogrid: a general repository for interaction datasets. *Nucleic Acids Res*, 34(Database issue):D535–D539.
- Stelzl, U., Worm, U., Lalowski, M., Haenig, C., Brembeck, F. H., Goehler, H., Stroedicke, M., Zenkner, M., Schoenherr, A., Koeppen, S., Timm, J., Mintzlaff, S., Abraham, C., Bock, N., Kietzmann, S., Goedde, A., Toksöz, E., Droege, A., Krobitch, S., Korn, B., Birchmeier, W., Lehrach, H., and Wanker, E. E. (2005). A human protein-protein interaction network: a resource for annotating the proteome. *Cell*, 122(6):957–968.
- Sutton, D. R. and Fox, J. (2003). The syntax and semantics of the proforma guideline modeling language. *J Am Med Inform Assoc*, 10(5):433–443.
- van der Lei, J., Duisterhout, J. S., Westerhof, H. P., van der Does, E., Cromme, P. V., Boon, W. M., and van Bemmel, J. H. (1993). The introduction of computer-based patient records in the netherlands. *Ann Intern Med*, 119(10):1036–1041.
- von Mering, C., Jensen, L. J., Snel, B., Hooper, S. D., Krupp, M., Foglierini, M., Jouffre, N., Huynen, M. A., and Bork, P. (2005). String: known and predicted protein-protein associations, integrated and transferred across organisms. *Nucleic Acids Res*, 33(Database issue):D433–D437.
- Wiebers, D., Whisnant, J., Huston, J., Meissner, I., Brown RD, J., Piepgras, G., Thielen, K., Nichols, D., O’Fallon, W., Peacock, J., Jaeger, L., Kassell, N., Kongable-Beckman, G., and Torner, J. (2003). Unruptured intracranial aneurysms: natural history, clinical outcome, and risks of surgical and endovascular treatment. *Lancet*, 362:103–110.

MY-HEART PROJECT: ANALYSIS OF SLEEP AND STRESS PROFILES FROM BIOMEDICAL SIGNAL

Sergio Cerutti, Anna M. Bianchi and Martin O. Mendez

*Department of Biomechanics, Polytechnic University of Milan, Piazza Leonardo da Vinci 32, Milan, Italy
sergio.cerutti@polimi.it, annamaria.bianchi@polimi.it, martin.mendez@biomed.polimi.it*

Keywords: Wearable Device, Heart Rate Variability, Sleep, Autonomic Nervous System, Stress, Arousal, Autoregressive Model, Hidden Markov Model.

Abstract: Advances in micro and nanotechnology, wireless technology, word-wide web networking, biomedical digital signal processing, textile tissue and implantable devices, etc. have permitted the development of alternative solutions for a better diagnosis of various pathologies, health care and prevention. These advances allow the remote continuous monitoring of persons, whenever and wherever they are. My-Heart FP6 EU Project integrates these components to develop a new concept of prevention and diagnosis for cardiovascular diseases. The present paper will mainly focus on a Concept of the Project which is called Take-Care and deals with the management, and integration of information for monitoring and personal motivation in the health care environment & applications. In its main goal, Take-Care Concept allows the learning in easy way about own self-body responses to different situations in the normal lifestyle: in particular, it offers an evaluation of life quality from sleep performance and stress management based on an integration of information between heart rate variability and respiration signals.

1 INTRODUCTION

My-Heart is a 6th Framework Programme EU Integrated Project for fighting cardiovascular diseases (CVD) by prevention and more objective diagnosis based upon vital signs processing in wearable devices. Cardiovascular diseases cause 50% of all deaths in the whole EU. It is well known that a healthy and active life style can significantly reduce the risk of cardiovascular diseases (primary prevention) and limits the recurrence rate of acute events (secondary prevention). Therefore, My-Heart aims at creating smart electronic solutions and appropriate services that empower users to take more control of their own health. The project addresses a spectrum of care, from prevention and the adoption of a healthier lifestyle to chronic disease management. The technological needs for My-Heart applications range from vital sign monitoring in daily life (ECG, respiration, skin impedance, etc), to body-worn low-power devices which run powerful and original detection algorithms for health status and acute cardiac event prediction, to low-power wireless links and server architectures for data handling at professional sites.

Inside My-Heart project a Take-Care concept is carried out that addresses people trying to find support to develop a healthier lifestyle. Today more and more people are facing problems: overweight,

inactivity, bad sleep and stress contribute to develop cardiovascular diseases and are affecting millions of European people.

Take-Care system provides people with a better understanding of their overall health, motivating them to become active in staying healthy and feeling well. It is the ambition of Take Care to provide easy-to-use technology and to develop solutions to manage bad sleep, stress, inactivity and overweight.

For the sleep management solution the focus is on analysing sleep stages and sleep fragmentation. For stress solution the focus is on providing the user a biofeedback tool to be used during relaxation exercises to more objectively measure the effect of the relaxation on the single subject.

In this context, the present paper describes the procedures for the evaluation of sleep quality and for stress management based on the analysis of the HRV and respiration signals. The analysis is thought for a device aimed to provide support to people who want to develop a healthier lifestyle, with major focus on cardiovascular disease prevention.

2 METHODOLOGY

Take-Care system (TCS) is integrated by different modules. Each module is developed by taking into account the analysis of a specific function. Sleep

fragmentation, sleep staging, obstructive sleep apnoea and stress are the functions singled out inside TCS. Firstly, a general scheme of TCS will be presented. Then, in subsection 2.1, the basic description of sleep fragmentation module is presented. Sleep staging module will be the argument of subsection 2.2. Sleep apnoea module is presented in subsection 2.3 and finally subsection 2.4 concerns with relaxation and stress module.

The main goal of TCS project is to improve life quality based on sleep performance. Sleep analysis is carried out by assessing different electrophysiological signals, which in general require well trained and specialized personnel for the signal interpretation and dedicated equipment. TCS takes advantage from peripheral signals of easy acquisition and sufficient signal to noise ratio. TCS is concentrated in the analysis of surface electrical activity of the heart (electrocardiogram, ECG), since it exhibits high sensitivity to different

physiologic and pathologic behaviours, and respiration. Figure 1 shows a TCS general scheme.

2.1 Sleep Fragmentation Module

Arousal from sleep (AS) has been one of the most studied sleep phenomena related to sleep fragmentation (SF). SF is associated with several symptoms, ranging from somnolence, excessive daytime sleepiness, impaired learning and memory capabilities, up to much more severe consequences such as cardiovascular diseases. When SF is associated with sleep-disordered breathings, such as obstructive sleep apnoea (OSA), there is high likelihood to develop arterial hypertension and other cardiovascular diseases.

AS is normally scored from either the central or occipital leads of the electroencephalogram (EEG) during standard polysomnographic studies. An AS consists in “an abrupt shift in EEG frequency, which may include theta, alpha and/or frequencies

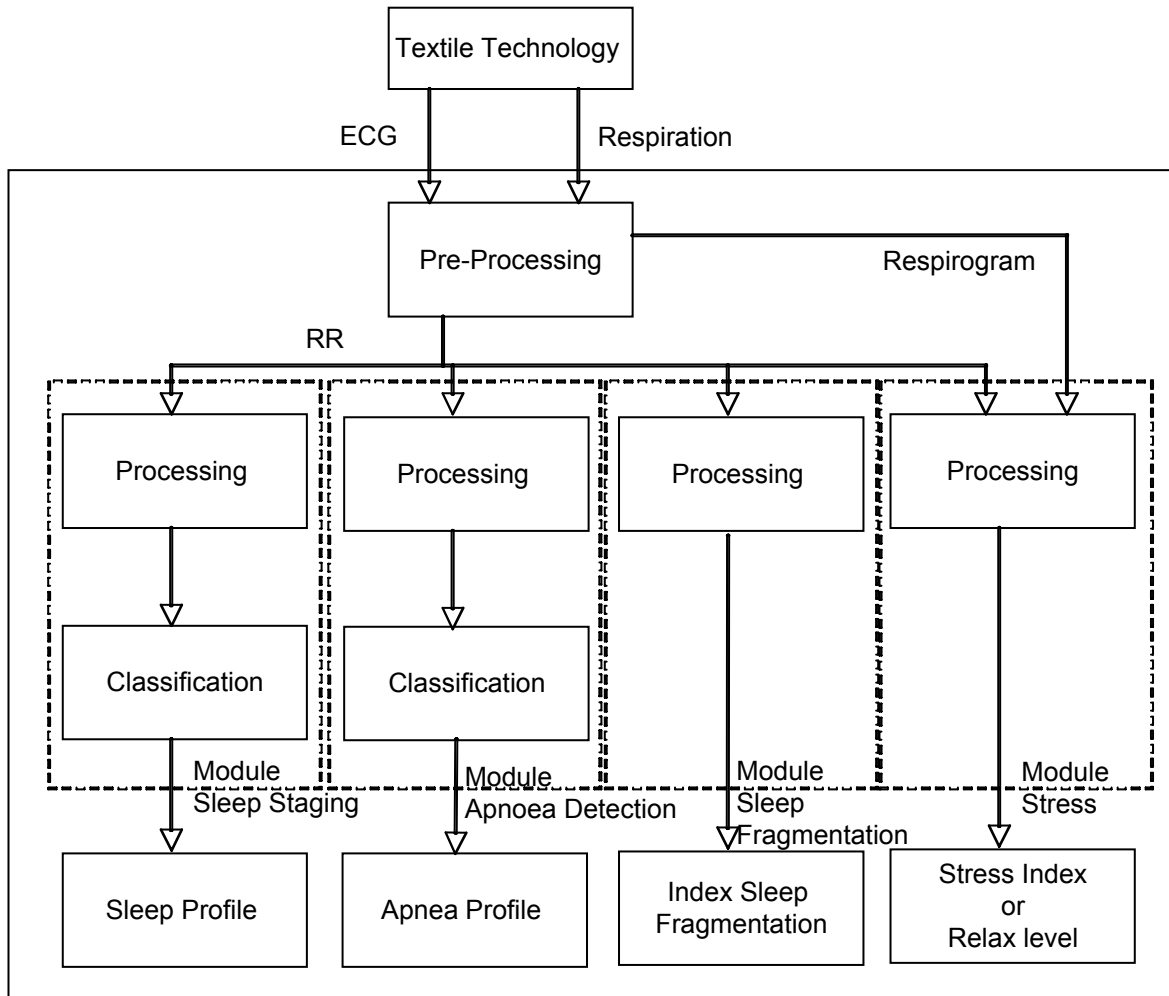


Figure 1: Take-Care System (TCS) scheme.

greater than 16 Hz but not spindles” (Atlas Task Force, 1992). AS generates a typical waveform in the heart rate (HR), which consists in an abrupt increment of the HR, followed few seconds later by a decrement in the HR (Sforza et al., 2000). Based on this knowledge, a search algorithm was implemented to identify the cortical arousal projection into the heart rhythm. This algorithm searches in the heart rate the pattern illustrated in Figure 2.

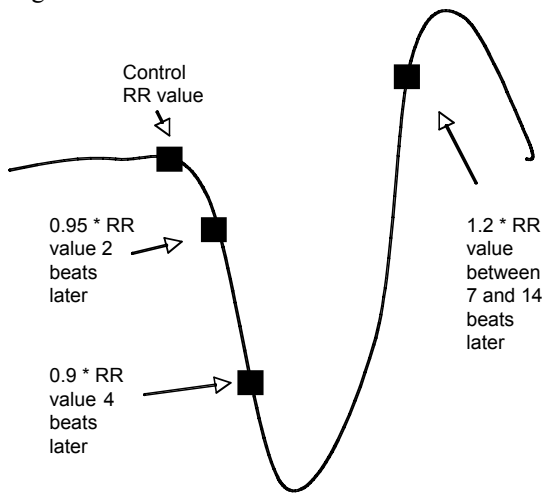


Figure 2: Arousal detection from RR intervals.

At each beat, in this case called control value, the algorithm observes if two and four beats later there is a decrement in the HR, this decrement must obey a predefined percentage. If this condition exists, the algorithm moves ahead and tries to find a HR value higher than the control one between the beats 7 and 14. If this condition is present, then an autonomic arousal is found and the process restart 20 beats later.

Sleep fragmentation is based on the number of arousal episodes during the whole night. However, the sleep process is affected in a different way if an arousal episode occurs during the first or last NREM-REM sleep cycle. A simple solution, that respects the physiologic impact of an arousal during the sleep process, is weighting differently arousal events depending of the time in which these occur. A proposed index (Sleep Fragmentation Index, SFI) is considered by splitting the total sleep time into three parts and is computed as follows:

$$SFI = 3 * (\#A's \text{ in the first segment}) + (\#A's \text{ in the second segment}) + 0.33 * (\#A's \text{ in the third (1) segment})$$

Where #A is the number of arousals in a specific time of the night. This equations gives high weight if an arousal episode occurs during the first sleep hours

and less weight if it occurs during the last sleep hours.

2.2 Sleep Staging Module

In the modern hospital Sleep Centers, the standard sleep evaluation procedure consists in the acquisition and analysis of three signals: electroencephalogram, muscular activity and electrooculogram. These signals present specific changes in time that characterize six different stages: wake, sleep stage 1, sleep stage 2, sleep stage 3, sleep stage 4 and REM. However, sleep stages 1 to 4 present similar characteristics that allow to regroup them all together in only one sleep stage, which is labelled NREM. With standard leads, at peripheral level, NREM and REM sleep present specific patterns, that with some grade of confidence, are significantly recognized. For instance, during NREM, heart rate is stable, eyes do not present movements, there is muscular tone and respiration is regular and deep. In contrast, during REM sleep, heart rate is instable, eyes present circular movements, there is no muscular tone and respiration exhibits instability (Guyton, 2000). Therefore, appropriate mathematical tools of signal processing and pattern recognition might possibly identify NREM and REM sleep periods on the basis of these characteristics.

2.2.1 Processing

This module is built up by two blocks. The first block is *processing*, in which extraction of the features that characterize REM and NREM sleep is carried out. During sleep, HR presents different behaviours as well as a series of non-stationarities. These characteristics in HR during sleep require necessarily the application of special mathematical tools. Wavelets, time-frequency and time-varying approaches have interesting properties to obtain the spectral features of the HR in the most diverse conditions. TCS uses a time-varying autoregressive model to extract important spectral features of the HR during sleep. The selection for this approach resides in its characteristic of real-time processing, high time and frequency resolution and very low computational cost.

Figure 3 shows an example where sleep stage 2, 4 and REM are analysed in the frequency domain. The power spectrum was obtained by a time varying autoregressive filter, which evaluates beat by beat the frequency content of a time series.

From a large amount of possible features to classify REM and NREM sleep, TCS uses only four features to discriminate between them:

- RR mean;
- Very low frequency component in RR;
- Modulus of the pole of the autoregressive model in the high frequency component;
- Phase of the pole of the autoregressive model in the high frequency component.

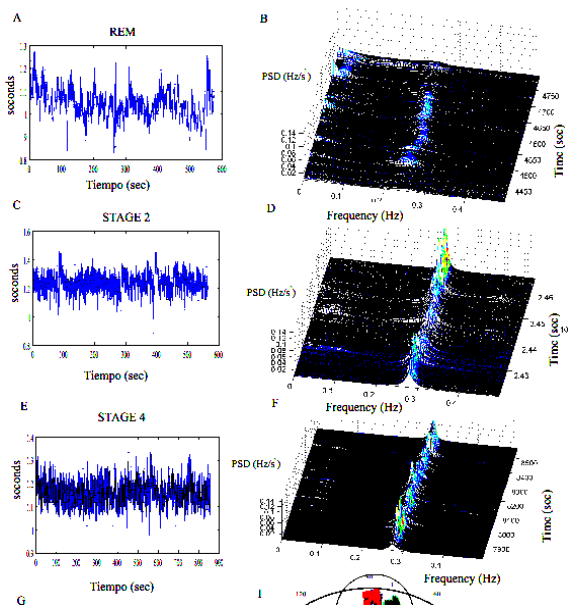


Figure 3: Time-varying spectral analysis of RR series in REM, stage 2 and stage 4 during sleep (A, C and E), with the respective power spectral densities (B, D and F).

2.2.2 Classification

A Hidden Markov Model classifier (HMM) integrates the second block, *classification*. HMM was selected since this classifier takes into account the temporal dynamic of a time series. This characteristic is very important for this study, since the length of the NREM-REM cycles varies in time. Figure 4 present an example where the selected features and the sleep profile are compared with the clinical hypnogram. The sleep stages are defined as: 0 = wake, 3 = NREM and 5 = REM. The classification was evaluated epoch by epoch of 30 sec according to the traditional clinical use.

2.3 Sleep Apnoea Module

Sleep apnoea is one the most common sleep pathologies. Only in the USA, this pathology affects sleep apnoea does not present evident symptoms, most of the time, sleep apnoea is undiagnosed. Sleep apnoea produces consequences that range from daily somnolence to heart failure. Sleep apnoea is divided into three different types: obstructive, central and mixed. Obstructive apnoea is the most common

sleep pathology and with more severe consequences and direct influences in the cardiovascular function.

Sleep apnoea is generated by an occlusion in the upper airways, which does not permit the air to enter to the lungs. Generally, this occlusion is produced by an inadequate sympathetic activation, decrement of the upper airways radio by obesity or exposition to agents as smog. The duration of a single obstructive apnoea episode is between 10 and 20 seconds. During this time, oxygen saturation decreases and respiratory efforts begin to increase in time in order to re-open the upper airways. If oxygen level in the blood decreases in such a way that the respiratory efforts are not enough to restore respiration, the central nervous system produces an arousal, which reactivates all the systems, and the respiration is restored. This process could occur hundreds of times during a single night: in this case a bad sleep quality and cardiac consequences are evidenced, giving as results social problems, accidents and heart attacks (Young et al., 2002).




Figure 4: Example of NREM-REM classification by Take-Care System. From the top to the bottom: RR intervals, very low frequency power, module and phase of the pole in high frequency, sleep profile evaluated by Take-Care System and clinical hypnogram. The sleep stages are defined as: 0 = wake, 3 = NREM and 5 = REM.

This process produces a defined bradycardia pattern in the heart rate. TCS explores the feasibility for obtaining an alternative detector of apnoeas using the same tools presented in the study for sleep staging. As autoregressive models present dedicated characteristic for evaluating features in dynamic time series, it was implemented as feature extractor. However, for a more correct detection of apnoeas, some new features were explored: i.e., an estimation of the respiratory effort was obtained from the calculation of the area of the QRS complexes. Correlation between both time series was used in order to extract more robust features to classify obstructive apnoea.

2.3.1 Processing

Also sleep apnoea module is built up from two blocks. The *processing* block helps to extract the features that separate apnoea and non-apnoea conditions. Again, an autoregressive model was used to extract the features. However, each problem requires the selection of its own features to solve it. In this case, the set of features is formed by:

- RR mean;
- Very low frequency component in RR;
- Very low frequency component in derived respiratory signal from ECG;
- Coherence between very low frequency components of the two signals.

2.3.2 Classification

In this block a K-nearest neighbour classifier is used to separate between apnoea and non-apnoea periods. The classification is evaluated in a minute-by-minute basis. From here it is possible to obtain an estimate of the time that a person spends in apnoea during the sleep time. Figure 5 shows an example of sleep apnoea classification for 25 subjects.




Figure 5: Class separation based on minutes per night calculated by the KNN classifier processing 4 features for 25 recordings of the testing group. Note that applying a threshold of 50 minutes per night, apnoea and normal classes are accordingly separated.

2.4 Stress Module

The stress concept employs a sensorised T-shirt which allows the continuous recordings of ECG and respiration. Previously, the subject is asked to fill in a questionnaire which allows to detect his/her level of stress, according to a clinical classification score (APA, 1994). A kind of personalised initial level of stress is hence obtained (IS). Then, according to the instructions delivered from the computer, the subject makes a rest-to-stand manoeuvre (from sitting to a standing position) which indicates the degree of responsiveness to a predominantly sympathetic stimulation.

Then, according to the computer indications, the subject makes some relaxation exercises, which consist in deep regular respirations cycles, trying to “synchronise” as much as possible cardiorespiratory activity. On the basis of Heart Rate Variability parameters (HRV) and respiration, as well as on bivariate magnitudes calculated from the signals, it is possible to measure the personalised physiological effects of training sessions after days or weeks of treatment and hence to re-position the subject possibly in another location of the stress level plane which started from IS.

It is advised that the subject could do this exercise on a regular base (i.e. once per day) and hence there is the possibility to monitor his/her level of stress from the responses of his/her vital signs. After a proper coaching it is believed that through such exercises the subject could monitor his/her level of stress and these objective measurements could be important elements for helping physicians in a better diagnosis and treatment follow-up of stress related cardiac pathologies.

3 CONCLUSIONS

Take Care Concept has developed original applicative tools, implemented through advanced technological implementations (textiles, microelectronics storing and controlling devices, modern wireless communication protocols, etc) in order to provide a precious instrument of prevention of cardiovascular pathologies. The basic philosophy is to detect from subject’s vital signs physiological and clinical parameters even in continuous recordings, by employing easy-to-use wearable devices which allow comfortable home or ambulatory applications.

REFERENCES

- Atlas Task Force of American Sleep disorders Association
1992. EEG arousal: scoring rules and examples. *Sleep*
15: 174-184.
- Sforza E., Jouny C., Ibanez V., 2000. Cardiac activation
during arousal in humans: further evidence for
hierarchy in the arousal response. *Clin Neurophysiol*
111:1611-1619.
- Guyton AC., 2000. Textbook of medical physiology, W.B.
Saunders Company 10th edition.
- Young T., Peppard PE., Gottlieb DJ., 2002. Epidemiology
of obstructive sleep apnea. *Am J Respir Crit Care*
Med. 165:1217-1231.
- American Psychiatric Society (APA) , 1994. Diagnostic
and Statistical Manual of Mental Disorders IV
(DSMIV). *American Psychiatric Society*. Washington,
DC, 4th edition.

A SEMANTIC GRID SERVICES ARCHITECTURE IN SUPPORT OF EFFICIENT KNOWLEDGE DISCOVERY FROM MULTILEVEL CLINICAL AND GENOMIC DATASETS

Manolis Tsiknakis, Stelios Sfakianakis

*Institute of Computer Science, Foundation for Research and Technology - Hellas
GR-71110 Heraklion, Crete, Greece
{tsiknaki, ssfak}@ics.forth.gr*

Stefan Rueping

*Fraunhofer IAIS, Schloss Birlinghoven, 53754 St. Augustin, Germany
stefan.rueping@iais.fraunhofer.de*

Oswaldo Trelles

*Computer Architecture Dept., University of Malaga, Bulevar Louis Pasteur 45, 29010, Malaga, Spain
ots@ac.uma.es*

Thierry Siestang

*Swiss Institute of Bioinformatics, Bâtiment GénopodeCH-1015 Lausanne, Switzerland
thierry.sengstag@isb-sib.ch*

Brecht Claerhout

*Custodix NV, Verlorenbroodstraat 120b14, Merelbeke, Belgium
brecht.claerhout@custodix.com*

Vasiliis Virvilis

*Biovista S.A. 34 Rodopoleos Street, Ellinikon, Athens 1677, Greece
vassilis.virvilis@biovista.com*

Keywords: Ontology-based Biomedical Database Integration, Semantic Mediation, Ontologies, Post-genomic Clinical Trials, Service Oriented Architectures.

Abstract: This paper presents the architectural considerations of the Advancing Clinico-Genomic Trials on Cancer (ACGT) project aiming at delivering a European Biomedical Grid in support of efficient knowledge discovery in the context of post-genomic clinical trials on cancer. Our main research challenge in ACGT is the requirement to develop an infrastructure able to produce, use, and deploy knowledge as a basic element of advanced applications, which will mainly constitute a Biomedical Knowledge Grid. Our approach to offer semantic modelling of available services and data sources to support high level services and dynamic services for discovery and composition will be presented. In particular, ontologies and metadata are the basic elements through which Grid intelligence services can be developed, and the current achievements of the project in this domain will be discussed.

1 INTRODUCTION

Life sciences are currently at the centre of an informational revolution. Dramatic changes are being registered as a consequence of the development of techniques and tools that allow the

collection of biological information at an unprecedented level of detail and in extremely large quantities.

The nature and amount of the information now available open directions of research that were once in the realm of science fiction. Pharmacogenomics (Roses, 2000), diagnostics (Sotiriou, 2007) and drug

target identification (Schuppe-Koistinen, 2007) are just a few of the many areas that have the potential to use this information to change dramatically the scientific landscape in the life sciences.

During this informational revolution, the data gathering capabilities have greatly surpassed the data analysis techniques. If we were to imagine the Holy Grail of life sciences, we might envision a technology that would allow us to fully understand the data at the speed at which it is collected. Ideally, we would like knowledge manipulation to become tomorrow the way goods manufacturing is today: highly automated, producing more goods, of higher quality and in more cost effective manner than manual production. It is our belief that, in a sense, knowledge manipulation is now reaching its pre-industrial age. The explosive growth in the number of new and powerful technologies within proteomics and functional genomics can now produce massive amounts of data but using it to manufacture highly processed pieces of knowledge still requires the involvement of skilled human experts to forge through small pieces of raw data one at a time. The ultimate challenge in coming years, we believe, will be to automate this knowledge discovery process.

This paper presents a short background section discussing the urgent needs faced by the biomedical informatics research community, and very briefly describes the clinical trials upon which the ACGT project is based for both gathering and eliciting requirements and also for validating the technological infrastructure designed. It continues with a presentation of the initial ACGT architecture defined, and presents its layers and key enabling services.

2 POST-GENOMIC CLINICAL TRIALS

In ACGT we focus in the domain of clinical trials on cancer. Cancer, being a complex multifactorial disease group that affects a significant proportion of the population worldwide, is a prime target for focused multidisciplinary efforts using currently available novel, high throughput and powerful technologies. Exciting new research on the molecular mechanisms that control cell growth and differentiation has resulted in a quantum leap in our understanding of the fundamental nature of cancer cells.

While these opportunities exist, the lack of a common infrastructure has prevented clinical

research institutions from being able to mine and analyze disparate, multi-level data sources. As a result, very few cross-site studies and multi-centric clinical trials are performed and in most cases it isn't possible to seamlessly integrate multi-level data.

It is well established that patient recruitment is often the time-limiting factor for clinical trials. As a result, clinical trials are gradually turning multi-centric to limit the time required for their execution (Sotiriou, 2007).

The ACGT project has been structured within such a context. It has selected two cancer domains and has defined three specific trials. These trials serve a dual purpose. Firstly, they are used for developing a range of post-genomic analytical scenarios for feeding the requirement analysis and elicitation phase of the project, and secondly they will be used for the validation of the functionality of the ACGT technologies.

The ACGT trials are in the domain of Breast Cancer and Wilm's Tumor (pediatric nephroblastoma). Specifically:

- The ACGT Test of Principle (TOP) study aims to identify biological markers associated with pathological complete response to anthracycline therapy (epirubicin), one of the most active drugs used in breast cancer treatment (Sotiriou, 2003).
- Wilms' tumour, although rare, is the most common primary renal malignancy in children and is associated with a number of congenital anomalies and documented syndromes (Graf, 2007).

In addition to these two clinical trials and on the basis of data collected for the purpose of their execution, and in-silico modelling and simulation experiment is also planned. The aim of this experiment is to provide clinicians with a decision support tool able to simulate, within defined reliability limits, the response of a solid tumour to therapeutic interventions based on the individual patient's multi-level data (Stamatakis, 2007).

2.1 Technical Challenges

ACGT's vision is to become a pan-European voluntary network connecting individuals and institutions and to enable the sharing of data and tools (see figure 1). In order to achieve its goals and objectives, ACGT is creating an infrastructure for cancer research by using a virtual web of trusted and interconnected organizations and individuals to leverage the combined strengths of cancer centers and investigators and enable the sharing of

biomedical cancer-related data and research tools in a way that the common needs of interdisciplinary research are met and tackled (Tsiknakis, 2006, Tsiknakis, 2007b).

Considering the current size of clinical trials (hundreds or thousands of patients) there is a clear need, both from the viewpoint of the fundamental research and from that of the treatment of individual patients, for a data analysis environment that allows the exploitation of this enormous pool of data generated.

As a result, a major part of the project is devoted to research and development in infrastructure components that are gradually been integrated into a workable demonstration platform upon which the selected (and those to be selected during the lifecycle of the project) clinical studies will be demonstrated and evaluated against user requirements defined at the onset of the project.

2.2 Scientific and Functional Requirements

The real and specific problem that underlies the ACGT concept is co-ordinated resource sharing and problem solving in dynamic, multi-institutional, pan-European virtual organisations. A set of individuals and/or organisations defined by such sharing relationships form what we call “an ACGT virtual organisation (VO)”. Simply stated, the participants in a multi-centric clinical trial form a VO, which exists for the duration of a trial or for any other period of time based on mutual agreements.

The task, therefore, of ACGT is to make data and tools securely available in this inter-enterprise environment where and when needed to all authorised users. As a result, the scientific and functional requirements for the ACGT platform can be summarised as follows:

- **Virtual Organisation Management:** support for the dynamic creation a VOs, defined as a group of individuals or institutions who share the computing and other resources of a "grid" for a common goal.
- **Data federation:** seamless navigation across and access to heterogeneous data sources, both private and public.
- **Data integration:** the capacity to pool data from heterogeneous sources in a scientifically, semantically and mathematically consistent manner for further computation.
- **Shared services:** the development, sharing and integration of relevant and powerful data exploitation tools such as tools for

bioinformatics analysis, data mining, modelling and simulation.

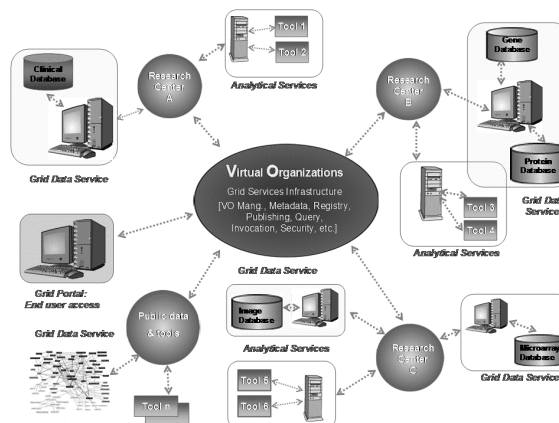


Figure 1: The vision of ACGT. Creating and managing Virtual Organisations on the Grid who are jointly participating in the execution of multicentric, post-genomic Clinical Trials.

The requirements elicitation process that has taken place in the project, based on input for a diverse range of users has resulted in the identification of the following key technical requirements.

- **Flexibility;** in other words modularity (supporting integration of new resources in a standardised way) and configurability (accommodating existing and emerging needs). This is required because (a) The a priori scientific and functional requirements are broad and diverse; (b) The data resources to be federated by the ACGT platform are characterised by deep heterogeneities in terms of source, ownership, availability, content, database design, data organisation, semantics and so on; and (c) the complexity of the underlying science, as well as the complexity of applicable knowledge representation schemas and applicable scientific algorithms;
- **Intuitive access to information;** From the user’s point of view, the ACGT knowledge management platform must provide relevant and simple access to information – both in terms of searching and navigation – and to services. In addition, it must provide a dynamically evolving set of validated data exploration, analysis, simulation, and modelling services.
- **Security;** Finally, it must be consistent with the European ethical and legal framework,

providing a high degree of trust and security to its users.

3 THE ACGT INITIAL ARCHITECTURE

In principle, the requirements for the ACGT platform can be met by designing a federated environment articulating independent tools, components and resources based on open architectural standards, which is customizable and capable of dynamic reconfiguration.

In order to fulfill the requirements imposed by scenarios identified in the ACGT project a heterogeneous, scalable and flexible environment is needed and the following technologies, which have gained momentum in the recent years, have been considered for adoption:

- Web Services technologies
- Grid technologies
- Semantic web technologies

Although initially separated, these technologies are currently converging in a complementary way.

Considering that the amount of data generated in the context of post-genomic clinical trials is expected to rise to several gigabytes of data per patient in a close future access to high-performance computing resources will be unavoidable. Hence, Grid computing (Foster, 2001) appears as a promising technology. Access and use of Grid-based resources is thus an integral part of the design of the infrastructure.

From the technical point of view, the requirements identified can be met using a distributed/federated, multi-layer, service oriented, and ontology driven architecture. The ACGT project decided to build on open software frameworks based on WS-Resource Framework (WSRF) and Open Grid Service Architecture (OGSA), the de facto standards in Grid computing.

Building on concepts and technologies from both the Grid and Web services communities, OGSA defines uniform exposed service semantics (the Grid service); defines standard mechanisms for creating, naming, and discovering transient service instances; provides location transparency and multiple protocol bindings for service instances; and supports integration with underlying native platform facilities. These standards are implemented in the middleware selected, namely Globus Toolkit 4 (GT4 - <http://www.globus.org/>).

An overview of the ACGT system layered architecture is given in Fig. 2, which is shortly presented in the sequel.

A layered approach has been selected for providing different levels of abstraction and a classification of functionality into groups of homologous software entities (Tsiknakis, 2007a, Rueping, 2007). In this approach we consider the security services and components to be pervasive throughout ACGT so as to provide both for the user management, access rights management and enforcement, and trust bindings that are facilitated by the Grid and domain specific security requirements like pseudonymization and anonymization.

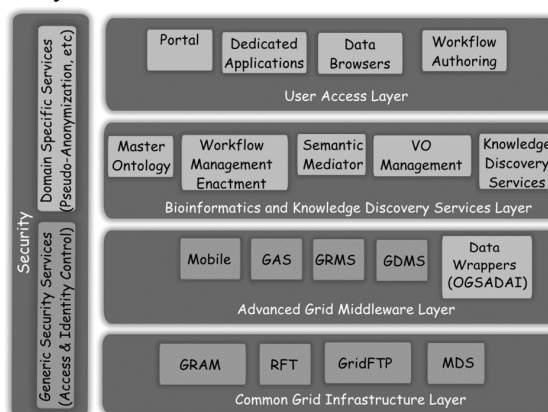


Figure 2: The ACGT layered architecture and its main services.

In specifying the initial architecture of the ACGT technological platform, architectural specifications of other relevant projects have been thoroughly studied. Of particular relevance are the Cancer Biomedical Informatics Grid (caBIG - <https://cabig.nci.nih.gov/>) in the US and the CancerGrid (<http://www.cancergrid.eu/>) project in the UK.

3.1 Heterogeneous Biomedical Databases Integration

Distributed and heterogeneous databases, created in the context of multi-centric post-genomic clinical trials on cancer, need to be seamlessly accessible and transparently queried in the context of a user's discovery driven analytical tasks. A central challenge, therefore, to which ACGT needs to respond, is the issue of semantic integration of heterogeneous biomedical databases.

The process of heterogeneous database integration may be defined as “*the creation of a single, uniform query interface to data that are collected and stored in multiple, heterogeneous databases.*” Several varieties of heterogeneous database integration are useful in biomedicine. The most important ones are:

- **Vertical integration.** The aggregation of semantically similar data from multiple heterogeneous sources. For example, a “virtual repository” that provides homogeneous access to clinical data that are stored and managed in databases across a regional health information network is reported in (Katehakis, 2007) and (Lesch, 1997).
- **Horizontal integration.** The composition of semantically complementary data from multiple heterogeneous sources. For example, systems that support complex queries across genomic, proteomic, and clinical information sources for molecular biologists are reported in (Stevens, 2000) and (Gupta, 2000).

The approach adopted in ACGT is based on the use of domain ontologies, acting as the global schema in a Local-as-View (LAV) integration methodology. Detailed presentation of the data integration architecture of the project and the tools and services utilized for this purpose is outside the scope of this paper. Such a detailed presentation of the data integration architecture can be found at (Anguita, 2007).

4 KNOWLEDGE DISCOVERY SERVICES

Once these multilevel clinical and genomic data are integrated, they can be mined to extract new knowledge that can be useful in topics such clinical diagnosis, therapy, prevention and, of course, the design of new studies (such as in the case of ACGT, clinico-genomic trials).

Knowledge discovery in clinico-genomic data presents a new array of challenges since it differs significantly from the original problems of data analysis that prompted the development of Grid technologies. The exploitation of semantics information in the description of data sources and data analysis tools is of high importance for the effective design and realization of knowledge discovery processes. Semantics are usually made concrete by the adoption of metadata descriptions and relevant vocabularies, classifications, and

ontologies. In ACGT these semantics descriptions are managed by the Grid infrastructure and therefore the knowledge discovery services build and operate on a Knowledge Grid platform (Cannataro, 2003).

4.1 Workflows

The Workflow Management Coalition (WFMC, <http://www.wfmc.org/>) defines a workflow as “The automation of a business process, in whole or part, during which documents, information or tasks are passed from one participant to another for action, according to a set of procedural rules”. In other words a workflow consists of all the steps and the orchestration of a set of activities that should be executed in order to deliver an output or achieve a larger and sophisticated goal. In essence a workflow can be abstracted as a composite service, i.e. a service that is composed by other services that are orchestrated in order to perform some higher level functionality.

The aim of the ACGT workflow environment is to assist the users in their scientific research by supporting the ad hoc composition of different data access and knowledge extraction and analytical services into complex workflows. This way the users can extend and enrich the functionality of the ACGT system by reusing existing ACGT compliant services and producing “added value” composite services. This reuse and composition of services is in some sense a programming task where the user actually writes a program to realize a scenario or to test a scientific hypothesis.

In order to support the ACGT users to build and design their workflows a visual workflow programming environment has been designed.

It is a web based workflow editor and designer that is integrated into the rest of ACGT system so as to take advantage of the Grid platform and the ACGT specific infrastructure and services. In particular, this workflow designer features a user friendly Graphical User Interface (GUI) that supports the efficient browsing and searching of the available ACGT services and their graphical interconnection and manipulation to construct complex scientific workflows. The choice of a graphical representation of the workflow and the support for ‘point-and-click’ handling of the workflow graph was made on the basis that this is more intuitive for the users and increases their productivity. Additional features that also take advantage of the metadata descriptions of services include the validation in the design phase of the workflows in order to reduce or even eliminate the

developed to provide a single point of reference for these concepts and to support reasoning of concept expressions.

5 THE ACGT SECURITY FRAMEWORK AND ITS SERVICES

We recognise that the sharing of multilevel data outside the walls of a hospital or a research organisation generates complex ethical and legal issues. It is also well known that the concerns around “security issues” have been one of the major obstacles that have inhibited wider adoption of information technology solutions in the healthcare domain. As a result we have devoted significant efforts in the study and analysis of the ethical and legal issues related to cross-institutional sharing of post-genomic data sets.

Based on such an approach we concluded that trust and security must to be addressed at multiple levels; these include (a) infrastructure, (b) application access, (c) data protection, (d) access control, which would be policy-governed, and (e) privacy-enhancing technology, such as de-identification.

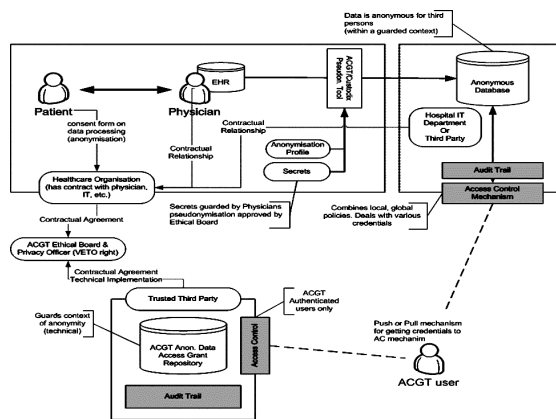


Figure 4: Overview of the ACGT security framework – actors, procedures and technological services.

The European Directive on Data Protection (http://www.cdt.org/privacy/eudirective/EU_Directive_.html) deals with the protection of personal data and imposes many restrictions on its use. In order to allow ACGT partners to handle and exchange medical data in conformance with the requirements of European Directive on Data Protection, an advanced Data Protection Framework has been designed. This framework (illustrated on figure 4)

achieves this goal through an integrated approach that includes technical requirements but also policies and procedures. Some of the aspects of the Data Protection framework are (a) Anonymization or pseudonymisation of the data, (b) a Trusted Third Party (TTP) pseudonymisation and a corresponding pseudonymisation tool, (c) technology supported measures to control the anonymity context, (d) an ACGT data protection board (acting as a Trusted Third Party) responsible for issuing credentials for data access to authorised users, and (e) definition of the necessary consent forms and legal agreements that need to be signed by all members of any ACGT Virtual Organisation.

Description of the technical details of the security architecture of ACGT (the data protection framework) goes beyond the scope of the current article. Nevertheless, the main message that we want to stress is the fact that a well designed set of both technological as well as procedural measures have been taken, so that a high degree of trust and security is build in the final infrastructure to be delivered.

6 CREATING AND SHARING ACGT COMPLIANT SERVICES

Achieving the level of automation, that is graphically depicted in Figure 3, requires the creation of highly interoperable services. In turn creating a service involves describing, in some conventional manner, the operations that the service supports; defining the protocol used to invoke those operations over the Internet; and operating a server to process incoming requests (Foster, 2005).

Although a fair amount of experience has been gained with the creation of services and applications in different science domains, significant problems do still remain, especially with respect to interoperability, quality control and performance. These are issues to which ACGT focuses, and these are briefly discussed in the next subsections.

6.1 Interoperability and Re-use

Services have little value if others cannot discover, access, and make sense of them. Yet, as Stein has observed (Stein, 2002), today’s scientific communities too often resemble medieval Italy’s collection of warring city states, each with its own legal system and dialect. Available technological (i.e. Web services) mechanisms for describing,

discovering, accessing, and securing services provide a common alphabet, but a true lingua franca requires agreement on protocols, data formats, and ultimately semantics (d. Roure, 2003). In the ACGT project we are paying particular attention on these issues, and especially on the issue of semantics (see section on metadata).

6.2 Management

In a networked world, any useful service will become overloaded. Thus, we need to control who uses services and for what purposes. Particularly valuable services may become community resources requiring coordinated management. Grid architectures and software can play an important role in this regard and ACGT is focusing on exploiting these opportunities made available by Grid computing.

6.3 Quality Control

As the number and variety of services grow and interdependencies among services increase, it becomes important to automate previously manual quality control processes—so that, for example, users can determine the provenance of a particular derived data product (Goble, 2004). The ability to associate metadata with data and services can be important, as can the ability to determine the identity of entities that assert metadata, so that consumers can make their own decisions concerning quality.

7 DISCUSSION AND CONCLUSIONS

In this paper, we consider a world where biomedical software modules and data can be detected and composed to define problem-dependent applications. We wish to provide an environment allowing clinical and biomedical researchers to search and compose bioinformatics and other analytical software tools for solving biomedical problems. We focus on semantic modelling of the requirements of such applications using ontologies.

The project has conceived an overall architecture for an integrating biomedical sciences platform. The infrastructure being developed uses a common set of services and service registrations for the entire clinical trial on cancer community. We are currently focusing on the development of the core set of components up to a stage where they can effectively support *in silico* investigation. Initial prototypes

have been useful in crystallizing requirements for semantics.

The project has set up cross-disciplinary task forces to propose guidelines concerning issues related to data sharing, for example legal, regulatory, ethical and intellectual property, and is developing enhanced standards for data protection in a web (grid) services environment.

In addition the project is developing

- standards and models for exposing web services (semantics), scientific services, and the properties of data sources, datasets, scientific objects, and data elements;
- new, domain-specific ontologies, built on established theoretical foundations and taking into account current initiatives, existing standard data representation models, and reference ontologies;
- innovative and powerful data exploitation tools, for example multi-scale modelling and simulation;
- standards for exposing the properties of local sources in a federated environment;
- a biomedical GRID infrastructure offering seamless mediation services for sharing data and data-processing methods and tools;
- advanced security tools including anonymisation and pseudonymisation of personal data according to European legal and ethical regulations;
- a Master Ontology on Cancer and use standard clinical and genomic ontologies and metadata for the semantic integration of heterogeneous databases;
- an ontology based Trial Builder for helping to easily set up new clinico-genomic trials, to collect clinical, research and administrative data, and to put researchers in the position to perform cross trial analysis;
- data-mining services in order to support and improve complex knowledge discovery processes;
- an easy to use workflow environment, so that biomedical researchers can easily design their “discovery workflows” and execute them securely on the grid.

A range of demonstrators, stemming from the user defined scenarios, together with these core set of components are currently enabling us to both begin evaluating and gathering additional and more concrete requirements from our users. These will

allow us to improve and refine the facilities of the ACGT services.

ACKNOWLEDGEMENTS

The authors would like to thank all members of the ACGT consortium who are actively contributing to addressing the R&D challenges faced. The ACGT project (FP6-2005-IST-026996) is partly funded by the EC and the authors are grateful for this support.

REFERENCES

- Anguita, A., et al, 2007. Solving semantic heterogeneities and integration between clinical and image databases in post-genomic clinical trials, Proc. of Personalised Healthcare (phealth2007) Conference (sponsored by the IEEE Engineering in Medicine and Biology Society (EMBS)), Chalkidiki, Greece.
- Cannataro M. and Talia, D., 2003. KNOWLEDGE Grid - An Architecture for Distributed Knowledge Discovery. CACM, vol. 46, no 1, pages 89-93.
- d. Roure, D., Jennings N. R. and Shadbolt, N., 2003. The Semantic Grid: A future e-Science infrastructure, in: Grid Computing: Making The Global Infrastructure a Reality. (Eds.) Berman, F., Hey, A. J. G. e Fox, G., John Wiley & Sons, pages 437-470.
- Foster, I., Kesselman, C., Tuecke, S., 2001. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. International Journal of High Performance Computing Applications, vol. 15, no. 3, pages 200—222.
- Foster, I., 2005. Service oriented Science. Science, vol 308, no. 5723, pages 814-817.
- Goble, C. , Pettifer, S., Stevens, R., 2004. The Grid: Blueprint for a New Computing Infrastructure, Morgan Kaufmann, San Francisco, ed. 2, pages 121–134.
- Graf, N., 2007. The importance of an ontology based clinical data management system (OCDMS) for clinico-genomic trials in ACGT (Advancing Clinico-Genomic Trials on Cancer). In Proc. of the International Society of Paediatric Oncology Conference 2007, Mumbai, India (to appear).
- Gupta, A., Ludascher, B and Martone, M. E., 2000. Knowledge-based Integration of Neuroscience Data Sources. In proc. of the 12th Intl. Conference on Scientific and Statistical Database Management (SSDBM), IEEE Computer Society, Berlin.
- Katehakis, D.G., et al, 2007. Delivering a Lifelong Integrated Electronic Health Record based on a Service Oriented Architecture. IEEE Transactions on Information Technology in Biomedicine (to appear), available at: <http://ieeexplore.ieee.org/xpl/tocpreprint.jsp?isnumber=26793&punumber=4233>)
- Leisch E., et. al, 1997. A Framework for the Integration of Distributed Autonomous Healthcare Information Systems. Medical Informatics, Vol. 22, No. 4, pages. 325-335.
- Roses A.D., 2000. Pharmacogenomics and the practice of medicine. Nature, 405, pages 857-865.
- Rüping, S. et al, 2007. Extending workflow management for knowledge discovery in clinico-genomic data. In Nicolas Jacq et. al., editor, Proceedings of HealthGrid 2007, volume 126 of Studies in Health Technology and Informatics, pages 184–193. IOS Press.
- Schuppe-Koistinen, I., 2007. The application of metabolic profiling technologies in biomarker discovery during drug R&D. In Proc. Pharmaceutical Science World Congress 2007 (PSWC2007), Amsterdam, The Netherlands.
- Sotiriou, C., et al., 2003. Breast cancer classification and prognosis based on gene expression profiles from a population-based study. Proc Natl Acad Sci USA, vol. 100, no. 18, pages 10393-8.
- Sotiriou C., Piccart, M.J., 2007. Taking gene-expression profiling to the clinic: when will molecular signatures become relevant to patient care?, Nature Reviews, Vol. 7, July 2007, pages 545-553.
- Stamatatos, G.S. et al, 2007. The “Oncosimulator”: a multilevel, clinically oriented simulation system of tumor growth and organism response to therapeutic schemes: towards the clinical evaluation of in silico oncology. In Proceedings of the 29th Annual International Conference of the IEEE EMBS, Cité Internationale, Lyon, France, August 23-26, pages 6628-6631.
- Stevens R., et al., 2000. TAMBIS: transparent access to multiple bioinformatics information sources. Bioinformatics, vol 16, no. 2, pages. 184–185.
- Stein, L., 2002. Creating a Bioinformatics Nation. Nature, 317, pages. 119-120.
- Tsiknakis M., et al, 2006. Building a European Biomedical Grid on Cancer: The ACGT Integrated Project. In Proc. HealthGrid 2006 Conference, Stud. Health Technol. Inform., vol.120, pages 247-58.
- Tsiknakis M., et al, 2007a. A Semantic Grid Infrastructure Enabling Integrated Access and Analysis of Multilevel Biomedical Data in Support of Post-Genomic Clinical Trials on Cancer, IEEE Transactions on Information Technology in Biomedicine (to appear), DOI: 10.1109/TITB.2007.903519, available at: <http://ieeexplore.ieee.org/xpl/tocpreprint.jsp?isnumber=26793&punumber=4233>
- Tsiknakis M, et al, 2007b. Developing a European Grid infrastructure for cancer research: vision, architecture, and services. ecancermedicalscience Journal, DOI: 10.3332/eCMS.2007.56.
- Wegener, D., et. Al., 2007. GridR: An R-based grid-enabled tool for data analysis in ACGT clinico-genomics trials. In proc. of the eScience Conference 2007, Bangalore, India, (accepted, to appear).

DECISION SUPPORT AND IMAGE & SIGNAL ANALYSIS IN HEART FAILURE

A Comprehensive Use Case

Sara Colantonio, Massimo Martinelli, Davide Moroni, Ovidio Salvetti
Inst. of Information Science and Technologies, ISTI-CNR, Pisa, Italy
{name.surname}@isti.cnr.it

Francesco Perticone, Angela Sciacqua
Univerity Magna Graecia, Catanzaro, Italy
angela.sciacqua@inwind.it; perticone@unicz.it

Franco Chiarugi
Institute of Computer Science (ICS), Foundation for Research and Technology – Hellas (FORTH)
Heraklion, Crete, Greece
chiarugi@ics.forth.gr

Domenico Conforti, Antonio Gualtieri, Vincenzo Lagani
University of Calabria, Cosenza, Italy
mimmo.conforti@unical.it; gualtieri.antonio@gmail.com; vlagani@deis.unical.it

Keywords: Decision Support Theory, Echocardiographic Imaging Analysis, Signal Processing.

Abstract: The European STREP project HEARTFAID aims at defining an innovative platform of services able to intelligently support clinical operators in the daily management of heart failure patients. The core of the platform intelligence is a Clinical Decision Support System, developed by integrating innovative knowledge representation techniques and hybrid reasoning methods, and including advanced tools for the analysis of diagnostic data, i.e. signals and images. Aiming at showing how all these issues are combined in the HEARTFAID platform, we present a comprehensive use case, centred on echocardiography workflow and covering the clinical course leading from visit scheduling to therapeutic choices, highlighting the intervention and the value added by the Clinical Decision Support System.

1 INTRODUCTION

Heart Failure (HF) is a complex clinical syndrome resulting from any structural or functional cardiac disorder which impairs the ability of the ventricle to fill with or eject blood. In its chronic form, HF is one of the most remarkable health problems for prevalence and morbidity, especially in the developed western countries, with a strong impact in terms of social and economic effects. All these aspects are typically emphasized within the elderly population, with very frequent hospital admissions and a significant increase of medical costs. Recent studies and experiences have demonstrated that accurate heart failure management programs, based on a suitable integration of inpatient and outpatient

clinical procedures, might prevent and reduce hospital admissions, improving clinical status and reducing costs.

The European project HEARTFAID (“A knowledge based platform of services for supporting medical-clinical management of the heart failure within the elderly population” — IST-2005-027107) aims at defining efficient and effective health care delivery organization and management models for the “optimal” management of the care in the field of cardiovascular diseases.

The HEARTFAID platform (HFP) has been conceived as an integrated and interoperable system, able to guarantee an umbrella of services that range from the acquisition and management of raw data to the provision of effective decisional support to

clinicians. All the functionalities and services supplied by the entire HFP can be further grouped into *data*, *decision* and *end-users* macro “contexts”. The former is devoted to the collection and management of information, which consists of biomedical data, acquired from biomedical devices and structured/unstructured information such as clinical reports collected during patient hospitalisation and outpatient visits and within a homecare program by *telemonitoring* patients’ conditions. The decision context includes a *knowledge-based Clinical Decision Support System* (CDSS) whose main goal is supporting the HF health care operators, by making more effective and efficient all the processes related to diagnosis, prognosis, therapy and health care personalization of the HF patients. The latter context provides the doorway to a multitude of end-user utilities and services, such as accessing an electronic health record, querying the CDSS, applying advanced models and methods for the extraction of new knowledge, and so forth.

The CDSS represents the core of HFP and has been carefully designed by combining innovative knowledge representation formalisms, robust and reliable *reasoning* approaches, based on *Machine Learning* and inference methodologies, and innovative methods for diagnostic images analysis and biomedical signals processing (VV.AA., 2007).

This paper aims at showing how all these issues are combined within a comprehensive use case, centred on an echocardiography workflow and covering the clinical course leading from visit scheduling to therapeutic choices, highlighting the intervention and the value added by the CDSS.

In the following sections, we will briefly review the related literature, then introduce the HEARTFAID CDSS and its main functionalities, and, finally, describe the use case and CDSS interventions.

2 DECISION SUPPORT AND DATA PROCESSING IN HEARTFAID

The development of computerized applications for supporting health care givers is an old but still alive quest, started more than 45 years ago, in the early 1960s, and with ascending and descending periods of interest and growth (Greenes, 2007).

A plethora of CDSS has already evolved with different platforms and architectures, encompassing a variety of services, from information retrieval and reporting, scheduling and communications, to cost-

effectiveness, error prevention, safety, and improvement of health care quality. The most common realizations include electronic medical/patient records (Poissant et al., 2005), computerized alerts and reminders, clinical guidelines formalizations (GEM, 2003), provider order entries (Park et al., 2005), diagnostic support, clinical result interpretation, adverse event monitoring, shared patient-doctor decision-making (Wirtz et al., 2006).

The primary task for developing effective CDSS is to select the corpus of *pertinent* knowledge and/or collect and process data to create pertinent knowledge which is relevant for bringing the health care to effect. Knowledge representation just concerns understanding, designing, and implementing ways of formally coding the knowledge necessary for deriving other knowledge, planning future activities, solving problems that normally require human expertise. Representing knowledge requires the selection of a suitable *language* or *formalism* and the definition of a *Knowledge Base* (KB) built by formalizing clinical experts’ know-how and guidelines. Usually, the formalism is *symbolic* and the KB contains *statements* or *expressions* of one of the following formalisms: (i) rule based; (ii) frame based; (iii) network based; and (iv) logic based (Helbig, 2006). Workflow based representation is also becoming well-known, especially for modelling guidelines (Boxwala et al., 2004). Moreover, in recent years ontologies are emerging as a powerful knowledge representation formalism which is conceptually equivalent to the frame based and to first order logic approach (Bayegan et al., 2002).

The KB is exploited by a *reasoning engine* which processes available information for formulating new conclusions and answering questions. *Inferential* reasoning is employed for inferring new knowledge from the KB by deduction, induction or abduction.

In some cases, making a decision requires an investigation on the hidden, complex, often non-linear correlations among data, together with high-level analytical processing functions. In such cases, the knowledge needed for the solution should be acquired directly from data (*inductive knowledge*) and stored in a model (e.g. *Artificial Neural Networks*, *Support Vectors Machines*), able to induce *sub-symbolic* knowledge by data-driven processing. Computational models are also useful for representing uncertain knowledge, as *Bayesian Networks* and *Fuzzy theory*.

HF routine practice presents several aspects in which an automatic, computer-based support could have a favourable impact. Some attempts to support HF clinical operators have been presented, such as

an Electronic Patient Record (Bosmans et al., 2000) or computerized guidelines (Dassen et al., 2003). More complex decision support systems have been developed for suggesting the most appropriate therapy (e.g. Perlini et al., 1990).

Within HEARTFAID, a careful investigation about the needs of HF practitioners and the effective benefits assured by decision support was performed: four problems have been identified as highly beneficial of HEARTFAID CDSS point-of-care intervention. They can be referred as *macro domain problems* and listed up as: (i) HF diagnosis, (ii) prognosis, (iii) therapy planning, and (iv) follow-up. Further detailed decision problems were identified for specifying these macro domains, focusing as much as possible on the medical users' needs; explicative examples are:

- severity evaluation of heart failure;
- identification of suitable pathways
- planning of adequate, patient's specific therapy;
- analysis of diagnostic exams
- early detection of patient's decompensation.

An accurate analysis highlighted that the needed corpus of knowledge mainly consisted of domain know-how. Nevertheless, the solution of some of these problems seemed still debated in the medical community, due to the lack of validated and assessed evidences. In such cases, computational models appeared the best solution for modelling the decision making extracting knowledge directly from available data. Moreover, specific processing algorithms were designed for analyzing diagnostic examinations. In this perspective, HEARTFAID CDSS was designed for combining different models of reasoning, as will be described in the next sections.

2.1 Significance of Signal Acquisition and Analysis in HF

Electrocardiography (ECG) is one of the very basic examinations performed in the evaluation and assessment of HF. According to ESC (2005) guidelines, the negative predictive value of normal ECG to exclude left ventricular systolic dysfunction exceeds 90%.

The most common ECG examinations are the "Resting ECG" and the "Holter ECG". While the latter is more commonly used for the discovery of rhythm abnormalities and the computation of the Heart Rate Variability (HRV), the former is more commonly used for the evaluation of morphological abnormalities in the PQRST shape.

In both examinations, the first step to be performed is the QRS detection with the identification of the time occurrences of each heart

beat. This series of data allows for the evaluation of the heart rate and is preparatory to the beat classification for the discrimination between normal and abnormal beats. This task is usually performed in the "Holter ECG" reading stations as a starting point for the arrhythmias' classification and for the evaluation of the NN series (time intervals between consecutive normal beats) that is the input for the HRV evaluation. In case of "Resting ECG" examination (typical duration is 10 seconds), the evaluation of the normal beats allows the normal beat averaging with the construction of a more noise-free reference beat that can be used for a better evaluation of wave durations and amplitudes.

Wave durations and amplitudes are paramount in the evaluation of the "Resting ECG" (usually with 12 leads) parameters of high significance for the HF patients like ST depression, QRS and QT durations, Sokolow-Lyon index for left ventricular hypertrophy, presence of left or right branch bundle block and presence of pathological Q waves.

2.2 Significance of Imaging Techniques and Image Processing in HF

Imaging techniques offer invaluable aid in the objective documentation of cardiac function, allowing for the computation of parameters relative to chamber dimensions, wall thickness, systolic and diastolic function, regurgitations and pulmonary blood pressure.

According to ESC (2005), chest X-ray and echocardiography should be included in the HF initial diagnostic work-up. Further, echocardiography will be regularly repeated to monitor in an objective way the changes in the clinical course of a HF patient. Additional techniques, like nuclear imaging and cardiac magnetic resonance, may be also considered for particular patients, since they have not been shown to be superior to echocardiography in the management of most HF population.

Thus, echocardiography—and in particular 2-D TransThoracic Echocardiography (TTE) for its portability and versatility—is the key imaging technique for the practical management of HF.

The most important measurement performed by TTE is Left Ventricle Ejection Fraction (LVEF), which permits to distinguish patients with cardiac systolic dysfunction from patients with preserved systolic function. LVEF is given by the normalized (non-dimensional) difference between left ventricle End-Diastolic Volume (EDV) and the End-Systolic volume (ESV). Among different models for the computation of such volumes, the American Society of Echocardiography (Lang et al., 2005) suggests the

use of the so-called Simpson's rule, by which the left ventricle is approximated by a stack of circular (or elliptical) disks whose centres lie in the major axis. Simpson's method, therefore, relies on left ventricle border tracing. It is well-known that manual border tracing, besides being time-consuming, is prone to inter- and intra- observer variability, and thus is unable to provide a satisfactory and reproducible measurement of LVEF.

Image processing techniques may relieve this problem, by providing automated or, at least, semi-automated methods for tracing contours of relevant structures found in an image, an issue called *image segmentation* in the specific literature. However, the segmentation problem for ultrasound images is by no means trivial, due mainly to low signal to noise ratio, low contrast, and image anisotropy and speckle noise (Noble and Boukerroui, 2006).

3 A SIGNIFICANT SCENARIO

HEARTFAID CDSS was designed after a careful analysis of the problems to be faced and the expectations of the medical users.

A complete use case was defined for guiding the development activity of CDSS by considering many of the integrated services of the platform.

More in detail, we are considering a 65 years old patient, already enrolled in the HFP, former smoker, suffering from hypertension for several years. The patient was enrolled in the HFP six months ago and, in particular, the telemonitoring services offered by the platform were activated. At the baseline visit, the patient referred a slight limitation of physical activity, since he felt comfortable at rest but ordinary activity resulted in fatigue and dyspnoea. For these reasons, the patient was assigned to NYHA class II. Anamnesis data were also collected, from which it is known that the patient had an acute myocardial infarction five years before and he underwent to aorto-coronary bypass. The patient had a post ischaemic dilated cardiomyopathy with associated systolic dysfunction.

The TTE test (performed at baseline evaluation) showed an LVEF equal to 40%, ESV and EDV being respectively 114 ml and 190 ml. The left ventricle end-diastolic diameter was 6.0 cm. The pharmacological treatment consisted in ACE-inhibitor, beta-blockers, spironolactone, aspirin and statin. Neither pulmonary nor systemic congestion signs were present. Blood examinations of renal function and electrolytes were normal. During these six months, the patient has been telemonitored. In particular, the pharmacological therapy has been

followed with care and no relevant changes have been detected by the platform.

Suddenly, the patient observes a worsening of his symptoms, with a marked limitation of physical activity. After he fills in a periodic questionnaire suggested by the platform based on Minnesota questionnaire, the changes in the symptoms are automatically detected and considered relevant. A medical visit is suggested by the CDSS, accepted by the referent physician and immediately scheduled.

At the visit, the NYHA class changes from II to III. No variations in the signs are observed by the cardiologist, apart from a slight worsening of blood pressure (150/90 mmHg) and an increase of 10 beats/min in the heart rate. An ECG is performed also to confirm the heart rate increment.

The cardiologist, supported by the CDSS, decides however to evaluate other parameters by echocardiography. During the TTE examination, the sonographer acquires images and images sequences according to a protocol specified by the platform. Finally the images and the parameters manually evaluated by the sonographer are stored in the platform image archive. The reviewing cardiologist visualizes the echocardiographic images and the estimated parameters. Left ventricle volume and ejection fractions are computed again by automatic methods, exploiting the available image sequences. These values are compared with the historical data of the patient. EDV increases to 210 ml, ESV increases to 145 ml, EF decreased from 40% to 30%.

Mild tricuspidal insufficiency is Doppler-detected by its regurgitation. By tricuspidal regurgitation extent, the pressure gradient (mmHg) between right ventricle and right atrium is measured. Pulmonary pressure is then estimated. With this aim, the subcostal view is taken into account, so as to determine Inferior Vena Cava (IVC) diameter and its collapsibility index. The pulmonary pressure is estimated to be 40 mmHg, by using a lookup table with entries consisting in the tricuspidal gradient, IVC diameter and collapsibility index. Since this value indicates a slight pulmonary congestion, the CDSS suggests the physician to integrate the pharmacological therapy with diuretics, for example loop diuretics or thiazides. Further, since there are no up-to-date information about the renal function and electrolytes, the CDSS suggests to start with a safe diuretic dosage and to perform blood examinations, which are scheduled for few days later. The physician opts for a loop diuretics therapy, for quicker beneficial effects.

Back to his home, the patient is monitored in the subsequent days. In particular control of weight, urine output, blood pressure, symptoms are scheduled daily. Blood examinations are scheduled seven days after the beginning of the new treatment.

The results of such blood examinations are uploaded to the platform.

An up-titration table for the diuretics is compiled by the CDSS, considering in particular symptoms and electrolytes balance, creatinine clearance, blood pressure, weight slope and urine output. The CDSS also suggests to control weight and urine output daily and to schedule blood examinations weekly. A visit is also suggested in one month, to appreciate the response to the therapy. The physician reviews this program and decides to approve it. After approval, the up-titration table for diuretics is automatically sent to the patient.

One week after, telemonitoring evidences persistence of symptoms; the patient is thus required to continue the up-titration program for diuretics. During the subsequent weeks symptoms get better until the visit. At that visit, the patient refers that symptoms are relieved. NYHA class is moved back to II. However, the CDSS suggests the physician to explore the possible origins of the change in the symptoms reported in the previous visit (i.e. the probable cause of heart failure decompensation). In particular, with the aim of controlling the ischemic disease, a stress test is scheduled.

3.1 Methods

The CDSS was defined for an overall support of HF management, facing the main decisional problems of diagnosis, prognosis, therapy and follow-up, by using patients' heterogeneous information (e.g., actual status, anamnesis, clinical history, diagnostic parameters, and clinicians' evaluation).

Ontologies combined with rules were chosen as representation formalism, because of the more suitable and up-to-date methodology for formalizing the declarative and procedural knowledge derived from the guidelines and the experts' know-how. Actually, ontologies constitute a logic-based representation which also assures easy re-use and sharing of knowledge. Moreover, the rule based approach appeared the more appropriate both to complete possible representation lacks of ontological model and to involve more effectively the experts in the elicitation process. An inference engine was then devised for the corresponding inferential reasoning processes, by induction and deduction on the formalized knowledge for assessing patients' status, formulating diagnosis and prognosis, assisting therapy planning, and monitoring patients.

Computational reasoning models were included for those difficult HF decision problems, such as prognosis assessment and early detection of patient's decompensation.

The HEARTFAID CDSS architecture was designed according to a multilevel conceptualization scheme for distinguishing among

- the *knowledge level*, corresponding to all the information needed by the system for performing tasks, e.g. data, domain knowledge, computational decision models;
- the *processing level*, consisting of the system components that are responsible for tasks accomplishment by using the knowledge level;
- the *end-user application level*, including the system components whose functionalities are specifically defined for interacting with the user.

More in detail, the CDSS architecture consists of the following components (Figure 1):

- *Domain Knowledge Base*, consisting of the domain knowledge, formalized from the guidelines and of the clinicians' know-how;
- *Model Base*, containing the computational decision models, signals and images processing methods and pattern searching procedures;
- *Meta Knowledge Base*, composed by the strategy knowledge, i.e. about the organization of CDSS tasks.
- *Brain*, the system component endowed with the reasoning capability, which is divided into the meta and object level;
- *Explanation Facility*, providing a means to explain conclusions taken.

The Brain was modelled by functionally separating a *meta level*, devoted to task accomplishment and organisation, and an *object level*, responsible for actually performing tasks, by reasoning on the computational and domain knowledge. A *Strategy Controller* was inserted for performing the meta level functionalities, by orchestrating the two components of the object level, i.e. the *Inference Engine* and the *Model Manager*.

Moving from the design to the development activity, the use case is being used as a real scenario for implementing the above architecture.




Figure 1: The general view of the HEARTFAID CDSS architecture – dashed arrows correspond to reference to the ontologies, while the others denote a direct communication.

The required CDSS interventions consist in the following services, listed in order of occurrence in the workflow:

- interpretation of telemonitored data and assessment of patient's status;
- visit scheduling;
- suggestion of new diagnostic examinations;
- analysis of imaging examinations;
- interpretation of diagnostic findings;
- suggestion of therapy changes.

The necessary pieces of knowledge have been identified as mainly symbolic and an elicitation process has been performed for their formalization in the Domain KB. Specific algorithms for extracting a number of useful parameters from the echo images have been developed and inserted into the Model Base.

In particular, the use case highlights the interventions of other components of HFP which hold important roles in assuring the effectiveness of the support services, e.g. the agenda for scheduling new visits or examinations. Actually, HFP was conceived for consciously distributing the work load among the various components. A sketch of the platform with the components that interact with the CDSS is shown in Figure 2.

An EHR module was inserted for suitably organizing, visualizing and managing patients' data, stored into the platform *Repository*. In particular, a dedicated *Repository* for storing examination images was conceived in accordance to the DICOM standard. An *Agenda* module was included for managing patients' care planning, e.g., scheduling new visits, prescribing new examinations and so forth. The *User Interface* was designed as a fundamental component, responsible for all the interactions and communications with the users.

The different components of the platform were seen as *resources*, by *virtualizing* the operations required for their management. When involved, the different components are dynamically integrated for supplying sophisticated but much flexible applications. The responsible for guaranteeing integration and interoperability among all the HFP components was defined as the platform *Middleware*, which includes all the adapters necessary for the virtualization. For simplifying the provision of different services, a *Service Controller* (SC) was comprised for managing platform events and invoking the other components.




Figure 2: A sketch of HFP with the components that interact with the CDSS.

In this perspective, the CDSS component was designed as a resource able to offer a number of functionalities and to interact with the other resources for performing its tasks. Each decision-making problem has to be translated into a *request* or a *class of requests* committed to the CDSS, which is then activated *on-demand*. The system handles every request according to a specific policy encoded in the Meta KB, interacting, when necessary, with the other platform components.

A brief (and partial), discursive description about how the scenario has been mapped onto a workflow of HFP services is useful for understanding its implementation.

Description of the Mapping into the HFP.

The patient answers a questionnaire through his web-based user interface and sends the information to the HFP that checks eventual missing values. Then the Service Controller stores this information into the repository, gets historical data and opportunely invokes the specific CDSS service responsible for handling the request.

The CDSS analyzes data and answers supplying the current patient's status, i.e. worsening of symptoms, and a set of suggested actions the clinician should undertake, i.e. schedule a new visit, change the NYHA class or change the therapy and so forth. Then the SC stores CDSS results into the repository.

When the doctor on duty logs in the HFP, the list of patients is displayed ordered by their severity status and the timestamp of the last related event. Once the patient is selected, the change in his status is shown along with the list of suggested actions, for instance as a list of operations that can be selected. He then approves the schedule of the visit and the SC forwards the request to the Agenda component that opportunely records it and informs the patient. During the visit, the physician inserts his observations into the patient's record and decides to approve the change of the NYHA class: he selects the corresponding action within the list and the SC

takes care of registering the change in the patient's record. An ECG is then performed acquired by the platform and processed with QRS detection and classification algorithms in order to produce noise-free reference beats for all the 12 leads. The availability of graphical tools for the signal display and measurements (ruler) allows the cardiologist for an accurate and reliable evaluation of the significant morphological parameters. Then the information obtained by the ECG is inputted and a request is sent by the SC to the CDSS, which suggests performing an echocardiography as displayed in the recommended actions list.

Inference on Patients' Data. In order to include information derived from the use case, starting from a preliminary ontology mainly corresponding to a structured terminology of the domain, we began to develop a new ontology by inserting relevant properties, classes and relations for a coherent and comprehensive formalization, also in accordance to standard medical ontologies, such as the Unified Medical Language System (UMLS, 2007). An excerpt of this new ontology is shown in Figure 3.

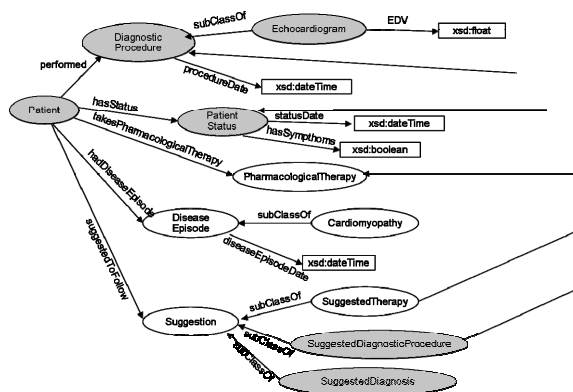


Figure 3: Some relevant classes and properties of the ontology.

A careful elicitation activity was performed for formalizing a set of rules, founded on the developed ontology and to be encoded in the KB. An elicited rule which is used for therapy suggestion is: *“If a patient has Left Ventricle Ejection Fraction <= 40% and he is asymptomatic and is assuming ACE Inhibitors or ARB) and he had a myocardial infarction then a suggestion for the doctor is to give the patient Betablockers”*.

Echocardiography Images Analysis. A prototype module for the computation of LVEF has been developed (Barcaro et al., 2007). The module is able to process an apical-view sequence of the heart (the

so-called two- and four-chamber views) in order to identify the left ventricle cavity in every frame of the sequence. This segmentation stage is accomplished augmenting a variational formulation of level set methods with *mimetic criteria* for contour initialization (see Figure 4).




Figure 4: Segmented left ventricle cavity at end-systole.

After segmentation, the left ventricle volume is computed as a function of time by applying monoplane Simpson's method. Then ESV and EDV are regarded as the minimum and maximum respectively of the volume time-course. Finally LVEF is obtained by the simple formula:

$$LVEF = \frac{EDV - ESV}{EDV} .$$

4 RESULTS AND CONCLUSIONS

A number of tools and instruments are available for developing the CDSS according to the design specifications. The key factors that were taken into account for defining an up-to-date system were accordance to standards, efficiency and robustness.

We investigated several technologies, with particular attention to the Semantic Web field, since it offers various tools for building ontological models, knowledge bases and reasoning on them. Moreover, the platform was conceived for web applications developed in Java. For selecting the appropriate tools, we carefully analyzed the W3C recommendations along with the performance, compatibility and maintenance of the same tools.

As to the knowledge representation formalism, we selected the Web Ontology Language (OWL, 2007) for defining the ontologies, since it is the actual de-facto standard semantic markup language for this task.

The ontology has been built using the two editors Protégé and Swoop. For defining the rules of the KB, we chose SWRL (2007), the Semantic Web Rule Language combining OWL and Rule Mark-up Language, which is a submission to W3C that extends the set of OWL axioms to include Horn-like

rules. For realizing the reasoning component, Jena (McBride, 2001) was selected as a Java programmatic environment that includes OWL, a language for querying ontologies, SPARQL (2007), and a rule-based inference engines. In particular, for improving the reasoning capability of the latter, we also used Bossam (2007) and Pellet (2007). An example of the rules we are developing in SWRL is shown in Figure 5 as it has been defined in Protégé.




Figure 5: A rule developed in SWRL.

Prototypical methods for processing the echo images were realized implementing the various procedures in Matlab.

Future activities consist in finalizing the platform implementation by concluding the realization of the Domain KB, the algorithms contained in the Model Base and the Brain, in particular of its meta level for integrating all the object models and the interface.

ACKNOWLEDGEMENTS

This research work is supported by the European Community, under the Sixth Framework Programme, Information Society Technology – ICT for Health, within the STREP project HEARTFAID (IST-2005-027107), 2006-2009.

REFERENCES

- Barcaro, U., Moroni, D., Salvetti, O. (2007) Left ventricle segmentation in ultrasound sequences for the automatic computation of ejection fraction. In *OGRW07 Open Gern Russ Work on Pat Rec & Im An*.
- Bayegan, E., Nytro, O., Grimsmo, A. (2002). Ontologies for knowledge representation in a computer-based patient record. In *ICTAI 2002, 14th IEEE Int. Conf. on Tools with Artificial*. IEEE.
- Bosmans, J.E., Dassen, W.R.M., van der Velde, E.T., Dijk, W.A., Spruijt, H.J., Baljon M.H. (2000). Towards ICT Support for Treatment of Congestive Heart Failure. *Computers in Cardiology*, 27, 641-644.
- Bossam (2007) <http://bossam.wordpress.com/>
- Boxwala, A.A., Peleg, M., Tu, S. et al. (2004) GLIF3: a representation format for sharable computer-interpretable clinical practice guidelines. *J. Biomed Inform.*, 37(3), 147-161.
- Dassen, W.R.M., Gorgels, A.P.M., Berendsen, A., Dijk, W.A., de Clercq, P.A., Hasman, A., Baljon M.H. (2003). Guideline Assessment and Implementation in Congestive Heart Failure. *Computers in Cardiology*, 30, 331-334.
- ESC (2005) Swedberg, K. et al., “Guidelines for the diagnosis and treatment of Chronic Heart Failure: full text (update 2005)”, *European Heart J*, 45 pages.
- GEM (2003) *E2210-02 Standard Specification for Guideline Elements Model (GEM)* Book of Standards Volume: 14.01. ASTM International.
- Greenes, D. (2007). *Clinical Decision Support: The Road Ahead*. Academic Press.
- Helbig, H. (2006) *Knowledge Representation and the Semantics of Natural Language*, Springer, Berlin, Heidelberg, New York.
- Lang, R.M. et al., (2005). Recommendation for Chamber Quantification.. *J Am Soc Echocardiogr*, 18, 1440-1463.
- McBride, B. (2001) Jena: Implementing the RDF Model and Syntax Specification. In *WWW2001, Semantic Web Workshop*.
- Noble, J.A., & Boukerroui, D. (2006). Ultrasound image segmentation: a survey. In *IEEE Trans. Med. Imag.*, 25, 987-1010.
- OWL (2007) <http://www.w3.org/TR/owl-features/>
- Park, R.W., Shin, S.S., Choi, Y.I., Ahn, J.O., Hwang, S.C. (2005). Computerized physician order entry and electronic medical record systems in Korean teaching and general hospitals: results of a 2004 survey. *J Am Med Inform Assoc*. 12(6):642-7.
- Pellet (2007) <http://pellet.owldl.com/>
- Perlini, S. et al. (1990) Treatment of chronic heart failure: an expert system advisor for general practitioners. *Acta Cardiol.*, 45(5), 365-78.
- Poissant, L., Pereira, J., Tamblin, R. Kawasumi, Y. (2005). The impact of electronic health records on time efficiency of physicians and nurses: a systematic review. *J Am Med Inform Assoc*. 12 (5), 505-16.
- SPARQL (2007) <http://www.w3.org/TR/rdf-sparql-query/>
- UMLS (2007). <http://www.nlm.nih.gov/research/umls/>
- SWRL (2007) <http://www.w3.org/Submission/SWRL/>
- VV.AA. (2007). Functional Specifications of Data Processing and Decision Support Services. Deliverable D15, HEARTFAID IST-2005-027107.
- Wirtz, V., Cribb, A., Barber, N. (2006). Patient–doctor decision-making about treatment within the consultation—A critical analysis of models. *Social Science & Medicine*. 62, 1, 116-124.

MYHEART

Fighting Cardio-vascular Diseases by Prevention and Early Diagnosis

Ralf Schmidt, Jörg Habetha and Matthew Harris

Philips Research Europe, Weisshausstrasse 2, 52062 Aachen, Germany

ralf.schmidt@philips.com, joerg.habetha@philips.com, matthew.harris@philips.com

Keywords: Personal Healthcare, cardio-vascular, prevention, vital body signs, wearables.

Abstract: MyHeart is an Integrated Project of the European Union aimed at developing intelligent systems for the prevention and monitoring of cardiovascular diseases. The approach of the MyHeart project is to monitor Vital Body Signs (VBS), to process the measured data and to give the user (therapy) recommendations. Using its broad base of technical and business expertise, four concepts addressing cardiac health have been developed and tested on a technical, business, realisability and usability level.

1 INTRODUCTION

Cardiovascular disease (CVD) is the leading cause of death in developed countries. Roughly 45% of all deaths in the EU, and 37% in the U.S. are due to CVD (Thom, 2006). Hundreds of billions of Euros are spent worldwide each year on the treatment of CVD. In order to maintain and improve the quality of health care without exploding costs, health care systems are undergoing a paradigm shift from patient care in the hospital to care at home.

A healthy and preventive lifestyle as well as early diagnosis of heart disease could save millions of life years annually, simultaneously reducing the morbidity and improving patient quality of life. Prevention offers the opportunity to systematically fight the origin of cardio-vascular diseases as well as to improve the medical outcome after an event. To enable a preventative health care system, a move is required from the current, event driven treatment to continuous and ubiquitous access to medical excellence. Innovative methods are needed that provide access to medical excellence in a cost-effective way.

The MyHeart consortium (MyHeart, 2004) involves 33 partners from 10 different countries. It is a balanced multidisciplinary consortium of industry (including Small and Medium Enterprises (SMEs)), research institutes, academia and medical hospitals. Prominent industrial partners are Philips, with its medical and technological expertise, Vodafone

(Foundation) as a leading service provider, and Medtronic, a world-leader in cardiac technology.

The project started in January 2004 and has a total duration of 45 months (until September 2007). It is one of the largest biomedical and healthcare research projects in the European Union with a budget of about 35 million Euros.

The project brings technical capabilities in functional clothing, on-body electronics, user interaction, professional interaction, and algorithmic development together with the business assessment and development capabilities necessary to bring new health technologies to the health care system.

The technological needs for MyHeart applications span a wide range covering: monitoring of vital signs (ECG, respiration, activity, etc.); body-worn, low-power, mixed-signal hardware which runs algorithms for detection of health status and prediction of acute cardiac events; user interfaces for citizens and medical professionals; low-power wireless links and server architectures for data handling at professional sites.

2 THE MYHEART APPROACH

It is the aim of the MyHeart project to fight CVD by prevention and early diagnosis. This is done by monitoring Vital Body Signs (VBS) with wearable technology, processing the measured data and giving (therapy) recommendations to the user of the system. Using the measured data to give user

feedback ‘closes the loop’ of measurement and therapy. As illustrated in figure 1, this closed loop can either consist of direct local feedback to the user or of professional help by a physician or nurse. The latter will typically be provided remotely, which implies that the MyHeart system also comprises a telemedical element. Data are transmitted to a remote server, where a professional can access the data and contact the patient subsequently.

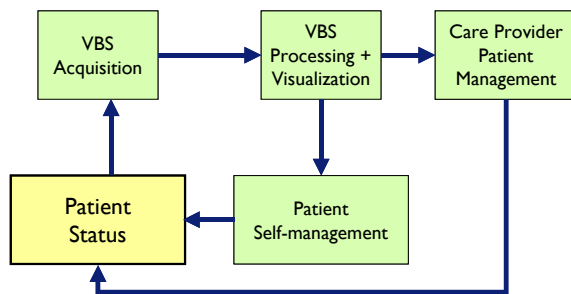


Figure 1: MyHeart disease management and prevention approach.

The system can be used for helping people to lead a healthier life as well as for the improved management of chronic diseases.

3 THE MYHEART CONCEPTS

MyHeart has taken a very innovative approach in ensuring the applicability of the project results in the real world. The consortium has started with a set of application ideas and only afterwards investigated the necessary technologies in order to serve these applications.

A concept is a concrete CVD application tailored to a specific user group or customer segment. The MyHeart project began with 16 concepts.

In the first one and a half years of the project, the 16 application concepts worked on answering the following questions in detail:

- **What** is the application/value proposition?
- **Who** are the customers and how to address them?
- **How** to do it technically?
- **Why** to believe in the concept (How to prove that it delivers what is claimed)?
- **Where** is the business?

In mid 2005, 4 of the 16 concepts were selected for further development in the remaining two years of the project. The criteria for selecting or combining concepts were:

- Medical credibility and feasibility
- Technical credibility and feasibility

- Business credibility and feasibility
- Critical project success factors (like size and excellence of the consortium)

The selected product concepts cover four major user segments: the healthy (Activity Coach), those at risk for developing CVD (Take Care), sufferers from a cardiac event (Neurological Rehabilitation), and chronically ill people (Heart Failure Management). In the following sections, the four product concepts are presented.

3.1 Activity Coach

The value proposition for the Activity Coach is to empower and allow the end user get maximum benefit from regular exercise sessions, both in terms of pleasure and health impact, anywhere, anytime through giving professional, easy to understand coaching which is tailored to the user’s profile, goals (Dunn, 1999) and personal performance.

The target group is people exercising for fitness and fun. The activity coach will help guide and motivate this group, both in the fitness studio and outdoors to give them optimal exercise result for the effort given.

As shown in figure 2, the system consists of four main components:

The **Body Signal Sensor (BSS)**, integrated into a textile garment, is responsible for monitoring the required vital signals. A one lead ECG is used to derive the heart rate. A stretch sensor is used to measure respiration rate. Furthermore, an accelerometer is used to measure the step rate while running.

The **Fitness Coach Bike (FCB)** is an indoor bike with integrated sensors measuring the pedaling rate, a processing and communication unit, and a user interaction device.

The **Personal Mobile Coach (PMC)** is a device for the outdoor scenario. It receives data from BSS via Bluetooth and generates appropriate feedback and interacts with the user and the service centre.

The **Fitness Coach Service Centre (FCSC)** is the professional platform that provides online services to the user. It receives all the data from the session, processes it using algorithms for fitness status assessment and performance analysis, and stores all the results. It also provides a web-based interface through which professional users are able to access different functionalities such as session results visualisation, messaging services, or the training program schedule.

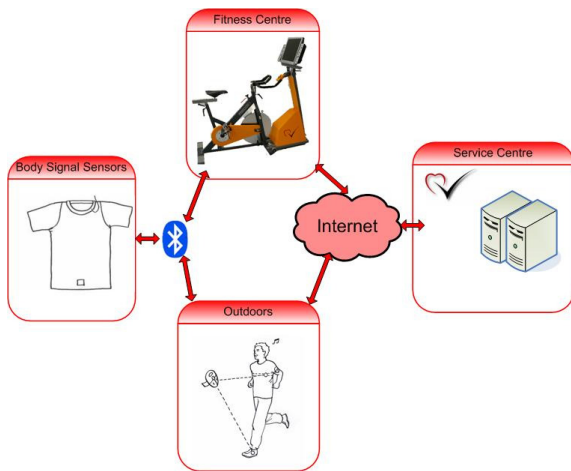


Figure 2: Architecture of the Activity Coach System.

During exercise, body and exercise equipment sensors measure heart rate, respiration rate, temperature and step or pedal rate. The data is processed by personalised algorithms, and user feedback is given on the FCB and the PMC. When using the FCB, the level of exercise can also be adapted automatically, guiding the user through the exercise. The system coaches and motivates the user to continue the trainings plan, and creates an immersive environment.

3.2 Take Care

The value proposition of Take Care is to empower the user to change her lifestyle by assessing CVD risk factors and providing appropriate improvement plans and personalised recommendations (Duchna, 2003; Euroaspire, 2001; Gordon, 2004).

By providing reliable and trustworthy education, monitoring and coaching, the Take Care system aims at supporting the user to learn and listen to his/her own body, reducing the risk factors for CVD. The Take Care system is aimed at healthy users that have risk factors for CVD that are willing to spend money out of their own pocket for help in adopting a healthier lifestyle.

As shown in Figure 3, the Take Care user interaction (UI) device is at the centre of the Take Care system. It is the platform for giving feedback and receiving input from the user, receiving input from sensors, and running personalised algorithms. To initialise the system, initial user data from a weight scale, a blood pressure meter and a cholesterol meter are inputted into the UI device and used to automatically generate a risk profile and lifestyle plan. The UI device then controls and communicates with the measurement devices, following a daily routine. It receives vital body signs

(heart rate, respiration rate and activity level) from the on-body electronics connected to textile sensors, and sleep quality data from piezo and textile electrodes integrated into the bed. The data is processed on the UI device, which then gives feedback and coaching to the user. The UI device can also forward the data to a professional centre for further examination.

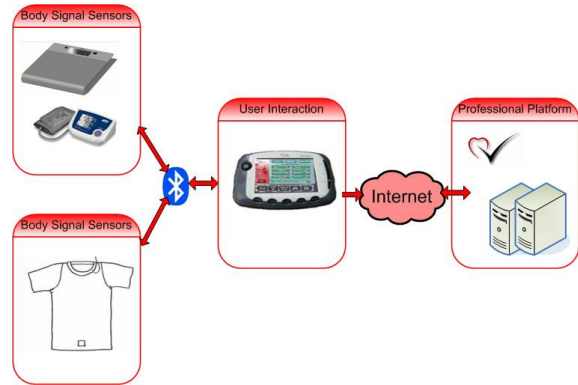


Figure 3: Architecture of the Take Care System.

3.3 Neurological Rehabilitation

The value proposition of the Neural Rehabilitation concept is to enable early intensive rehabilitation for patients following a cerebrovascular event (Sulch, 2001; Micieli, 2002) by using a telemonitoring system, using wearable technology, speech therapy tools, learning tools and communication tools.

The main users are patients with stroke symptoms, physicians, physiotherapists, and occupational therapists.

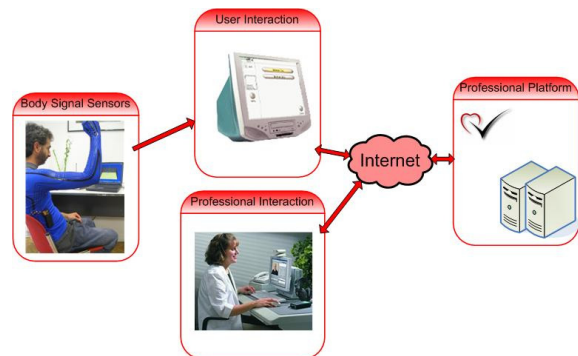


Figure 4: Architecture of the Neurological Rehabilitation System.

The **patient station**, connected to the therapist station, the server site and the user sensors, is the user interface with the patient, giving feedback on the exercises carried out. Wearable electronics integrated into an upper torso garment is used to

monitor the patient movement during rehabilitation exercises. A speech therapy unit is used to carry out and evaluate speech exercises.

The **therapist site** is used by the physician to monitor the patient's exercises and progress. At present, the therapist can only monitor the patient exercises offline. In the future, online monitoring could be possible.

The **server site** is a central server that hosts a database of configurations, exercises, session recordings, demographic data, and the rehabilitation protocol. The physician may access the server through the therapist site to configure the treatment for a particular patient, or to view the recorded data.

The system consists of three main stations (the patient station, the therapist station, and the server site), and the communication structure between them.

3.4 Heart Failure Management

The main objective of the Heart Failure Management concept is to improve the outcome of heart failure patients with respect to mortality, morbidity and quality of life (Swedberg, 2005; Steward, 2001). This objective is achieved by monitoring vital body signs that are relevant for heart failure on a daily basis (currently these parameters are only measured at infrequent visits to the physician) using easy to use equipment in the patient's home. The data is automatically analysed in order to detect changes in the patient's health status early enough to allow early therapy intervention, thus avoiding severe deterioration and hospitalisation.

The end users of the system are patients with heart failure (NYHA classes II-IV), and the physicians and nurses caring for the patient. Typically, the system would be funded by disease management organisations (DMO) and health care insurances.

The user interface for the heart failure concept is a PDA. Like in the Take Care concept, it is the platform for giving feedback and receiving input from the user, receiving input from sensors, and running personalised algorithms. The PDA controls and communicates with the measurement devices. A textile vest with integrated textile sensors and wearable electronics is used to measure vital body signs relevant for heart failure management. ECG sensors incorporated into the bed sheet and pillow, and a piezo sensor under the sheet capture ECG, breathing and movement data during the night. A weight scale and blood pressure cuff send their measured values to the PDA using Bluetooth. The PDA uses personalised algorithms to process the measured data, and to detect a possible deterioration

in the patient's health status, triggering action by the patient or medical professional.

The PDA also communicates with a professional platform which receives preprocessed patient data and gives health care professionals access to the application configuration, and the patient's data.

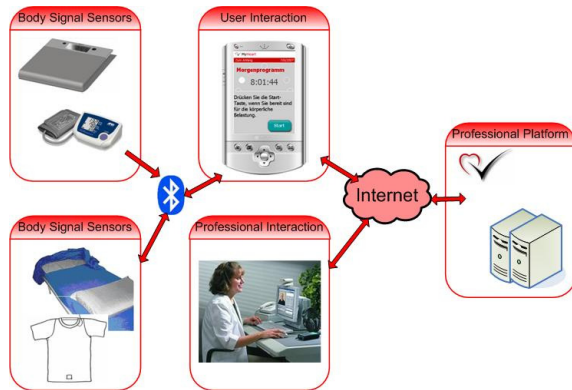


Figure 5: Architecture of the Heart Failure Management System.

4 VALIDATION

Medical and technical validation and business assessment are important aspects to be addressed by each MyHeart concept. In each concept studies are being carried out with prototype systems with end users to assess usability and medical effectiveness. For the Heart Failure concept an observational study with 200 heart failure patients will be carried out in Germany and Spain. In this one year study, the heart failure system will be used to make daily measurements of vital body parameters. At the end of the study, medical incidents will be correlated with the measured data to deduce which (combination of) parameters can be used to give warnings of a forthcoming decompensation.

ACKNOWLEDGEMENTS

This work was supported in part by the IST-2002-507816 MyHeart Project.

REFERENCES

- Thom, T. et al., 2006. *AHA Statistical Update. Heart Disease and Stroke Statistics - 2006 Update*. Circulation. 113:e85-e151.
- MyHeart IST-2002-507816 project, 2004. Information available on <http://www.cordis.lu/ist>.

- Dunn A et al., 1999. *Comparison of Lifestyle and Structured Interventions to Increase Physical Activity and Cardiorespiratory Fitness*. JAMA 28:327.
- Duchna HW et al., 2003. *Sleep Disordered Breathing and Cardio and Cerebrovascular Diseases: Update of Clinical Significance and Future Prospectives*. 2003 Somnologie 7: 101-121.
- Lifestyle and Risk Factor Management and Use of Drug Therapies in Coronary Patients from 15 Countries – Principal Results from EUROASPIRE II*, 2001. European Heart Journal 22: 554-572.
- Gordon NF et al., 2004. *Getting Risk Factors to Goal: Lifestyle Intervention is Worth the Effort in Patients With Hypertension, Hyperlipidemia and/or Hyperglycemia*. 53rd Annual Scientific Sessions of American College of Cardiology. New Orleans.
- Sulch D et al., 2001. *Randomized Controlled Trial of Integrated (Managed) Care Pathway for Stroke Rehabilitation*. Stroke 31:1929-34.
- Micieli G et al., 2002. *Guideline Application for Decision Making in Ischemic Stroke (GLADIS) Study Group. Guideline Compliance Improves Stroke Outcome: A Preliminary Study in 4 Districts in the Italian Region of Lombardia*. Stroke: 33(5) 1341-7.
- Swedberg K et al., 2005. *Guidelines for the Diagnosis and Treatment of Chronic Heart Failure: Executive Summary (Update 2005)*: The Task Force for the Diagnosis and Treatment of Chronic Heart Failure of the European Society of Cardiology. Eur Heart J:26 1115-1140.
- Stewart S and Blue L (Ed.), 2001. *Improving Outcomes in Chronic Heart Failure*, BMJ Books.

ONTOLOGY BASED INTEGRATION OF DISTRIBUTED AND HETEROGENEOUS DATA SOURCES IN ACGT

Luis Martín, Alberto Anguita, Víctor Maojo

*Biomedical Informatics Group, Artificial Intelligence Laboratory
School of Computer Science, Universidad Politécnica de Madrid, Campus de Montegancedo S/N
28660 Boadilla del Monte, Madrid, Spain
{lmartin, aanguita, vmaoj}@infomed.dia.fi.upm.es*

Erwin Bonsma, Anca Bucur, Jeroen Vrijnsen

*Phillips Research, Healthcare System Architecture, High Tech Campus 37, 5656 AE Eindhoven, The Netherlands
{erwin.bonsma, anca.bucur, jeroen.vrijnsen}@philips.com*

Mathias Brochhausen, Christian Cocos, Holger Stenzhorn

*IFOMIS, Universität des Saarlandes, Postfach 151150, 66041 Saarbrücken, Germany
{mathias.brochhausen, holger.stenzhorn, cristian.cocos}@ifomis.uni-saarland.de*

Manolis Tsiknakis, Martin Doerr, Haridimos Kondylakis

*Institute of Computer Science
Foundation for Research and Technology - Hellas, GR-71110 Heraklion, Crete, Greece
{martin, kondylak}@ics.forth.gr*

Keywords: Ontology-based Biomedical Database Integration, Semantic Mediation, Ontologies, Post-genomic Clinical Trials, Service Oriented Architectures.

Abstract: In this work, we describe the set of tools comprising the Data Access Infrastructure within Advancing Clinico-genomic Trials on Cancer (ACGT), a R&D Project funded in part by the European. This infrastructure aims at improving Post-genomic clinical trials by providing seamless access to integrated clinical, genetic, and image databases. A data access layer, based on OGSA-DAI, has been developed in order to cope with syntactic heterogeneities in databases. The semantic problems present in data sources with different nature are tackled by two core tools, namely the Semantic Mediator and the Master Ontology on Cancer. The ontology is used as a common framework for semantics, modelling the domain and acting as giving support to homogenization. SPARQL has been selected as query language for the Data Access Services and the Mediator. Two experiments have been carried out in order to test the suitability of the selected approach, integrating clinical and DICOM image databases.

1 INTRODUCTION

Data integration across heterogeneous data sources and data aggregation across different aspects of the biomedical spectrum is at the centre of current biopharmaceutical R&D. A technological infrastructure supporting such a knowledge discovery process should, ideally, allow for:

- Data to be searched, queried, extracted, integrated and shared in a scientifically and semantically consistent manner across

heterogeneous sources, both public and proprietary, ranging from chemical structures and omics to clinical trials data;

- Discovery and invocation of scientific tools that are shared by the community, rather than repeatedly developed by each and every organisation that needs to analyse their data and
- Both the sharing of tools, and their integration as modules in a generic framework, applied to relevant dynamic

datasets. We refer to this process as “discovery driven scientific workflows” which ideally would also execute fast and in an unsupervised manner.

Needless to say that our current inability to efficiently share data and tools, in a secure and efficient way, is severely hampering the research process. The objective of the Advancing Clinico-Genomic Trials on Cancer (ACGT) project is to contribute to the resolution of these problems through the development of a unified technological infrastructure which will facilitate the seamless and secure access and analysis, of multi-level clinical and genomic data enriched with high-performing knowledge discovery operations and services in support of multi-centric, postgenomic clinical trials.

Integrated access to heterogeneous biomedical data is at the core of the problems that need to be resolved. This paper presents the main methodological and technological challenges addressed in the implementation of an ontology-based data integration architecture within the context of the ACGT project. Emphasis is given to the description of the ACGT Data Access Architecture which is comprised by a set of key services, namely the ACGT-Data Access Services, the ACGT-Semantic Mediator, and the ACGT-Master Ontology, as well as additional dedicated tools. While the first two services provide the means to resolve syntactic and semantic heterogeneities when accessing integrated databases, the latter acts as a core resource supporting the data integration process.

2 BACKGROUND

Database Integration aims at facilitating users in querying sets of heterogeneous sources of information in an intuitive and transparent way. The research community has been dealing with different kinds of methods during the last decade, namely *Data Warehousing* (Kimball, 1996), *Federated Database Systems* (Sheth, 1990), *Mediator-based approaches* (Wiederhold, 1992), and other hybrid approaches. From the technical point of view, three categories can be differentiated, namely data translation, query translation and information linkage.

In Data Translation, data from the different databases are integrated in a centralized repository. Before the integration, these data must be modified in order to fit the requirements of the unified

schema—the central repository has its own schema different from the ones belonging to the underlying databases. The most popular example of a DT-based technology is Data Warehousing, which is now in its industrial exploitation phase.

By contrast, Query Translation does not perform actual integration of data, but transformation of a query when it is launched. A mediation software offers a representation of a virtually integrated set of databases to the users. The user is able to build and launch a query based on this representation. The mediator receives the query and transforms it into a set of dedicated sub-queries for the underlying databases. After their actual execution in the corresponding databases, the results are integrated by the mediator software to be presented to the user. On the other hand, Information Linkage just defines cross-reference links between databases to perform database integration. Some examples of usage of IL are MEDLINE, GENBANK, OMIM, and the World Wide Web itself.

There exist two main ways to deal with Query Translation: *Global as View* and *Local as View*. In both approaches the system has a description of the domain. In *Local as View*, views representing the databases are described using the knowledge contained in the global schema. In *Local as View* no additional work apart from defining a single view is necessary when a new database needs to be integrated in the system. However, translation of queries becomes leads to performance problems (Abiteboul, 1998) (Ullman, 1997). Conversely, in *Global as View* (Cali, 2001) a global model is built using information from the underlying databases and from the domain model. Query translation in *Global as View* is straightforward, since the links are actually stored in the schema, but it needs of a global revision when new sources are added.

During the last years, ontologies have been used as global domain models in database integration, obtaining promising results, mainly in the fields of biomedicine and bioinformatics. Biomedical ontologies have adopted the role of domain homogenizing tools in the last decade. We distinguish three major classes of biomedical ontologies: Generic Medical Ontologies—dealing with the entire domain of Medicine—, Specific Medical Ontologies—describing a single domain within Medicine—, and Specific Biomedical Ontologies—supporting a specific biomedical domain. Some examples of these three categories are:

- Generic Medical Ontologies: SNOMED CT (SNOMED, 2007), UMLS (Lindberg, 1990).

and GALEN (GALEN, 2007), HL7 RIM (HL7, 2007). Both SNOMED CT and UMLS have been proved to be theoretically unsound (Ceusters, 2003). HL7 RIM, even though widely used, has been subjected to a number of criticisms that also question its theoretical soundness (Smith, 2006).

- **Specific Medical Ontologies:** The Foundational Model of Anatomy (FMA, 2007) is a highly stable and rigorously developed ontology. One system frequently mentioned when talking about the state of the art in ontology-based cancer research and management is caCORE (caCORE, 2007), a highly developed environment making use of UMLS and NCI Thesaurus and Metathesaurus representations (NCI, 2007).
- **Specific biomedical Ontologies:** Gene Ontology (GO, 2007) is one example. Other examples can be found at the OBO Foundry (OBO, 2007). There is a high number of Specific Biomedical Ontologies, and they follow a variety of different standards. This shows the importance of quality assessment in ontology development within this domain.

The following section describes in detail the database access architecture adopted to integrate clinical trials databases including image information.

3 THE ACGT DATA ACCESS INFRASTRUCTURE

The ACGT platform is comprised by a set of services and resources supporting the different needs of clinicians and researchers involved in a post-genomic clinical trial. The ACGT platform architecture follows a layer based design, as can be seen in Figure 1.

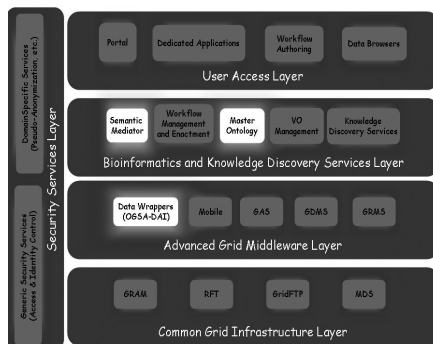


Figure 1: The ACGT platform architecture.

The ACGT data access infrastructure forms part of this architecture. This infrastructure is comprised by three core resources, together with other satellite tools that give support to the complete data access task. These core resources are, namely: the ACGT Master Ontology on Cancer (ACGT-MO), the ACGT Data Access Services (ACGT-DAS) and the ACGT Semantic Mediator (ACTG-SM). Figure 2 shows the architecture of the ACGT Data Access Infrastructure.

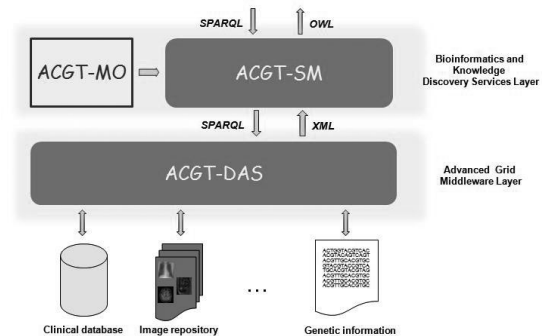


Figure 2: The ACGT data access infrastructure.

The next sections give a detailed description of these three components.

3.1 ACGT-MO

ACGT deals with the integration of data from a variety of heterogeneous sources. There exists a lack of standardization among data from different clinical trials, which leads to a loss in the possible knowledge exchanging power. Ontology based data management becomes then a major advantage in the way to achieve consistency in data collection and processing policies.

The ACGT-MO employs the resources of a Top Level Ontology, called Basic Formal Ontology (BFO, 2007). This choice is based on its proven high applicability to the biomedical field (Grenon, 2004). The ACGT Master Ontology inherits BFO's foundational principles:

- realism
- perspectivalism
- fallibilism
- adequatism

Figure 3 shows the BFO structure.

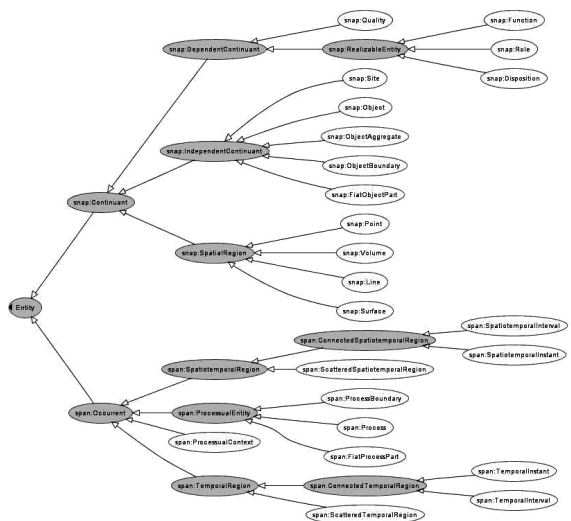


Figure 3: The Basic Formal Ontology.

The ACGT-MO has been developed using the OWL-DL language, achieving the maximum level of expressivity to describe the domain of post-genomic clinical trials on cancer. For its development and maintenance, the Protégé editor (Protégé, 2007) has been used.

The ACGT-MO basically contains two sets of elements, namely i) Classes and ii) Properties. The former group contains the concepts of the ontology (the so-called universals) structured in a taxonomy using *is_a* type relations to establish links between classes—e.g., CanonicalBodySubstance is_a BodySubstance. The latter represents the set of relations connecting the classes of the taxonomy. In order to fit the requirements of data integration in biomedical reality, and to express the truths of Medicine and Biology, a wide variety of relations (besides from mere *is_a*) has been included. A few examples of structure in the tree of relations are hasBloodPressure is a child of hasPressure which, in turn, is a child of hasMagnitude, or hasFunction is a child of implements. An important part of the relations list has been imported from the Open Source Relation Ontology (RO) (RO, 2007).

3.2 ACGT-DAS

The ACGT-DAS provide a means to solve syntactic heterogeneities—i.e. they provide uniform data access interface. ACGT-DAS are required also to export the data schema of each individual source, in order to aid the clients in building queries.

The ACGT-DAS offer a web service interface. They have been implemented using the Open Grid

Services Architecture Data Access and Integration (OGSA-DAI) services (Antonioletti, 2005)

SPARQL (SPARQL, 2007) has been chosen as the query language. More expressive than its predecessor RDQL, the language used by an early version of the mediator, SPARQL offers new features, becoming an intermediate level (in terms of expression) language, appropriate for being used as common query language. It is less expressive than Structured Query Language (SQL), due to the lack of support of any form of aggregation. SQL is a relational specific language, so it cannot be used as common language by the ACGT-DAS (mainly because of the selected query translation approach). On the other hand, it is more expressive than DICOM.

```

SELECT ?patientId ?studyId ?seriesId
WHERE {
  ?patient dicom:PatientID ?patientId ;
  dicom:PatientsName "Huge, Lurch" .
  ?study dicom:Patient ?patient ;
  dicom:StudyInstanceUID ?studyId .
  ?series dicom:Study ?study ;
  dicom:SeriesNumber "3" ;
  dicom:SeriesInstanceUID ?seriesId .
}

```

Figure 4: Example of DICOM query in SPARQL.

A relational database and a DICOM wrapper have been developed so far. We used D2RQMap (Bizer, 2004) for the implementation of the relational databases wrapper. This was a straightforward process, due to the technology choices. The development of the wrapper for querying DICOM image databases was not as direct. DICOM uses a four-level hierarchical information model, not a classical relational model. This structure caused difficulties in the query transformation process, since SPARQL is more expressive than DICOM, so the final queries may not be able to represent the view that is expressed in the original one. Figure 4 shows an example of a DICOM query expressed in SPARQL. A set of special functionalities had to be implemented to support retrieving of DICOM images as well.

3.3 ACGT-SM

The ACGT-SM aims at solving the semantic heterogeneities present in databases to support interoperability and integration. The ACGT-SM is supported by a set of satellite tools, like the mapping tool and the data cleaning module among others.

The approach selected to perform database integration has been *Local as View*. This decision is based on the nature of data in the biomedical

domain, and more concretely in post-genomic clinical trials. *Local as View* based techniques require less amount of effort when the structure of data sources changes or when new ones need to be integrated in the system. However, as stated before, some performance issues are associated to this kind of approaches. In order to overcome these difficulties, the domain representing the integrated set of databases is constrained. This restriction is based on requirements specified by the end users.

Semantic heterogeneities are tackled following an ontology based approach. The ACGT-MO acts as a semantic framework supporting homogenization. In *Local as View*, the ACGT-MO acts as global schema. The goal of this global schema is twofold: 1) provides a means to build the local views of the underlying databases, and 2) represents the set of queries that can be formulated by the users.

SPARQL has been selected as query language for the ACGT-SM. As said previously, SPARQL is used as query language by the ACGT-DAS. This homogeneity allows saving time and memory resources. When a query is launched through the ACGT-SM, the software divides it into a set of dedicated queries for the underlying data access services, wrapping the actual databases. No interface is needed to translate these queries, given that the same query language is used by the ACGT-SM and the ACGT-DAS.

The results of a query are returned by the wrappers in XML SPARQL Result format. The ACGT-SM builds an integrated set of results as an ontology instance file. These instances are represented using the OWL ontology description language. Other formats, such as CSV (Comma Separated Variables) are supported to fit the requirements of the Data Analysis Tools in ACGT.

Parallel to the ACGT-SM, an API for creating mappings between the ACGT-MO and RDF Schemas of data sources to be integrated has been developed. This API offers a flexible and generic approach for creating mappings, and is based on path mapping. Paths from the ACGT-MO are mapped to paths from an RDF Schema, providing a way to translate queries to the ACGT-MO—which are basically sets of paths—into queries to the RDF Schema. A graphical interface on top of this API is being developed as well.

3 EXPERIMENTS AND RESULTS

We have tested the first version of our tools in a case study including a set of three different sources,

including two clinical relational databases—SIOP nephroblastoma database and TOP breast cancer database— and a DICOM repository of images. We carried out two experiments integrating DICOM source and each one of the clinical trials databases. Our tool integrated the sources successfully, and the generated schemas were validated by experts in the domain.

The experiments were performed using a dedicated web interface. This interface was built for demonstration purposes within the ACGT project, but it is not the final query interface—this interface is now in its design phase, and is going to be able to support more complex queries. Figure 5 shows the result of the execution of a query combining SIOP and DICOM data.

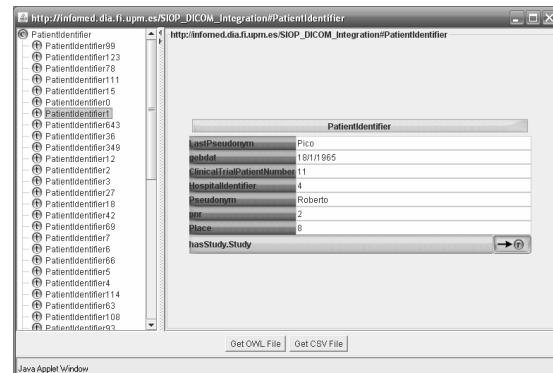


Figure 5: Instances retrieved: SIOP and DICOM integration.

As can be seen in Figure 5, the user can request the DICOM studies related to the patient selected by clicking on the relation button.

The experiments showed the suitability of the approach adopted to cope with data access and integration in the post-genomic clinical trials domain. This prototype was developed using Java and HTML languages.

4 CONCLUSIONS

In this work, we present a set of tools to provide data access, allowing seamless access and integration of heterogeneous databases. To this end, a clinical trials on cancer domain ontology has been developed, the ACGT-MO, and two core services to overcome syntactic and semantic heterogeneities, namely the ACGT-DAS and the ACGT-SM.

The ACGT-MO covers the domain of clinical trials on cancer, and has been built using Clinical Report Forms from SIOP and TOP trials. The

ACGT-MO follows the recommendations of the OBO Foundry.

The ACGT-DAS resolve syntactic heterogeneities present in disparate sources of information. They provide a homogenous query language, SPARQL, and a web service interface developed using the OGSA-DAI middleware.

The ACGT-SM is able to process user queries formulated by means of a global model—i.e. the ACGT-MO—, and to retrieve information from a set of integrated heterogeneous databases. The ACGT-SM is supported by a set of satellite tools tackling with problems such as mapping and instance homogenization.

The results obtained in the carried out experiments prove that this approach can properly integrate relational and image databases.

In the second phase of the project, we plan to add new types of sources, such as public web databases, different file formats—e.g. plain text, Excel spreadsheets, XML, etc— and microarray data.

ACKNOWLEDGEMENTS

The authors would like to thank all members of the ACGT consortium who are actively contributing to addressing the R&D challenges faced. The ACGT project (FP6-2005-IST-026996) is partly funded by the EC and the authors are grateful for this support.

REFERENCES

- Abiteboul S., Duschka O., 1998. Complexity of answering queries using materialized views. In *Proc. of the 17th ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS'98)*, pages 254-265.
- Antonioletti, M., et al. 2005. The Design and Implementation of Grid Database Services in OGSA-DAI- In: *Concurrency and Computation: Practice and Experience*, Volume 17, Issue 2-4, Pages 357-376.
- BFO, The Basic Formal Ontology. Available at: <http://www.ifomis.uni-saarland.de/bfo> [13 oct 2007]
- Bizer, C., Seaborne, A. 2004. D2RQ – Treating Non-RDF Databases as Virtual RDF Graphs. In: *Proc. of the 3rd International Semantic Web Conference (ISWC2004)*, Hiroshima, Japan. Poster presentation.
- caCORE: A Common Framework for Cancer Data Management. Available at: http://ncicb.nci.nih.gov/infrastructure/cacore_overview [13 oct 2007]
- Cali, A., De Giacomo, G., Lenzerini, M., 2001. Models for information integration: Turning local-as-view into global-as-view. In: *Proc. of Int. Workshop on*

Foundations of Models for Information Integration (10th Workshop in the series Foundations of Models and Languages for Data and Objects).

- Ceusters W, Smith B, Kumar A, Dhaen C. 2003. Mistakes in Medical Ontologies: Where Do They Come From and How Can They Be Detected?. in: Pisanelli DM (ed.): *Ontologies in Medicine: Proceedings of the Workshop on Medical Ontologies*, Rome, October 2003. IOS Press, Amsterdam.
- FMA, The Foundational Model of Anatomy. Available at: <http://fme.biostr.washington.edu:8089/FME/index.html> [13 oct 2007]
- GALEN, Common Reference Model. Available at: <http://www.opengalen.org> [13 oct 2007]
- GO, Gene Ontology project. Available at: <http://www.geneontology.org> [13 oct 2007]
- Grenon, P., Smith, B., Goldberg, L., “Biodynamic Ontology. 2004. Applying BFO in the Biomedical Domain,” in: *Ontologies in Medicine*, D. M. Pisanelli, Ed., Amsterdam: IOS Press, pp. 20-38.
- HL7, Reference Information Model. Available at: http://www.hl7.org/Library/data-model/RIM/modelpage_mem.htm. [13 oct 2007]
- Kimball R., 1996. *The Data Warehouse Toolkit: Practical Techniques for Building Dimensional Data Warehouses*, John Wiley.
- Lindberg C. The Unified Medical Language System (UMLS) of the National Library of Medicine. *Journal of the American Medical Record Association* 1990; 61(5):40-2.
- NCI Enterprise Vocabulary Services. Available at: <http://www.nci.nih.gov/cancerinfo/terminologyresources> [13 oct 2007]
- OBO Foundry ontologies. Available at: <http://obofoundry.org> [13 oct 2007]
- Protégé, The Ontology Editor and Knowledge Acquisition System. Available at: <http://protege.stanford.edu> [13 oct 2007]
- RO, OBO Relation Ontology. Available at: <http://www.obofoundry.org/ro> [13 oct 2007]
- Sheth A. P., Larson J. A., 1990. Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases, *ACM Computing Surveys*, 22(3): pp.183-236.
- Smith, B, Ceusters, W. HL7 RIM: An Incoherent Standard, *Stud Health Technol. Inform.* 2006; 124: 133–38
- SNOMED Clinical Terms Core Content. Available at: <http://www.snomed.org/snomedct/coreterms.html> [13 oct 2007]
- SPARQL Query Language for RDF. Available at: <http://www.w3.org/TR/rdf-sparql-query/> [13 oct 2007]
- Ullman J.D., 1997. Information integration using logical views. In *Proc. of the 6th Int. Conf. on Database Theory (ICDT'97)*, volume 1186 of *Lecture Notes in Computer Science*, pages 19-40. Springer-Verlag.
- Wiederhold G., 1992. Mediators in the Architecture of Future Information Systems, *IEEE Computer*, 25(3): pp. 38-49.

MEDICAL KNOWLEDGE REPRESENTATION WITHIN HEARTFAID PLATFORM

Dragan Gamberger, Marin Prcela, Alan Jović, Tomislav Šmuc
Rudjer Boskovic Institute, Bijenicka 54, Zagreb, Croatia
dragan.gamberger@irb.hr, marin.prcela@irb.hr, alan.jovic@irb.hr, tomislav.smuc@irb.hr

Ganfranco Parati, Mariaconsuelo Valentini
Cardiology II, S.Luca Hospital, Istituto Auxologico Italiano, Milano, Italy
parati@auxologico.it, mcvalent@hotmail.com

Kalina Kawecka-Jaszcz, Katarzyna Styczkiewicz, Andrzej Kononowicz
I Cardiac Department, Jagiellonian University Medical College, Cracow, Poland
mckaweck@cyf-kr.edu.pl, kstyczki@cm-uj.krakow.pl, mykonono@cyf-kr.edu.pl

Antonio Candelieri, Domenico Conforti, Rosita Guido
Department of Electronics, Informatics and Systems, University of Calabria, Rende (Cosenza), Italy
acandelieri@deis.unical.it, mimmo.conforti@unical.it, rguido@deis.unical.it

Keywords: Heart failure, knowledge representation, ontology, production rules.

Abstract: The paper presents the results of the development of the knowledge base for the HEARTFAID platform. The means and methods used to collect, systematize and formalize medical knowledge, as well as to test the developed knowledge representation are described. The descriptive part of the knowledge base is realized as an ontology which conceptualizes the heart failure medical domain. The procedural part of the knowledge base is realized through sets of production rules. The procedural knowledge covers the tasks of heart failure diagnosis, severity assessment, treatment process, medication prescription and dosage, medication contraindications, prognosis estimation, and acute decompensation detection. Finally, medical plans are used to present medical actionable knowledge. Currently they are used only to systemize procedural knowledge development process but they present a challenge for the future work in the field of medical knowledge representation.

1 INTRODUCTION

Medicine is the field characterized by the enormous amount of existing expert knowledge and at the same time there is a need for constant and reliable decision making. This is an ideal scenario for building and using automated knowledge based decision systems. Building an effective knowledge base is a challenge relevant not only for the Heartfaid platform, but for all artificial intelligence applications as well. It is also known as a hard problem with possibly many different solutions among which none can be selected as ideal or optimal for all situations.

The knowledge representation is actually a very lively research field, especially in medical

application. Heartfaid platform is a good example of a real medical environment for which automated intelligent decision making is necessary. In our work we first tried to test different knowledge representation options and then to select and use modern and most appropriate technology for solving concrete decision making problems.

2 MEDICAL MOTIVATION

Heart failure is increasingly frequent in western world, carrying a high mortality rate and being responsible for a consistent increase in healthcare costs related to the multiple therapeutic interventions

and the high frequency of hospital admissions required by these patients. There is therefore an increasing need for a better care, that might be provided not only by highly specialized centers, but also by small hospitals and by field cardiologists, a need that has to be matched with a policy of cost containment. Progress in technology may offer an important support to make this possible, allowing adequate knowledge to be made available to all health care providers in this field. It might also offer new methods for regular and accurate collection of biological signals in patients living in their home environment, making use of sensors, either traditional or wearable, able to provide a continuous monitoring also through telemedicine facilities. More recent progress might further offer an advanced platform of services for the automated integration between the signals collected both at home and in the clinic environment on one side, and the available state-of-the-art knowledge on the other side, providing an artificial intelligence support to clinical decision. The proper adoption of these tools might help improving the daily care of chronic heart failure patients, through a prompt titration of treatment in response to early detection of even minor changes in clinical conditions, as well as through a reduction of diagnostic and therapeutic errors, by reinforcing the implementation of the most advanced recommendations provided by international clinical guidelines. Such an approach might also help improving the cost/effectiveness of heart failure care facilitating the implementation of a disease management approach, in which therapy, education and follow-up are tailored for each patient by a multidisciplinary team constantly supported by an advanced platform of computerized services guiding the clinical decision through a continuous update of patient's clinical conditions allowed by advanced wireless telecommunication technology. The Heartfaid project is aimed at developing such a tool and at testing its feasibility and usefulness in the management of chronic heart failure patients, focussing in particular on those with advancing age.

3 RELATION WITH DECISION SUPPORT SUBSYSTEM

Decision support subsystem (DSS) is the part of the platform responsible for its intelligent behaviour and the knowledge base (KB) is the representation of the medical knowledge necessary for the DSS operability. In order that this knowledge can be used

by the DSS, it must be presented in a formally sound way. The task of building the knowledge base consists of collecting the relevant medical knowledge, its systematization, and technical formalization.

User services are not supposed to directly access the knowledge available in the KB. They can only ask for the assistance of the DSS, which can then decide to use the KB for its decision making process. It means that during normal platform operation, the KB, with exception of DSS, is isolated from other platform parts. In contrast to that, during the platform development the KB is perhaps the most relevant integrative part between medical and technical partners. Building it presents the challenge of transferring all aspects of relevant medical knowledge into the platform. The success in this work significantly determines the overall performance of the system.

On the other side, DSS completely determines the requirements set on formal aspects of the knowledge representation task. The final goal is effective interplay between the DSS and the KB.

Already in the project definition it has been determined that the ontological form of knowledge representation is the most appropriate for the concrete task. In the DSS development phase the semantic web OWL ontology form with integrated rules in the SWRL form have been selected as optimal solution. The main reason for these decisions is that only relative simple representation forms without explicit time component and with deterministic logic can be effectively handled by available open source interpreters and reasoning systems.

In our work we have also experimented with some other systems, especially those specifically designed for medical applications and guideline modelling. We have worked with Arden syntax, GLIF, Asbru, and PROforma.

Arden syntax (Hripessak et al, 1993) is a rule-based system adopted in the year 1992 and it is now part of HL7 standard. The main idea of the Arden syntax is to add as much as possible human-readable information to machine-readable rules. Each rule is stored in a single file and is called a Medical Logic Module (MLM). The drawback of Arden syntax is that it does not refer to any kind of domain description (ontology). Due to this, the system needs to interact directly with a clinical database in order to provide alerts and reminders, which strongly hinders knowledge sharing. Also, the execution components are not freely available (mostly because of Arden's institution dependency). However, unlike

the vast majority of other systems, it has found a practical usage in real clinical environments.

GLIF (Peleg et al, 2000) provides a framework for developing medical guidelines that are both easily understandable by humans (medical experts) and interpretable by machines. Each GLIF guideline is modelled in the form of a flowchart (directed graph). GLIF is suitable for describing logic sequence of actions. Within the HEARTFAID platform GLIF may be used to represent the logical flow of actions, e.g. sequence of tests performed for diagnosing disease or prescribing therapy but the problem is that there exists only commercial execution engine (Glee).

Asbru (Shahar et al, 1998) is a guideline modelling tool which focuses on representing medical plans. It is highly aware of the time dimension in the medical procedures and actions. A plan in Asbru is a set of actions that are performed when certain preconditions hold. Each plan is decomposed into more sub-plans that are performed sequentially, concurrent (parallel execution) or cyclical. Within the HEARTFAID platform, Asbru can be used in situations where actions are taken in a predefined order, e.g. to describe the procedure at the baseline evaluation or additional patient visits to the clinics. However, there are no freely available execution engines that may be integrated into HEARTFAID platform.

PROforma (Sutton et al, 2003) is a knowledge composition language that aims to assist patient care through active decision support and workflow management. Similar to the GLIF model, it represents also guidelines as a directed graph in which nodes represent instances from the PROforma task ontology. PROforma contains a number of tools for developing guidelines. A major focus point is on guideline safety by defining additional safety-related operators such as integrity and safety constraints. Considering the execution engines, Arezzo is a commercial version of PROforma, while Tallis is a version available for educational and research purposes (under license agreement).

4 DESCRIPTIVE HEART FAILURE KNOWLEDGE

The first step in the development of the knowledge base for the Heartfaid platform has been development of the heart failure (HF) ontology. It presents the formalized description of concepts for the whole heart failure domain. It includes basic HF

concepts, properties that characterize patients, all relevant diagnostic examinations and tests, and treatment procedures. The ontology also includes other cardiovascular system related concepts as well as concepts related to other organs when they are connected with HF. The information presented in the ontology has been obtained by human interpretation of guidelines for congestive and acute heart failure (<http://www.escardio.org/knowledge/guidelines/>), Heartfaid reports, as well as from other medical knowledge sources, including, but not limited to UMLS (Unified Medical Language System), Mayo clinic web site and Open Clinical web site.

In its current form the ontology presents the detailed taxonomic overview of the HF domain with around 200 classes describing HF related concepts. Examples are "Cardiac_hypertrophy", "Blood_pressure_signs" or "Heart_murmurs". These concepts are interconnected with super-class and sub-class properties into a hierarchical tree-like structure. At the basic level there are five relevant super-classes: "HF_concept", "Patient_characteristic", "Patients", "Testing", and "Treatment". Figure 1 presents the Protégé tool displaying these five super-classes with some of their most relevant sub-classes.

Individuals or instances are members of the classes and typically present exhaustive list of concrete concepts relevant for the class. For example, the "Cardiac_hypertrophy" class has following six instances: "Cardiomegaly", "Combined ventricular hypertrophy", "Left atrial hypertrophy", "Left ventricular-hypertrophy", "Right atrial hypertrophy", and "Right ventricular hypertrophy". The ontology includes more than 2000 individuals. When possible, classes are specified with their CUI number (Concept Unique Identifier according to UMLS) and with a list of synonyms. For example, for the class "Heart_diseases" its CUI is C0018799 and its synonyms are "Disorder_of_heart", "Cardiac_diseases", "Cardiopathy".

Finally, the ontology contains properties that connect individuals in different classes. These properties are relevant because they enable introduction of relations among concepts. For example, individual "Valvular_heart_disease" from the class "Heart_valve_diseases" is indicated by the individual "Dyspnea" from the class of "Signs_and_symptoms". Or that "Hyperkalemia" from the class "Potassium_disorder" may be caused by medications like "Potassium_sparing_diuretics" or "Spironolactone". The names of these properties are "Indicated" and "MayBeCausedByMedication".

The HF ontology includes definitions of more than 100 properties.

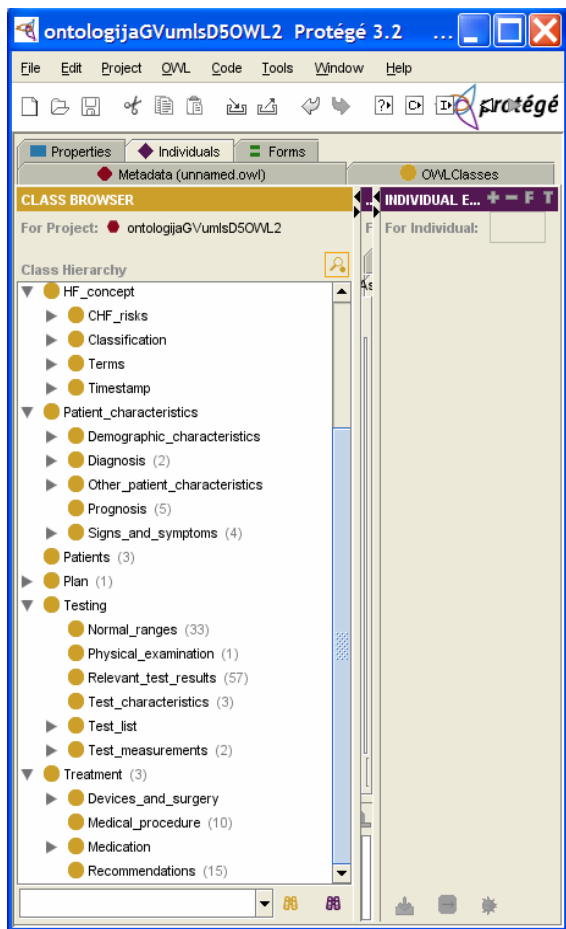


Figure 1: Protégé tool used to display a part of the HF ontology.

The ontology presents descriptive domain knowledge. This knowledge is of two types. The first is defined by the generality relations among instances and classes, as well as by the generality relations among sub-classes and super-classes. In this way for each concept presented by some instance there is a series of *is-a* relation. For example, it means that "Cardiomegaly" *is-a* "Cardiac_hypertrophy" while "Cardiac_hypertrophy" *is-a* "Heart_disease". The second type of descriptive knowledge contained in the ontology comes from properties that define relations between classes, such as "Indicated" or "Maybe CausedByMedication", mentioned before.

Descriptive HF ontology is publicly available in the web form from the project web site at <http://www.heartfaid.org/links.php>.

5 PROCEDURAL KNOWLEDGE

The HF ontology presents the detailed taxonomic overview of complete heart failure domain including relevant relations among concepts. It represents *descriptive knowledge* about the domain. Platform should also be able to perform some actions, typically in the form of suggestions for patients and medical personnel. The knowledge representing sufficient and necessary conditions that some actions can be done is the so called *procedural knowledge*. Descriptive and procedural knowledge together present the knowledge base of the HF platform.

5.1 Production Rules

Production rules in a form "IF *some condition is true* THEN *make some action*" are a widely used approach for the presentation of procedural knowledge. Their advantages are natural interpretation by humans and modularity during construction. It is also relevant that production rules are also a formal way of presenting knowledge and in this way a good starting point for practical realization of the decision support system. For the integration of descriptive and procedural knowledge it is important that production rules can use only the concepts defined in the HF ontology.

At the knowledge presentation level it is very important that production rules can be easily understood and corrected by medical experts. In this way the major advantage of presenting procedural knowledge in the form of production rules is that they present formal enough way to present knowledge that can be used by the platform and that at the same time medical experts can easily control the expected performance of the platform. The correction of rules or adding them should only be performed by the authorized medical personnel.

The HF procedural knowledge has been divided into 10 functional subtasks in order to enable easier human control of the completeness and consistency conditions. They are:

1. HF diagnosis
2. Alternative or additional diagnosis
3. Heart failure severity assessment – specifying NYHA class for patient, which is required for general patient treatment approach
4. HF general treatment process based on severity assessment
5. HF medications contraindications, adverse effects & additional treatment rules
6. Prognosis estimation for HF patients

7. Non-pharmacological management and recommendations
8. Specific medication prescription and dosage
9. Acute decompensation of congestive heart failure
10. Heart failure cause and CAD risk factors

An example for a rule from the diagnosis subtask is:

Diagnosis: **Systolic heart failure**

IF

Patient has either heart failure signs or heart failure symptoms

AND

ECG abnormal (left bundle branch block AND anterior Q waves)

AND

patient has (ischemic heart disease)

AND

Chest X-ray abnormal (cardiothoracic ratio > 0.5)

AND

natriuretic peptides abnormal (BNP > 100 pg/ml)

5.2 Soft Computing

Intelligent medical applications require the ability to work with imprecise or only partially true data. The goal is to ensure robustness and efficiency of the decision making process in a real world environment. Soft computing techniques including fuzzy systems and probabilistic reasoning can be used to solve these problems. These techniques typically lead to relative complex systems whose performance and final decisions are rather difficult to predict. For the platform we have decided to solve the problem by a) a mixture of deterministic production rules with fuzzy output values (consequences) and b) complex but deterministic computation of some input values that mimic fuzzy inputs.

5.2.1 Fuzzy Consequences from Deterministic Rules

The approach means that we have deterministic rules, deterministic rule inputs, and deterministic outputs which may have different, but in advance predefined levels of reliability or probability.

An example of the deterministic rule with fuzzy consequence is that *Heart failure is possible* IF *Patient has either heart failure signs or heart failure symptoms* AND *cardiothoracic ratio > 0.5*. Such rule has precisely defined conditions but the level of the reliability of the consequence is rather low. Higher level of reliability is the rule with the consequence *Heart failure is probable*, while the

highest level is that *Heart failure diagnosis is suggested*.

The advantage of this approach is that decision making process is deterministic and consequently relatively simple. The application of this approach is possible for the HF platform because decisions are directed to humans who must decide upon their acceptability and they are not automatically executed. Medical doctors are the only ones who can confirm and follow these suggestions. In this framework we do not have closed loop decisions. Suggestions with a predefined level of reliability are completely acceptable. Moreover, they are easier to interpret by humans than the numerical values produced by probabilistic reasoning and completely fuzzy systems.

In addition to different levels of the diagnosis reliability in which we have four levels (suggested, probable, possible, unlikely), we use fuzzy conclusions for the prognosis (good, worse, very poor), for medication recommendation (suggested, consider) etc.

5.2.2 Computation of Complex Patient Descriptors

Some deterministic rules require inputs that describe patient status that are difficult to define by absolute values. Such inputs can be described as fuzzy because the same value can in different situations have different meaning. Good examples are all values that should be interpreted relative to some other, current or previous, patient characteristics and measured values.

We avoid using fuzzy values by implementing complex computation in the process of preparing the inputs for decision making. This means that in this computation we must take into account, besides basic patient characteristic, also all other properties that significantly determine its previous or current status relevant for the interpretation of the basic characteristic. For example, for the input *significant arterial drop* we have to look into the complete patient history, compute mean blood pressure values, and then based on the current value that is more than 30 mmHg lower conclude on significant arterial drop.

The computation of complex descriptors can be effectively realized inside the factual knowledge building block. It has access to the complete set of patient data and this enables that all data necessary for complex computations can be acquired. If some data can not be found in the patient record then the final patient descriptor will have unknown value and

this will prevent the respective rule to fire. For the sake of decision reliability, the rules should have also simple security cut-off points, like present systolic blood pressure below 90 mmHg in the previous example that will fire also if data necessary for complex computations are not available.

The main disadvantage of this approach is that rather complex computations relevant for the decision making process are built into fixed programmed logic with the consequence that they can not be changed easily. The advantages are simplicity of procedural knowledge and the reliability of the DSS process.

5.3 Ontological Representation of Procedural Knowledge

Presentation of HF procedural knowledge in the form of production rules does not mean that their practical realization for the decision support purposes must be in the same form. It is true that these rules may be used to build a rule based expert system, but these rules can also be used for representing procedural knowledge in other forms. For the HF platform, the integration with descriptive knowledge presented in the HF ontology is relevant. In this situation, an appropriate form for procedural knowledge is SWRL (Semantic Web Rule Language), which is a logical extension of the OWL (Web Ontology Language).

But there is also another possibility. The unique property of OWL is the ability to represent logical operations between classes and between classes and individuals using the so called *concept constructors*. It enables that logical relations contained in production rules can be presented in the ontological form. The result is the ontology that contains both descriptive and procedural knowledge. The advantage of the approach is a tight connection and conceptualized representation of complete domain knowledge which may potentially lead to a more intuitive representation of medical knowledge. In the future this can also enable web based distributed decision support, but this is not relevant for the current platform realization.

By integrating all of the ten sets of production rules, we have built the *procedural HF ontology*. It is different than the descriptive HF ontology because it has two root classes: "Patient" and "Patient_characteristics". The "Patient_characteristics" class contains descriptive knowledge necessary for logical relations in production rules. All of its subclasses and individuals, including the class hierarchy, are based on and can be thought of as a

subset of the HF ontology. Theoretically the complete descriptive HF ontology could be integrated here but this was not done because of reasoning efficiency. So the "Patient_characteristics" class contains only descriptive knowledge necessary for reasoning with current version of procedural knowledge.

The class "Patient" contains the complete procedural knowledge. In order to be compatible with production rules organization it has ten subclasses, each of them representing one rule set. They are further divided into many subclasses. Every class with no subclasses has necessary and sufficient conditions defined, and this is where procedural knowledge is stored. Each of the class definitions can be found as a rule in one of the rule sets. Fulfilling conditions for being in the class is perceived as a suggestion for a particular patient, e.g. class Patient_Severity_assessment has subclass NYHA_IV. Conditions for this class are: patient has dyspnoea, fatigue or palpitations at rest and patient has heart failure. The patient with these characteristics fulfils the conditions for being in this class.

Testing is always the significant part of the KB development process. By presenting production rules in the ontological form we have enabled that the problem can be solved by developing a java-based OWL interpreter with added concept "negation-as-failure" into the logic semantics of OWL. The hybrid *instance checking* process obtained in that manner aims to combine the OWL syntax with the *closed world assumption* semantics. The developed Protégé plug-in integrates the interpreter into the Protégé with a simple user interface. In this way, the interpreter is introduced directly into the knowledge base building facility, which eases the process of building and maintaining the knowledge base.

6 MEDICAL PLANS

Medical plans are textual and visual presentations of procedures that take place after detection of some events. Events can be any type of health disorder including signs, symptoms, and diagnosis. The main characteristic of medical plans is that they are event driven and because of that they present actionable view of the medical knowledge. This actionable view is a special case of the more general procedural knowledge.

In the Heartfaid platform medical plans are only a middle step between experts and the guideline modelling tools which persuade the experts to clearly state the procedure they would normally perform when facing a specific problem. At the same time they enable technicians to understand it and encode it in a machine readable form (de Clercq et al, 2004). They are similar to medical pathways but in contrast to medical pathways which are designed to be used by medical doctors in order to systemize and standardize their work, medical plans are designed for technical people in order to better understand medical concepts. In the HF domain we have used medical plans as an auxiliary tool for systemizing procedural knowledge development and to enable some verification of the implemented knowledge base.

The syntax of the medical plans highly resembles the traditional workflow management. The difference is that the medical plans will not be executed by machines; they are written in an almost-free graph/text form with main purpose to be fully understandable by humans. Their main characteristic is that they are event driven and their main advantage is a clear systematization of the medical procedures and interconnections among them. Additionally, their visual presentation facilitates understandability by medical experts.

For the heart failure platform, medical plans describe the disorders that can occur as events to the patient who is treated by the platform. These disorders have assigned urgency levels which correspond to the type of response needed from the medical team. For example, pulmonary edema has the highest urgency level, requiring immediate admission to the hospital. An example of a symptom that has a low urgency level is cough. Cough does not require for the patient to report to the hospital, but rather if it is persistent, he should contact his general practitioner. The urgency levels solve the problem of entering the appropriate medical plan in situations when more than one triggering event occurs. The plans of the lower urgency level can be interrupted if another event of the higher urgency level happens.

At the moment, the heart failure system has 38 interconnected plans for signs, symptoms and diagnosis assessment and treatment and 15 plans for medications prescription and dosage. Most of them have been presented in both graphical and textual form. Figure 2 presents the medical plan for handling heart failure patients with increased body temperature.

The usual way of designing medical plans is in close resemblance to the medical doctor's way of thinking when handling a patient. The most common procedure would be to ask for other symptoms in relation to the one the patient complains about and to do the examination and find the appropriate signs. These other signs and symptoms can either confirm the initial suspicion, or request that some tests should be taken, or completely disprove the existence of the disorder. Usually, the next thing a medical doctor would do is to order a series of tests. These tests can also confirm the suspicions or give rise to a new possible diagnosis. The next course of action is to prescribe the appropriate treatment or to give recommendations.

6.1 Further Research Topics

In the HF project the medical plans have been developed by technical people in the phase of procedural knowledge development. The goals were to demonstrate medical domain understanding and to systemize acquired knowledge. Their significance is in the fact that they present a middle step between the experts and their expert knowledge and technical people that formalize the knowledge.

Development of medical plans opened some potentially interesting questions that might be very relevant as further research topics. The first is whether all types of useful medical procedural knowledge (or at least its major part) can be described by medical plans. If the answer is positive, then it would be interesting to think about the possibility to make medical plans executable directly without their transformation to other forms (rules, ontologies, workflows) or to try to enable their automatic conversion without human intervention (de Clercq et al, 2004). Based on the work and results in the HF project, these options seem interesting because the approach based on medical plans as the first and potentially the only creative part requiring human intervention, could significantly change the traditional way of designing procedural knowledge.

7 CONCLUSIONS

The paper presents the main results of the work related to collection, systematization, and formalization of the knowledge related to the heart failure domain. The main results are: descriptive HF

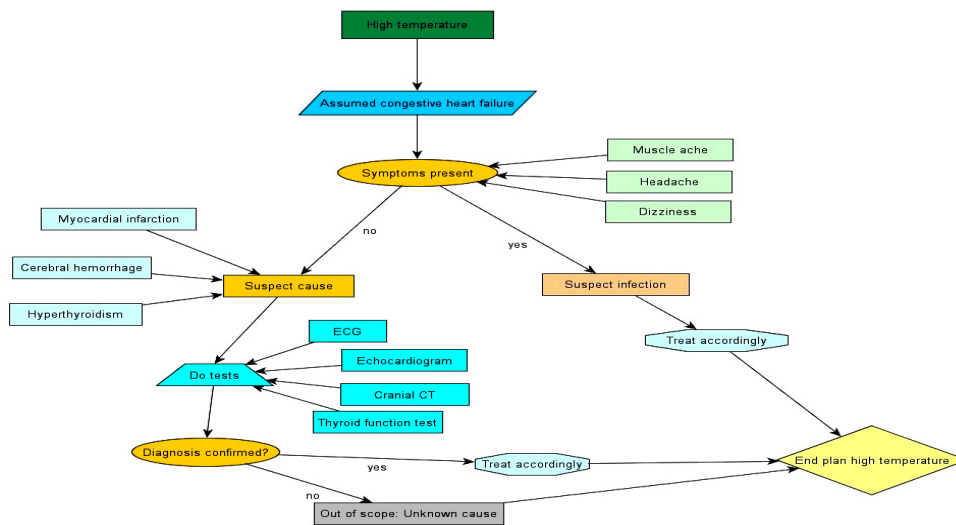


Figure 2: Medical plan for heart failure patients with increased body temperature.

ontology, procedural knowledge base, HF medical plans, and ontological presentation of procedural knowledge.

Although the constructed knowledge base has been partially verified and improved by medical doctors, the current version presents technical formalization of medical guidelines and starting point for Heartfaid platform implementation. It can be expected that tests with prototypes will demonstrate deficiencies in the form and content of the knowledge base. By these improvements we expect to be able to collect and formalize also the tacit medical knowledge related to HF. Long and detailed experimental work with the platform is the necessary condition for the success of this process. Moreover, even in the operational life of the platform it can be expected that continuous improvements in the knowledge base will be necessary.

ACKNOWLEDGEMENTS

The work was done for EU FP6 project "Heartfaid: A knowledge based platform of services for supporting medical-clinical management of the heart failure within the elderly population". It was supported by Croatian Ministry of Science, Education and Sport project "Machine Learning Algorithms and Applications".

REFERENCES

- Hripesak, G., Clayton, P.D., Pryor, T.A., Haug, P., Wigertz, O.B., Van der lei, J., 1993. *The arden syntax for medical logic modules*. Forteenth Annual Symposium on Computer Applications in Medical Care, 200-204. IEEE Computer Society Press.
- Peleg, M., Boxwala, A., Ogunyemi, O., 2000. *GLIF3: The Evolution of a Guideline Representation Format*. Proceedings of the American Medical Informatics Association (AMIA) Annual Symposium. (20 Suppl):645-649. Washington DC: AMIA.
- Shahar, Y., Miksch, S., Johnson, P., 1998. *The Asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines*. Artificial Intelligence in Medicine, 14(1-2), 29-51. Springer, 1998.
- de Clercq, A., Blom, J.A., Korsten, H.M., Hasman, A., 2004. *Approaches for creating computer-interpretable guidelines that facilitate decision support*. Artificial Intelligence in Medicine 31(1), 1-27. Springer, 2004.
- Sutton, D.R., Fox, J., 2003. *The Syntax and Semantics of the PROforma guideline modelling language*. Journal of the American Medical Informatics Association: JAMIA, 2003 Sep-Oct;10(5):433-43. Hanley & Belfus.

AUTHOR INDEX

Aalst, W.	118	Dusek, L.	168
Abaydulla, B.	196	Faria, F.	182
Abreu, F.	182	Fensli, R.	3
Aguiar, F.	182	Ferreira, A.	182
Andreão, R.	186	Freund, J.	134
Andresen, H.	59	Friedrich, C.	265
Anguita, A.	301	Funyu, Y.	196
Antunes, L.	182, 215	Fuxreiter, T.	176
Arisawa, H.	258	Gaffney, P.	81, 142
Aspray, W.	17	Gamberger, D.	307
Ávila, F.	28	Gomes, F.	182
Bakker, P.	118	Gomes, H.	72
Baldacchino, J.	172	Gottberg, H.	205
Bianchi, A.	273	Grad, J.	209
Binotto, A.	28	Grillo, A.	43
Boisen, E.	3	Grimson, J.	81, 142
Bonsma, E.	301	Gualtieri, A.	288
Brabec, P.	168	Guelfi, N.	172
Brochhausen, M.	301	Guido, R.	307
Bucur, A.	301	Habetha, J.	296
Buitelaar, P.	134	Han, D.	110
Burdescu, D.	192	Hanke, S.	176
Cabitza, F.	9	Harris, M.	296
Candelieri, A.	307	Hata, S.	258
Cascio, M.	102	Hayes, B.	17
Cerutti, S.	273	Hofmann-Apitius, M.	265
Chiarugi, F.	288	Hoogendoorn, M.	148
Chronakis, I.	265	Hsiao, C.	67
Claerhout, B.	279	Huang, C.	102
Clarke, F.	81, 142	Huang, L.	67
Cocos, C.	301	Ikeda, T.	242
Coenen, F.	34	Inoue, Y.	242
Colantonio, S.	288	Iochpe, C.	28
Conforti, D.	288, 307	Ion, A.	192
Contreras, D.	159	Jang, E.	254
Costa, T.	227	Johnson, W.	126
Costa-Pereira, A.	215, 219	Jović, A.	307
Cruz-Correia, R.	182, 215, 219	Jung, S.	110
Cunha, A.	215	Kao, T.	67
Cunha, J.	72	Kawecka-Jaszcz, K.	307
Damij, N.	209	Khan, M.	34
Darlington, K.	223	Kim, D.	254
Das, A.	87	Kim, T.	254
Doerr, M.	301	Klein, M.	148
Dunlop, R.	265	Ko, I.	110

AUTHOR INDEX (CONT.)

Kondylakis, H.	301	Pereira, C.	28
Kononowicz, A.	307	Perticone, F.	288
Kurumatani, K.	242	Pinho, C.	182
Lagani, V.	288	Pinto, V.	227
Lanza, V.	102	Pisa, I.	205, 227, 231, 235
Lederman, H.	235	Plumer, P.	172
Lee, J.	254	Prcela, M.	307
Lee, M.	110	Pudil, P.	246
Lentini, A.	43	Risselada, R.	265
Lien, C.	67	Rubio-Sánchez, M.	94, 239
Liu, H.	126	Rueping, S.	279
Locatelli, M.	9	Ruiz-Rico, F.	94, 239
Lopes, P.	231, 235	Rulli, G.	43
Luglio, M.	200	Sá, C.	182
Macedo, F.	182	Sachpazidis, I.	28
Mans, R.	118	Sakas, G.	28
Maojo, V.	301	Salvetti, O.	288
Martín, L.	301	Santos, E.	182
Martinelli, M.	288	Sanz, F.	265
Martins, S.	87, 182	Sanz-Bobi, M.	159
McGrory, J.	81, 142	Sarcinelli-Filho, M.	186
Me, G.	43	Sasaki, J.	196
Memon, Z.	148	Sashima, A.	242
Mendes, E.	182	Sato, T.	258
Mendez, M.	273	Schaffer, M.	51
Millward, D.	164	Schartner, P.	51
Möller, M.	134	Schmidt, R.	296
Moroni, D.	288	Schonenberg, M.	118
Mota, F.	182	Schor, P.	227, 235
Moturu, S.	126	Schwarz, D.	168
Mukherjee, S.	134	Sciacqua, A.	288
Muyeba, M.	34	Seil, M.	172
Naveira, M.	227	Sfakianakis, S.	279
Neira, R.	235	Shankar, R.	87
Nicholls, W.	164	Siestang, T.	279
Novovicova, J.	246	Sigulem, D.	227, 235
O'Connor, M.	87	Silva, F.	182
Ohta, M.	242	Silveira, M.	172
Oliva, B.	265	Simone, C.	9
Oliveira, L.	186	Sintek, M.	134
Ortolani, C.	231	Šmuc, T.	307
Parati, G.	307	Somol, P.	246
Park, S.	110	Song, M.	118
Pascual, J.	159	Sousa, F.	182
Penka, M.	168	Stanescu, D.	192

AUTHOR INDEX (CONT.)

Stanescu, L.	192
Stenzhorn, H.	301
Stoica-Spahiu, C.	192
Sturkenboom, M.	265
Styczkiewicz, K.	307
Tanaka, M.	196
Tenório, J.	231
Trelles, O.	279
Treur, J.	148
Tsiknakis, M.	279, 301
Valentini, M.	307
Vicedo, J.	94, 239
Vieira-Marques, P.	215
Virvilis, V.	279
Vrijnsen, J.	301
Wienecke, A.	172
Yamada, K.	196
Yamashita, T.	242
Zampognaro, F.	200
Zavitoski, L.	235
Zhou, X.	134
Zúquete, A.	72

