HEARTFAID

Models and Methods for Signals and Images Processing

Submission date: 15/03/2008
Due date of document: 31/01/2008

HEARTFAID

A KNOWLEDGE BASED PLATFORM OF SERVICES FOR SUPPORTING MEDICAL-CLINICAL MANAGEMENT OF THE HEART FAILURE WITHIN THE ELDERLY POPULATION

Project summary	
Project acronym:	HEARTFAID
Project identifier:	IST – 2005 – 027107
Duration of the Project:	01/02/2006 – 31/01/2009
Project Co-ordinator Organization:	UNICAL University of Calabria (Italy)
Project Co-ordinator name:	Domenico Conforti
Thematic Priority:	Information Society Technology
Instrument:	Specific Targeted Research or Innovation Project

Consortium

- UNICAL- Università della Calabria (Italy)
- UNICZ- Università degli studi Magna Graecia di Catanzaro (Italy)
- UNIMIB- Università degli studi di Milano Bicocca (Italy)
- JUMC- Jagiellonian University Medical College (Poland)
- VMWS- Virtual Medical World Solutions Ltd (United Kingdom)
- FORTHNET S. A.- Hellenic Telecommunications and Telematic Applications Company S. A. (Greece)
- SYNAP- Synapsis s.r.l. (Italy)
- CNR- Consiglio Nazionale delle Ricerche (Italy)
- FORTH-Foundation for Research and Technology Hellas (Greece)
- RBI- Rudjer Boskovic Institute (Croatia)
- AUXOL- Istituto Auxologico Italiano (Italy)

Models and Methods for Signals and Images Processing

Document summary	
Document title:	**Models and Methods for Signals and Images Processing**
Document classification:	Derivable D30 Report
Dissemination level:	PU (public)
Submission date:	15 March 2008
Due date:	31 January 2008
Authors:	Sara Colantonio – CNR Davide Moroni – CNR Gabriele Pieri – CNR Ovidio Salvetti – CNR Franco Chiarugi – FORTH Dimitra Emmanouilidou – FORTH
Work package:	WP5 – Data Processing and Decision Support Services
Report version:	1.0

Short description
This deliverable reports the activities carried out in the Task T5.2. In particular, after some recall of the diagnostic resources relevant to HEARTFAID platform, the prototypical algorithms developed for image and signal processing are described in great detail. Then, deployment of these algorithms for the extraction of clinical significant parameters is explained. In parallel, the IT infrastructure needed for the seamless integration of the HEARTFAID Image and Signal Processing toolkit is introduced and the already available modules are described. A discussion about the value added to the CDSS (along with some illustrative use-cases) ends the deliverable, clarifying the possible role of image and signal processing in the management of heart failure.

Change record		
Version number	Changes	Release date
0.1	First draft of the table of contents	04/01/2008
0.2	First version of the report (CNR)	25/01/2008
0.3	Contribution by FORTH (Signal Processing, Standards & Integration issues)	08/02/2008
0.4	Contribution on Value Added to Decision Support	29/02/2008
0.5	Revision of the Signal processing Section	05/03/2008
0.6	Final Draft	08/03/2008
0.7	Final Draft Revised Version	10/03/2008
1.0	Final Version	11/03/2008

Table of contents

Executive Summary

HEARTFAID aims at defining efficient and effective health care delivery organization and management models for an "optimal" patients' management in the field of cardiovascular diseases. An informative and decision support platform of services is in phase of development for improving all the processes related to diagnosis, prognosis, treatment and personalization of health care of the Heart Failure (HF) in elderly population.

Since signal and imaging investigations are a basic step of the diagnostic, prognostic and follow-up processes of heart diseases, the development of methods for computer-aided analysis of such diagnostic data has been included in the goals of HEARTFAID project. The basic idea is to improve the accuracy and consistency in the clinical interpretation of the diagnostic images and signals inherent to the management of heart failure, thus reducing the subjectivity in the findings of diagnostic examinations and improving the reproducibility of the results. In addition computer-aided methods are able to process in a more complete manner the huge datasets produced by some diagnostic resources. As it is obvious that we all already beneficiate of such automated methods (just think how much time is saved with the automatic generation of an Holter report or about the refined algorithms involved in the generation of tomographic images), it is also clear that computer-aided methods will have in general an increasingly important role in health care.

This deliverable is intended to cover the activity on the development of methods for signal and image processing carried out during WP5. Since several and complex diagnostic resources are involved in the management of heart failure, a careful analysis has been conducted (together with the medical partners) to single out the most significant and useful algorithms for the processing of diagnostic data, also in accordance with the suggestions we got after the first review of the project in 2007. The choices we made both for signal and image processing are motivated in Chapter 1 (Introduction), while a more extensive treatment is available in the specific chapters.

Chapter 2 is devoted to image processing. After a general recall of the selected imaging modalities in Section 2.1, we address the segmentation problem in section 2.2. Having the general theory of segmentation being reported in deliverable D15, we describe the algorithms developed for the segmentation of the left ventricle both in echocardiographic and cardiac magnetic resonance image sequences and for the segmentation and classification of chest X-ray images. This is respectively the content of Sections 2.2.1, 2.2.2 and 2.2.3, which are of a scientific nature.
Segmentation, that is roughly speaking the identification of the boundary of some object in an image (here the left ventricle or the lungs), is of course an intermediate results of image analysis. Actually, the boundary of the left ventricle, which can be traced by the developed algorithms, may be used to compute clinical parameters (such as dimensions of the left ventricle) by the process of *feature extraction*. This is the content of Section 2.3 in which methods for performing common linear, area and volume measurements

are described. Specific subsections are accorded to particularly important or complex measurements (such as ejection fraction and wall thickness). Section 2.4 ends the chapter with a discussion about the use, in research workflows, of shape modeling methods for the assessment of heart deformation pattern.

Chapter 3 is devoted to signal processing and is organized in a similar way to Chapter 2. In particular, after a recall of the non-imaging resources involved in the management of heart failure in Section 3.1, a scientific contribution is presented in Sections 3.2 and 3.3, which describe respectively the algorithms developed for QRS detection and the construction of the dominant heart beat from ECG recordings. Extraction of significant clinical parameters is treated instead in Section 3.4, where the role of the signal processing is mainly dedicated to the provision of reliable supporting tools for a more objective evaluation of the significant parameters. This is obtained through a graphical ECG viewer endowed with specific tools like zoom and caliper, so that each ECG once processed could be displayed, including the average dominant beat that was estimated by the signal processing chain.

In the technical Chapter 4, the issues related to the seamless integration of diagnostic data processing are treated. After an analysis of the echocardiography workflow (Section 4.1.1), of DICOM servers (Section 4.1.2) and of IHE profiles (Section 4.1.3), an Image Archive for HEARTFAID platform is introduced in Section 4.1.4. The current installation is based on a JAVA open source implementation of DICOM. Some features of its native interface are described together with DICOM and IHE compliant functionalities for web visualization. The integration of the image processing results into the platform is mediated through the use of the image archive and standard-compliant network services, as described in Section 4.2. In particular we adopted DICOM Secondary Capture images as a way to store the processed images, thus keeping the personal information of the patient as well as the reference of the original images. Interfaces for the deployment of image processing tools are described next (Section 4.3) together with an extendable web viewer interface for easy and quick image visualization and selection. A discussion about the selected standard for ECG storing and transmission ends the chapter (Section 4.4).

Finally, we discuss the high level integration of signal and image processing tools in the platform, by linking them to the Clinical Decision Support System (CDSS) in Chapter 5. After a general discussion about the value added by diagnostic data processing (Section 5.1), we describe representative examples taken from a comprehensive use case that has been ideated in collaboration with the clinical partners (Section 5.2). The actual integration of signal and image processing functionalities in the CDSS is reported in Section 5.3.

1. Introduction

Signal and imaging investigations are a basic step of the diagnostic, prognostic and follow-up processes of heart diseases. Not by chance, in the last decades, the development of *Computer-Aided Diagnosis* (CAD) schemes has attracted a lot of interest and effort within medical imaging and diagnostic radiology, becoming in some cases a practical clinical approach. The basic concept of CAD is to provide a "second opinion" or a "second reader" that can assist clinicians' by improving the accuracy and consistency in the clinical interpretation of the diagnostic images and signals. Actually the findings largely depend on the reader's subjective point of view, knowledge and experience. Hence, computer-aided methods, able to make this interpretation reproducible and consistent, are fundamental for reducing subjectivity while increasing the accuracy in diagnosis. As such, they are likely to become an essential component of applications designed to support physicians' decision making in their clinical routine workflow. Another important motivation relies on the possibility to perform more complete measurements on the vast amount of data generated by some devices (e.g. on image sequences obtained by echocardiography or magnetic resonance imaging) where the detection of potential disease is a burdensome task and may cause oversight errors. Further, signal and image processing may extend the usual set of extracted parameters, by adding novel representation features. Although such kind of features may be not suitable for the routine workflow, it may play an important role in research workflows.
Motivated by these issues, the signal and image processing activity in the HEARTFAID project was focused on the development of algorithms for the extraction of clinical parameters from the data produced by the available diagnostic devices.
Since heart failure is a very complex syndrome, whose management requires the use of several diagnostic devices, a careful analysis has been conducted to single out the most significant and useful algorithms for the processing of diagnostic data, also in accordance with the suggestions we got after the first review of the project in 2007. Actually, the development of a comprehensive integrated suite of signal and image processing tools, covering all the diagnostic resources, would require an effort far beyond the capability of the project. Nevertheless, a set of significant and representative tools has been selected after some interviews with the HEARTFAID clinical partners.
In particular, the signal processing applied in the HEARTFAID platform was driven by the type of raw data acquired in an interoperable format available in the HEARTFAID platform. At this purpose an accurate inventory of the medical devices available in the 3 clinical sites (JUMC, AUXOL/UNIMIB and UNICZ) (and that were going to be used in the HEARTFAID clinical validation) was performed in the context of WP2 and for each of them it was analyzed its capability of transferring the raw data in an interoperable format to a computerized system. In case of the signal processing the attention was devoted to all non-imaging medical devices in case of the routine workflows and of the research workflows.
We considered only the "time series" medical devices or, in other words, the medical devices that acquire one (or more) vital sign for a certain observation period T at a sampling rate of ΔT providing in output a time series of that (or those) vital sign. In fact, the other medical devices (non time series) provide only a single value

measurement that is not suitable for any signal processing but simple criteria for the rejection of measurement outliers.
With the view to "time series" medical devices with raw data in interoperable format, we identified in the routine workflow only the Archimed 4210 electrocardiograph that adopts a dialectal implementation of the SCP-ECG CEN standard for the storage of the 12-lead diagnostic (but non-interpretive) ECG.
Considering that the ECG processing is one of the most important examinations in the diagnosis and follow-up of heart failure patients, it was immediately judged important to design and implement some basic, robust and scalable algorithms for ECG processing that could be immediately applied to the raw data acquired by the Archimed 4210 cardiograph and at the same time used on the data acquired by other ECG devices (with different lead number and different acquisition period), if such ECG devices with raw data in interoperable format were made available in the clinical sites.
After some interviews with the clinical partners, it has been identified a significant operative scenario, where the ECG acquired with the non-interpretive electrocardiograph is transferred to the hospital gateway and from there processed in order to detect the QRS complexes, to identify the dominant beats and to evaluate the averaged dominant beat (for all the leads) that can be used by the cardiologists (with the help of a graphical ECG viewer), for the evaluation of all the measurements of interest for the diagnosis or the follow-up of heart failure patients. Consequently, the processing was based on QRS detection, QRS morphological classification (for the dominant beat) and dominant QRS averaging with special care to build-up an algorithm able to work with a scalable number of ECG leads, in order to be adaptable to ECG acquired by ECG devices with different "space dimensions".
The correctness of this strategic approach was further confirmed by the fact that, for research workflow, it was considered the use of wearable devices able to record single-lead ECG together with other parameters. Also for this kind of device, a simpler version of the QRS detector could be applied in order to estimate the heart rate time series.
For what regards instead the image processing applied in HEARTFAID platform, it was early realized that echocardiography is the most important imaging modality for the practical management of heart failure. In this case the interoperability problems were significantly less than for non-imaging diagnostic resources, by the large adoption of DICOM standard, and were solved in the context of WP2. However an echocardiographic study contains a very large collection of images, not all suited for post-processing. Since echocardiographic devices run proprietary software, the only possibility to perform image processing is in a post-processing workflow. This consideration somewhat limits the ability to perform comprehensive image analysis on echocardiographic data.
Nevertheless, the routine analysis of echocardiographic data involves the delineation of some anatomical structures (e.g. the heart chambers), which is needed for the computation of fundamental and basic parameters such as left ventricle ejection fraction.
Such task is accomplished in the HEARTFAID validation sites by manually drawing the contours of the anatomical structure with a mouse or a trackball, using tools provided in the proprietary software of the echocardiography device. This manual procedure has two drawbacks. First it is time consuming. Then, and more importantly

perhaps, the manual delineation of contours is prone to strong intra- and inter- observer variability. Such variability propagates to the derived measures (e.g. to ejection fraction), giving results which are poorly reproducible. Image processing tools may try to alleviate these problems, by providing automated or semi-automated methods for the delineation of the border of an anatomical structure. This image processing task is known as *segmentation* in the specific literature and it has a long history which starts from the very beginning of medical imaging. However, automatic image segmentation is still a *challenging problem*, especially when applied to complex, noisy and low contrast images. Echocardiographic images are precisely in this class of difficult images. Nevertheless, we developed some algorithms for their segmentation which may be exploited at different level of automatism, depending on the user's choice.

Apart from echocardiography, other discussed imaging modalities are chest X-ray and magnetic resonance imaging, for which some other segmentation algorithms have been developed. In particular, we report an algorithm for the segmentation of the left ventricle from cardiac magnetic resonance image sequences, by which 3D reconstruction of this anatomical structure are obtained. According to the generally accepted terminology, the algorithms discussed above solve *low level* image analysis tasks, in the sense that some further image processing should be applied to their output in order to obtain clinically meaningful parameters. Such kind of parameters extraction is also treated in this deliverable and tools for performing common simple measurements (such as linear measurements) as well as more sophisticated ones (such as Simpson estimation of volumes and wall thickness) are also introduced. In particular, methods for fully exploiting the 3D model of the left ventricle (obtained by segmentation of cardiac MRI) by extracting complex geometrical features are reported.

In parallel, the IT infrastructure for the seamless integration of signal and image processing has also been taken into account. In particular the need of a repository (in which to store in a structured way the original images and the processed ones) has leaded us to the study, choice and installation of a HEARTFAID image archive. In accordance to the general "interoperability" philosophy of the platform, the choice was driven by the level of conformance to the DICOM standard and by the implemented IHE integration profiles. In this way, the image archive constitutes a smoothly integrated and coordinated piece of the global HEARTFAID repository.

In addition, a web interface for quick visualization of the images is reported. More importantly for the scopes of WP5, tools for network communication between the image archive and the image processing toolkit were also analyzed and the prototypical algorithms were equipped with them.

If the IT infrastructure offers a "technical" integration of the modules for signal and image processing in the platform, the Clinical Decision Support System represents their integration at a higher level. Indeed, the results of the image and signal processing, besides being visualized and interpreted by the physician, have also the CDSS as end-consumer. Actually, the CDSS is able to exploit the results of processing and to formulate suggestions about diagnosis, prognosis and treatment that should then be submitted to the physician. The value added by signal and image processing to the CDSS is discussed together with examples drawn from a comprehensive use case prepared in collaboration with the clinical partners of the consortium.

2. Image processing

This chapter starts with a critical review of the imaging modalities relevant to the management of chronic heart failure. Segmentation, a basic but important image processing task, is introduced next. The methods developed for achieving segmentation of echocardiographic and cardiac magnetic resonance images are reported in detail. Then, methods for the computation of several typologies of clinical parameters are described, while Section 2.4 concludes the chapter with a discussion about shape modelling methods for the assessment of heart deformation pattern.

2.1 Imaging Modalities in the HEARTFAID

Imaging techniques offer an 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 deliverable D5 and ESC (Swedberg et al. 2005), chest X-ray and echocardiography should be included in the HF initial diagnostic work-up. Further, echocardiography will be repeated regularly 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 groups of patients, since they have not been shown to be superior to echocardiography in the routine 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.
However, in this section, we will also present some image processing results relative to other imaging modalities besides echocardiography. The reason is twofold. First of all, tomographic techniques like cardiac magnetic resonance imaging and fast computed tomography are often the golden standard for the computation of clinical parameters (this is the case for example of left ventricle ejection fraction which is one of the fundamental parameter for the documentation of heart function at rest). Thus, processing of tomographic images is an essential tool for the validation of methods for features extraction from echocardiography. Then, tomographic techniques allows for the computation of a bunch of parameters regarding shape and motion. These parameters may be eventually used for the analysis and characterization of heart deformation pattern in research workflows.
Finally, it is worthwhile to recall that tomographic techniques are less sensitive than echocardiography from the actual way in which images are acquired. For example, during an echocardiography examination, the position of the transducer greatly influences the acquired images. The transducer is usually positioned by trial and error, using the echocardiography device screen for immediate feedback. This somewhat

limits the ability to perform image processing in a post-processing workflow on echocardiographic data (that is after the images are acquired and the patient has left the examination room), since a misplaced or misoriented transducer cannot be corrected anymore. Although dependencies form the acquisition protocol is always strong, this kind of problems is mitigated in tomographic techniques.

2.1.1 TransThoracic Echocardiography (TTE)

According to ESC guidelines (Swedberg et al. 2005), echocardiography is the preferred method for the documentation of cardiac dysfunction at rest; TTE is encouraged for the diagnosis of heart failure, being a non-invasive, rapid and widely available technique.
Actually, when compared with other promising techniques like MRI and ultra fast CT, echocardiography is essentially inexpensive and, in a single routine examination, it provides a bunch of information ranging from heart chambers dimension to semi-quantitative assessment of the valvular functions. Furthermore, being real time, it is suitable for deformation analysis and assessment of dynamic functional features of the organ. One of the most important measurements is the left ventricular Ejection Fraction (EF) that permits to discriminate between patients with preserved or impaired systolic function. In latest years, the assessment of LV diastolic function has been recognized as in important examination in the heart failure domain. In particular, in combination with conventional Doppler echo, Tissue Doppler Imaging (TDI) is currently employed in some HEARTFAID validation sites to get a precise characterization of the LV filling pattern.
Thus, echocardiography is recognized as a valuable imaging modality for the qualitative and quantitative assessment of heart health. Nevertheless, the quality of information content is often reduced by noise, speckle, attenuation in the far field, suboptimal or insufficient acoustic windows and the like. Even worse, due to small impedance difference between blood and endocardium, the endocardial border may be broken or ill-defined. For these reasons, accurate assessment of cardiac parameters depends tightly from the expertise of the sonographer.
TTE is not the only ultrasound examination carried out in the cardiovascular domain. Actually, there exist other related techniques, including contrast-enhanced ultrasound or transesophageal echocardiography (TEE). In the first case gas-filled microbubbles are injected into the venous system to improve tissue and blood delineation and improve flow-related measurements. In TEE examination the heart is imaged from a transducer introduced in the esophagus using an endoscope. It conveys clearer, higher contrast images, since lung tissue does not intervene between the heart and the esophagus. Further a shorter field of view is enough in TEE, thus permitting the use of higher frequency transducers. However, TEE is an invasive, unpleasant examination. Since its use is limited to those patients having inadequate echo windows, TEE is not addressed directly in this report, although some of the methods developed for TTE may be adapted to the analysis of TEE images and image sequences.

Available Echo techniques and Echo windows

We briefly recall the available echo techniques and the echo windows relevant for the imaging of chronic heart failure patients, referring to deliverable D15 for further information and related bibliography.

As it is well known, ultrasound may be used to obtain different kind of images, the most important being two dimensional echo, M-mode, conventional Doppler and TDI.

Two dimensional echo gives a snapshot of a planar cross section of the heart. These snapshot may be produced and recorded in quick succession, giving a 2D image sequences, showing the heart chambers, vessels and valves in real time during the deformation cycle of the organ.

M-mode refers instead to ultrasound data acquired along a single direction and acquired over time. An M-mode image consists in a 2D image in which the Y-direction is a spatial coordinate (corresponding to a virtual chord across some anatomical structure) and the X-direction corresponds to time. Thus the intensity value of the pixel (x,y) is related to the ultrasound response at time x in the point y of the chord. M-mode echocardiograms are used mainly for linear measurement, for example for quantification of chambers size (e.g. Teichholz estimation of left ventricle ejection fraction), valvular orifice and thickness of heart walls and for semi-quantitative assessment of regional wall motion. Notice that the choice of a suitable transducer orientation is crucial in most applications; actually the beam should be oriented perpendicularly to the structure to be analyzed. In modern echocardiographic devices, a complimentary small 2D image is acquired to provide real time guidance in placing the transducer (this is the so-called 2D targeted M-Mode method).

Doppler echo uses the Doppler shift principle to determine the velocity and direction of blood flow or tissues in TDI. On modern devices, the sonographer sets manually a small region of interest in a 2D echo image and then the device use an appropriate time delay to acquire data on the selected region. Although, Doppler imaging conveys a lot of information about the heart and, in particular, about valve regurgitations and pressure, in the following we will only consider 2D and M-mode echo methods, since they are more adequate for performing post-processing tasks.

Three dimensional echocardiography, a promising and still under development technique, has not been addressed either, since it is not available in HEARTFAID validation sites. See deliverable D15 for a discussion about the current state of the art of 3D echocardiography.

Besides the previous methods for acquiring ultrasound data, different positions for the transducer on the chest may be used to obtain different *views* of the heart. Each of these transducer positions, corresponding to suitable *acoustic windows,* should be manually selected by the sonographer with no guidance but direct feedback on the device screen.

The main windows are 1) the left parasternal window which is used to record long-axis views and short axis views of the heart, 2) the Apical window which gives the fundamental 2-chamber and 4-chambers views and 3) the subcostal view, which is useful for the assessment of inferior vena cava diameter and its collapsibility index.

Proposed Acquisition Protocol

A protocol, first proposed by JUMC, has been discussed and adopted for the acquisition of TTE data. The protocol is enough rich to allow for the application of the developed

methods for image processing. It includes also Doppler images that may reveal useful in the future and which are already considered by the HEARTFAID image archive (see Section 4.1). The protocol prescribes the acquisition of the following images and image sequences:

- Parasternal long-axis:
  - 2 D of LV
  - M-mode cursor perpendicular through left ventricle just below the level of the mitral leaflet tips
  - M-mode cursor perpendicular through aorta and left atrium : record M-mode 5 beats
  - Colour flow Doppler recordings for mitral and aortic regurgitation
- Parasternal short-axis:
  - 2D of LV at level of mitral valve (basal)
  - 2D of LV at level of papillary muscles (middle)
- Apical four-chamber view :
  - 2D of LV
  - Colour flow Doppler recordings of mitral valve
  - Colour flow Doppler recordings of aortic valve
  - Colour flow Doppler recordings in tricuspidal valve and CW Doppler recordings for detection of tricuspidal regurgitation
  - PW Doppler transmitral flow recordings during diastole
  - CW Doppler aortic valve recordings
- Apical two-chamber view
  - 2D of LV
  - Apical three-chamber view
  - 2D of LV during quiet respiration
- Subcostal view
  - 2D of inferior vena cava
  - 2D of inferior vena cava during inspiration

Notice that all 2D acquisitions are image sequences during at least 3 heart beats to be stored in DICOM format.

2.1.2 Chest X-ray

According to ESC guidelines, the chest radiography should be part of the initial diagnostic work-up of heart failure. Notice, however, that the predictive value of biplane chest X-ray is only obtained by the interpretation of images in the context of clinical findings of ECG anomalies.
Chest X-ray is useful for detecting cardiomegaly by the computation of linear parameters and the estimation of the cardio-thoracic ratio. Also signs of pleural effusion, pulmonary congestion, interstitial and alveolar edema and other diseases contributing or causing dyspnoea can be detected.

2.1.3 Cardiac Magnetic Resonance Imaging (MRI)

According to ESC guidelines, MRI is a versatile well-established imaging technique that has become the *de facto* gold standard in accuracy and reproducibility for the assessment of heart volume, mass, motion, thickening and deformation. It is less operator-dependent when compared to TTE and it is more suitable for post-processing workflows.
As already stated, however, MRI has not been shown to be superior to TTE in the practical management of HF patients. So, being expensive, it is not routinely performed; its use is limited to patients with inadequate acoustic windows that require specific parameters determination or in research workflows.
Nevertheless, it is likely that, in the following years, MRI will become a less rare resource. The insertion of these imaging modalities in the HEARTFAID platform would be a perspective value added to the platform, since medical imaging and pattern recognition methods could be successfully applied to MRI imaging sequences in order to convey a unified picture of the state of the heart, suitable for precise classification by the HEARTFAID CDSS.

2.1.4 Other imaging modalities

For a discussion of other imaging modalities we refer to deliverable D15.

2.2 Image Segmentation

Image segmentation is a basic though, in general, non trivial image processing task. Its aim is to separate the image domain into regions that are, in a broad sense, meaningful for a further image analysis or interpretation task. For example, image segmentation could be used to delineate the borders of some structure of interest. In turn, these borders may be used to compute geometric features of the structure (e.g. length and area to cite the simplest ones) or densitometric features (i.e. relative to the distribution of gray level values inside the delineated region).
Although the ultimate goal is to achieve automatic segmentation of images, performing automatic segmentation is very often extremely difficult because of low contrast, poor image quality and noise. For this reason, methods of semi-automated or assisted segmentation have their own interest if they come equipped with suitable tools for user interactions. Having the general theory of segmentation being described in deliverable D15, we move next to three specific tasks (namely segmentation of echocardiographic images, segmentation of MRI images and segmentation and classification of Chest X-ray images) detailing the developed algorithms.

2.2.1 Segmentation of Echocardiographic Images

Among various segmentation problems that arise in the analysis of echocardiographic images, the research activity within HEARTFAID has been focused on the problem of Left Ventricle (LV) segmentation. Indeed, as we will see, the delineation of LV is crucial for the estimation of fundamental functional parameters and, as such, it is routinely manually performed by sonographers.

Related works

A number of methods for segmentation of echocardiographic data have been recently introduced. They can be divided into two categories: methods that analyze each frame independently, and spatio-temporal methods which exploit the fact that echocardiographic data are actually movies, given by a succession of images (frames). Generally, it was found that commonly applied approaches to edge detection by means of gradients were not helpful. As an interesting example of single-frame methods, we can consider that proposed by Mignotte and Meunier (2001) for the analysis of short-axis parasternal images. They applied active contour algorithms and introduced suitable statistical hypotheses for determining the initial contour, for modeling the gray level distribution, and for minimizing the energy function. Another example is given by the work by Chen et al. (2003). They used prior profiles for region shape and for intensity and determined the region borders by means of an optimized balance between the actual image information and the prior estimations. Among the methods based on artificial neural networks, we can mention the two-layer back-propagation network applied by Binder et al. (1999) to parasternal short-axis images. An example of spatio-temporal approach is given by the method applied by Chalana et al. (1996) to short axis views. This method focused on the motion between consecutive frames, which was described as due to an attracting force, which was quantitatively expressed by an external energy term. A modification of the active appearance model was proposed by Bosch et al. (2002) in order to obtain the motion of the endocardium. A recent careful survey of methods for echocardiographic image segmentation has been made by Noble and Boukerroui (2006).

Proposed methodology

We devised a procedure that exploits a variational formulation of level set methods, i.e. methods in which the evolving contour is represented implicitly as the zero level set of a suitable 2D scalar function (Osher, 1988; Sethian, 1999). With the aim of solving the difficult initialization problem, we augmented the level set method by a computational stage that provides a suitable initialization contour: in fact, level set methods, as other active contour schemes based on energy minimization, may become trapped in undesirable local minima. The initial contour is obtained by means of the introduction of mimetic criteria, based on the idea of mimicking the processing steps applied by an expert observer for the identification of LV cavity (see also Barcaro et al. 2007a).
We applied this two-stage procedure to apical-view images of the heart (the so-called two- and four-chamber views), in order to identify the LV cavity in every frame of the processed image sequence.

Identification of the Region of Interest

The identification of the Region of Interest, i.e. the LV cavity, was carried out by applying a mimetic criterion. Figure 1 shows an outline of this stage. First of all the segmentation process was carried out applying *local thresholding*. Then a careful analysis of the connected components produced by thresholding was used for identifying the mitral valve leaflets. In turn, knowing the position of the leaflets, a closure on the mitral level plane is obtained, leading finally to the delineation of the LV cavity.

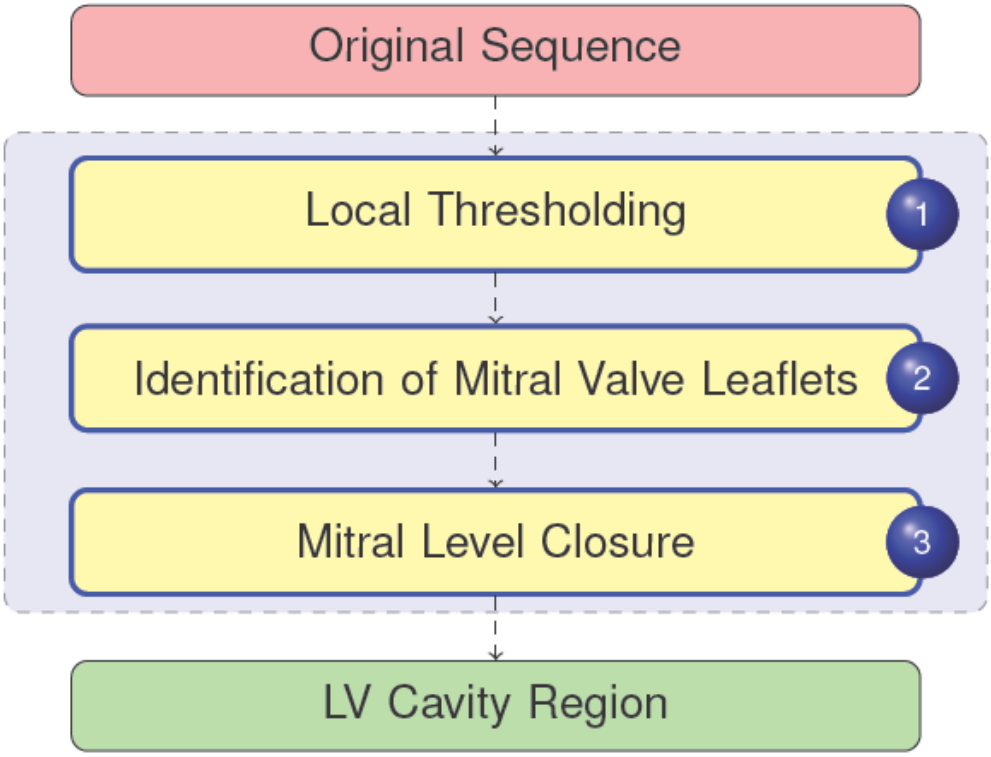

Figure 1. Outline of the mimetic criterion for the identification of the LV cavity

More in detail, local thresholding was chosen as segmentation process in this stage on the base of a preliminary comparison of the results provided by thresholding with the results provided by edge detection, region growing, and region splitting and merging. The last two methods provided results very close to those given by thresholding, but the computation was slightly complicated by the necessity of repeatedly applying morphologic operations. Edge detection provided interesting results, but there were too many interruptions in the contour. For this reasons, we decided to apply thresholding. Figure 2 shows a comparison between the results provided by edge detection and thresholding, respectively.

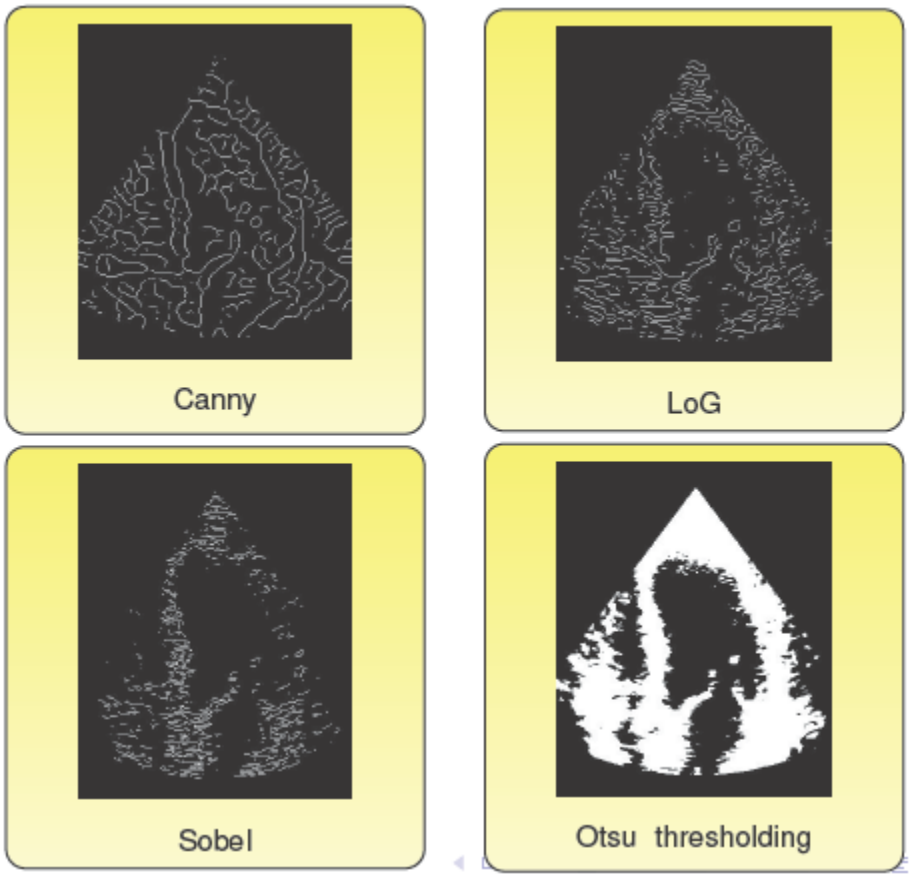

Figure 2. Edge detection vs. thresholding

For the choice of the threshold, the automatic and visual contours were compared by varying the percentages (from 10% to 200%, with 10% increments) of Otsu's threshold. The best agreement was for a percentage equal to 50%. The percentage applied was the same for all frames, while Otsu's threshold obviously changed.

Since the application of a global threshold provided poor results in the apical region of the image, the global Otsu's threshold was locally corrected in the apical region. Again in agreement with visual assessment, we found that a local correction to a percentage of 70% was appropriate. For a gradual correction, a local sinusoidal (vertically varying) function was applied to the threshold. The formula applied for threshold T in the sub-image including all columns and the rows i in the range $[i_0 - \Delta i_0; i_0 + \Delta i_0]$ centered on i_0 with width Δi_0 was:

$$T(i) = \{\alpha \cos(i - i_0 / \Delta i_0) + \beta\} T_0$$

where $\alpha = 0{,}1$, $\beta = 0{,}6$ and T_0 is Otsu's threshold. Figure 3 shows how the original frames were filtered in order to apply local thresholding.

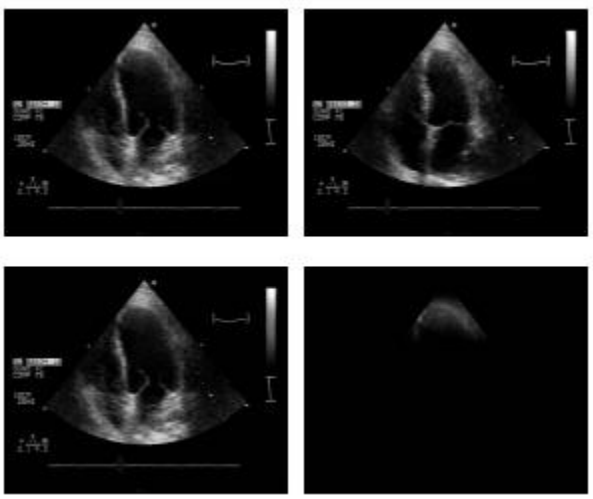

Figure 3. Top: End- diastole (left) and end-systole frames (right); bottom left: end diastole frame modified in order to apply local thresholding; bottom right: absolute difference between the original and modified frame

An effective method for the identification of the mitral valve leaflets was based on finding the minimum distance pixels between two separate regions, each containing one leaflet, in a suitable portion of the image.
The few lateral contour interruptions were eliminated in the following way: a set $\{P_n\}$ was built including the pixels placed at the extremes of each contour portion; then Matrix $D_{n,m}$ was defined as:

$$D_{n,m} = \text{distance}\ (P_n, P_m)$$

At last, each pixel P_n was connected by a segment to P_z such that:

$$D_{n,z} = \min_m \{ D_{n,m}\}$$

The ventricle contour was closed at the mitral valve level simply by considering a sub-image that contained the leaflets and then applying the threshold to this sub-image. The number of L separate connected components was always either 1 or 2. If $L = 1$, the contour was already closed. If $L = 2$, the following matrix was built:

$$D_{n,m} = \text{distance}\ (P_n, Q_m)$$

where P_n belonged to the first contour and Q_m to the second: a segment was then traced connecting P_w and Q_z such that

$$D_{w,z} = \min_{n,m} \{D_{n,m}\}$$

At last, the ventricle "base" was given by the segment whose extremes were the left and right inferior pixels of the region. This procedure was only applied to the initial frame,

because we found that for the successive frames the base could be easily traced by imposing that its extremes should be close to the previous extremes.
Figure 4 shows the segmented LV regions selected from eight consecutive frames, together with their respective major axes.

Figure 4. Sequence of segmented LV cavities

Level Set Segmentation
The mimetic stage provided a slight underestimation of LV cavity which suffered from spurious oscillations in the contour. These issues were solved using a level set method which achieved both contour regularization and better adherence to image edges. Among different formulations of level sets, we decided to work in the geodesic active contour framework (Caselleset al. 1997), exploiting a quite recent variational formulation without contour re-initialization (Li et al. 2005).
As usual, a 2D *edge indicator* function g was computed for each frame, according to the following definition:

$$g = \frac{1}{1 + \|\nabla(G_\sigma * I)\|^2}$$

where I indicates the gray-level matrix for the image and G_σ the Gaussian kernel with variance σ^2.
Notice that $g \approx 1$ in homogeneous regions of the image, since there the gradient is approximately zero. On an ideal edge, instead, $g = 0$ since the gradient of the intensity function diverges. On real edges, g assumes small but anyway finite positives values. An example of the computed edge indicator function is shown in Figure 5.

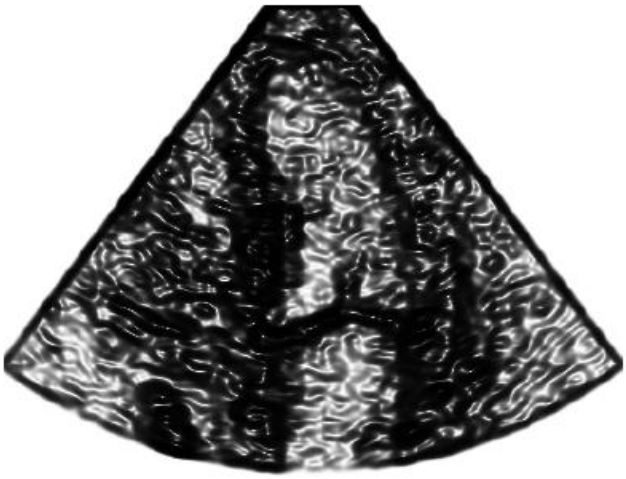

Figure 5. Plot of the edge indicator function. Black areas correspond to edges, while brighter areas correspond to homogeneous regions

The evolving contour is represented as the zero level set of a function Φ, i.e. the contour at time *t* is given by $\delta(\Phi)$, where δ is the Dirac function.
The initial value Φ_0 at time t=0 was set equal to the Signed Distance Function (SDF) from the LV cavity determined as in the previous section. More in detail, if Ω is the region corresponding to the LV cavity:

$$\Phi_0(x) = \begin{cases} \min_{y \in \partial\Omega}(|x-y|) & \text{if } x \in \Omega \\ -\min_{y \in \partial\Omega}(|x-y|) & \text{otherwise} \end{cases}$$

The initialization of the level set function to a SDF is suitable for different reasons, and especially for numerical computation. Indeed, if ε is a sufficiently small perturbation, letting $\Psi = \Phi + \varepsilon$, the contours $\delta(\Phi)$ and $\delta(\Psi)$ are close each other. Notice that this wouldn't be the case if we chose Φ to be the *distance*, instead of the *signed distance function*, because in that case also small perturbation may drastically change the contour, including its topological properties. The main reason for this nice behaviour is that a SDF crosses monotonically each of its level sets. Further a SDF is almost everywhere differentiable (actually $|\nabla\Phi| = 1$), reducing numerical inaccuracy that may show up during contour evolution.
A composite energy term was associated with Φ:

$$E(\Phi) = \alpha \int g \cdot \delta(\Phi) \, \|\nabla\Phi\| \, dxdy - \beta \int g \cdot H(\Phi) \, dxdy + \gamma \int (\|\nabla\Phi\| - 1)^2 dxdy$$

where $H(\Phi)$ denotes the region bounded by $\delta(\Phi)$ and the parameters α, β, γ are real positive numbers. The first two terms favor contour smoothness and penalize contours

that are distant from image edges. In the corresponding integrals both the contour length and region area are weighted by the edge indicator function g. To understand the behavior of the third term, recall that any signed distance function Φ satisfies $\|\nabla\Phi\| = 1$ almost everywhere. Conversely, it is not hard to show that any function Φ satisfying $\|\nabla\Phi\| = 1$ is a signed distance function up to a constant. Thus, the effect of the third term is to constrain Φ to be approximately a signed distance function.
In this way the *shocks* encountered usually in traditional level set evolution are avoided without any need of numerical remedies or subtle contour re-initialization procedures.
The minimization of the functional E was performed by gradient descent, using the initialization provided by the mimetic stage. We refer to (Barcaro et al. 2007b) for the details of the computation of the first variation of the energy term.

Results
Our method has been tested on 2D image sequences, recorded from the apical window (2-chamber and 4-chamber views). The echocardiographic device was GE Vivid 7, which is available at UNICZ and JUMC validation sites. The data consisted of image sequences acquired at the rate of 25 frames per second. Three full cardiac cycles have been imaged for each patient. The various procedures have been implemented in the Matlab environment, exploiting the Image Processing Toolbox (IPT). In pre-processing, standard corner detection algorithms have been employed for restricting the analysis to the region corresponding to the ultrasound wedge in the images. This restriction, besides reducing computation time, is essential for a meaningful determination of Otsu's threshold, otherwise biased by black pixels outside the ultrasound wedge. Standard IPT routines have been used for handling connected components during the mimetic stage. For the second stage, consisting in level set segmentation, the parameters α, β, γ have been selected heuristically, by visually comparing the results. Further, the number of iterations and the time-step in the discretization of the level set equation have been determined experimentally. In particular, for the choice of the time-step, a trade-off between speed of evolution and stability has been looked for. Indeed, a too small time-step may lead to very slow convergence of the contour to the desired one. Bigger values of the time-step may instead produce incorrect behaviour, since the evolving contour may go beyond edges belonging to the left ventricle. The domain of the level set function was set equal to the ultrasound wedge and it was computed on an equally-spaced rectangular grid, having as spacing the pixel size. Although it is usually necessary to compute with high precision the level set function only near its zero level set, we used the same precision in its computation throughout the entire domain. In the future, for computational advantages, it will be possible to use different resolutions and narrow-band algorithms. As already stated, a simple finite difference scheme was used. In this scheme, central approximations are adopted for spatial derivatives, instead of the more complicated up-wind approximation which are very often adopted in numerical discretization of the level set equation. Indeed, experimental results and qualitative analysis of the behaviour of the level set function during evolution guarantee that the adopted numerical scheme is suitable for our purposes. In the future, we hope to perform a more careful analysis (both experimental and theoretical) of the stability of this numerical scheme, and, in case, upgrade it. Figure 6 and Figure 7 show the evolution process of the level set in an end-diastole and end-systole frame. The

segmentation provided by the mimetic stage is used in both case for initialization of the signed distance function. Figure 8 and Figure 9 show instead the final result of segmentation procedure in an end-diastole and end-systole frame respectively.

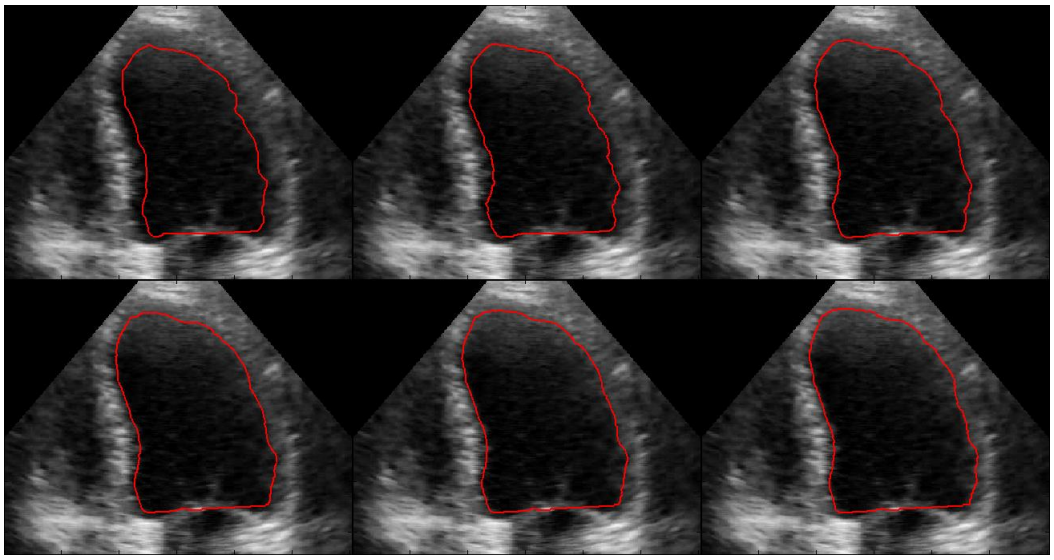

Figure 6. Evolution of the level set in an end-diastole frame. In the upper left corner the red lines corresponds to the segmentation provided by the mimetic stage, which is used for initialization of the evolution equation. Then from left to right, from top to bottom the figures show the level sets after 20, 40, 60, 80 and 100 iterations.

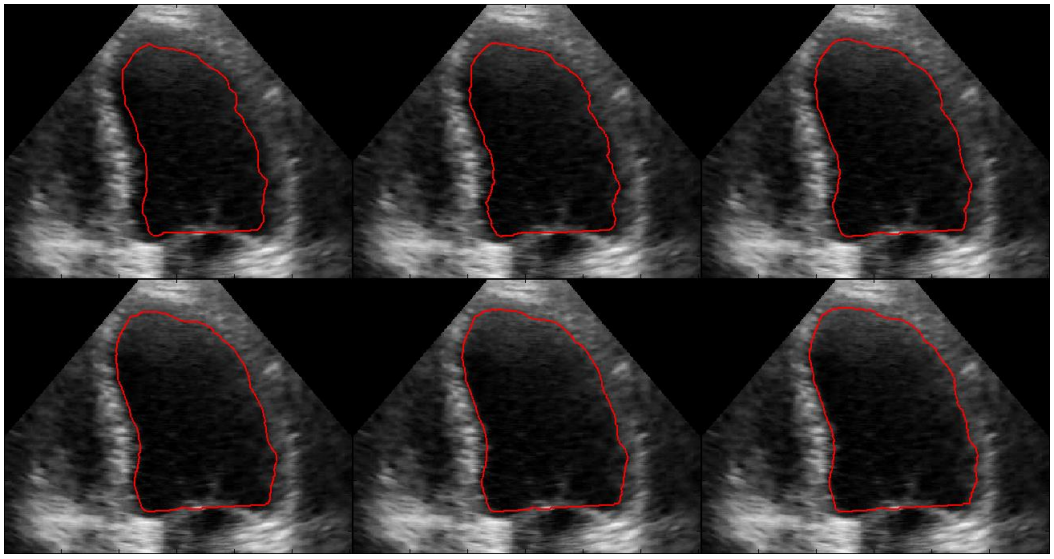

Figure 7. Evolution of the level set in an end-systole frame. In the upper left corner the red lines corresponds to the segmentation provided by the mimetic stage, which is used for initialization of the evolution equation. Then from left to right, from top to bottom the figures show the level sets after 20, 40, 60, 80 and 100 iterations.

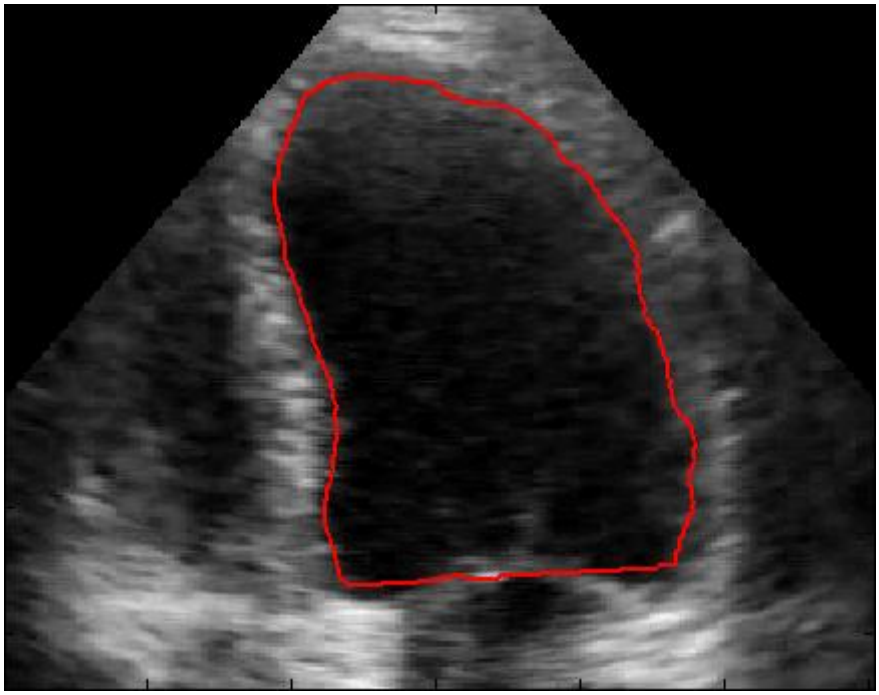

Figure 8. Final result of the level set segmentation in an end-diastole frame

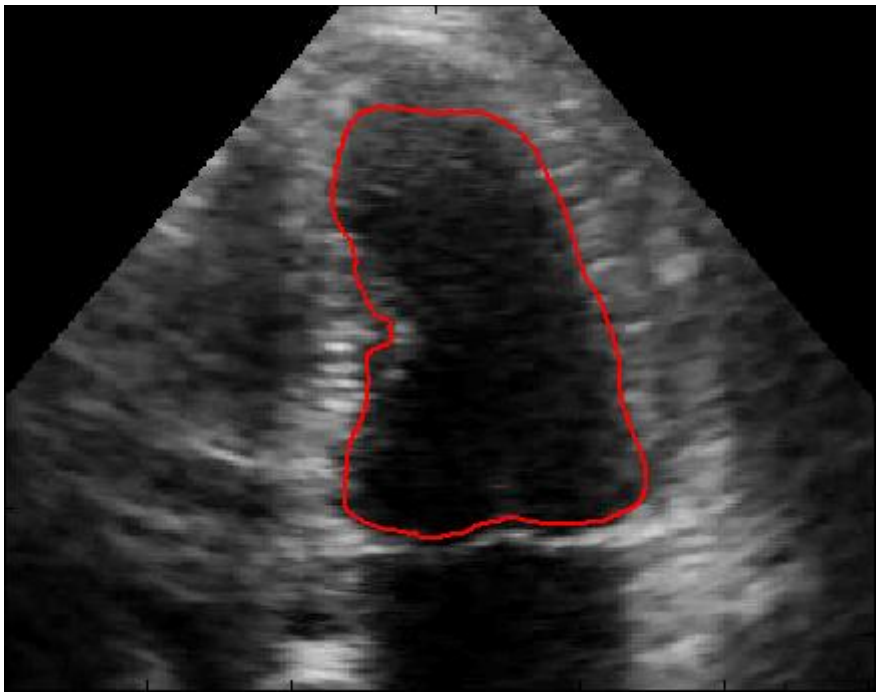

Figure 9. Result of the level set segmentation in an end-systole frame.

2.2.2 Segmentation of Cardiac Magnetic Resonance Images

In this section, we present current results regarding the automatic three-dimensional LV segmentation and reconstruction from MRI scans. The method proposed consists in a two-stage approach, based on fuzzy clustering and Artificial Neural Networks (ANN), for the identification of the LV surfaces.

Related works

As already stated, accurate segmentation of cardiac cavities is fundamental in assessing cardiac function and determining quantitative parameters. Owing to poor contrast and complex shape, this task is not trivial at all and there is still a need for fully automated tools, required to speed up diagnosis and eliminate intra-observer variability.

Magnetic Resonance Imaging (MRI) is a high quality and well-established imaging modality in analysing heart diseases and has proved to be more reliable than other techniques, both in supplying accurate morphological information and in assessing heart functions. However, due to noise or acquisition artefacts, visual information can be corrupted or ill defined. Moreover, in MR images, the presence of papillary muscles and trabeculations causes problems in detecting the LV contours. In such cases, only expert knowledge can help: the exact location of the contours cannot be based only on image evidence, but should be learned from examples provided by expert observers.

Usually, researchers have tried to design ad hoc algorithms able to incorporate *a priori* information about the LV shape. Model based surface detector have been widely used: for example, Declerck et al. (1997) employed a Canny-Deriche edge detector in a 3D polar map to segment endocardial and epicardial surfaces, while Faber et al. (1991) defined a hybrid spherical-cylindrical coordinate system.

Snakes, since their introduction in the seminal paper Kass et al. (1988), have been a powerful tool in cardiac images analysis for segmentation and motion tracking. Recent improvements in this field include works by Jolly et al. (2001), who reduced sensitivity to initial contour through Dijkstra algorithm, and by Paragios (2002) and Huang et al. (2004) who introduced deformable models influenced by forces derived from image region information. A biomechanical volumetric model of the heart, and not only of its boundary surfaces, is proposed in Sermesant et al. (2003) where segmentation of time series of images is obtained by means of a non-rigid deformation procedure.

Mitchell et al. (2002) used the concept of *active appearance model* (AAM). An AAM is a technique of analysis by synthesis, which, in principle, could describe any heart through a set of learned 3D shapes and a set of allowed variations.

Finally, also neural networks approaches have been proposed. In Stadilis et al. (1999) a *Generating-Shrinking Neural Classifier* is used to distinguish among lung, blood and myocardium points. This classification allows to extract a set of points on myocardial surfaces and, then, to assess parameters for a wavelets-based model. Two dedicated neural networks are presented in Coppini et al. (1995). The first is used to select from an edge map boundaries belonging to the ventricle. To get a meaningful and unbroken surface, these edges are further processed: a thin plate model for the left ventricle is introduced and a stable configuration of minimum potential is found by means of an analog neural-network implementation.

The methods reported above depend tightly from the choice of model parameters and initial conditions. The last problem has traditionally been solved with a manual intervention by an expert observer, but this contrasts the need of a fully automatic segmentation.

Proposed methodology

The LV segmentation in MRI can be viewed as a bi-modal problem: the structure of interest is the *myocardium*, which can be identified and extracted seeking the separation among the *endocardium*, the *epicardium* and the other cardiac structures. We address this problem with a two-stage method, which, firstly, identifies the left ventricle cavity (LVC) and then extracts the endocardial and epicardial surfaces.

The heart is imaged at the various phases of the cardiac cycle as a stack of slices, perpendicular to the left-ventricular long axis, by means of a short axis gradient echo MRI modality. More precisely, n 3D scans, S_0, S_1,..., S_{n-1}, are acquired for any cardiac cycle. Each scan S_t depicts the heart at its deformation phase t and consists of a set of k parallel 2D slices. The pixels of each slice are identified by their position (x,y) in the slice plane and by a third coordinate, z ($z = 1, \ldots, k$), which refers to the index of the slice itself in the stack. In the following, the three coordinates (x,y,z) will be referred to as a *voxel* and $S_t(x,y,z)$ will denote its intensity value.

To each scan S_t, a two-stage procedure is applied:

1. *LVC automatic localization*: a cluster analysis, based on the *fuzzy c-means* algorithm, is applied to identify and label homogeneous regions in each slice. Through a region tracking procedure, the behaviour of these regions is analysed over an entire cardiac cycle, in order to determine LVC.
2. *Extraction of the LV contours*: LVCs are used to compute the LV long axis, which, in its turn, is used to extract three-dimensional features processed by a dedicated ANN, in order to complete the segmentation, by identifying the endo- and epi-cardium contours.

A sketch of the method is shown in Figure 10.

Automatic localization of left ventricle cavity

In short axis gradient echo MRI sequences, LVCs are entirely depicted and appear as high intensity round-shaped regions, whose areas change periodically in time.

Clustering: Homogeneous image regions are first labelled using an unsupervised clustering method, based on the *fuzzy c-means* algorithm (FCM). This algorithm groups a set of data in a predefined number of regions so as to iteratively minimize a criterion function, namely the sum-of-squared-distance from region centroids, weighted by a cluster membership function. A membership grade $p \in [0,1]$ is associated to each element of the data set, describing its probability to be in a particular cluster.

The FCM algorithm is applied to each 3D scan S_t to produce a number of clusters: for any voxel $\mathbf{x}$, a features vector $(I_0(\mathbf{x}), I_1(\mathbf{x}), I_2(\mathbf{x}),\ldots, I_r(\mathbf{x}))$ is computed so that $I_0(\mathbf{x})=S_t(\mathbf{x})$, and for $d=1,\ldots,r$, $I_d(\mathbf{x})=S_t*\Gamma^d(\mathbf{x})$, where Γ^d is a Gaussian filter with $\Gamma=d$. This induces a partition of each slice plane $z=1,\ldots,k$ into a set $P_t(z)=\{R_t^1(z),R_t^2(z),\ldots\}$ of disjoint connected regions, where the upper indices 1,2,… are region labels. In the following, $\rho_t(z)$ will denote the generic region in $P_t(z)$.

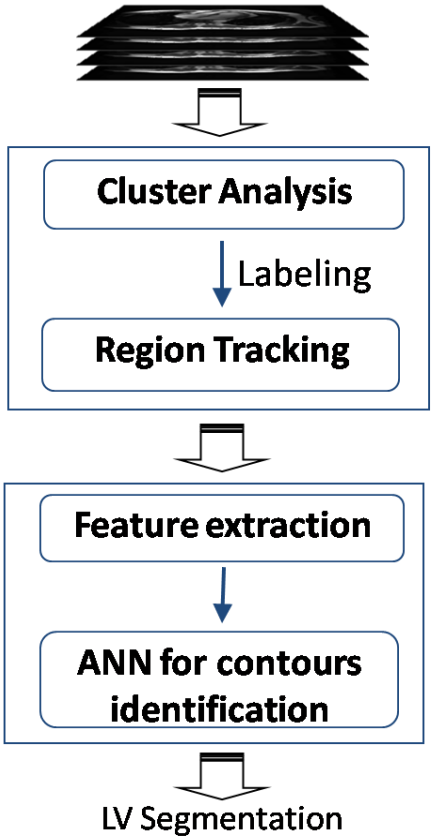

Figure 10. Sketch of the two-stage methods for LV segmentation

Region tracking: Once eliminated regions of negligible area (thresholding), an *intra-cycle* tracking procedure is performed. A simple centroid-based tracking algorithm associates, to any region $\rho_t(z) \in P_t(z)$ in the phase t, its correspondent region $T(\rho_t(z)) \in P_{t+1}(z)$ in the subsequent phase $t+1$. The procedure may be iterated, thus producing a region sequence

$$\rho_t(z) = T^0(\rho_t(z)) \rightarrow T^1(\rho_t(z)) \rightarrow T^2(\rho_t(z)) = T\,(T(\rho_t(z))) \rightarrow \ldots,$$

which can be thought as the evolution of the starting region $\rho_t(z)$ in the different phases. Considering the end diastole as reference phase ($t=0$), for each $\rho_0(z) \in P_0(z)$ the regions appearing in its evolution are collected in a list $Ev(\rho_0(z)) = (T^t(\rho_0(z)))_{0 \le t \le n-1}$ (See Figure 11).

Features extraction: For any region $\rho_0(z) \in P_0(z)$, the behavior in time of a geometric property G (such as area, eccentricity,…) can be estimated by evaluating G for every element in the list $Ev(\rho_0(z))$, thus obtaining a vector $G.Ev(\rho_0(z)) = (G.T^t(\rho_0(z)))_{0 \le t \le n-1}$.

To detect the oscillatory behavior of $G.Ev(\rho_0(z))$, it is effective and convenient to switch to frequency domain and consider its power spectrum density (PSD) function. In more detail, after eliminating the 0^{th}-harmonic (DC component), the DFT of the vector

$G.Ev(\rho_0(z))$ is normalized and the 1st-harmonic coefficient $\nu^G(\rho_0(z))$ in the periodogram is selected.
For a fixed set $\{G,H,\ldots\}$ of geometric properties, a features vector $\Im(\rho_0(z))=(G.\rho_0(z), \nu^G(\rho_0(z)), H.\rho_0(z), \nu^H(\rho_0(z)),\ldots)$ is associated to each $\rho_0(z)\in P_0(z)$.
Region classification: Let $LVC_t(z)$ denote the region corresponding to the LVC in the slice plane z at the phase t. At first, the reference phase is considered and $LVC_0(z)$ is searched among regions $\rho_0(z)\in P_0(z)$, taking into account their features vectors $\Im(\rho_0(z))$. More precisely, a set of learning examples is used to estimate the conditional probability density function $p(\Im|LVC_0(z))$ of the features vector $\Im$, given it pertains to $LVC_0(z)$. Then, for any new case, $LVC_0(z)$ is selected among candidate regions $\rho_0(z)\in P_0(z)$ according to the maximum likelihood estimator:

$$LVC_0(z) = \arg \max_{\rho_0(z)\in P_0(z)} \log p\left(\Im(\rho_0(z)) \mid LVC_0(z)\right)$$

In subsequent phases, the LVC is singled out using the tracking algorithm, namely $LVC_t(z)$ is defined as $T^t(LVC_0(z))$.
The results of the first stage are the volumes LVC_t obtained, for each scan S_t, as the stack of the extracted $LVC_t(z)$, for $z=1,\ldots,k$.

Region Tracking

for $z=1,2,\ldots,k$
for each $\rho_0(z)\in P_0(z)$ **do**
$T^0(\rho_0(z))= \rho_0(z)$
for $t=1,2,\ldots,n-1$
$c = centroid(T^{t-1}(\rho_0(z)))$
$T^t(\rho_0(z))=select\{\rho_t(z)\in P_t(z)$ s.t. $c\in \rho_t(z)\}$
end
return $Ev(\rho_0(z))=(T^t(\rho_0(z)))_{0\le t\le n-1}$
end
end

Figure 11. A sketch of the tracking algorithm

Automatic localization of left ventricle cavity
The boundary of the LVC identified at the previous stage supplies an approximation of the endocardial contour, which may suffer from poor intensity contrast or the presence of the papillary muscles. A refinement of this contour is then necessary, as well as the identification of the epicardial surface.

The set up is as follows. Let Ω, a subset of R^3, be the image domain of the scan S_t. First we define a features function $F_t : \Omega \rightarrow R^s$, that assigns to each point $\mathbf{x} \in \Omega$ a vector $F(\mathbf{x})$ of local features extracted from the image data S_t . Then we use an approach base on Multi-Level Artificial neural networks (MANN) to find functions Φ_α:

$$\Omega \times R^s \rightarrow R \ (\alpha = 1 \text{ or } 2)$$

s.t. the level sets:

$$V_\alpha = \{\mathbf{x} \in \Omega | \ \Phi_\alpha(\mathbf{x}, F(\mathbf{x})) = 0\} \quad \alpha = 1 \text{ or } 2$$

correspond to the endocardial and epicardial surfaces respectively. The functions Φ_α are learned using a training set of segmented images and they can be used subsequently to segment new instances.

To compute the feature vector F, the long axis l of the LV is computed for each scan S_t, as the principal inertia axis of the 3D LVC and used to define an intrinsic coordinate system suitable for characterizing the LV geometry. Since the LV is approximately *bullet-shaped*, a hybrid cylindrical/spherical reference system is chosen along the axis l. Given a point $O \in l$, let H be the plane through O orthogonal to l and H^+, H^- be the two half-spaces in which 3-dimensional space is divided by H. Points $P \in H^+$ are then described by cylindrical coordinates (r,φ,z), whereas spherical coordinates (r,φ,θ) w.r.t. O are associated to points $P \in H^-$

Selection of the point O is based on the request that the hybrid reference system should fit cardiac geometry or, more precisely, that the field $\hat{r}$, consisting of unit vectors pointing in the direction of increasing radial coordinate r, should be, as far as possible, orthogonal to endocardial surface. This property can be estimated by means of a cost function $D(O, l, A.LVC(z))$, where $A.LVC(z)$ denotes the area of the region $LVC(z)$; in this way the point O can be chosen as a global optimum for D.

The hybrid reference system is used to associate to each *voxel* $\mathbf{x} \in S_t$ a feature vector F consisting of the following features:

- *Position* with respect to the reference system, expressed as a quadruple (r,φ,θ,z). For $\mathbf{x} \in H^-$, the entries r, φ, θ represent its spherical coordinates, whereas z is set to 0. Similarly, for $\mathbf{x} \in H^+$, the entries r, φ, z represent its cylindrical coordinates whereas θ is set to $\pi/2$. With this choice both definitions agree for points in H;
- *Intensity* value $S_t(\mathbf{x})$;
- *Mean intensity* $S_t * F(\mathbf{x})$, computed applying an averaging filter F;
- *Gradient norm* $\|\nabla S_t(\mathbf{x})\|$ and its mean;
- *Radial derivative* $G_{r,t}(\mathbf{x})$, defined as the gradient component in the radial direction $\hat{r}$ and its mean;

The use of the radial information is motivated by the expected orientation of LV edges: for a LV boundary point, radial gradient component is likely to be a high fraction of total gradient magnitude.

The set of selected features are processed to accomplish the voxel classification by means of a Multilevel Artificial Neural Network (MANN), which assures several computational advantages (Di Bona et al., 2003).
For each voxel **x**, its computed features vector F_t (**x**) is divided into vectors F^i_t (**x**), each one containing features of the same typology and/or correlated.
Then each F^i_t (**x**) is processed by a dedicated classifier based on an unsupervised Self Organizing Maps (SOM) architecture. The set of parallel SOM modules constitutes the first level of the MANN which aims at clustering each portion of the feature vector into crisp classes, thus reducing the computational complexity.
Cluster indexes, in turn, are the input of the final decisional level, operated by a single EBP network. The output of this last module consists in a vector of membership grade of the voxel **x** to the two surfaces V_α (α=1,2). The SOM modules are trained according to Kohonen algorithm (Kohonen, 1997). For the EBP module, a set of 3D scans should be preclassified by an expert observer and used for supervised training, performed according to the Resilient Back-Propagation algorithm (Riedmiller and Braun, 1993).

Results
The two-stage method described above has been applied to short axis gradient echo MR images, acquired with the FIESTA, GENESIS SIGNA MRI device (GE medical system), 1.5 Tesla, TR = 4.9 ms, TE = 2.1ms, flip angle 45° and resolution ($r_x \times r_y \times r_z$) = (1.48 × 1.48 × 8 mm). Sets of n = 30 3D scans, consisting of k = 11 2D slices, were acquired at the rate of 30 ms for cardiac cycles [diastole-systole-diastole]. Various clinical cases were considered, for a total of 360 scans, corresponding to 12 cardiac cycles.
With the previous notation, fuzzy c-means was applied separately to each scan to produce two cluster using 2 as fuzziness parameter; we considered as a feature vector (I_0, I_1, . . . , I_r) where I_d is obtained from the original scan applying a Gaussian filter with standard deviation d times the inslice resolution r_x. Experimental testing showed that setting r = 2 is sufficient to get a good partition of the image domain. The result of the tracking procedure on a middle slice is shown in Figure 12.
The convex-hull area and the inertia moments were considered as geometric properties. The use of convex area (instead of the simpler area) reduces the effect of papillary muscles that sometimes move towards the boundary of the region corresponding to the LV. Processing was performed only on middle slices, thus eliminating the apical cap and the basal segments of the LV.
Analysis of various clinical cases has been used to introduce the Mahalanobis distance D; for simplicity, the covariance matrix S has been assumed to be diagonal.
For what regards the second stage, an example of computation of the radial derivative is presented in Figure 13. After having computed the feature vector described above, voxels are classified on its basis as belonging or not to the epi- and endocardial surfaces using the 2-level ANN. More in detail, the set of extracted features is divided into two vectors F_1, F_2 containing respectively 1) position and intensity and 2) position, gradient norm and radial derivative.

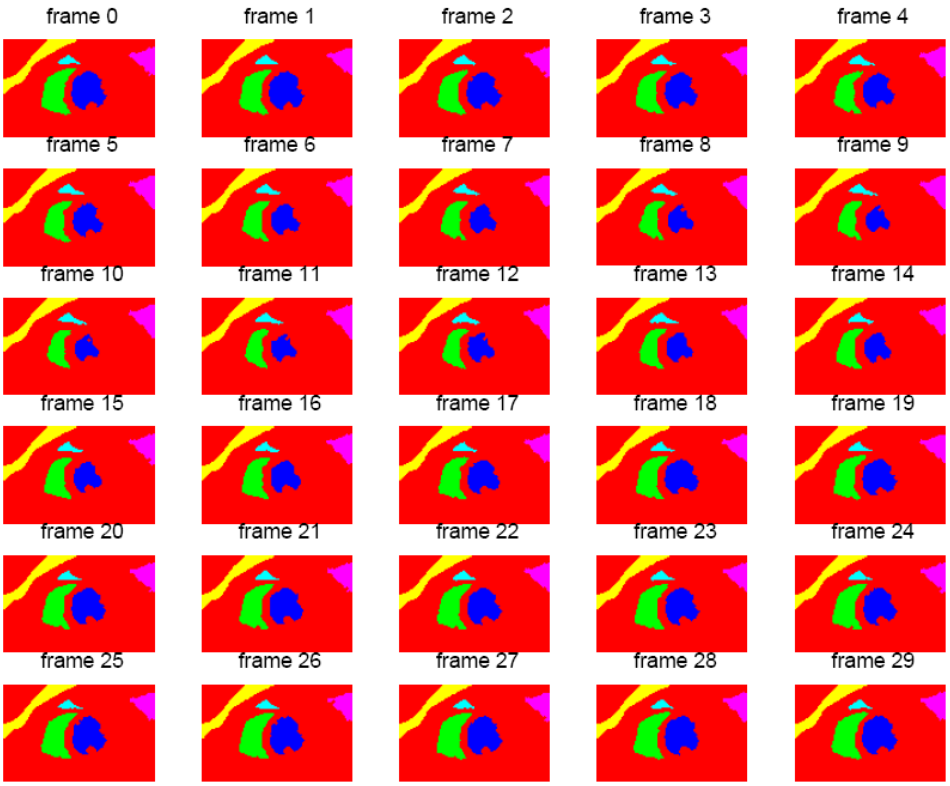

Figure 12. Visualization of the results of clusterization and tracking algorithm (slice number 6)

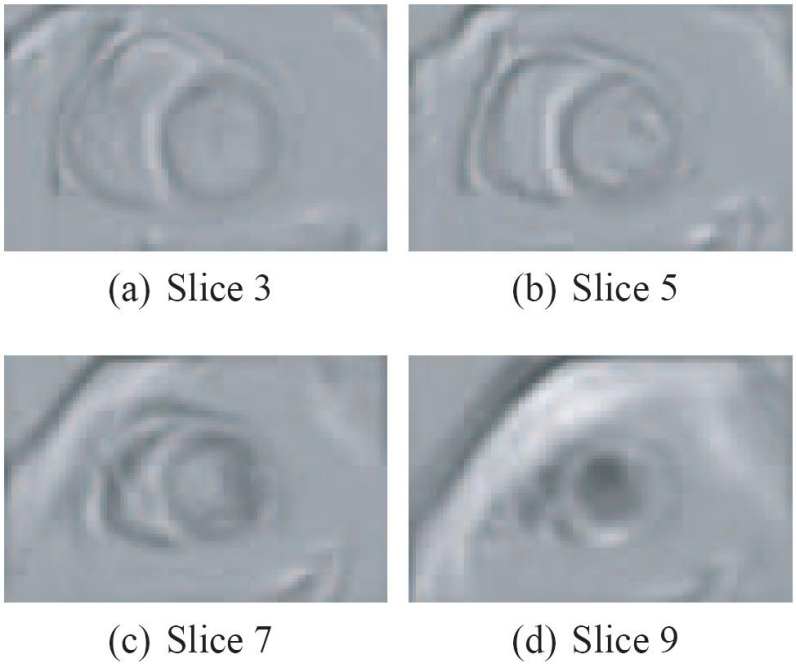

Figure 13. Example of computed features: radial derivative in different slices

The position w.r.t. the hybrid reference system is replicated in both vectors because it reveals salient for clustering both features subsets. Then, the first level of the MANN consists of two SOM modules, which have been defined as 2D lattice of neurons and dimensioned experimentally, controlling the asymptotic behavior of the number of excited neurons versus the non-excited ones, when increasing the number of total neurons (Di Bono et al., 2004).

An 8×8 lattice SOM was then trained for clustering the features vector F_1, while F_2 was processed by a 10×10 lattice SOM.
A single EBP module has been trained to combine the results of the first level and supply the final response of the MANN. The output layer of this final module consists in two nodes, which are used separately for reconstructing the epicardium and the endocardium. Since each cardiac surface divides the space into two connected regions (one of which is bounded), each output node can be trained using the signed distance function with respect to the relative cardiac surface.
In this way, points inside the surface are given negative values, whereas positive values are given to points in the outside. Henceforth the surface of interest corresponds to the zero-level set of the output function. Different architectures have been tested, finding the best performance for a network with only one hidden layer of 15 units. A manual segmentation was performed with expert assistance on the available data. A set of 240 scans was used for network training, while the remaining ones were used for network performance test. The voxel classification, supplied by the MANN, may be directly used for visualization purposes by using an isosurface extraction method, as shown in Figure 14. Figure 15 shows the intersection of the two cardiac surfaces with a slice plane.

Figure 14. Different views of the rendered left ventricle at end diastole. The surfaces are obtained applying marching cubes on the two output functions of the network. To eliminate satellites, a standard island removing procedure is applied

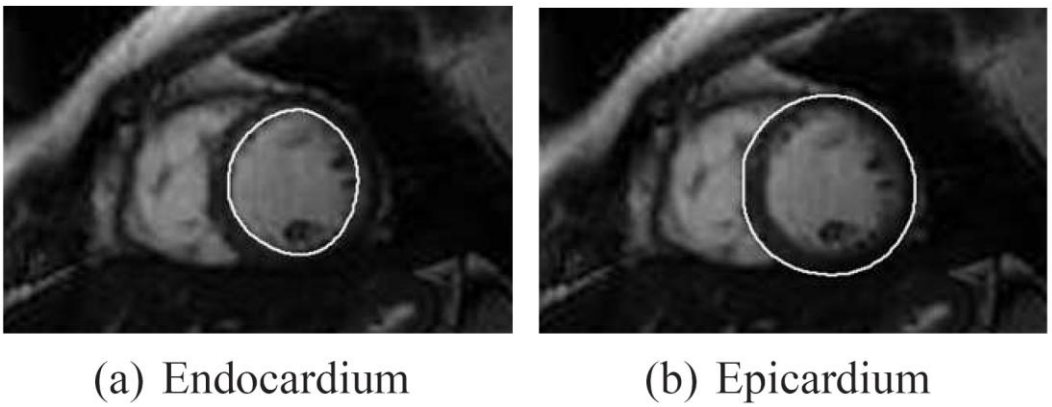

Figure 15. Intersection of cardiac surfaces with a slice plane.

2.2.3 Segmentation and Classification of Chest X-ray images

Chest X-ray is the most commonly used non-invasive modality for assessing a number of pathological conditions relevant to CHF diagnosis and management, such as cardiomegaly, pleural effusion, congestion, and interstitial and/or alveolar oedema. Within HEARTFAID, the work concentrated on supporting the identification of the presence, severity, or changes in pulmonary oedema by suitable image analysis methods.

Related works

Interpreting a chest radiograph is extremely challenging, since superimposed anatomical structures make the image complicated. Even experienced radiologists have troubles distinguishing infiltrates from the normal pattern of branching blood vessels in the lung fields (Quekel *et al.*, 1999). When radiologists rate the severity of abnormal findings, large interobserver and even intraobserver differences occur (Yerushalmy, 1969). The clinical importance of chest radiographs, combined with their complicated nature, explains the interest in developing computer algorithms to assist radiologists in reading chest images.

Recently, some works, aimed at fully automating chest X-ray examination, have been reported (Katsuragawa and Doi, 2007); however the general agreement is that the focus should be on making useful computer-generated information available to physicians for decision support rather than trying to make a computer act like a diagnostician. Three main areas can be distinguished in the literature on computer analysis of chest radiographs:

- general processing techniques, e.g., enhancement;
- algorithms for segmentation of anatomical structures, e.g, lung field or rib cage;
- analysis aimed at solving a particular task or application, usually an attempt to detect a specific kind of abnormality.

A review of them can be found in (van Ginneken *et al.*, 2001); however works that address the problem of oedema detection are difficult to find.

Proposed methodology

Pulmonary oedema refers to extravasation of fluid from the pulmonary vasculature into the interstitium and alveoli of the lung. The cardiogenic type is due to increased capillary hydrostatic pressure secondary to elevated pulmonary venous pressure. Cardiogenic pulmonary oedema (CPE) reflects the accumulation of fluid with low-protein content in the lung interstitium and alveoli, when pulmonary veins and left atrium venous return exceeds left ventricular output.

The progression of fluid accumulation in CPE can be identified as three distinct physiologic stages.

- In the first stage, elevated left atrial pressure causes distention and opening of small pulmonary vessels. At this stage, blood gas exchange does not deteriorate, or it may even be slightly improved. No fluid accumulation is recorded.
- In the second or *interstitial* stage, fluid and colloid shift into the lung interstitium from the pulmonary capillaries, but an initial increase in lymphatic outflow efficiently removes the fluid. The continuing filtration of liquid and solutes may overpower the pumping capacity of the lymphatics. In this case, the fluid initially collects in the relatively compliant interstitial compartment, which is generally the perivascular tissue of the large vessels, especially in the dependent zones. The accumulation of liquid in the interstitium may compromise the small airways, leading to mild hypoxemia. Hypoxemia at this stage is rarely of sufficient magnitude to stimulate tachypnea. Tachypnea at this stage is mainly the result of the stimulation of stretch receptors or J-type receptors.
- In the third or *alveolar* stage, as fluid filtration continues to increase and the filling of loose interstitial space occurs, fluid accumulates in the relatively noncompliant interstitial space. The interstitial space can contain up to 500 mL of fluid. With further accumulations, the fluid crosses the alveolar epithelium in to the alveoli, leading to alveolar flooding. At this stage, abnormalities in gas exchange are notable, vital capacity and other respiratory volumes are substantially reduced, and severe hypoxemia is usually associated with hypocapnia.

Detecting interstitial lung diseases on chest radiographs is one of the most difficult tasks for radiologists, because the contrast of lesion opacities is low and the pattern of opacities is very complex. Therefore, if the lung texture can be quantified objectively, it is expected that the accuracy and reproducibility in the detection of interstitial diseases could be improved.

The aim of the activities carried out within HEARTFAID was the possibility of aiding radiologists in their interpretation of chest radiograms by detecting and highlighting the lung area corresponding to an oedema, thus allowing the estimation of its extension, and track its changes in time.

A pattern recognition approach has been selected for defining a semi-automatic method that classifies a region selected by the radiographer on a Posterior-Anterior (PA) chest radiogram. The functioning of the method for lung region categorization can be summarized as consisting of four steps (Figure 16):

- *Initialization*: the operator selects a set of landmarks along the contour of the lung that should be analysed;

- *Lung extraction*: an interpolation method is applied for tracing the lung contour curves;
- *Region characterisation*: a vector of features is associated to each pixel of the region;
- *Region classification*: a MANN model is applied for classifying the features vectors and then label the different types of regions hence detecting eventual anomalies.

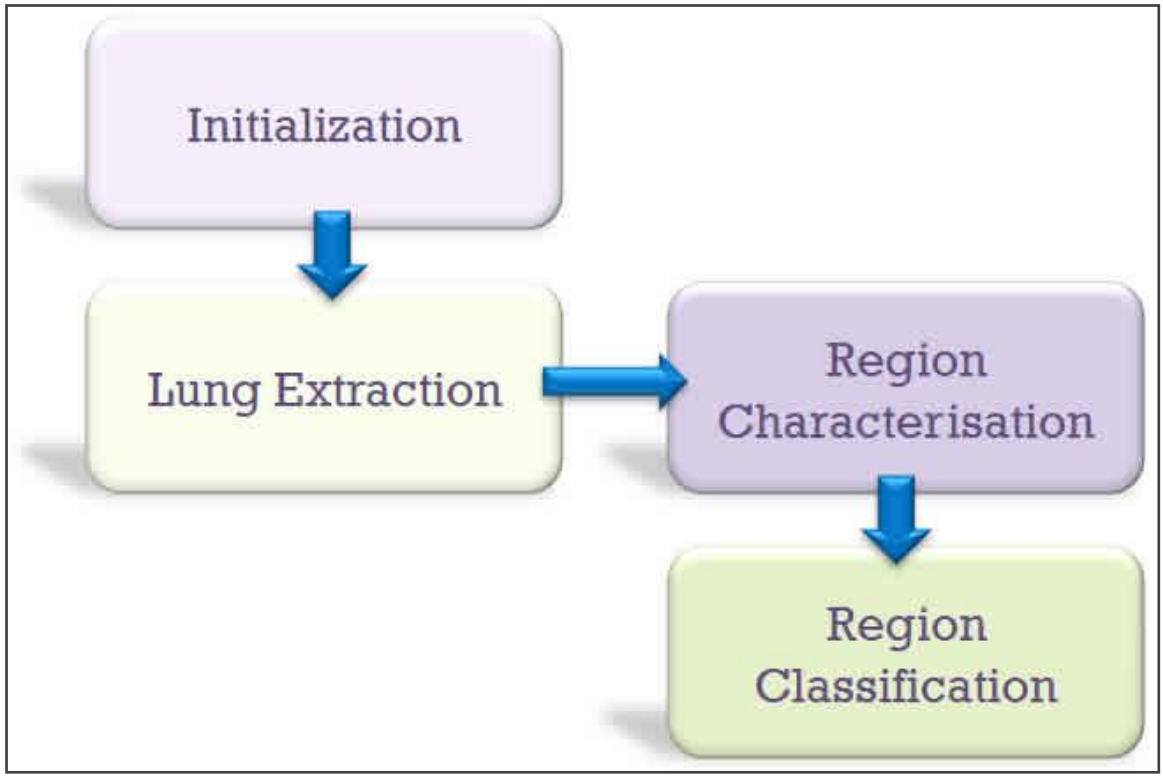

Figure 16. Four processing steps for the lung region categorization

For the first step, the image is interactively displayed for the selection of a series of points along the lung contours; this way, a vector of landmarks, *knots*, is obtained for each lung (Figure 17):

$$K_R = [k_{1R},\ldots,k_{nR}] \text{ and } K_{sx} = [k_{1L},\ldots,k_{mL}].$$

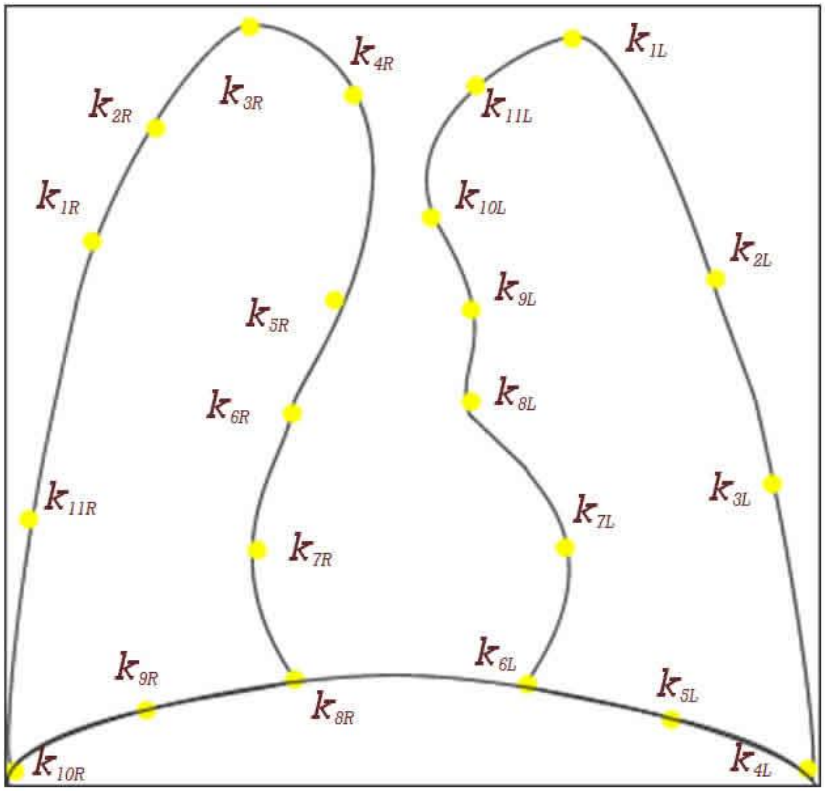

Figure 17. Sketch of landmarks selection along lung borders

Such vectors are used in the second step for extracting the lungs by tracing their contours. The segmentation step is easily performed by interpolation, i.e., by fitting two bicubic splines, C_R and C_L, to the two vectors of knots (see Figure 18).
Each pixel of the extracted regions are then characterised by a vector of features that take into account three types of information:

(i) intensity level based measures;
(ii) local difference measures;
(iii) local texture measures.

More in detail, the grey level of a pixel is primarily determined by the amount of X-ray attenuation of the tissues and two features are included in this group F_I:

- the grey level of the pixel,
- the grey level of the pixel expressed as a percentage of the histogram of the image (to account for the shape of the histogram).

Local difference measures are used to characterize the amount of homogeneity in a local neighbourhood F_D:

- total Sobel gradients (first-order difference) in a 7x7 neighbourhood;
- the Laplacian (second order difference) in a 7x7 neighbourhood;
- the standard deviation of pixels in a 7x7 neighbourhood.

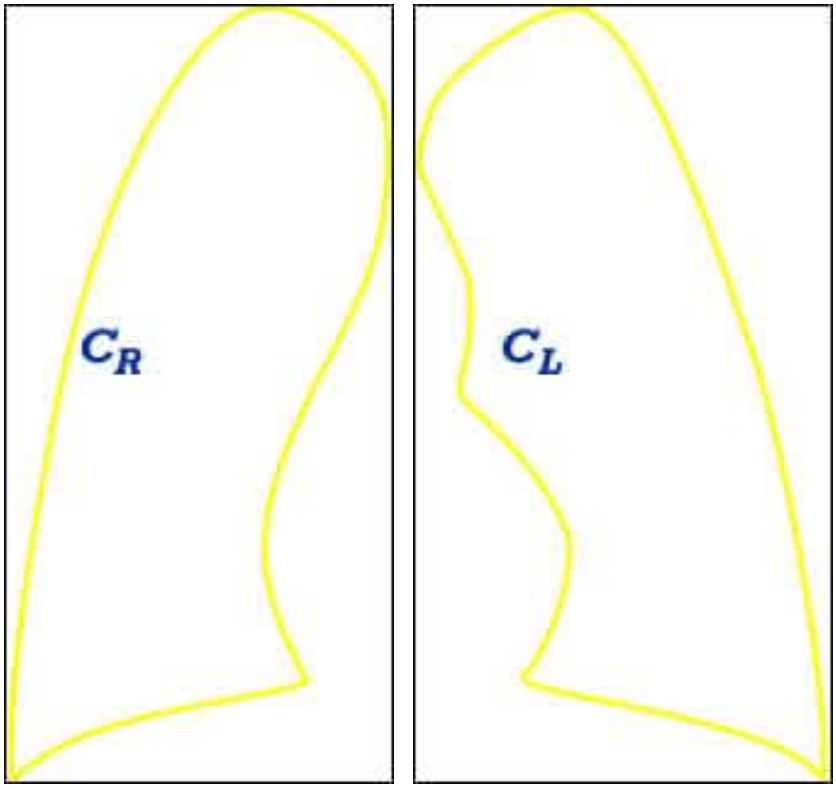

Figure 18 Lung contours: C_R and C_L are the curves drawn for segmenting the right lung (on the left) and the left lung (on the right), respectively.

Local texture measures are used to quantify the nature of local differences and often take into account both distance and direction of the differences. The features considered are F_T:

- microtexture features computed with four Laws' masks (Laws, 1980);
- spatial grey level dependence features, consisting in energy and entropy which measure of uniformity, homogeneity which measures the similarity between neighbours, and inertia for dissimilarity.

The extracted features are grouped according to their types and processed by a MANN model which consists of two levels: a first level of three SOM modules is introduced for classifying each of three groups of features, while an EBP module further elaborate the first level output and supplies the final classification of each region pixel. Six classes have been considered for training the network: (i) lung (including ribs); (ii) heart and upper mediastinum; (iii) subdiaphram; (iv) oedema; (v) other. Figure 19 shows the MANN architecture.

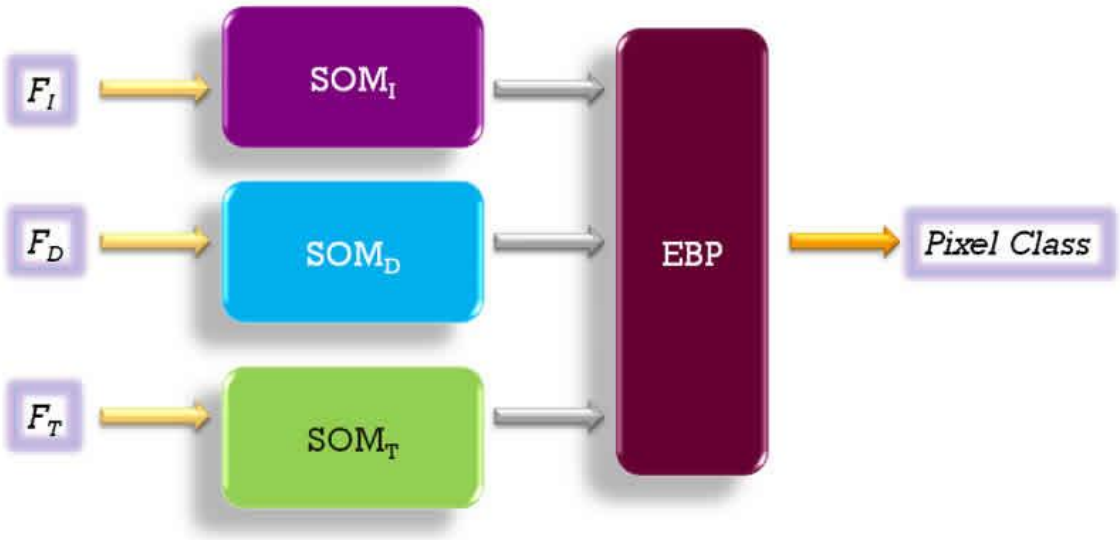

Figure 19. MANN architecture

Results

The method described was developed by using a dataset of 50 chest radiograms acquired with Thorax 2000 (IMIX, Finland), at a standard 2-m focus-to-detector distance with the patients upright, holding their breath at full inspiration. Each radiograph was 2000 × 2000 pixels (198 μm per pixel), with a dynamic range of 12 bits. Some examples of normal and oedematous cases are shown in Figure 20, while
Figure 21 depicts an example of landmarks selection and the corresponding lung contours drawn by interpolation.

The dataset has been divided in two subsets of 30 and 20 images used, respectively, for training and testing the MANN model. In particular, for each radiograph selected for training, some regions have been selected and pre-classified by an expert observer. The networks were then trained on a dataset of 2000 vectors and used for classifying the entire images of the test dataset.

For processing the input features vector, three SOM modules of 5×5, 10×10 and 12×12, were trained for intensity, difference and texture feature groups respectively. Outputs of such modules were used for a supervised training of an EBP module of two layers, the first consisting of 15 and the second of five neurons (one per each pixel class). Once classified, the regions are grouped and distinguished into two macro-categories: (a) *normal*, and (b) *suspicious*. An example of the final classification can be seen in Figure 22.

Once classified the lung region and detected the oedema portion, some measurements can be computed for aiding clinicians' analysis. Dimension, location and ratio with respect to the lung area are the main parameters that can be easily extracted.

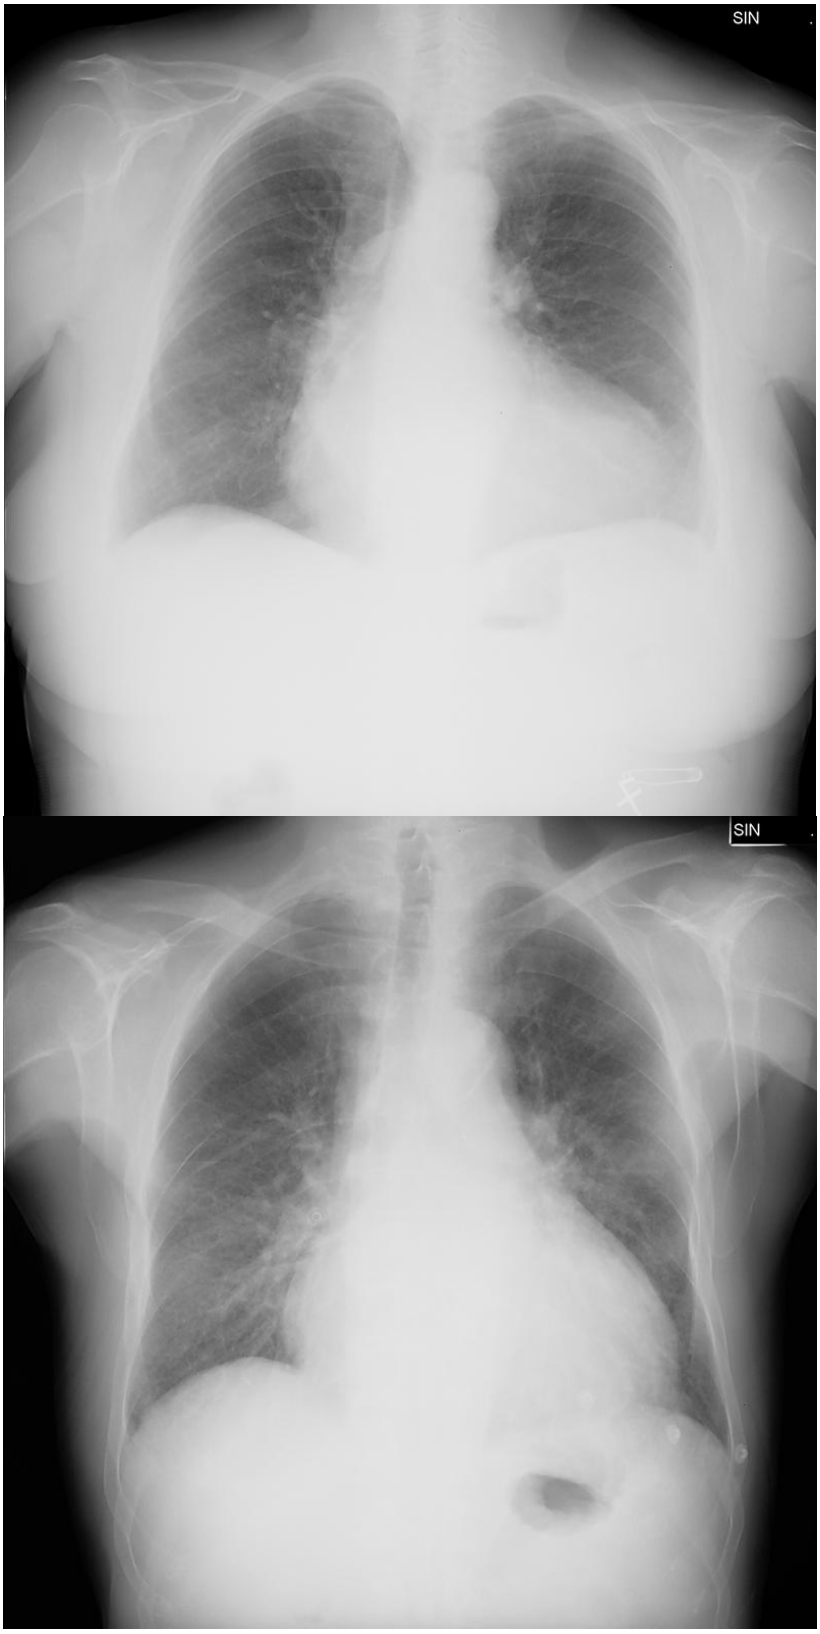

Figure 20. Example a normal (up) and oedema (low) cases of patients affected by cardiomegaly.

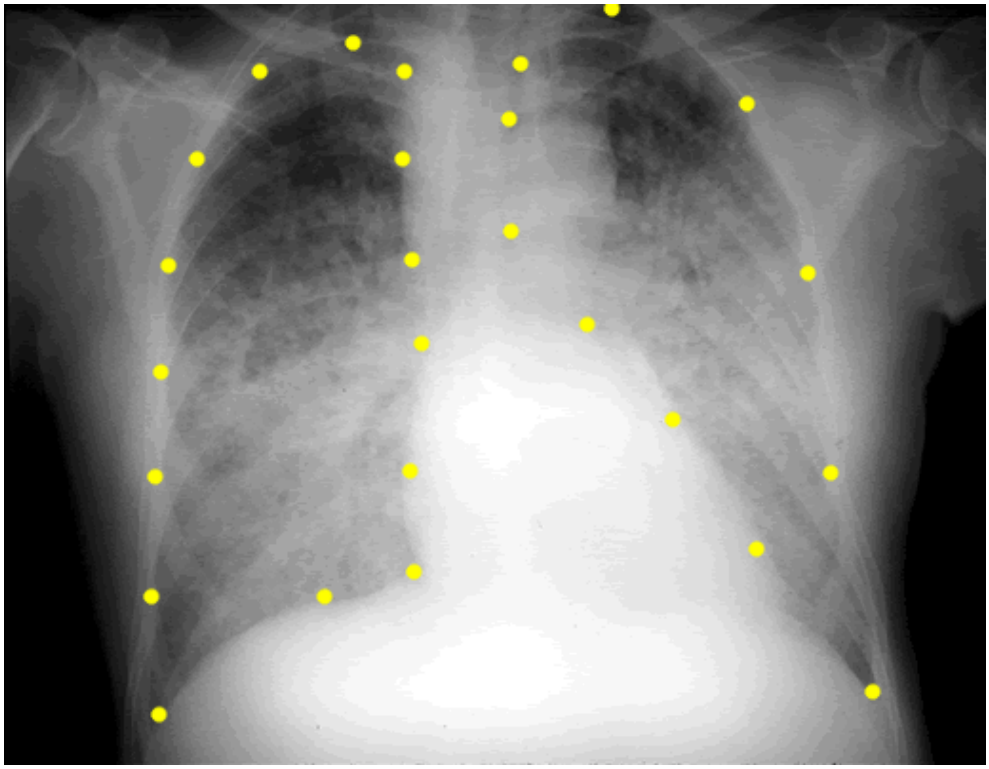

(a)

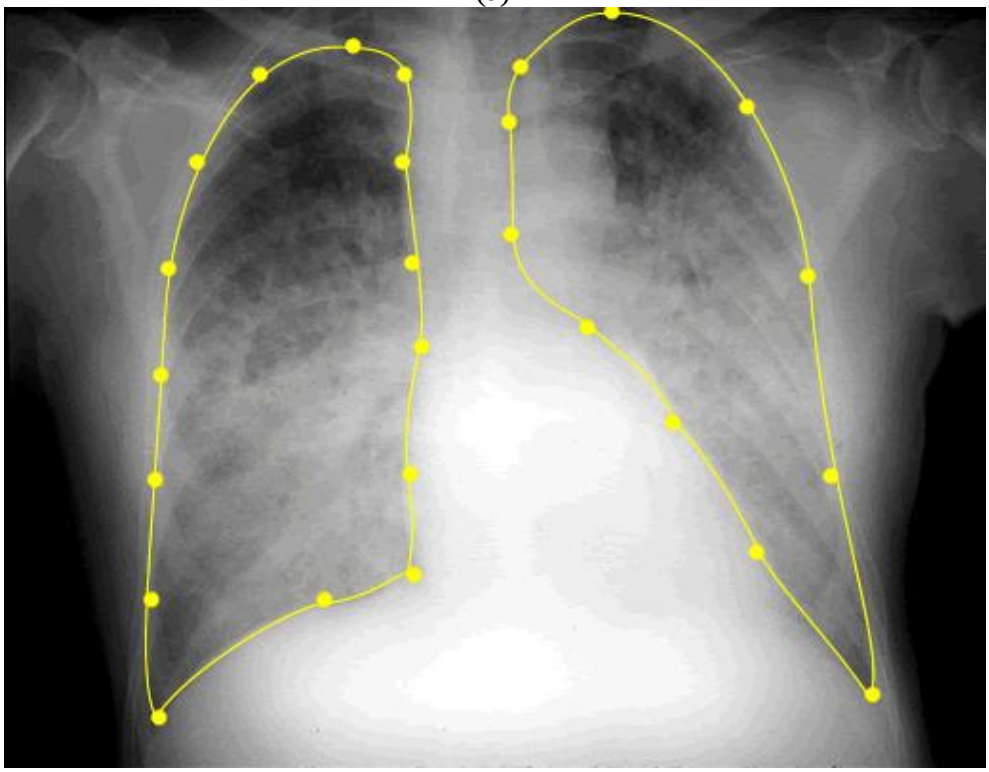

(b)

Figure 21. Results of the first two steps: (a) landmarks selection along the lung borders, and (b) knots interpolation and lung contours drawing.

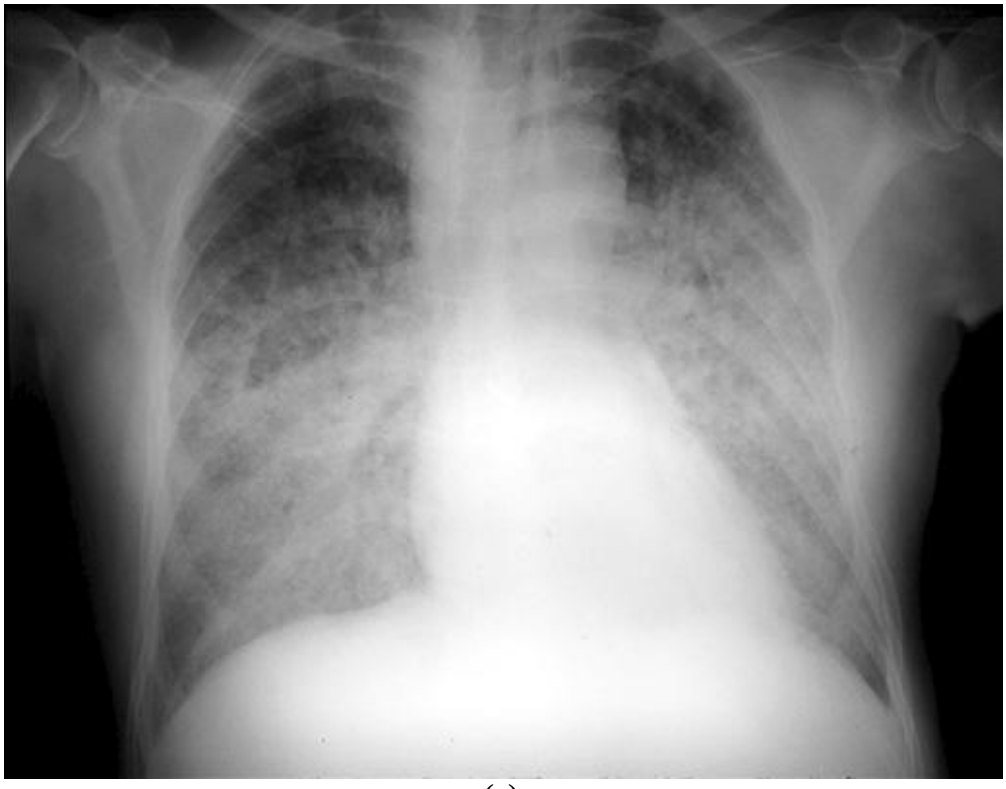

(a)

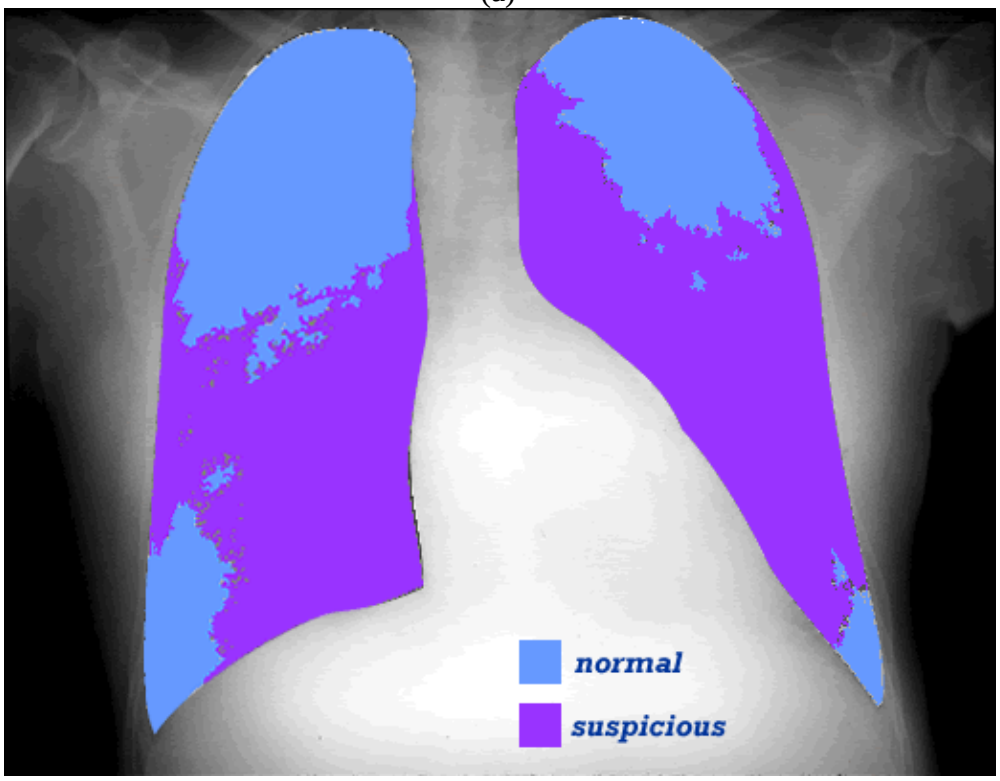

(b)

Figure 22. Examples of lung classification: (a) original image (b) lung final classification.

2.3 Computation of Image Derived Clinical Parameters

In this Section we focus on the extraction of clinically meaningful parameters from images. In particular we single out a set of parameters whose computation ends the image processing phase. Then, the computed value will be stored, visualized by the physician to perform a diagnosis or transferred to the CDSS for triggering decision support tasks (see Chapter 5). General linear, area and volume measurements are discussed along with examples taken from ultrasound, MRI and chest X-ray imaging modalities, while treatment in independent sections is accorded to particularly complex or important parameters, such as left ventricle ejection fraction, vena cava collapsibility index and wall thickness.

2.3.1 Linear Measurements

Linear measurements are by definition measurements operated on an image which involve the computation of the length of suitable segments. Thus, in our context, the value of such a measurement is a length expressed in the units of the image reference system. For example in a raw image such value is in pixel units, whereas if additional information is available (for example by the DICOM header of an image) the value in pixel units can be converted into physical units. Although the most common physical units are millimetres, some of the images relevant to HFP have different physical units on the image axes. In particular, ultrasound M-Mode images have values in mm in the vertical direction and in ms in the horizontal direction (see Section 2.1.1). In such cases, the only meaningful linear measurements are the one performed along segments parallel to one of the axes. Nevertheless, it is worthwhile to notice that other non strictly speaking linear measurements may be obtained from general segments traced on the image. For example, the slope of a segment in an M-mode image physically represents the *velocity* of the tissue.
After the measurement has been performed, some simple algebraic manipulation may be used to extrapolate (usually according to heuristic or approximate formulas) other parameters in different units. This is the case for example of Teichholz method for the estimation of left ventricle volumes; loosely speaking this method raises to the cube the result of a linear measurements of the left ventricle to obtain a value in mm^3, which is then multiplied by some corrective factor.
From an operational point of view the process of linear measurements can be described as follows:

1. Selection of the image on which the measurements should be performed
2. Selection of two anatomical landmarks on the image
3. Measurement of their distance in pixel units
4. Conversion into physical units (if DICOM header infos are available)
5. Association of the computed value to a tag describing the performed measurement

6. Extrapolation of related parameters by algebraic manipulation on the available tags.

Steps from 2 to 5 may be repeated to produce multiple tags, which can then be exploited in the last step.
The image processing toolkit needed to accomplish this process includes:

- General visualization interface, with pan & zoom and contrast adjustment functionalities
- A "ruler" tool. This tool allows landmark selection, real-time visualization of the produced segment and distance computation both in pixel and physical units. Further features required for this tool are:
    - Ruler constraints. Default ruler permits to draw segments with any orientation. For particular tasks (see the example below) an options should be activated to select horizontal or vertical segments only
    - Tag association. The user should have the ability to associate a tag (drawn from a predefined list or user-defined) to any of the performed measurements
- A small medical calculator, able to accomplish step 6 in the in the linear measurement process

Such kind of an interface for linear measurement is quite standard and open source software (e.g. ImageJ) provide some of these tools.
We thus prefer to describe some examples pertinent to image processing for the management of HF.

Computation of Left Ventricle diameter from M-mode images
In a TTE study, left ventricle diameter may be computed either by 2D targeted M-Mode or 2D echo. Figure 23 shows a typical 2D targeted M-mode image, where the cursor has been placed across the ventricle at the mid-papillary level (look at the chord in the small wedge in the top right of the figure, where a low quality 2D echo image is shown). The big rectangle in the figure represents instead the truly M-mode image. Each of the column of this rectangle corresponds to a temporal instant; in each column, the grey level values are the ultrasonic echoes of the tissues along the chord at the corresponding instant. In particular the darker area in the middle of a column corresponds to the left ventricle cavity. By using the ruler tool (with the "constrain to vertical" option activated), it is then easy to compute the maximum and minimum diameters of the left ventricle cavity. The maximum diameter is obtained when the heart is fully expanded, that is in end diastole, and is associated to the tag EDD (End Diastolic Diameter). Similarly the tag ESD (End Systolic Diameter) is associated to the minimum diameter. Sometimes end systole and end-diastole are identified by a simultaneous ECG, with a single-lead directly connected to the echocardiographic device. The typical diameter of the LV cavity at the mid-papillary level in end-systole is 3.1cm and in end-diastole 4.7cm.

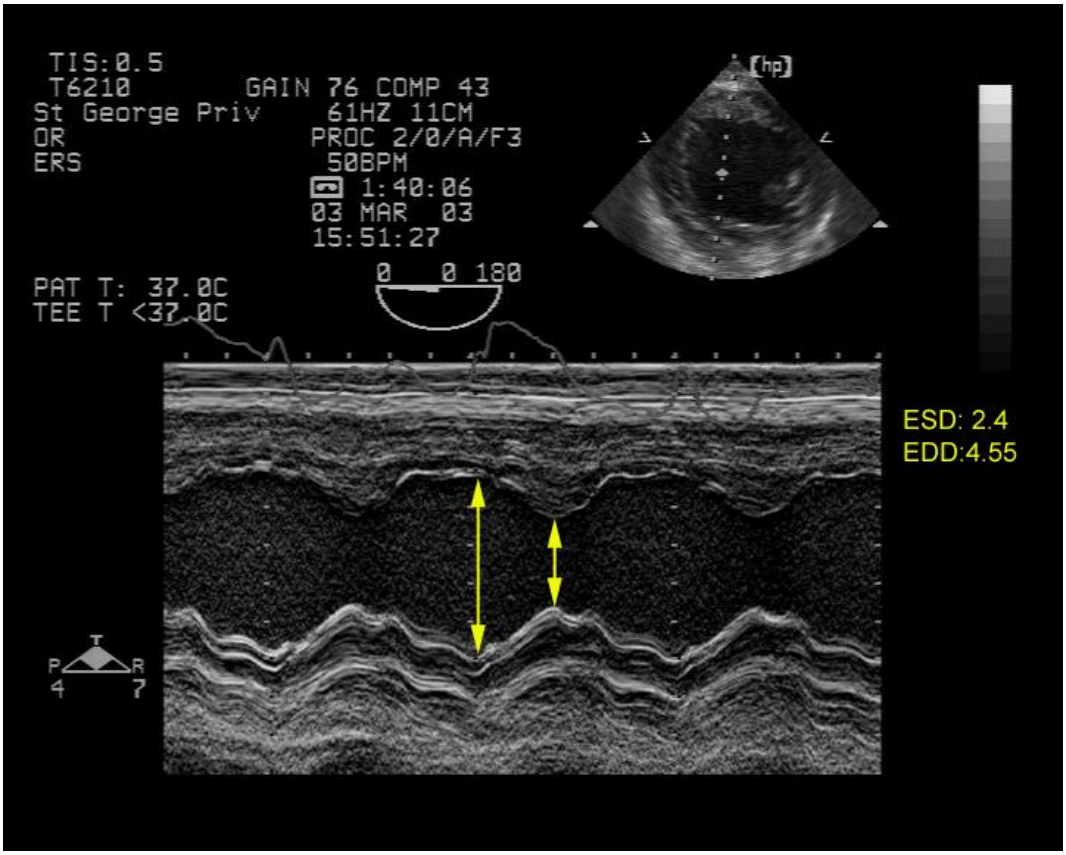

Figure 23. Estimation of end systolic and end diastolic LV diameters by 2D targeted M-Mode

Computation of Inferior Vena Cava Diameter from 2D echo

The inferior vena cava can be appreciated form the subcostal view; the computation of its diameter is important in the cardiovascular field because it correlates to pulmonary pressure. As in the example above, this linear measurement may be accomplished either in M-mode or in 2D echo. Figure 24 shows a 2D echo image depicting the inferior vena cava. For the computation of its diameter, the user selects a landmark by clicking with the mouse in the image region (denoted by a + sign in the figure). Moving then the mouse, the interface renders a segment (here in green) and shows its length in real-time. Clicking again the mouse selctes the second and last landmark and the measure is accomplished. The tag IVC (or IVCi/IVCe if the user wants to consider the respiratory state, that is inspiration/expiration) is associated to the measure.

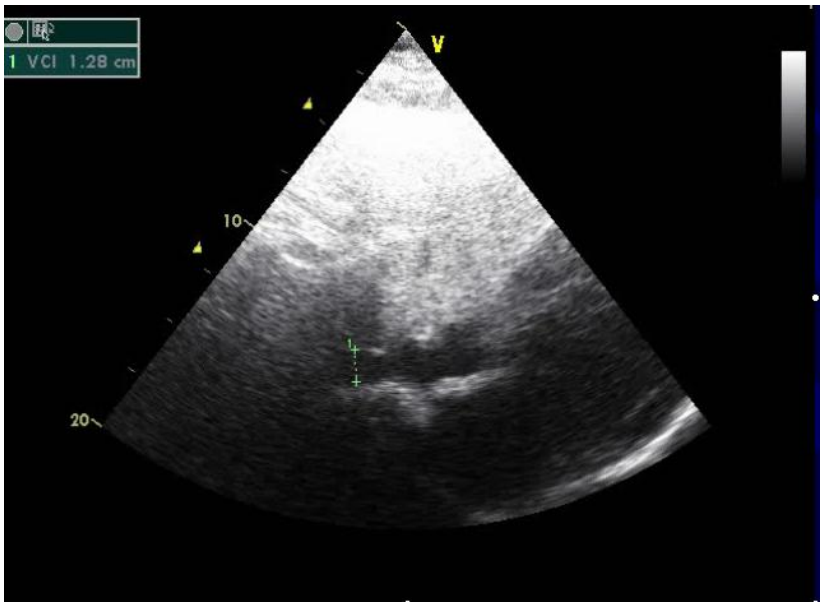

Figure 24. Estimation of IVC diameter during quite expiration by 2D echocardioghraphy

Cardio-Thoracic Ratio Estimation

Chest X-Ray is the preferred investigation in routine practice for estimating the cardio-thoracic ration (CTR) and, hence, detecting an increased cardiac size. Since its first description by Danzer in 1919, it has been a widely followed procedure to consider a CTR greater than 0.5 as representing cardiac enlargement, i.e. the presence of cardiomegaly which, even if frequently absent in acute HF, is considered an evidence of some heart disease.

In the standard practice, CTR is easily computed based on linear measurements on a PA chest X-ray image, i.e. by estimating heart diameter and the transverse chest diameter as sketched in Figure 25.

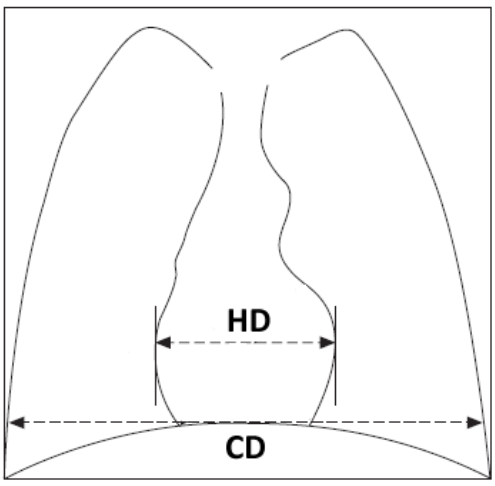

Figure 25. Sketch of the linear measurements, i.e., heart diameter (HD) and chest diameter (CD), used for estimating the cardio-thoracic ratio.

More precisely, CTR is obtained by measuring the distances from the midline to the most lateral points of the left and right heart borders (left apical and right atrial silhouettes, respectively) and dividing their sum by the maximum horizontal measurement of the thorax, from left pleural surface to right pleural surface (generally taken at the level of the diaphragmatic apices) on the PA chest radiograph.
For aiding and making more reliable CTR estimation, the interface supplies the users with orthogonal rulers for assessing the axial directions of the radiograph and then easily drawing straight horizontal or vertical segments when required. The procedure should consist in tracing a first segment along the spine midline, i.e., segment AB in Figure 26 drawn once selected landmark A and B; this is then used for tracing and measuring the distance from the midline to the most lateral point of the right atrium, i.e., segment dxAt in Figure 26, and in a similar fashion the distance from the midline to the most lateral point of the cardiac apex, i.e., segment CA in Figure 26. The largest horizontal width of the chest, from left to right pleural surface, is then drawn at the level of the left hemidiaphragmatic silhouette, i.e., segment CW in Figure 26. This way, CTR is computed as the ratio:

$$\mathbf{CTR} = \frac{\mathbf{dxAt} + \mathbf{CA}}{\mathbf{CW}}$$

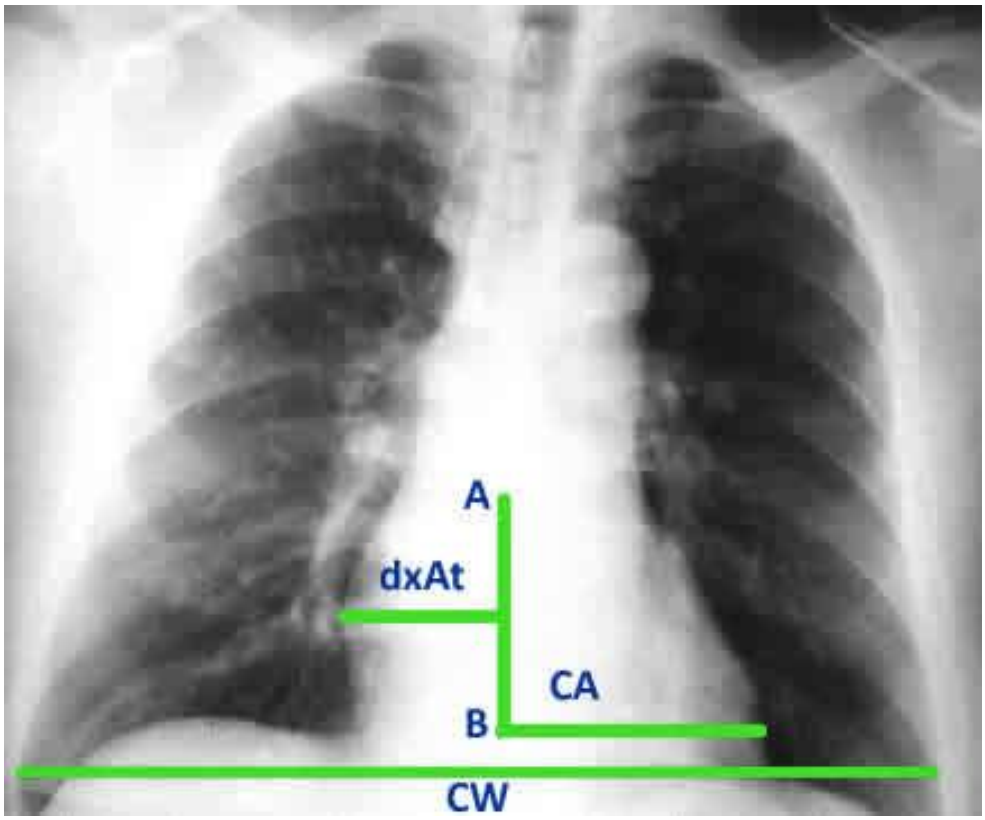

Figure 26. Linear measurements used for computing CTR: dxAt is the distance from the midline to the most lateral point of the right atrium; CA is the distance from the midline to the most lateral point of the cardiac apex, and CW is the horizontal width of the chest. A and B are landmarks used for drawing the spine midline used as reference.

2.3.2 Area Measurements

Area measurements are by definition measurements operated on an image which involve the drawing of the contour of suitable regions. Similarly to linear measurements, values in pixel^2 units are normally available for the area of the region and may be converted into physical units by exploiting the DICOM header. The most common physical unit is mm^2, although some area measurements with Doppler techniques may lead to different units (since roughly speaking they correspond to the integral of a flow). However, we restrict here to non-Doppler techniques.
From an operational point of view the process of area measurements involves only minor changes with respect to the one described for linear measurements:

1. Selection of the image on which the measurements should be performed
2. Selection of a region of the image
3. Measurement of the area in pixel units or measurement of other more general geometric features of the region
4. Conversion into physical units (if DICOM header infos are available)
5. Association of the computed value to a tag describing the performed measurement
6. Extrapolation of related parameters by algebraic manipulation on the available tags.

Steps from 2 to 5 may be repeated to produce multiple tags, which can then be exploited in the last step.
Thus, the *basic* image processing toolkit needed to accomplish the area measurement process needs, in addition to the tools already described for linear measurements, only a region selection tool with automatic computation of the areas. This basic functionality is implemented in several open source image processing interface (e.g. ImageJ).
However during WP5 a major effort has been devoted to the design of more advanced assisted region selection methods, which have been described in Section 2.2. Briefly, this work is motivated by the fact that manual selection of regions is time consuming and prone to inter- and intra-observer variability, especially in echocardiographic images. The algorithm developed (at the moment MATLAB prototypes) will be integrated in the standard interface for area measurements to offer additional services.
We discuss again some examples. The first one is basic and consists in the computation of the left ventricle area from a short-axis view. The second one involves instead the computation of more general features of a region and leads to the estimation of left ventricle volume by the fundamental Simpson rule.

First Example: Computation of Left Ventricle Area from a short-axis view
Figure 27 and Figure 28 show a short axis view of the left ventricle at end-systole and at end-diastole respectively. The image is obtained by 2D echo in a conventional TTE study. The goal is to compute the area of the left ventricle cavity which is clearly visible in both figures where it is enclosed by a white contour.
The user selects the “Select region” tool with the “free hand option”. Clicking on an image point sets the starting point of the contour. Moving the mouse draws the contour.

Returning to the starting point and clicking again closes the contour. The area inside the contour is automatically estimated in mm^2 if calibration is available or in $pixel^2$ units otherwise. After associating tags to the area measurements in the two images, fractional area change (see Section 2.3.4) can be computed.

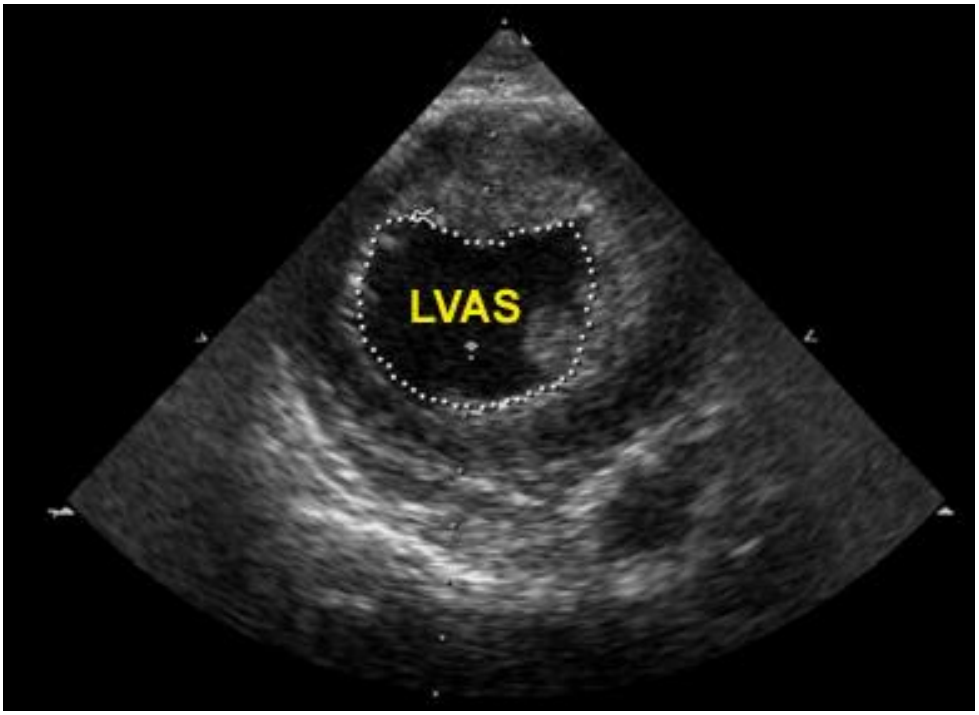

Figure 27. Estimation of left ventricle area at end systole. The user should be able to select in a manual or assisted way the internal boundary of left ventricle cavity in a mid-papillary level view.

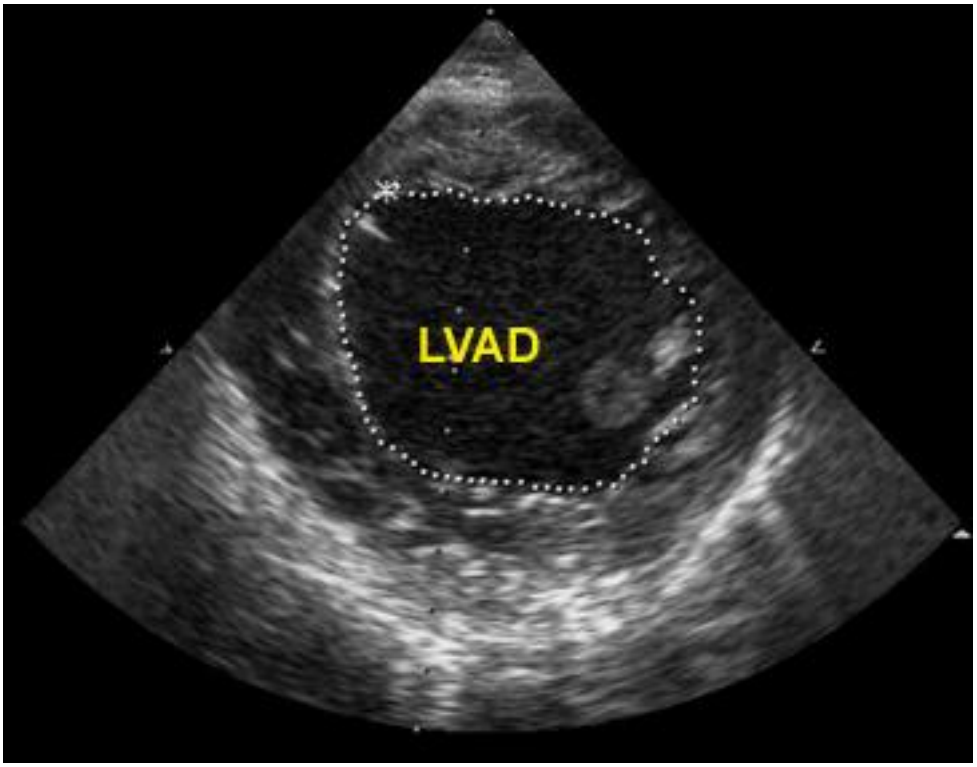

Figure 28. Same estimation of Figure 27 but at end diastole

Second Example: Computation of Left Ventricle Volume by Simpson rule

The computation of left ventricle volumes by Simpson rule is an important routinely-performed step in a conventional TTE study, since it is one of the preferred methods for the computation of left ventricle ejection fraction (see Section 2.3.4 below).

Having acquired image sequences taken from a 2-chamber, 4 chamber apical view or both, a first step towards the computation of the left ventricle volume is tracing the contour of the left ventricle at the end-systole and end-diastole frames of the available image sequences. Some examples of the traced contour are shown in Figure 29.

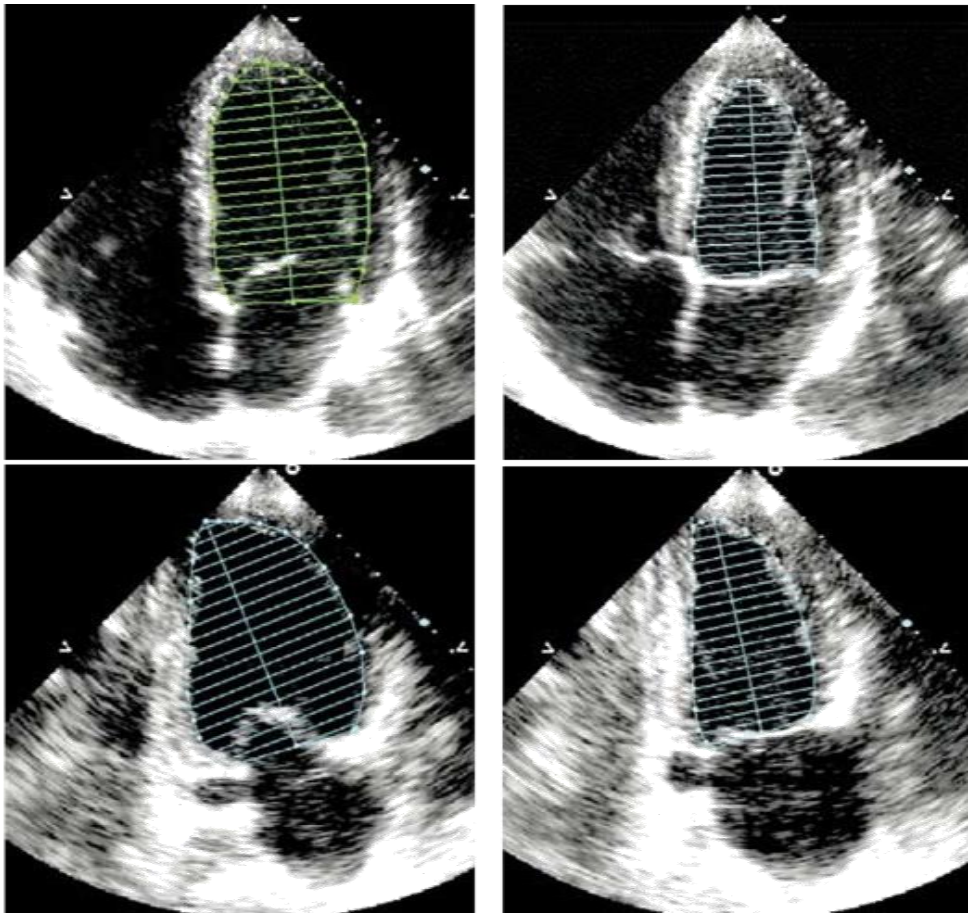

Figure 29. Segmentation of LV cavity. Top left: 4-chamber view at end-diastole. Top right: 4-chamber view at end-systole. Bottom left: 2-chamber view at end-diastole. Bottom right: 2-chamber view at end-systole

The contour tracing step may be accomplished manually (as in the example above) *or* in an assisted way exploiting the developed algorithms described in Section 2.2.1. More in detail, we may foresee three possible way to employ these algorithms, which we describe according to the automatism level, starting from the less automatic one.

Case A) Manual selection of the end-diastolic and end-systolic frames and rough manual contour tracing

In this case, the algorithm provides a refinement of the manually traced left ventricle contour in the manually selected frames. Instead of a "free hand" selection, the user may select just a polygonal region approximating the left ventricle cavity (as the one shown in Figure 30). The level set segmentation step described in Section 2.2.1 is then triggered. In particular, the manually drawn contour is used for the initialization of the contour evolution equation.

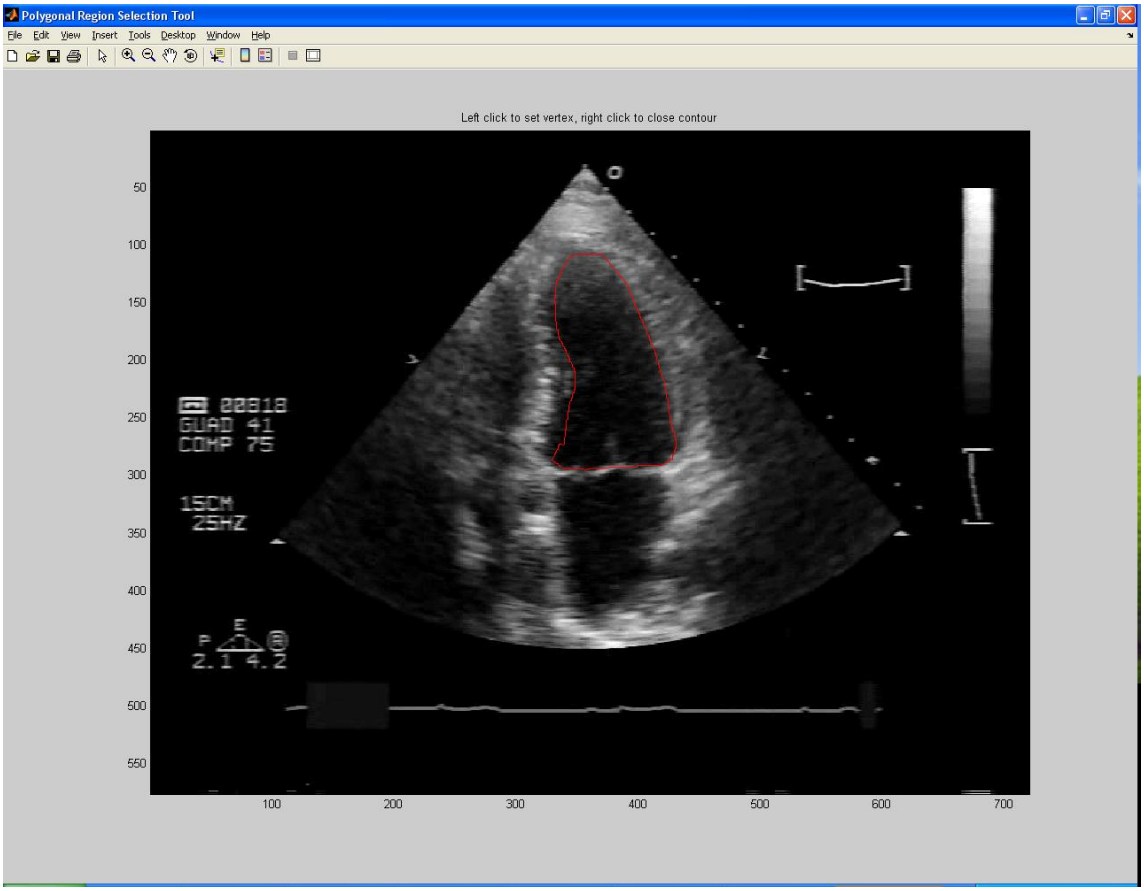

Figure 30. The Polygonal Region Selection Tool may be used to draw a first approximation of the left ventricle cavity, which is then automatically refined by the level set segmentation step. The prototype of the tool is implemented with the Image Processing Toolbox in Matlab

Case B) Manual selection of the end-diastolic and end-systolic frames and automatic contour tracing

In this case, the algorithm traces automatically the contour of the left ventricle in the manually selected frames. The mimetic stage is used to find an approximate left ventricle contour by local thresholding and manipulation on the connected components. Then the contour is refined by the level set segmentation step.

Case C) Automatic selection of the end-diastolic and end-systolic frames and automatic contour tracing

In this case the algorithm takes in input the whole image sequence and applies the mimetic stage to every frame in order to obtain a rough segmentation of the left ventricle. Then the volume of the cavity is computed on this rough segmentation (see below how this computation is accomplished), obtaining as a result the plot in Figure 31, from which we can appreciate the behaviour of the volume over time. The indices of the frames corresponding to the extremal values (i.e. maximum and minimum) of the volume are found and stored. Then, the level set method performs refinement of the contours in the frames which are near to those of extremal values. Computing again

volumes on the basis of the refined contours leads to the identification of the end-systole and end-diastole frames.

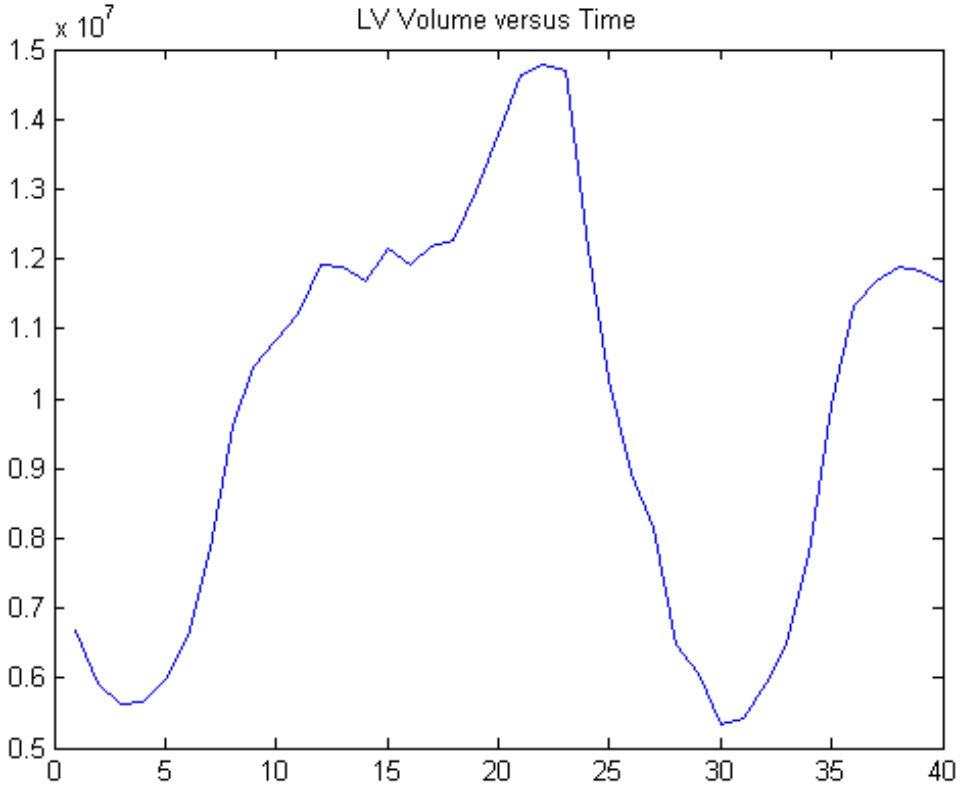

Figure 31. Plot of LV volume in pixels units versus time. The volume is computed by Simpson monoplane method on an apical 2-chamber image sequence

However it is important to notice that the algorithm for the left ventricle segmentation requires some input parameters. Although default value may be sufficient for some studies, in general a fine tuning step is necessary. Thus, an interface for easy setting of the parameters will be developed during integration.

After left ventricle border tracing step has been accomplished, the volume should be computed according to Simpson rule. The general principle underlying this rule is that the left ventricle can be approximated by a stack of circular or elliptical disks (elliptical disks are used if the 2 chamber and 4 chamber views are available both). Thus, the application of this rule requires the computation of several geometric features on the selected region. In particular the inertia axis of the regions and the long axis length are computed. Then the long axis is divided into equal parts (generally 20) and the radius of each part is computed (see again Figure 29). After having collected these geometric features, the volume is obtained by elementary algebraic manipulations.

This final volume computation stage (but not the border tracing step, up to the best of our knowledge) is performed automatically on several echocardiographic devices and is replicated with minor changes in the algorithms developed during the WP5.

2.3.3 Volume Measurements

As it is well known, 3D heart chambers are incompletely characterized by linear and area measurements, since volumes extrapolations are based on geometric assumption. Although some of these assumptions are mild (as in the case of Simpson biplane method), they may be no longer valid in case of misshapen heart chambers. In this context, magnetic resonance imaging permits the evaluation of truly 3D features like volumes and masses, which can be used as reference values for the corresponding measures obtained by conventional TTE. Actually, after having performed the segmentation of endocardium and epicardium (for example by the methods described in Section 2.2.2), a bunch of such kind of measurements can be obtained. The most important volumetric measurement is the evaluation of LV volume during the heart deformation cycle and especially at end-diastole and end-systole. Such measurements lead straightforwardly to the determination of LV ejection fraction (see Section 2.3.4 below). After having represented the endocardial surface as a mesh (see Figure 14), simple geometric computations allow to determine the volume inside the surface. Conversion from image voxel units to real world physical units is performed exploiting the information in the DICOM header. Notice that since the apical cap is not very well appreciable in the studies at our disposal, it is excluded from the computation of the volume.

Determining instead the volume between the endocardial and epicardial surfaces gives an approximation of the volume occupied by the myocardium and in turn an approximation of the LV mass, by multiplication with the standard approximate density for the myocardium (1.05 g/ml).

Three dimensional assessment of volumes and masses can also be obtained by 3D echocardiography. Since native 3D transducers are not available in HEARTFAID validation sites, we restrict to offline reconstructions from 2D cross sections of the heart obtained by conventional TTE. After acquisition of raw data (e.g. multiple apical views of the left ventricle or several short axis views), a segmentation step is required with the aim of tracing the endocardial and epicardial borders. The segmentation of endocardial border may be accomplished by using the assisted methods described in Section 2.2.1 while, for epicardial border tracing, stronger manual intervention is required (indeed no dedicated and robust automatic methods exists up to the best of our knowledge). In any case, at the end of this stage, a silhouette of the LV is obtained for each of the considered views. Then the silhouettes corresponding to the various views are aligned in the 3D space. The alignment algorithm assumes that the orientation of the considered views are known, thus introducing some dependence on plane positioning errors. After alignment has been performed, a 3D reconstruction of the LV is obtained by interpolation of the various silhouettes.

A part from dependency from plane positioning, this sort of offline echocardiography could only be applied to patients with regular rhythm, since the various views cannot be acquired altogether, but they necessarily refer to different heartbeats. Such different heartbeat are considered to be identical by the reconstruction algorithm. Although these limitations, the produced 3D reconstruction is less biased by geometric assumptions, also with respect to biplane Simpson methods. Actually the main source of bias in the 3D reconstruction consists in the interpolation step used to transform the various

silhouettes into a surface. However such an interpolation is based on a richer number of data points that, at least in principle, could be further enriched adding more and more views useful to describe a particular misshapen chamber.

2.3.4 Computation of Left Ventricle Ejection Fraction

Left Ventricle Ejection Fraction is the most fundamental parameter for systolic function derived from echocardiography. After having computed the left ventricle ESV and EDV by linear, area or volume measurements (as described in the previous subsections) on the basis of a TTE or MR study, the computation of LV EF is straightforward. Actually the simple definition of this parameter is as follows:

$$LV\ EF = \frac{EDV - ESV}{EDV}$$

As such it is a number in the range [0,1], which is more commonly expressed as a percentage.
The golden standard for the estimation of LV EF is cardiac magnetic resonance imaging. In alternative, ESC suggests to employ biplane Simpson method on TTE or TEE images, which lead to almost comparable results.
The reference range for LV EF when estimated by TTE is >55% while LV EF is severely abnormal if it is <30%. The commonly applied cutoffs are the same for men and women, although comparison with the results of MRI suggests that LV EF is somewhat higher in apparently healthy women than in men.

2.3.5 Computation of Fractional Shortening and Fractional Area Change

Similarly to left ventricle EF, there are other two parameters of systolic function normally used in clinical practice. They are the so-called Fractional Shortening (FS) and Fractional Area Change (FAC).
The FS parameter is computed on the basis of linear measurements of the LV diameter, performed at end-diastole and end-systole. Such measurements can be made by M-mode ultrasound as in Figure 23. After having computed the end diastolic diameter (EDD) and the end-systolic diameter (ESD), FS is given by:

$$FS = \frac{EDD - ESD}{EDD}$$

The FAC parameter is computed instead on the basis of area measurements of the LV performed again at end diastole and at end-systole. Such measurements should be made on a mid-papillary level view of the LV by tracing the internal boundary of the LV (as in Figure 27 and Figure 28). Semi-automated methods will be available, very similarly

to the one developed for the segmentation of apical views (see section 2.2.1). After having computed the end diastolic area (LVAD) and the end-systolic area (LVAS), the FAC parameter is given by:

$$FAC = \frac{LVAD - LVAS}{LVAD}$$

Both parameters are usually expressed as a percentage.

2.3.6 Computation of Vena Cava Collapsibility Index

The importance of Inferior Vena Cava Collapsibility Index (IVCCI) parameter is related to the estimation of pulmonary pressure. Actually, since the diameter of IVC decreases during inspiration when the negative intrathoracic pressure leads to an increasing in RV filling from the systemic veins, it is known that IVC diameter and its percent decrease correlate to RA pressure.
For the computation of IVCCI, the maximum anteroposterior diameter of the Inferior Vena Cava both in deep inspiration and expiration should be estimated.
Such values can be obtained by two linear measurements on a 2D or 2D-targeted M-mode image of the heart taken from the subcostal window.
Then, the IVC collapsibility index is defined as the normalized difference beetwen the diameter during expiration IVCe and the diameter during inspiration IVCi, that is:

$$IVCC = \frac{IVCe - IVCi}{IVCe}$$

2.3.7 Computation of Wall Thickness (WT)

Wall thickness is an important parameter for the analysis of LV regional function. Roughly speaking it is defined as the thickness of the myocardium in every point of the heart and it is computed as the distance between the epicardial and endocardial surfaces. To understand heuristically the meaning and the importance of wall thickness, notice that the heart could be assumed in first approximation as incompressible, since the percentage of water inside its tissues is high. Thus the thickness of the myocardium is expected to *increase* during contraction. Therefore thickening of the wall reflects the regional activity of the heart. For example a less drained ischemic region is expected to have a lower thickening with respect to the physiological behavior. Although regional wall motion may also convey similar information, in a comparison study by Azhari et al. (1990) in which they analyzed the power of wall thickening and wall motion in the detection of dysfunctional myocardium, it was concluded that WT is a more sensitive indicator of dysfunctional contraction. Further, when endocardial incursion and wall thickening are being used to judge the extent of ischemic myocardium, it is intuitive that

the impairment of wall movement tends to extend beyond the area of ischemia while by contrast the failure of wall thickening approximates more closely the actual area of ischemia. This finding has triggered several researchers to define refined methods for the quantification of wall thickness, especially from tomographic images (MRI, fastCT). One of the first proposed methods consists in dividing the myocardium into small elements (generally cubes) and then defining the local wall thickness as the ratio between the volume of the particular element and the average area of its endocardial and epicardial surfaces. However the most famous and widely employed method is the so-called *centerline method.* In its original formulation it is based on a 2D slice-wise approach. Starting with the endocardial and epicardial contours at each short-axis slice, the centerline method measures WT in chords drawn perpendicular to a line that is equidistant to both contours (the centerline). This method is more accurate than the previous one because it does not rely on a fixed coordinate system. However it makes some basic geometric assumptions (namely that the contours are perpendicular to the long axis of the LV) which are not satisfied in practice, especially in the apical region. Actually, the centerline method leads to an overestimation of WT in the apical region. A truly 3D approach may relieve this problem. Instead of treating each slice separately, one may indeed consider the epicardial and endocardial surfaces directly, by representing them explicitly as meshes or implicitly as zero level set of suitable signed distance functions. Then, instead of defining the centerline, one may compute the *center-surface*, which is a surface equidistant from the endocardial and epicardial surfaces. Methods for the computation of the center-surface range from classical medial axis transform to some type of computation involving signed distance functions, depending on the chosen representation of the surface.

Using the results of Section 2.2.1, it is not difficult to compute wall thickness by using short axis cardiac magnetic resonance imaging. Actually, since the output of the neural network achieving segmentation is a signed distance function, we adopt a variant of the center-surface method based on a suitable manipulation of this kind of functions. The result of this method is a dense estimation of wall thickness that can be visualized and navigated in 3D by plotting wall thickness as an attribute of the epicardial surface. Figure 32 and Figure 33 show respectively the estimated wall thickness at end diastole and at systole.

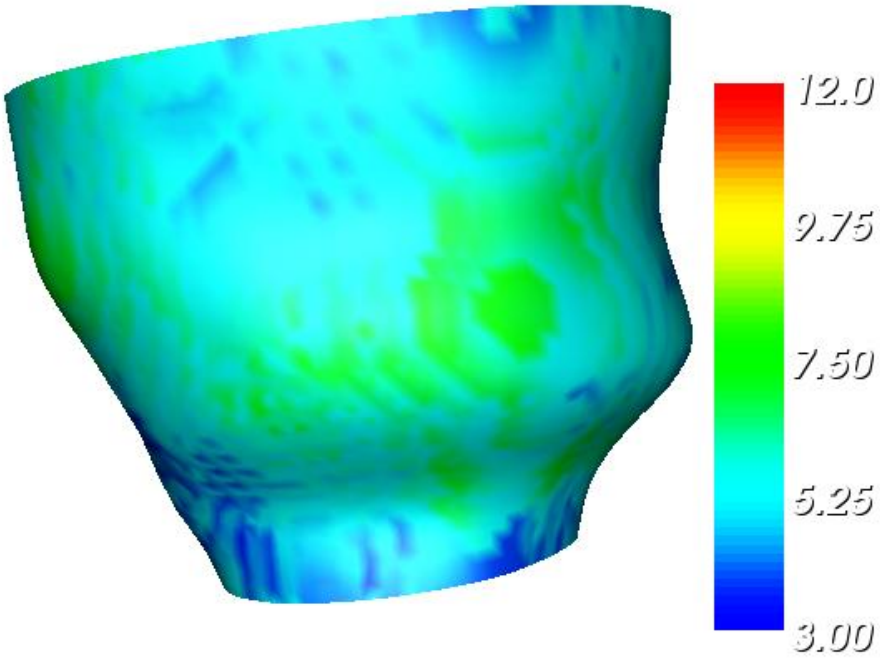

Figure 32. Estimation of wall thickness at end diastole

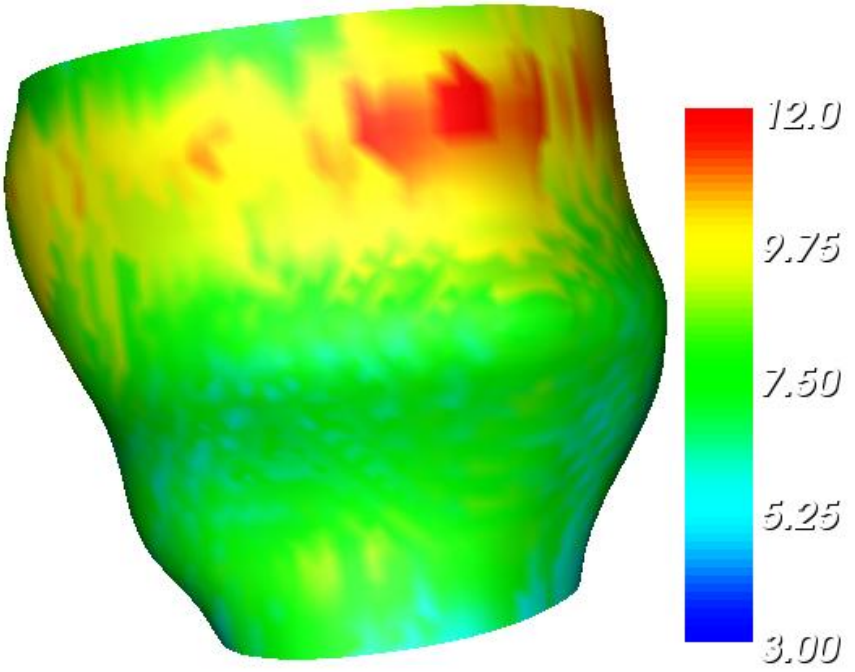

Figure 33. Estimation of wall thickness at end systole

Segmental analysis of wall thickening can also be performed by echocardiography. In this case, the analysis relies on a bunch of linear measurements computed on 2D targeted M-mode images or directly on 2D images for which we refer to the previous remarks on linear measurements (See Section 2.3.1). As for the quantification of wall motion from echocardiography, one may rely on a standardized model of LV, e.g. the 16 segment model described in D15 (see references therein) proposed by ASE (1989). See Figure 34 for a sketch of this model. Denoting by WT_S and WT_D the linear measurements corresponding to wall thickness at end systole and end diastole, wall thickening is defined by:

$$\text{Wall Thickening} = \frac{WT_S - WT_D}{WT_S} \times 100$$

In an echo study in which segmental analysis of wall thickness is performed, wall thickening is estimated for each of the 16 segments of the LV ventricle; these computed values are then used to associate a score of regional function. In a normal state, wall thickening is expected to be more than 50%, while by contrast it can decrease to less than 10% in case of akinesia.
An example of computation of wall thickening by targeted M-mode is given in Figure 35.

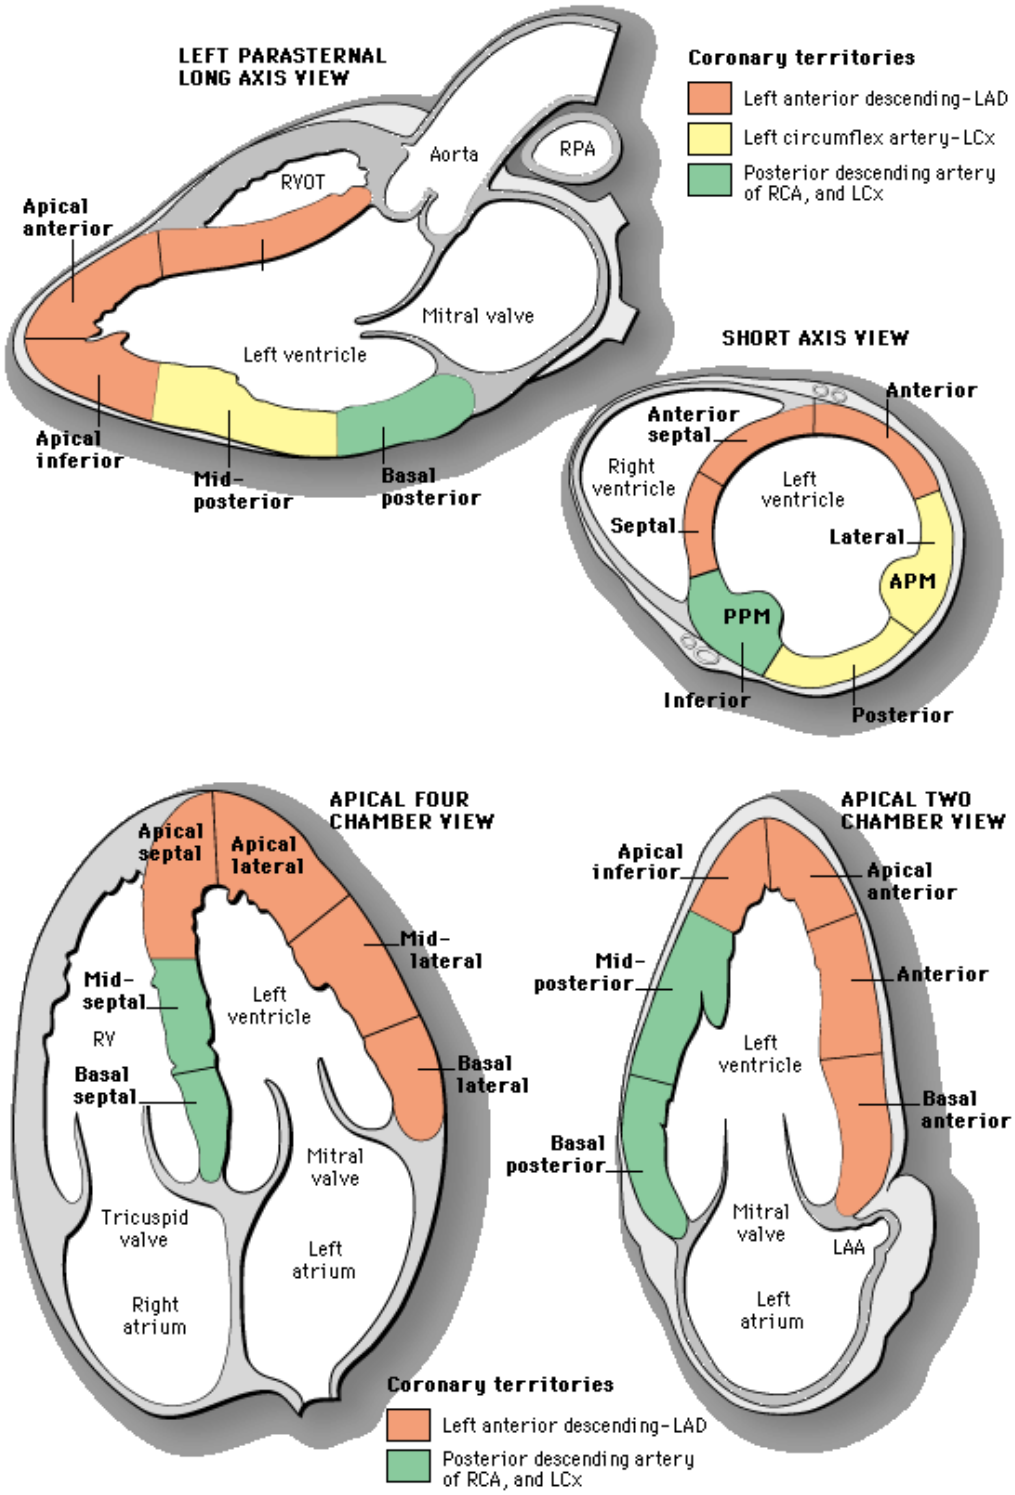

Figure 34. Sketch of the 16 segments model for segmental analysis of LV in echocardiography (source: Atlas of Echocardiography, Yale University)

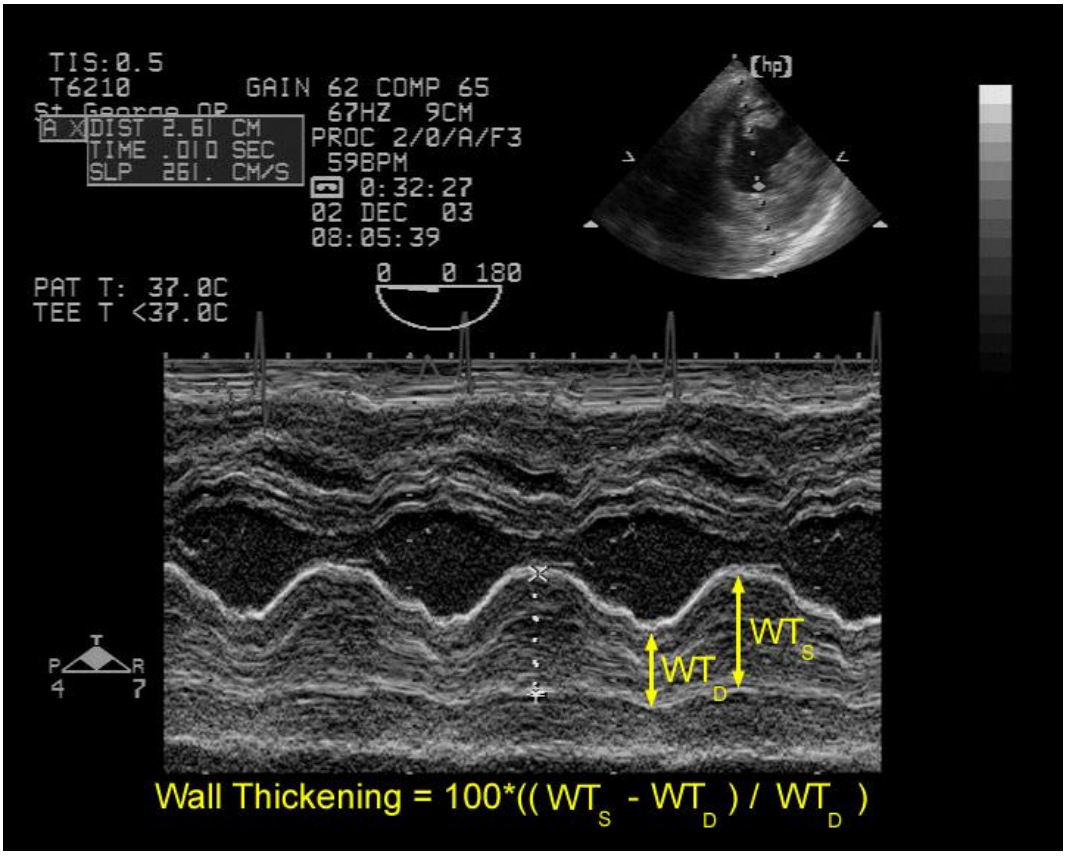

Figure 35. Example of estimation of wall thickening by echocardiography.

2.4 Towards Heart Deformation Pattern Assessment

As we have seen, imaging modalities provide an invaluable aid in analyzing the heart, a very complex structure whose deformations are of clinical importance for the assessment of its functional properties.
However the image sequences, involved in such kind of "deformation analysis", contain a huge amount of high dimensional data (2 or 3 spatial dimensions plus time) which cannot be fully exploited unless with the help of suitable tools for image processing, pattern recognition and image understanding. Actually, in cardiac analysis, many imaging techniques (like MRI, fastCT, PET, SPECT, echocardiography…) allow acquiring video sequences of the heart, from which its dynamical behavior can be inferred.
However, in daily practice, sometimes, physicians extract only the most salient frames from the video sequence (end diastole and systole) and perform direct comparison among images in the selected subset. Actually many of the clinical parameters we have defined so far (LV diameters and volumes, ejection fraction, wall thickening…) involve just the processing of the end-diastole and end systole frames. Considering the full video sequence, more precise and rich information about the state of the heart can be actually discovered.

In this frame, cardiac modelling provides powerful methods for pattern analysis. An abstract representation of the heart is built and can be instantiated to the particular anatomy under examination, with the aim of extracting shape and functional parameters. In addition, cardiac modelling can enable sophisticated quantitative assessment of heart pathologies, for which up-to-now only semi-quantitative evaluation is used in clinical practice. For example, this is the case of ventricular dyssynchrony, a complex phenomenon whose origins are to be tracked back to electrical conduction disturbances that affect both regional and global function of the heart; dyssynchrony results into uncoordinated ventricular wall motion due to activation delay. Despite its relevance, the only dyssynchrony marker that has received some consensus is an ECG-derived parameter, which however is poorly correlated to the outcome of resynchronization therapy. It may turn out that cardiac modelling may offer new insight into the problem of dyssynchrony characterization, by conveying novel representation features and suitable tools for their scientific visualization. Ultimately, dyssynchrony characterization may be translated via the extracted features into a statistical pattern recognition problem, thus allowing for new methods of quantification.

The main goal of cardiac modelling is to compactly but faithfully describe deformable structure in such a way to allow for deformation pattern characterization and assessment. Such an encoding would be useful to build up a reference database for similarity searches or data mining procedures.

Motivated by these problems, we would like to define a reference dynamic model of the heart: this model should be understood as an encoding of morphological and functional properties during the full heart deformation cycle. In particular, shape changes and evolution of local structure properties should be depicted in a concise form in the reference dynamic model, thus allowing for deformation analysis and deformation pattern classification.

One possible approach towards the definition of a reference dynamic model for the left ventricle is as follows. Suppose we have already accomplished the task of LV reconstruction, for example by the segmentation methods described in Section 2.2.2 for magnetic resonance imaging. Then, we may consider the reconstructed structures (epicardium and endocardium) as geometrical objects and start to perform *characterization* of them by computing *local attributes*. Formally, in this case, local attributes are mathematical functions defined on the epicardial and endocardial surfaces. We may distinguish among three types of local attributes:

– intensity based properties;
– local shape descriptors;
– local dynamic behavior descriptors.

Examples of attributes of the first type are gray level value, gradients, textures and the like. They could be extracted from the image sequence which leads us to LV reconstruction. If data collected from other imaging modalities are available, after performing registration, we can fuse this information to further annotate the structure (for example, information regarding perfusion and metabolism, obtained e.g. by means of PET imaging, can be referred to the reconstructed myocardium).

Geometric based properties, belonging to the second type, are extracted directly from the surfaces and are essential to describe locally the shape of the LV. We may distinguish between context independent features (automatically computable for every

surface such as Gaussian and mean curvature) and problem-specific properties (such as wall thickness). Finally, the local dynamic behavior may be described by properties borrowed from continuous mechanics (such as velocity field and strain tensor); they, however, require, at least, local motion estimation, that we haven't pursued yet. The chosen attributes for LV characterization are summarized in Figure 36.

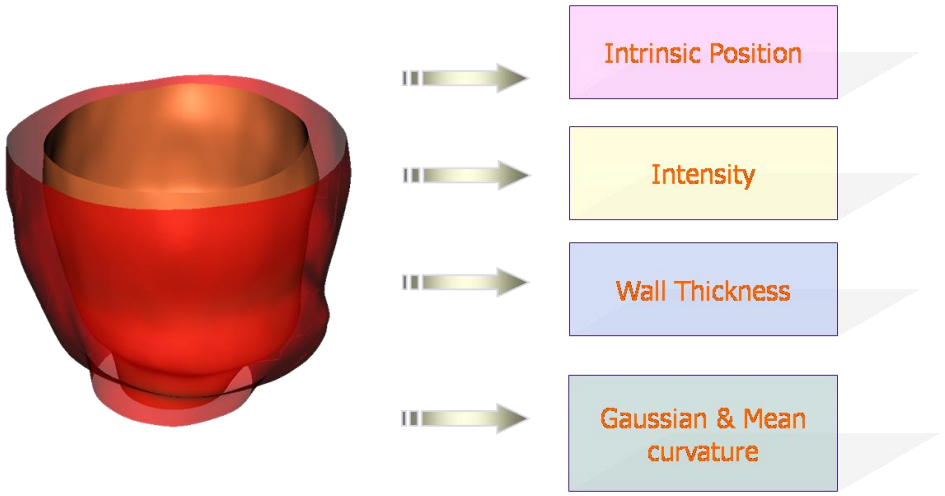

Figure 36. Local attributes for the characterization of the LV

Then we may use these local attributes to obtain a concise representation of the LV suitable for similarity searches and data mining procedures. The main idea is to combine the well established feature vector paradigm for 3D objects (see e.g. Bustos et al (2005) for a recent survey) with a modal analysis, able to cope with periodically deforming structures. Indeed, although the LV with the local attributes obtained above can be used in principle to assess the dynamic behavior of the structure and identify its deformation pattern, the voxel-wise characterization of the reconstructed endocardial and epicardial surfaces is not suited for similarity searches. The reason is twofold. First, the given description of the whole structures (two surfaces described by functions) has a dimensionality far too high to make the problem computationally feasible. Even worse, the voxel-wise characterization does not permit, at least in a straightforward manner, the comparison of anatomical structures belonging to different patients or relative to different phases in the cycle.

We address this issue using a deformable model and normalizing every instance of anatomical structure to that model: in this way the involved anatomical structures are uniformly described and can be then compared according to the feature paradigm.

To recap, we should look for a new set of 'more intrinsic' features that should be enough simple and, at the same time, capturing essential information about the structures. To obtain these new kinds of features, global information about the structures can be extracted from the local attributes, without introducing any problem-specific model. For example, one may consider the attribute spectrum, that is, by definition, the probability density functions (PDF) of a given attribute. This function captures how the attribute is globally distributed; thus, comparison of different attribute

spectra is directly feasible; to reduce dimensionality, moreover, it could be effective to compute the momenta of the PDF (mean, variance...).
However, attribute spectrum does not convey any information at all about regional distribution of the attribute. In clinical applications, this is a drawback which cannot be ignored: actually a small highly abnormal region may not affect appreciably the attribute spectrum, but its clinical relevance is, usually, not negligible. Hence, spatial distribution of attributes has to be analyzed. One approach would be to estimate multidimensional attribute spectrum. In this way, we may implicitly encode spatial relationship between different kinds of features. For example, considering the chords going from the endocardium to the epicardium orthogonal to the center-surface (see the previous section for the definition), we may use the chord length and orientation as a multidimensional attribute function defined on the center-surface. Then, the associated multidimensional PDF implicitly codifies the local distribution of wall thickness. However, a major issue in dealing with such sort of multidimensional shape distributions is the accurate estimation of the PDF. Some methods, based on the fast Gauss transform, have been reported. Although this approach may be conceivable for general-purpose 3D structure indexing and retrieval, its relevance in medical applications is not clear, for the too implicit encoding and the scarce characterization capabilities of local abnormal regions. In the same vein, approaches which do not need a refined model of the LV (e.g., Gaussian image, spherical harmonics, Gabor spherical wavelets and other general purposes shape descriptors used for content-based image retrieval) may be suitable.
However, in general one should define a model of the LV (whose primitives - elementary bricks- are regions, patches or landmarks) and then propagate it to the set of instances to be analyzed by using matching techniques. It is then possible to consider the average of a property on regions or patches (or its value in a landmark) as a good feature, since comparisons between averages on homologous regions can be performed straightforwardly.
Following this recipe, a vector of features with the desired properties could be obtained for each phase of the cycle. The deforming LV is then described by the dynamics of the temporal sequence of feature vectors obtained at different phases of the deformation cycle.
A further fruitful feature transformation may be performed. Indeed, the deformations of the LV could be assumed to be “smooth” to some extent. This implies that the LV has mainly low frequency excited deformation modes.
With these assumptions, an obvious choice is given by the Fourier transform, followed by a low pass filter, which supplies a new features vector.
The evaluation of the above mentioned parameters, at each phase, implicitly codifies information regarding structure dynamics. The overall procedure for the construction of the reference dynamic model is depicted in Figure 37.

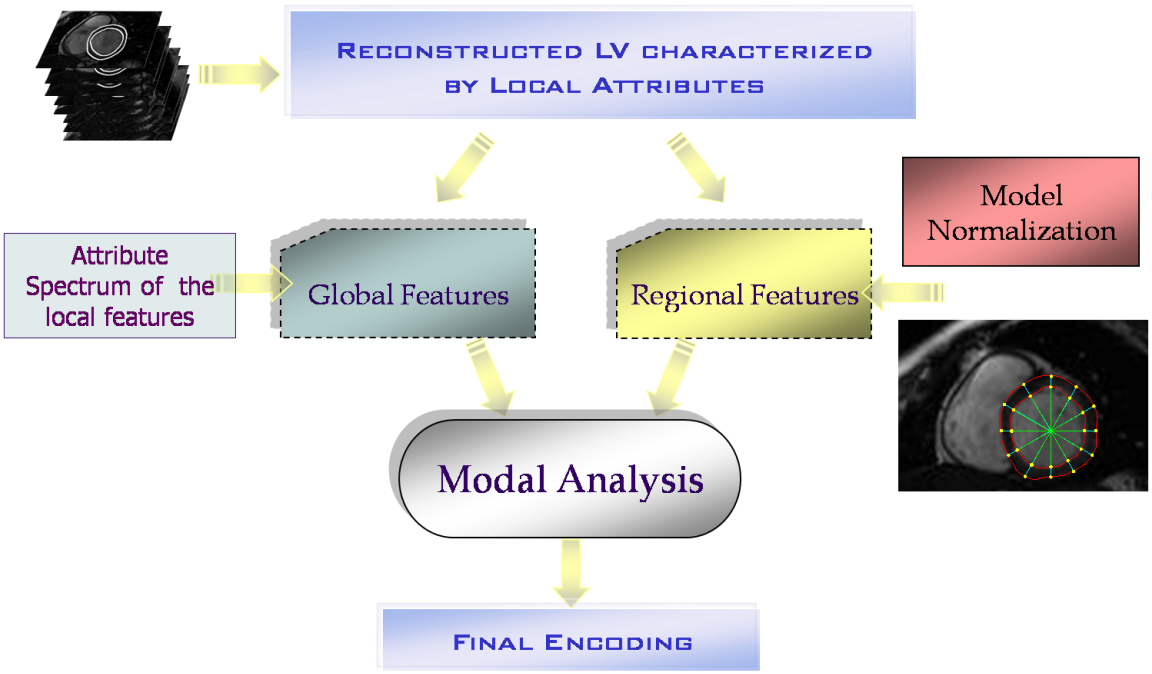

Figure 37. Sketch of the construction of the reference dynamic model

3. Signal Processing

3.1 Signal processing in the HEARTFAID: Diagnostic Resources

Due to the typical lack of interoperability in the domain of non-imaging medical devices (standards are not adopted by manufacturers) it was a serious problem to find medical devices able to transfer the non-imaging examinations to a host computer in an interoperable format containing the raw data.
Once it was decided to design some multi-purpose ECG processing algorithms, no data were available from the clinical sites and the medical device inventory was still in progress. We decided to use some publicly available annotated database and PhysioNet was a fundamental resource.

Real data (surface ECGs) have been used from the MIT-BIH Arrhythmia Database for a total of 48 records. The records are half-hour excerpts of two-channel ambulatory ECG recordings obtained from 47 subjects studied by the BIH Arrhythmia Laboratory between 1975 and 1979. The recordings are digitized at 360 Hz with 11-bit resolution over a 10 mV range. Each beat of each record has a reference annotation identifying the QRS position and type. Thus this database is very suitable for the design and evaluation of both QRS detection and QRS classification algorithms.

3.2 QRS Detection

QRS detection in electrocardiograms (ECGs) is the basic step for any further processing. Usually the limited number of available leads can be an obstacle to the attainment of high performances for a QRS detection algorithm especially when there is a high noise in one or more of the available leads (Chiarugi et al. 2007a). Noise in ECGs can appear due to several different sources like poor contact between the electrode and the skin, patient movements or breathing, etc. All these different sources can produce different types of noise like a) baseline wandering, b) powerline interference, c) muscle artifacts, d) spikes, e) sudden baseline shifts. In several circumstances the noise can appear only in one or few leads, while the others have a good-quality signal. A common approach to time-based methods for QRS detection is through the QRS enhancement achieved in a QRS enhanced signal (QeS) based on the derivatives of pre-filtered leads. The contribution of the noisy channel to the QeS can strongly deteriorate the performances of the overall algorithm. Thus, the estimation of the noise level in each channel with a criterion for the best channel selection (excluding the noisy channels from the algorithm) can improve the total performances of the QRS detector.

3.2.1 Methods

Several different techniques can be applied for QRS detection (Kohler et al. 2002). The selected approach belongs to the time-domain techniques (Chiarugi et al. 2007b). The first step consists in a signal pre-filtering using a moving-average linear filter in order to reduce the baseline wandering and the high-frequency noise, and to select the typical frequencies contained in the QRS complexes. Several different bands have been tested and the most appropriate results 5-15 Hz (Hamilton 2002).
The QeS is built as the sum of the absolute derivatives of each pre-filtered channel. The filter for the generation of the derivatives has been chosen trying to reduce the effect of the high frequency residual noise. In practice a pass-band filter is used with a derivative behaviour in the band of interest.
An adaptive threshold is initially set up as 40% of the average QeS peaks in windows of 2 sec discharging the cases out of the 98% percentiles. The average QeS peak (QeSap) is continuously updated after each QRS detection using the QeS peak (QeSp) detected in the current QRS with the formula (n is the progressive number of the detected QRS):

If
QeSp(n) >= 1.5*QeSap(n-1)
Then
QeSap(n) = 0.97*QeSap(n-1) + 0.045*QeSap(n-1)
Else
QeSap(n) = 0.97*QeSap(n-1) + 0.03*QeSp(n)
End

The beginning of a QRS is detected when the QeS overcomes the threshold (0.4*QeSap(n-1)) while the end of the QRS is revealed when the QeS goes down to the threshold and remains down for a sufficient number of consecutive samples.
To avoid False Positives (FP) due to high T-waves detection, a dead-time zone of 200 msec is set up in order to reject any QRS detection too close to the previous one. Furthermore, the QRS detection threshold is artificially increased after detecting a QRS peak and linearly decreased, with the time-distance from the previous QRS, to its base value. The QeS and the detection threshold in an excerpt of record 100 are shown in Figure 38.

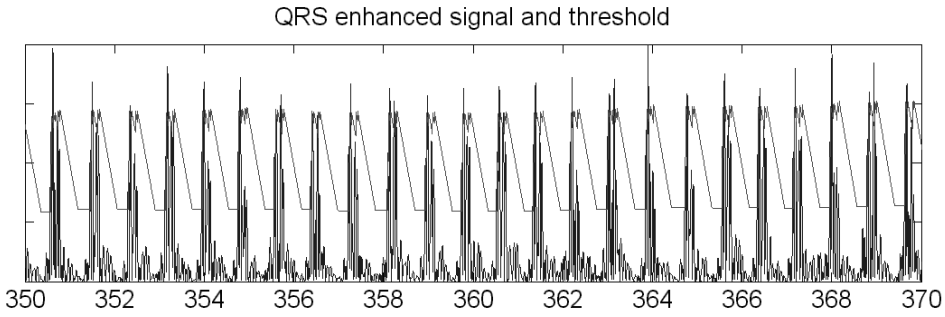

Figure 38. QeS and detection threshold for record 100 (on the abscissa there is the time expressed in seconds)

Using only the above algorithm the QRS detection results are good enough especially in recordings with low or medium content of noise. However, when the noise in one or both leads is high, the performances of the detector are significantly reduced. Therefore, we introduced a further technique in order to improve the detection performance when the noise is present only in one channel. Of course, this technique cannot be applied in case of single-lead ECG acquisitions (typical of wearable devices) but can be used in all the routine workflows of the HEARTFAID projects.
It has been decided to take into account the impact of the noise in the QRS detection as already done in other studies (Christov 2004), but with a different approach. In fact, it has been observed that, when the noise is present only in one channel, the exclusion of this noisy channel from the QeS improves the QRS detection results.
The noise level of each ECG lead is estimated with the following procedure. For each detected beat, the QRS average power is estimated as the square average of the samples in a 100 msec interval located around the detected R-peak. The T-P interval power is evaluated as the square average of the samples in an interval obtained by rough estimation of the end of T-wave and the onset of the following P-wave (Talmon 1983).
For each detected QRS a noise index (NI) is defined as the T-P interval average power divided by the QRS average power. The NI is quantized in three levels: NI<0.1 (low); 0.1<NI<0.2 (intermediate); NI>0.2 (high). The weights 0, 1, 2 are respectively assigned to each of these levels and, for any interval of an ECG; the Noise Score (NS) is estimated averaging the weights of the QRSs detected in that interval.
In Figure 39, related to record 103, the diagrams with the NI for the detected QRSs with QeS_1+2 (QeS built with channels 1+2) are shown respectively for channel 1 and channel 2.

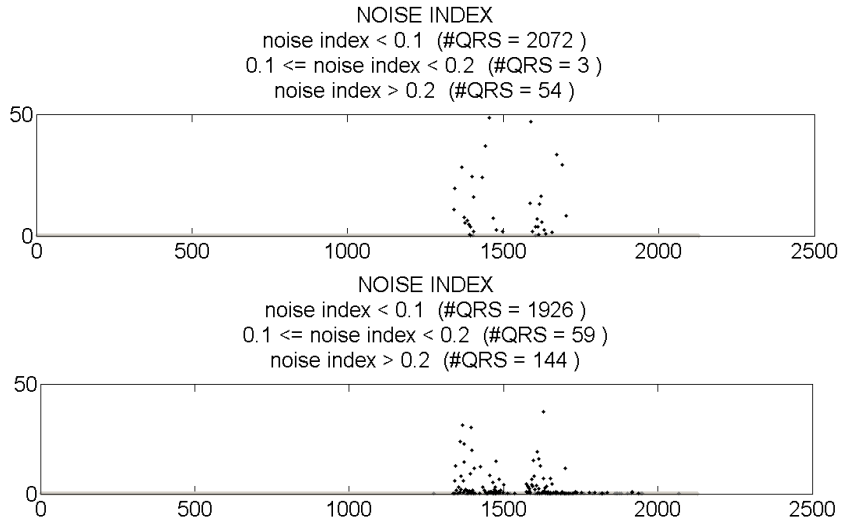

Figure 39. From top to bottom the NI in channel 1 and 2 of record 103 after QRS detection with QeS_1+2 (in the abscissa there is the #QRS). Around QRS #1500 it is evident an area (larger in channel 2) with high density of NI greater than 0.2

The presence of high noise in channel 2 for the QRSs from #1270 to #1710 induces several False Positives (FP) and False Negatives (FN) in that interval, but the high noise is revealed by the NI that indicates 3 and 59 QRSs with intermediate noise and 54 and 144 QRSs with high noise respectively in channel 1 and 2.

In Figure 40 an excerpt of record 103 from the area with high noise (identified in Fig. 2) shows how the noise is most contributed by channel 2.
If for building up the QeS in record 103 only channel 1 is used (QeS_1), then the NI for channel 1 is improved to a great extent as shown in Figure 41.

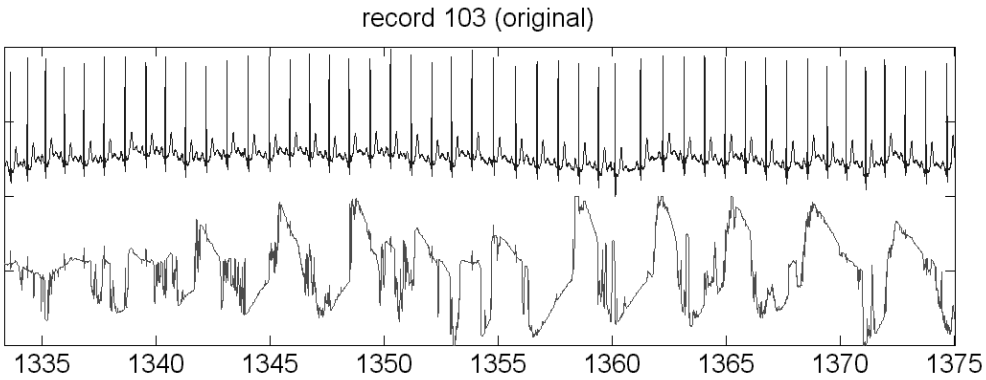

Figure 40. An excerpt of record 103 from the area with high noise identified in Figure 39. It is evident how the noise is mostly contributed by channel 2

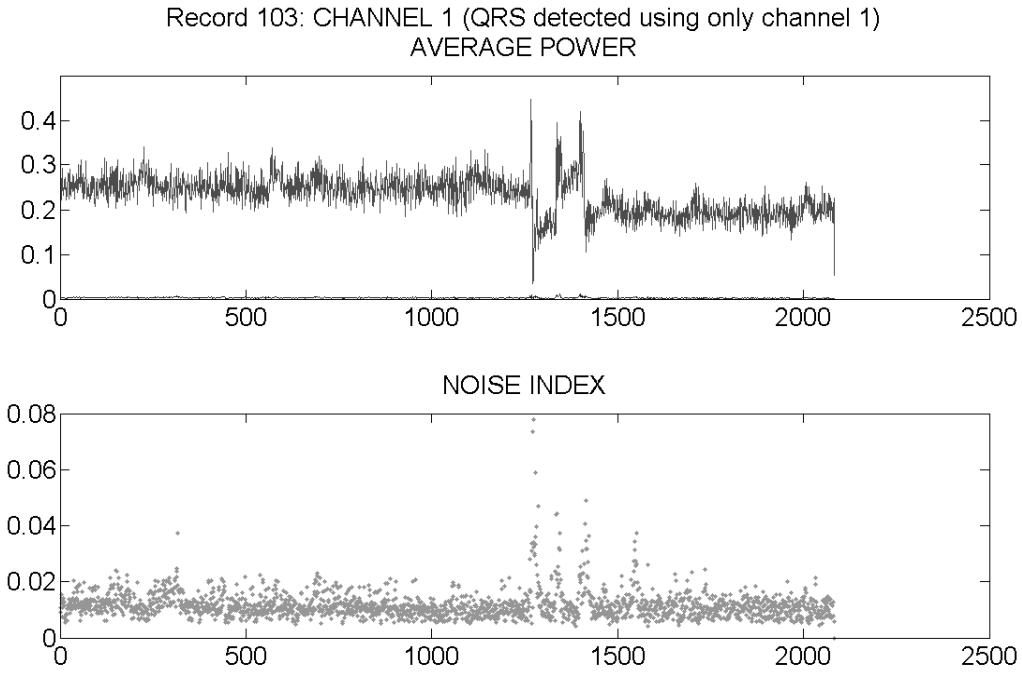

Figure 41. The (QRS and T-P interval) average power and NI in channel 1 of record 103 after QRS detection with QeS_1 (in the abscissa there is the #QRS). NI is always below 0.1

Since the NI can be used as an indicator of the noise in the two different channels and of good QRS detection, the following algorithm has been implemented for the best channel selection.

The appearance of a number of consecutive noisy QRSs (noise index greater than 0.1) determines the beginning of a Noisy Interval (No.In.), which ends once a few consecutive non-noisy QRSs appear. Based on the detected QRSs with QeS_1+2, the No.In. of channel 1 (No.In._1) and the No.In. of channel 2 (No.In._2) are calculated, and the total No.In. are evaluated as the union of No.In._1 and No.In._2. For each unified interval, the overlap percentages with channel 1 and channel 2 are calculated. For example, if a No.In. comes exclusively from channel 1, then the overlap percentage is 100% for channel 1 and 0% for channel 2.

For the record samples not belonging to any unified No.In., the QRS detection results with QeS_1+2 are used (best channel 1+2).
For the samples of each unified No.In., QRS detection with QeS_1 and QeS_2 are performed as well. Then, the following criterion is applied (j is the channel with the lowest overlap percentage and k is the other one):

1) If the NS (QeS_1+2) in channel j is 0, the QRS detection results with QeS_1+2 are used (best channel 1+2).
2) If the NS (QeS_j) in channel j is more than 0.75*NS (QeS_1+2) in the same channel, the QRS detection results with QeS_1+2 are used (best channel 1+2).
3) If the NS (QeS_j) in channel j is less than 0.75*NS (QeS_1+2) in channel j, the QRS detection results with QeS_j are used (best channel j) unless the NS (QeS_k) in channel k is less than the NS (QeS_j) in channel j and less than 0.1*NS (QeS_1+2) in channel k. In such case the QRS detection results with QeS_k are used (best channel k).

Furthermore, the database includes intervals with ventricular flutter waves that may be erroneously detected as QRSs increasing the number of FP. Although these VFL intervals might be manually excluded from the QRS detection evaluation, a first implementation of an automatic detection of the VFL intervals has been performed in order to obtain a fully automatic algorithm.
Similarly to ventricular fibrillation (VF) (Jekova 2000), the analysis of VFL can be performed with several different techniques. However, being the flutter waveforms almost sinusoidal, a frequency domain approach has been considered more appropriate.
The raw signals (both channels) are filtered with a moving average 1-30 Hz pass-band linear filter, with a Hamming window. Each pre-processed ECG signal of a record in the database (both channels of the record) is divided into periods of 5 seconds with a 50% overlap with the next one. Fourier transform of these intervals is computed and spectral characteristics as power peak amplitude and location are used to mark these 5-sec intervals as VFL or not. Each 5-sec interval marked as VFL is not considered in the QRS detection algorithm.

3.2.2 Results

In the overall database only record 207 contains VFL intervals. It has 6 VFL intervals for a total of about 142.5 seconds. The proposed algorithm for VFL interval

identification has satisfying performance in all the records of the database resulting in no FP and only 12.5 seconds of FN intervals.

In the MIT-BIH Arrhythmia Database the total number of annotated beats results 109494, with 109288 TP. FN and FP are respectively 266 and 210. The sensitivity TP/(TP+FN) is 99.76% while the positive predictive value (PPV) TP/(TP+FP) is 99.81%. In 15 records a perfect detection without any FN and FP has been obtained. 12 records have more than 10 FP+FN and only 5 records more than 30 FP+FN.

When the data from the clinical sites started to be available, a set of 75 ECGs (10 seconds, 12-lead ECG, 500 Hz) was collected from the University of Magna Graecia and the algorithm was tested on this subset showing very satisfactory results with only 6 FP and no FN. Considering that the main purpose of this task is to detect the QRS for a subsequent QRS classification and averaging, it is quite normal to reject from the averaging the first and the last QRS because the beat might not contain all its parts. For this reason the FP or FN as first or last QRS has not a real significance. With this view in mind, 4 FP were referred to the first QRS and 2 FP to the last QRS, thus no real FP were detected by our developed algorithm on the data set provided by the University of Magna Graecia.

3.3 Construction of the Average Heart Beat

A prerequisite for the construction of the average dominant beat is the morphological classification of each detected QRST. In fact, it is necessary to avoid the introduction of extrasystoles or non-dominant beat in the averaging process because they would alter the quality of the averaged beat.

Normally the evaluation of the heart beat type is performed considering its morphology and its occurrence compared to the previous and following beats (rhythm). If the requirement is to obtain a complete rhythm evaluation, then it is necessary an accurate classification of each heart beat based on both morphological and rhythm criteria. However, significant clinical information can be obtained from the analysis of the dominant beat morphology. In such respect, the identification of the dominant beats and their averaging can be very helpful because it allows the doctor to perform the measurement of amplitude and intervals on a beat that is cleaner from the noise than a generic beat selected from the entire ECG recording. Furthermore, an algorithm that allows the identification of the dominant beat if iterated on the remaining beats can provide a complete morphological classification identifying all the morphological classes present in the entire recording.

The algorithm implemented focuses on morphologically distinguishing dominant from non-dominant heartbeats of a patient's ECG recording, with feature extraction from the ECG analysis. For the classification algorithm, only the basic morphological parameters were taken into consideration, trying to limit as much as possible the complexity of such a system. For our purpose we have taken input from and tested our system on the records of MIT-BIH Arrhythmia Database (four records are from patients with pacemaker).

3.3.1 Initialization and pre-processing

High-level language and interactive environment, MATLAB, was used for the development and testing of our algorithm and code.
The two-lead ECG signals was first sub-sampled at 1/5 of the original frequency in order to pass it through a chain of median filters of 200 msec and 700 msec respectively. The first median filter has the goal of eliminating the QRSs while the second eliminates the T waves. At the end of the filtering we obtain a signal representing the course of the baseline. This signal low-pass (1 Hz) filtered and re-sampled at the original frequency is then subtracted from the original signal providing the original two- leads where the baseline wandering has been cancelled.
Then the signal is processed with a pass-band filter (1-40 Hz) in order to eliminate the residual noise.
Both ECG signals are considered in the algorithm.
The algorithm is described in two phases: we firstly calculate a possible classification of all beats, and then we re-estimate the created classes and if necessary redefine them.

3.3.2 Phase 1

In order to begin the first phase of our grouping algorithm, we need the positions (samples) of all R peaks of the signal. For the MIT-BIH Arrhythmia Database all R peaks have been annotated by cardiologists and we use them as input.
The problem is a typical problem of unsupervised clustering because information about the number of the existing classes and the morphology of the dominant QRST complexes is not available at the start of the algorithm.
The following algorithm is applied for each record of the database. We firstly calculate a rough centroid (template) for all beats of the signal, one for each lead, with duration approximately 0.6 seconds. Each sample of this template is built by the mean values of all beats for the specified lead (excluding the distribution tails). After the template is built, we then compare it with every beat of the signal and, thus, obtain a distance vector (one for each lead). However, each single QRST complex in the recording can have a shorter duration due to the premature occurrence of the next QRST complex. Thus, an estimation of the duration of each QRST complex is necessary and it is performed based on the occurrence of the next R peak.
For this comparison we only use for the centroid as many samples as the beat under consideration has (the minimum between 600 msec and the estimated QRST duration) and we, then, evaluate the L1 distance, normalized with a (p2pXbeat/p2pXcentr) factor; where p2pXbeat stands for the peak-to-peak value of the beat under consideration, and the p2pXcentr term stands for the peak-to-peak value of the centroid.
We extract the most similar beats according to L1 distance from both leads, and use them to calculate a new centroid, which will be the centroid of the dominant beat type with duration, again, approximately 0.6 seconds.
In Figure 42, Figure 43 and Figure 44, it is shown an example (from record 101) of a rough calculated centroid, of a centroid of the dominant beats after the initialization

phase, and a comparison between the dominant centroid (centroid after first initialization) and a generic dominant beat, for both channels.

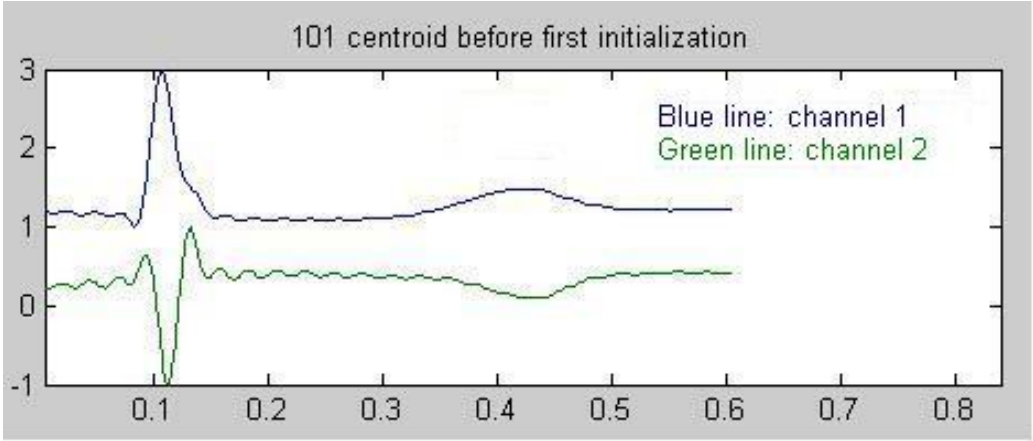

Figure 42. A rough centroid of record 101 built by all beats in the signal.
Blue line: channel 1, Green line: channel 2

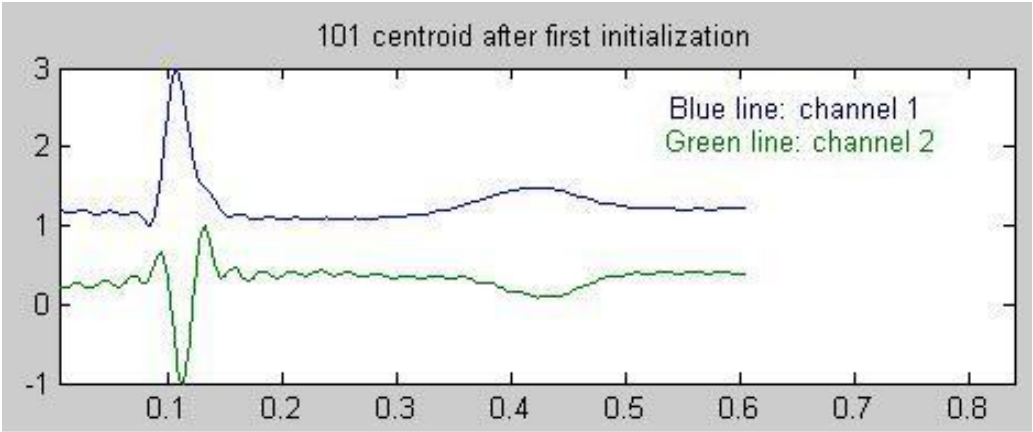

Figure 43. A centroid of dominant beats, of record 101, after initialization of the algorithm, built by the beats that are similar enough, according to L1 distance, with the rough centroid.
Blue line: channel 1, Green line: channel 2

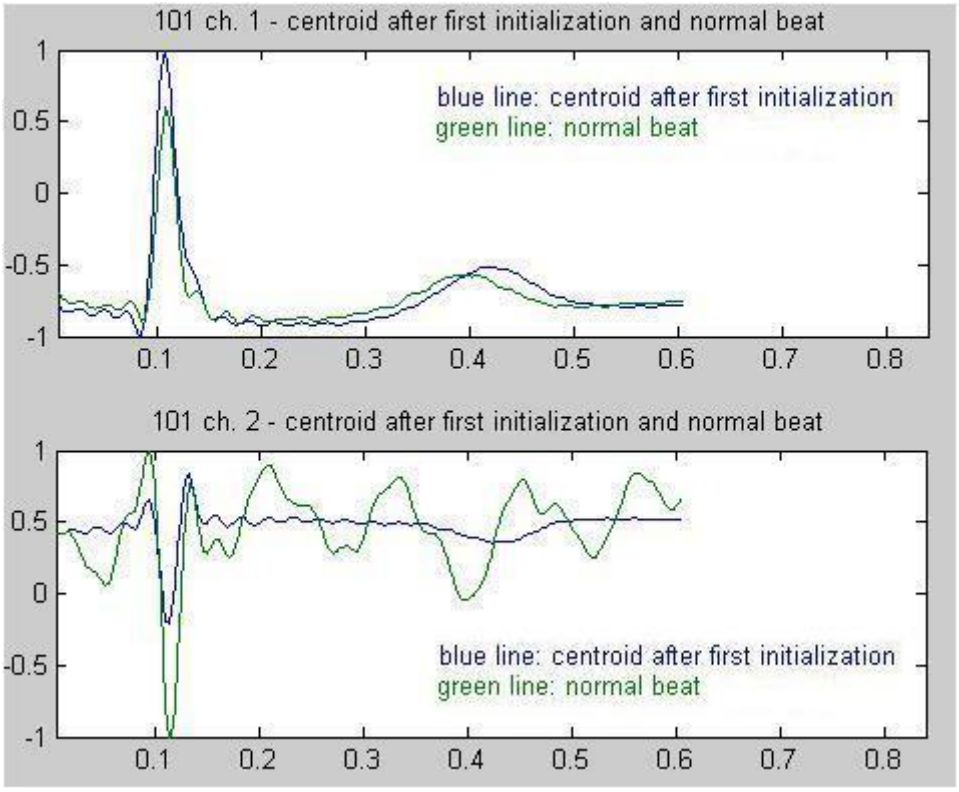

Figure 44. With blue line the centroid after first initialization and with green line a normal beat of record 101. Above for channel 1, Below for channel 2

As we have a new robust centroid we need to rebuild the L1 distance matrix, by the L1 distance of each complex with the new centroid, with the difference that this time we will build it using horizontal and vertical wiggling as we need it to be more accurate than before and a better alignment between the QRST complex under investigation and the estimated centroid is necessary. Along with L1, we calculate L2 distance matrix, also normalized on the factor (p2pXbeat/p2pXcentr). These two features combined with the centroid-to-beat correlation coefficient, are the main attributes for the "similarity" evaluation of the QRST complex under investigation. Using these 3 features and the p2pXbeat value, we determine whether a complex should be considered dominant or not. Beats in this class, formed by the dominant beats of the signal, are not processed again; only new beats may be added during the second phase.

On all the remaining beats we apply again the same algorithm, initialized this time only by the remaining beats (not by all complexes of the signal). We build a rough centroid and compare it with the remaining beats in order to find the dominant class of the remaining beats. We, then, built the new robust centroid by the dominant class, and compare every other beat with this centroid in order to classify them. The beats that do not fall in this dominant class of the remaining beats are the new remaining beats and are reprocessed in the same way. The algorithm stops when all beats have been classified or no further classification can be made (from the remaining beats) and thus, a final group with all remaining unclassified beats is additionally formed. This completes the first phase of the algorithm.

3.3.3 Phase 2

The output of the first phase constitutes the input of the second phase. Here we reprocess all groups except the one with the dominant beats of the signal. The groups containing non-dominant beats (according to first phase) that are large in number are split into smaller ones and reconsidered for misjudgment of being non-dominant. The validation criteria are two of the ones used in the first phase, but the thresholds are even stricter. We do not examine beats one by one as in the first phase also for reasons of performance (except from the group with the unclassified beats, see below). The groups are divided into smaller ones with respect to their L1 similarity with the dominant group. Centroids are calculated for every one of the reformed groups and are compared with the centroid of the dominant group using L2 distance, and centroid-to-centroid correlation, with horizontal and vertical wiggling. The beats of the group, whose comparison criteria are satisfied, are put with the dominant ones. All other group remains non-dominant.
Regarding the group which contains the remaining unclassified beats, being this group more a collection of the remaining groups than a group of similar complexes, its treatment is different and performed at the end: we compare (L2 distance and correlation with wiggling) each beat separately with the centroids of all formed groups as far; if a satisfactory match is not possible, then a new group is formed consisting of the beat that could not be put in any other class. Following beats are also compared with the centroid (at the beginning a single beat) of this new single class. This phase is the last step of the algorithm and produces the final separation between dominant and non-dominant beats.

3.3.4 Noise Contribution

This clustering algorithm, though, could not be as reliable if we didn't take into consideration noise issues. It is well known that noise could cause great inconvenience when it comes to processing ECG signals. Noise factor could not be ignored here, either. For this reason, before processing the two-channel signal to obtain a beat classification, we located every noisy segment in the signal. We calculated the square power of a QRS segment and divide it with the square power of the following TP segment. If this ratio is below 10, we consider the complex not to be noisy; if the ratio is between 10 and 20 we consider it to be low in noise and if it exceeds 20, then is considered noisy. According to this Noise ratio for every complex, we often lighten the criteria for the clustering algorithm of the first phase, and for the process of the last group with the remaining beats, during second phase. To better estimate the starting and ending sample of a QRS and ST segment we used an adaptive threshold for the PQ and QT estimation, considering the next R peak and avoiding T wave's contribution. Figure 45 illustrates a noisy QRST complex.

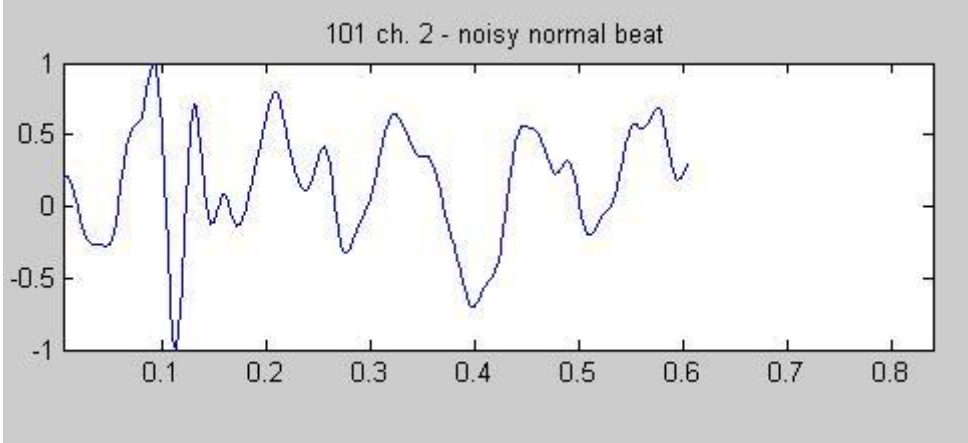

Figure 45. Noisy dominant beat in channel 2 of record 101

3.3.5 Results

The performances of the algorithm on all the records of the MIT-BIH Arrhythmia Database are resumed in Table 1.
The first column, Record, contains info of the record of the two-channel signal used from the database. The second column, TP (True Positives), represents the Dominant beats classified as dominant. The third column, FP (False Positives), represents the non-dominant beats classified as Dominant. The fourth column, FN (False Negatives), represents the Dominant beats that were classified as non-dominant and the final column, TN (True Negatives), represents the non-dominant beats that were correctly classified as non-dominant.

Table 1: The results in terms of true positive, false positive, false negative and true negative for each record of the MIT-BIH Arrhythmia Database.

Record	*TP*	*FP*	*FN*	*TN*
100	2271	0	1	1
101	1860	0	3	2
102	1970	16	58	143
103	2074	0	10	0
104	1328	95	52	754
105	2406	11	120	35
106	1500	3	7	517
107	2078	0	0	59
108	1494	0	251	18
109	2483	2	9	38
111	2076	0	47	1
112	2539	0	0	0
113	1789	1	0	5
114	1830	0	2	47
115	1950	0	3	0
116	2301	1	2	108

117	1523	0	12	0
118	2259	2	3	14
119	1543	0	0	444
121	1861	0	1	1
122	2476	0	0	0
123	1512	0	3	3
124	1566	1	1	51
200	1666	42	107	786
201	1665	20	1	277
202	2094	2	3	37
203	2359	22	170	429
205	2569	2	5	80
207	1559	34	5	262
208	1558	60	28	1309
209	2994	0	10	1
210	2416	31	7	196
212	1825	17	0	906
213	2662	100	4	485
214	2003	11	0	248
215	3188	0	10	165
217	1540	45	2	621
219	2086	10	3	55
220	2047	0	1	0
221	2028	0	3	396
222	2473	0	10	0
223	2104	5	13	483
228	1688	0	3	362
230	2096	0	159	1
231	1254	0	1	316
232	1774	0	6	0
233	2235	5	2	837
234	2750	0	0	3
Sum	***97322***	***538***	***1138***	***10496***

In order to evaluate our system we calculated commonly used parameters for validation of such systems; Sensitivity is the rate of TP/ (TP + FN), Specificity is the rate of TN/ (FP + TN), Positive Predictive Value (PPV) is the rate of TP/ (TP + FP) and Negative Predictive Value (NPV) is the rate of TN/ (TN+FN).
Our system achieved a Sensitivity of 98.84%, a Specificity of 95.12%, a Positive Predictive Value of 99.45% and a Negative Predictive Value of 90.22%.
In the tests performed on the data set provided by the University of Magna Graecia we had only 6 FN and 1 FP, but 5 FN were referred to the first QRST of the recording that was only partially contained in the registration. How specified during the discussion of the QRS detector it is a common rule to exclude the first and the last QRST from the averaging process because such complexes might not be completely recorded and they might alter the averaging process. Thus, for the purposes of our algorithm 5 FN have

not to be considered real FN and the total errors are just 1 FP and 1 FN. The final results are therefore very satisfactory.
The averaged dominant beat is represented by the class centroid of the dominant class evaluated on all the QRST assigned to the dominant class. The averaged dominant beat is then saved in section 5 of the SCP-ECG file format and can be shown by the ECG viewer.

3.4 Extraction of Clinical Parameters

Extraction of significant clinical parameters is a task reserved to the clinicians. Thus, due to the variety of different parameters and to the fact that clinicians usually like to check and confirm the measurements, the role of the signal processing was mainly dedicated to the provision of reliable supporting tools for a more objective evaluation of the significant parameters. The clinical sites were provided with a graphical ECG viewer endowed with specific tools like zoom and caliper, so that each ECG once processed could be displayed (and printed) including the average dominant beat that was estimated by the signal processing chain. The average dominant beat is more clean from the noise than the original signal and the clinicians can perform more accurate measurements using the caliper on this average beat reducing the inter and intra observer variability and properly filling the CRF for the ECG examination
Some screenshot of the ECG viewer are shown in the next figure pages.

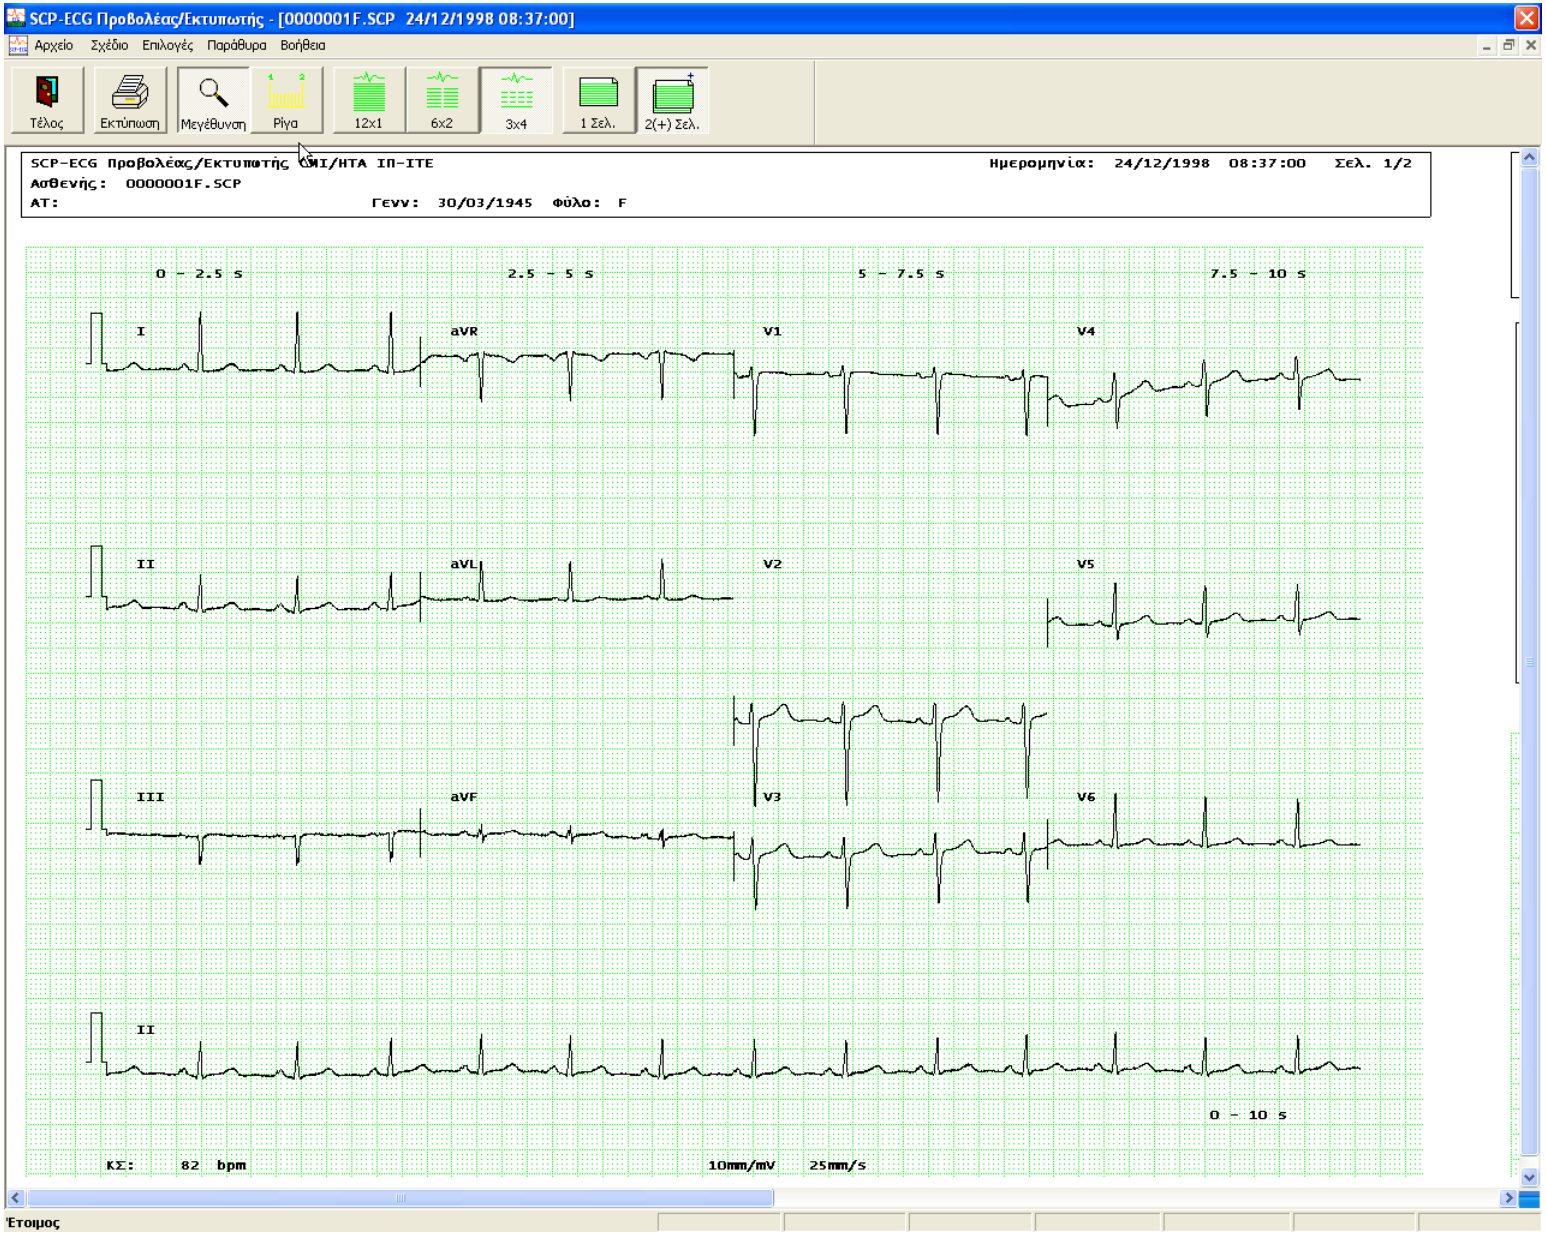

Figure 46. A screen of the ECG viewer displaying in format 3x4+1 the original ECG signal

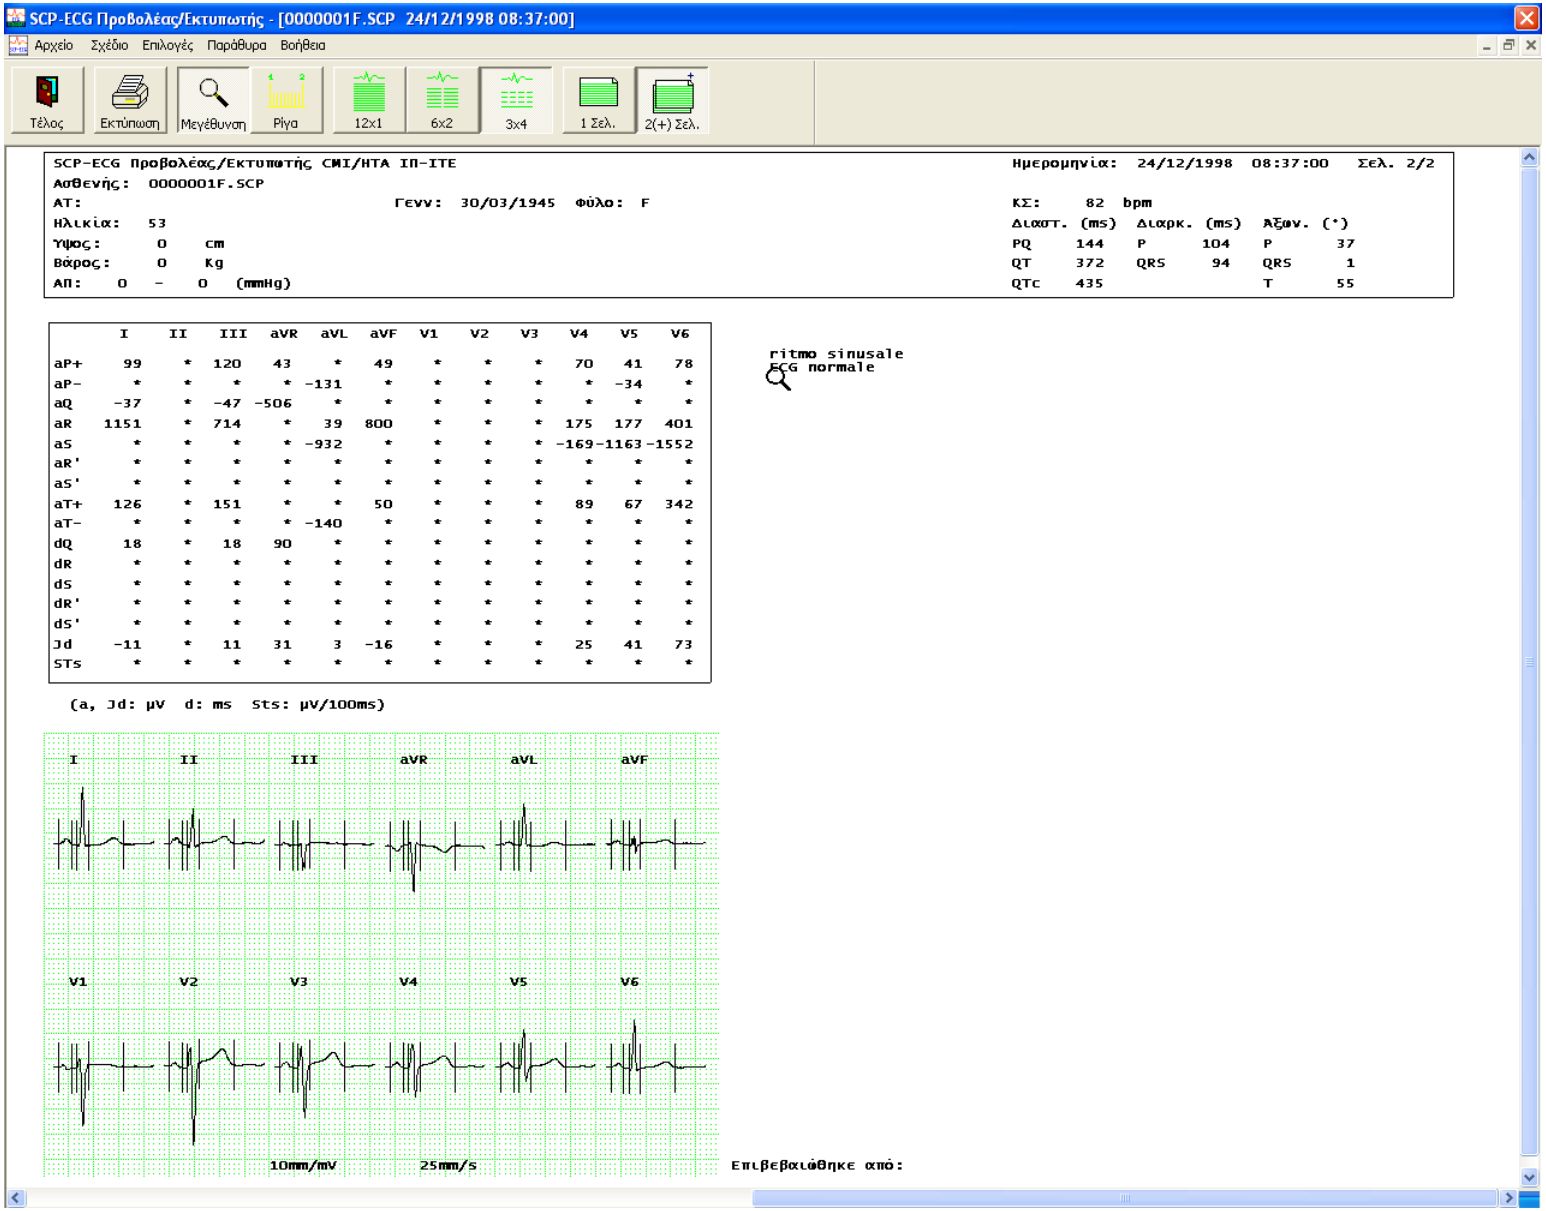

Figure 47. A screen of the ECG viewer displaying additional information including the reference (average) beats

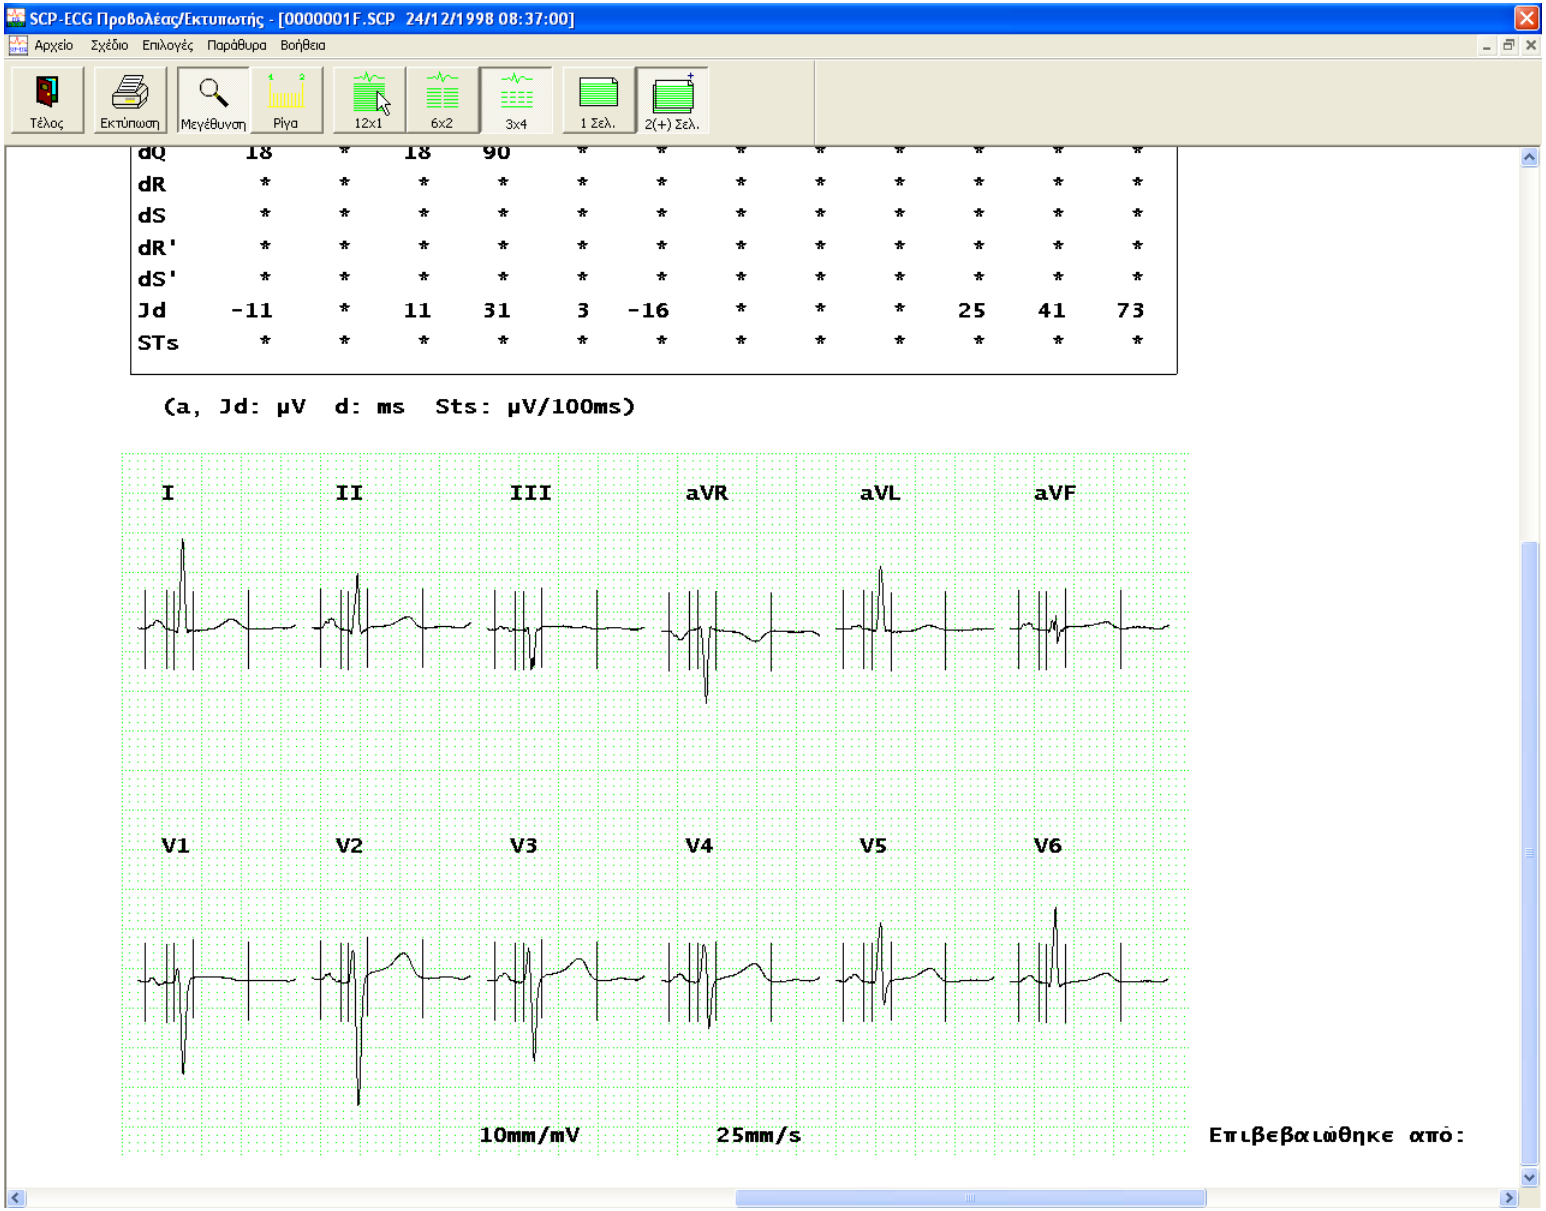

Figure 48. A screen of the ECG viewer displaying additional information including the reference (average) beats

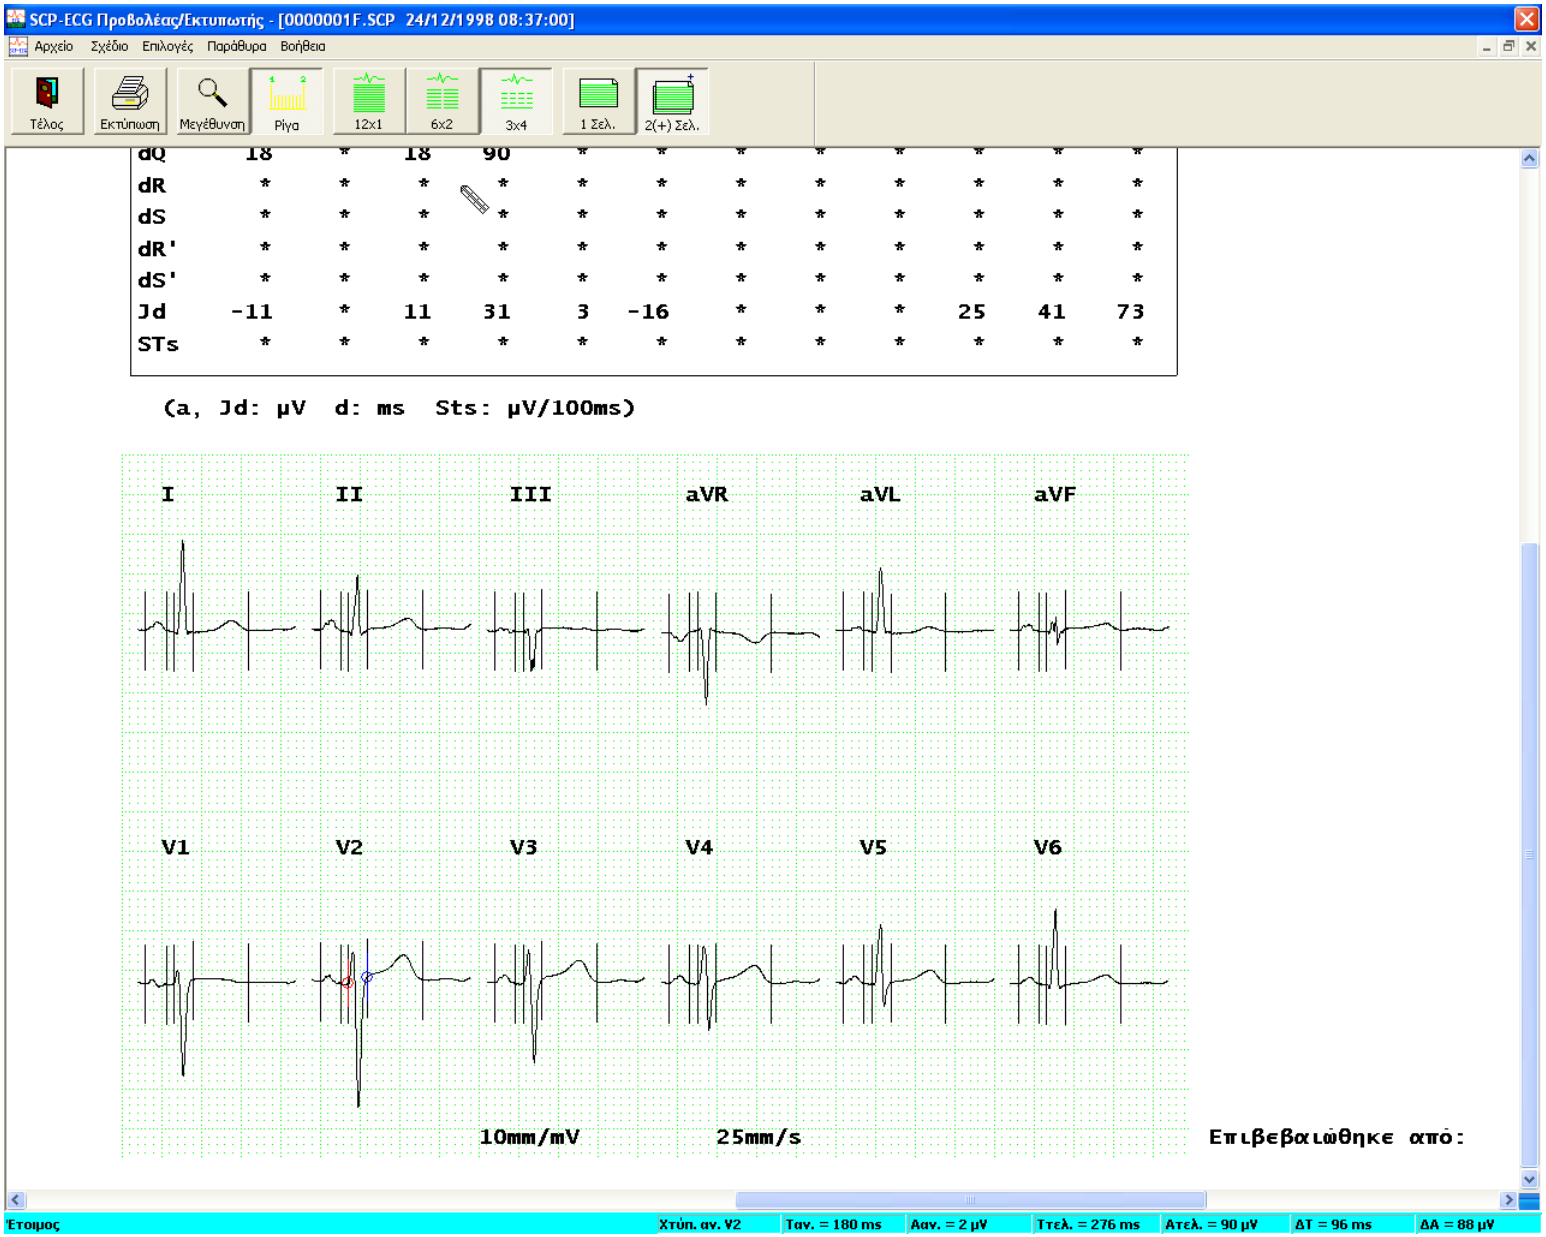

Figure 49. A screen of the ECG viewer displaying (in zoom mode) additional information including the reference (average) beats. The calliper (ruler) is active and the amplitude and intervals can be accurately measured

4. IT infrastructure for the integration of signals and images in the Platform

In this chapter we describe several issues related to the integration of the tools for signal and image processing. In particular, starting from an analysis of the echocardiography workflow, we introduce the need of a structured image archive and discuss the current installation (Section 4.1). The integration of the image processing results into the platform is mediated through the use of the image archive and standard-compliant network services, as described in Section 4.2. Interfaces for the deployment of image processing tools are described next (Section 4.3) together with an extendable web viewer interface for easy and quick image visualization and selection. A discussion about the selected standard for ECG storing and transmission ends the chapter (Section 4.4).

4.1 Image Archive and Manager

4.1.1 Image archive for echocardiography workflows

Nowadays echocardiography is a digital modality, offering the opportunity to coordinate its workflow in an IT framework. When considering echo workflow, we are mainly interested in TransThoracic echo (TTE), for its versatility and portability and for its fundamental importance in the management of heart failure patients.
An echo study generally consists of digital images (single- and multi-frame), measurements and an interpretive report.
Images are obtained by a sonographer who may make preliminary measurements and preliminary observations. According to IHE, the over-reading echocardiographer must have access to all of this data in discrete, structured format to synthesize a final report.
In an echo lab reached by HEARTFAID, it would be optimal to upload the original images, the annotated ones and the final report to the integrated HFP of services. However, the treatment of digital images, both original and annotated, poses several problems, due to the discrepancy between the ideal hospital (just from an IT point of view, of course) and the situation usually encountered in a real one.

Ideal echocardiography workflow
In the ideal situation, the hospital is equipped with a Hospital Information System (HIS) and a PACS dedicated to cardiology and, finally, echocardiography devices are persistently connected to the hospital network. For a patient, pre-admitted and registered in the HIS before undergoing an echocardiographic examination, a visit is scheduled and demographics and procedure information (for example why the visit is required, which parameters should be estimated during examinations,...) are transmitted accurately to the echo device. After the examination, images are securely stored to the PACS and can be displayed at any imaging workstation. Echo measurements, performed anywhere, are correctly associated and securely stored with the study as

discrete, structured data which can be interpreted by another workstation and finally incorporated into a report.
In particular, interoperability is guaranteed among HIS, PACS, Echo devices and various radiology workstations.

The actual workflow
In a real-world example, instead, echo devices are not connected to any network. Sometimes, a workstation provided by the echo device vendor and running proprietary software, is associated. This workstation has one (or several) storage units to setup a local picture archive for the echo lab. Although potentially connectable to the hospital or global network, the local archive often can export images only to physical devices (e.g. by CD-rom burning). In most cases, luckily, images are exported according to DICOM standard and, thus, there exists no image format problem.
For what regards echo measurements and findings (usually printed to paper and inserted into a patient's folder) cardiologists typically have to retype the information into a separate reporting system, since cross document sharing (XDS) seems not to be feasible (via HL7 export, for instance). Further, review of already annotated images is not possible, since annotation is made through the proprietary software of the echo device.

Discussion
Given such a discrepancy between ideal and real world, the HEARTFAID image archive should try to cope with the nowadays IT structures of its validation sites, with a view towards the future that -as far as we can see- will resemble more and more the ideal situation described above. In particular, it is essential for the success and diffusion of the platform to offer:

- State of the art interoperability. In this way, HFP could easily be connected to an ideal hospital IT infrastructure, allowing for the exchange of images, measures, reports, demographics and procedural information directly from the HIS/PACS to HFP and *vice versa*. Notice indeed that whereas original images have of course as source only the echo lab, procedural information and reports will be mediated by the CDSS. For example, CDSS may suggest which kind of parameters is more useful to be evaluated for a given patient during a TTE session. This list could be directly sent to the echo device. In such a way the sonographer has directly on his screen, besides patient demographics (reducing the risk of misspelling in his name), the purpose of the examination. Further some of the measures and annotated images will be produced through image analysis tools provided by the platform; this data may be of interest also for the HIS/PACS.

In such a situation only (though by no means trivial) security/privacy/non-repudiation concerns are left out.

- Customizable modules to solve IT infrastructure lacks. In particular, for what regards the image archive, this means to develop methods and interfaces for uploading echo images to HFP and query/retrieve images from the archive.

The rather obvious answer to these two items, at least for echocardiography workflow, is to adhere to standards (namely DICOM) and integration profiles (provided by IHE). The methods offered by the standards, suitably inserted in an interface, allow for image

uploading, query/retrieve and review of reports from any workstation, either in the case of IT infrastructure lacks.

4.1.2 DICOM Servers

Since the 1970s, the emerging idea of a digital image archive (PACS) and electronic image distribution in a hospital created the need to exchange digital images between medical devices of different manufacturers. The American College of Radiology (ACR) and the National Electrical Manufacturers Association (NEMA) formed a working group in order to develop an image exchange standard. After some first failure, DICOM was released in 1993 to offer an open vendor-independent platform for the communication of medical images and related data, supporting PACS networks and guaranteeing interoperability. DICOM standard, which is composed of several parts, is continuously being extended and updated. The last part "WADO" was released in 2003.
It is important to notice that DICOM is not just an exchange format for medical image data. Actually DICOM goes far beyond, defining among several others:

- Data structures (formats) for medical images and related data
- Network oriented services, e.g.:
  - Image transmission
  - Query of an image archive (PACS)
- Web-access to images
- Print (hardcopy)
- RIS - PACS - modality integration
- Formats for storage media exchange
- Requirements for conforming devices and programs.

For these reasons, any implementation of a DICOM server can't be merely a database of images in DICOM format, since storing, querying, retrieving, web-access are prescribed by the standard.

Network Services
DICOM network services are based on the client/server concept. In case two DICOM applications want to exchange information, they must establish a connection and agree on the following parameters:

- Who is client and who is server
- Which DICOM services are to be used
- In which format data is transmitted (e.g. JPEG compressed or uncompressed)

Only if both applications agree on a common set of parameters, the connection can and will be established. In addition to the most basic DICOM service image transmission (the so-called *Storage Service Class* such as C-STORE), several advanced services are available. The following two, for example, may reveal useful in HFP:

- The DICOM image archive service ("Query/ Retrieve Service Class") allows to search images in a PACS archive according to user-defined certain criteria (patient, time

of creation of the images, modality etc.) and to selectively download images from this archive

- The DICOM modality worklist service allows to automatically downloading up-to-date worklists including a patient's demographic data from an information system (HIS/RIS) to the modality.

Web Services

Users of medical information systems may benefit from rapid and reliable access to reports and images. For example it would be important during an examination to retrieve on the fly reports or even images of the previous examination. Within HFP it is likely that such access will be based on web technologies also for image retrieval and visualization. The access to relevant DICOM *persistent objects* (i.e. images and reports, not logs of workflows) should be either in native DICOM format for advanced use or rendered into a generic format (e.g. JPEG, PDF) that can be presented with off-the-shelf applications. DICOM standard offers the so-called "WADO" (Web Access to DICOM persistent Objects) to answer these needs. It is a web-based service for accessing and presenting DICOM persistent objects, consisting in a simple mechanism for accessing a DICOM persistent object from HTML pages or XML documents, through HTTP/HTTPs protocol, using the DICOM UIDs (study, series and instances).
In particular, according to RFC2396 and subsequent versions, an URI should be composed as:

http://<authority><path>?<query>

In our context <authority> is the provider of the Dicom services (e.g. HEARTFAID platform or a hospital), <path> is the path to the DICOM WADO services and <query> specifies the DICOM query, e.g.:

http://heartfaid.isti.cnr.it/wadoRequest.m?<wado query>

Notice that this address and the ones below do NOT correspond to the actual addresses used in the implementation of the platform but are used here for explanation purposes.
The WADO standard defines then the <wado query> syntax. A common WADO request looks like:

http://heartfaid.isti.cnr.it/wadoRequest.m?studyUID&seriesUID&objectUID[&frameNumber]

where:

- studyUID: UID of the study containing the object(s)
- seriesUID: UID of the series containing the object(s)
- objectUID: UID of the single object (Service Object Pair SOP)
- frameNumber: optional parameter number of the selected frame (for multiframe image objects only)

For example:

http://heartfaid.isti.cnr.it/wadoRequest.m?studyUID=1.2.840.113619.2.98.3456.1193039503.0.26&seriesUID=1.2.840.113619.2.98.3456.1193039503.0.27&objectUID=1.2.840.123619.2.98.3456.1193039503.0.28.512&frameNumber=12

This request retrieves frame number 12 from a multiframe image sequence specified by the unique identifiers of the study, series and object (that is, in this case, the image itself).
The MIME type of the retrieved object can also be specified by the contentType. Supported format are image/jpeg, application/dicom. text/html, application/pdf and others. For example, the URI:

http://heartfaid.isti.cnr.it/wadoRequest.m?contentType=application/dicom&studyUID=1.2.840.113619.2.98.3456.1193039503.0.26&seriesUID=1.2.840.113619.2.98.3456.1193039503.0.27&objectUID=1.2.840.123619.2.98.3456.1193039503.0.28.512

allows to download to the local file-system and visualize the multiframe image identified by its UIDs as before.
We refer to (Cordonnier 2003) for a discussion of WADO and the presentation of several related scenarios.

Security Concerns and Current limitations of WADO
Clearly the information contained within DICOM objects may be considered as Protected Healthcare Information (PHI). Thus the protocol used, that is HTTP, can be replaced by HTTPs for that purpose. Further, DICOM standard defines two optional parameters, *anonymize* and *annotation*, which control respectively the absence of patient identification in a retrieved DICOM object and the presence of patient identification burned into the pixel data of images. It is likely, however, that for patient enrolled into HFP personal information will be erased and replaced with an ID (the same used in eCRF) before storing images in the HEARTFAID image archive.
When dealing with echocardiography, one usually deals with multiframe images (i.e. image sequences). It would be interesting to be able to retrieve and visualize multiframe images via a web-based service. However, considering the size of such kind of images, some sort of *streaming* is necessary. Currently, up to the best of our knowledge, no kind of streaming is suggested or prescribed by DICOM standard. In particular, WADO gives access to the first frame of a multiframe image by default. Subsequent frames can be obtained by using the option frameNumber in the composition of the WADO URL.

4.1.3 Beyond DICOM

As we saw, DICOM defines a complex standard to guarantee interoperability among different DICOM compliant devices ranging from DICOM server to radiology workstation, from imaging devices (US, CT, X-ray...) to HIS. Any DICOM compliant device comes equipped with a DICOM conformance statement in which the portion of implemented DICOM features is described. In practice, however, conformance

statements are only comprehensible by experts and they are frequently inadequate since often only a minimum set of features is documented. In some cases interoperability problems tend to occur because some inconspicuous details do not go together.
Further, DICOM standard alone is not able to cope with the complex echocardiography workflows described above.
To answer these needs, IHE (an initiative by healthcare professionals and industry) tries to improve the sharing of healthcare information by promoting the coordinated use of established standards such as DICOM and HL7 to address specific clinical needs in support of optimal patient care.
IHE Technical Framework defines specific implementations of the afore-mentioned established standards to achieve seamless transmission of information among physicians, medical specialists, nurses, administrators and other stakeholders.
In particular IHE provided a Cardiology Technical Framework (written by the American College of Cardiology (ACC), the Radiology Society of North America (RSNA) and Healthcare Information and Management Systems Society (HIMSS)). Echocardiography workflow is included with the aim of:

- Providing echo measurements interoperability
- Ensuring images and measurements are securely stored
- Uploading echo reports to a repository
- Reconciling patient demographics

In IHE terminology, this is obtained implementing several *actors* (ADT/HIS, Scheduler, Acquisition modality, Image Manager, Evidence Creator). Since there is an obvious relationship between IHE actors and various modules in the HFP related to the image archive, it is at least reasonable to explore IHE Cardiology Technical Framework in the close future. For the moment, it is timely to use a DICOM server implementation which integrates IHE actors, that is, according to IHE terminology, *Image Manager/Archive* and optionally *Report Manager* and *Report Repository*.

4.1.4 Current Installation of the Image Archive

A) General Description

From the previous discussion, the good features expected from HEARTFAID Image Archive may be summarized as follows:

- DICOM network services
- Web access to DICOM objects
- Easy development of web interfaces for Image Archive Management
- Easy development of web interfaces for image uploading
- Implementation of IHE actors
- Extendibility to meet HFP needs (interaction with CDSS and Image Analysis Tools)
- Multi-platform or platform independent

Among different open-source implementation (CONQUEST, DCMTK, DCM4CHE), DCM4CHE, a Java implementation of DICOM, has been chosen according to the

previous requirements list. For sake of completeness, an overview of DCM4CHE components is presented in
Figure 50.
Besides being an image archive, DCM4CHEE provides a toolkit of standalone applications and methods to make network communication and interface development easier. Unfortunately, the documentation of DCM4CHE is not completed yet, so it is somewhat difficult to grasp its architecture.
In the current installation, DCM4CHEE has an underlying MySQL DB, though other choices (e.g. PostgreSQL) are conceivable. Both the Dicom and MySQL servers accept network connection. Thus it is possible to query directly the DB for low level platform management purposes (e.g. user name, group, password harmonization).

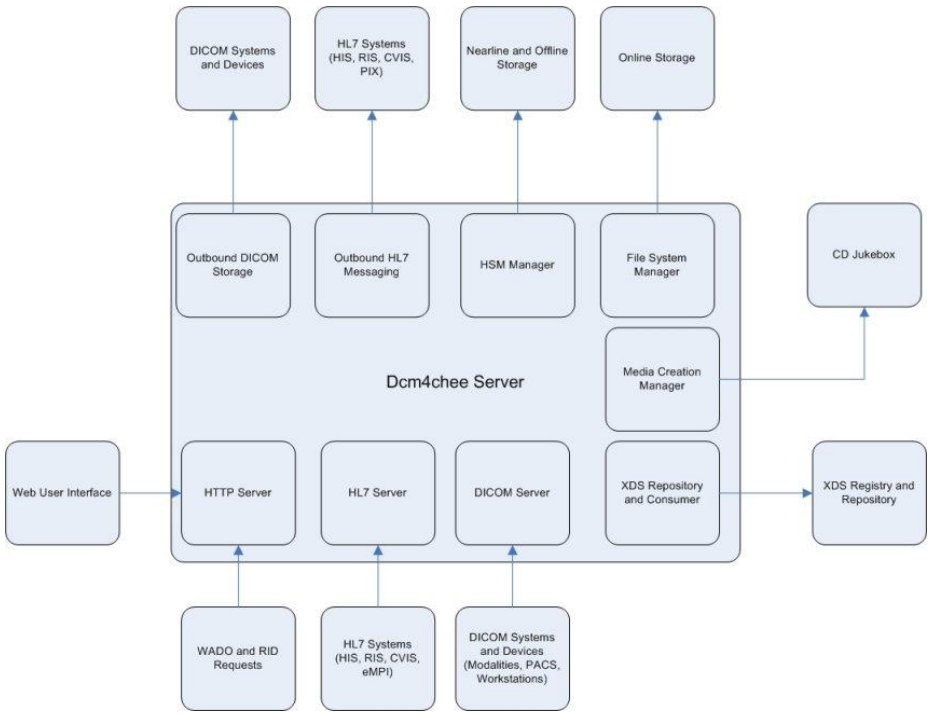

Figure 50. Overview of DCM4CHEE components

B) Interface for Image Archive Management

DCM4CHE comes equipped with two web-interfaces for Image Archive Management. The first interface allows among others for

- Making queries about patient studies in the DB
- Editing Patient information e.g. for anonymization purposes
- Downloading of images in native DICOM format
- Management of study permissions among the several user roles
- Visualization of images by WADO
- User administration for the management of user name, role, password
- Management of recognized DICOM entities (that is for example name and IP address of DICOM clients that are allowed to retrieve information)

- Export of studies for clinical trials as DICOM structured reports, with automatic on the fly generation of reports in PDF format (using Apache FOP)
- Audit Record, to monitor the jobs required to the Dicom server

This first interface could be easily modified, customized and integrated into HFP. Harmonization of Patients demographics and user data may be, for example, directly obtained through queries to the DB or by using more refined single sign on middleware applications (e.g. Shibboleth or openSSO to name a few). Some screenshots of the current interface are presented in Figure 50, Figure 51, Figure 52, Figure 53, Figure 54 and Figure 55.

However, it is important to point out that this kind of interface (although being powerful and complete) seems to be too complicated for standard management of image data by the medical staff. This conclusion has been drawn also from the feedback by our clinical partners. For this reason, a simpler and interoperable interface for image web-viewer has also been considered and integrated. See Section 4.3 below.
The second interface is a standard jmx-console *Mbean view*, that allows for invoking Mbean operations for internal DICOM server configuration, WADO services, HL7 messages routing. In addition it provides archive management and maintenance procedures. As such, it is not necessary to integrate it into the platform, since it is mainly addressed to technical archive administrators.

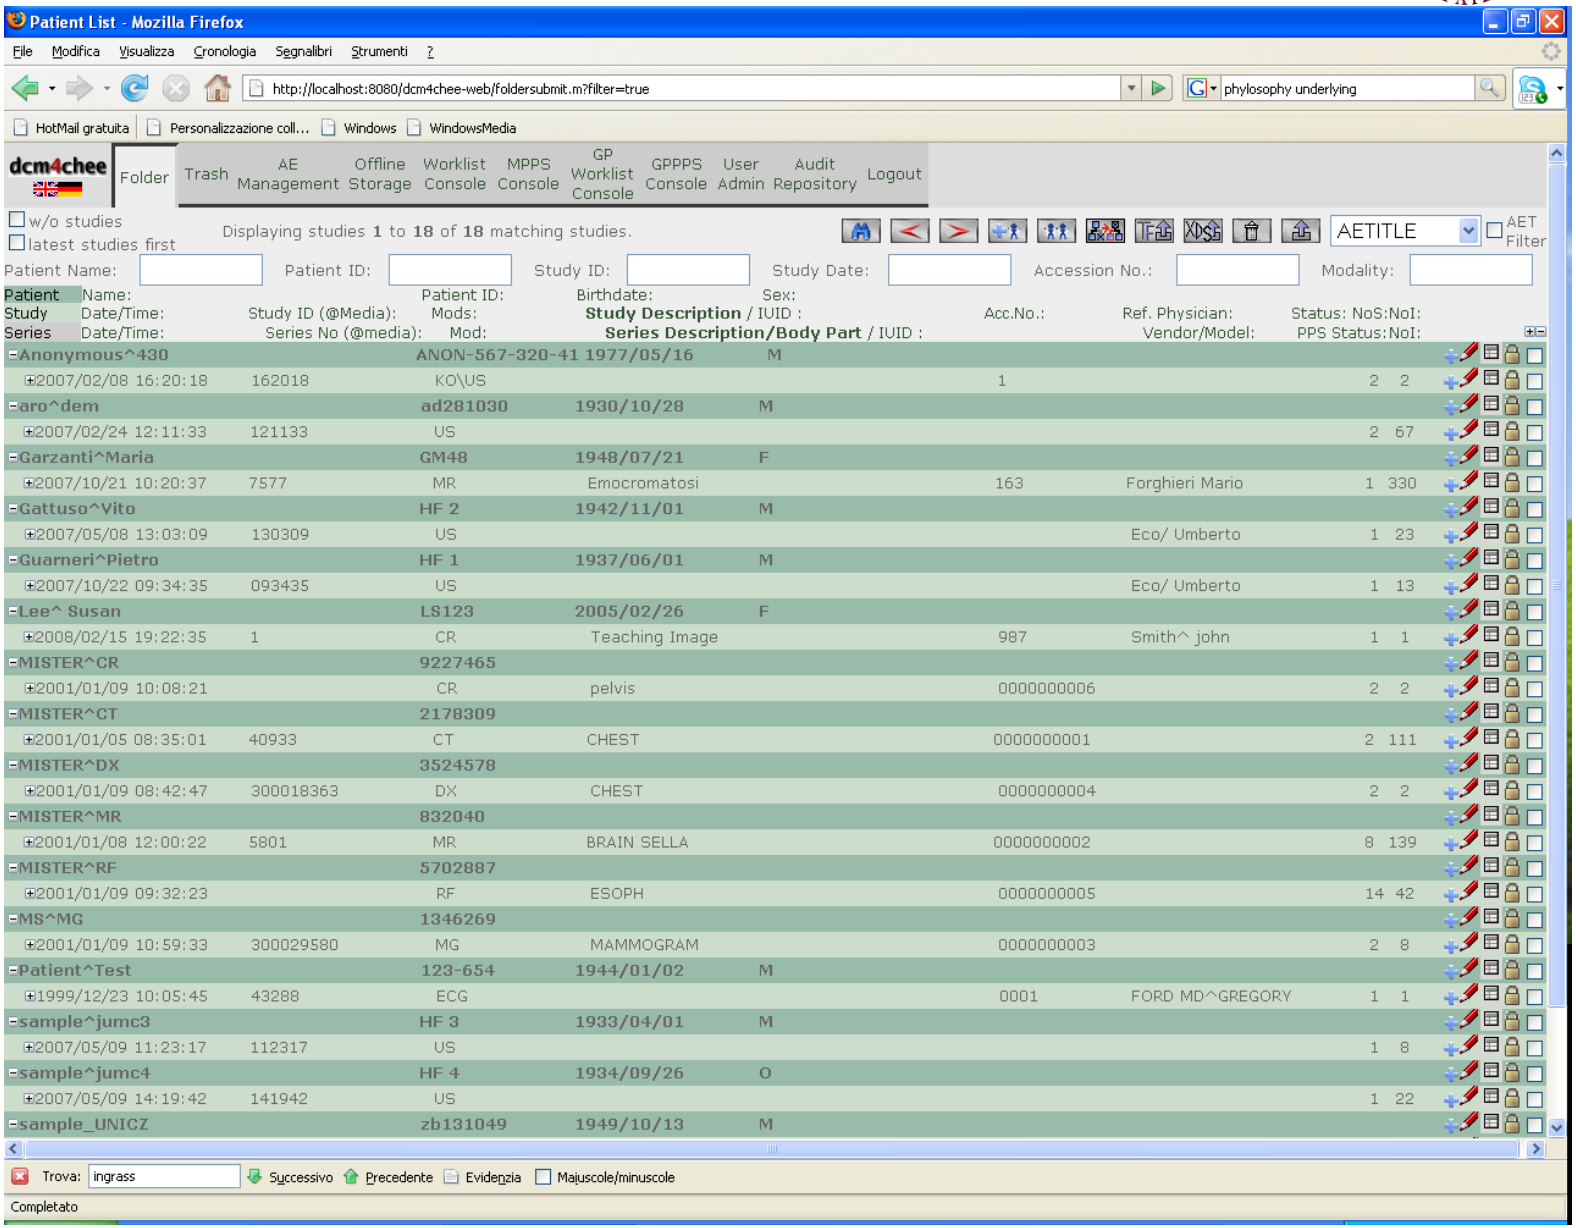

Figure 51. Main search window of the current installation. Apart from querying a specific patient or study, links to info editing functionalities are provided.

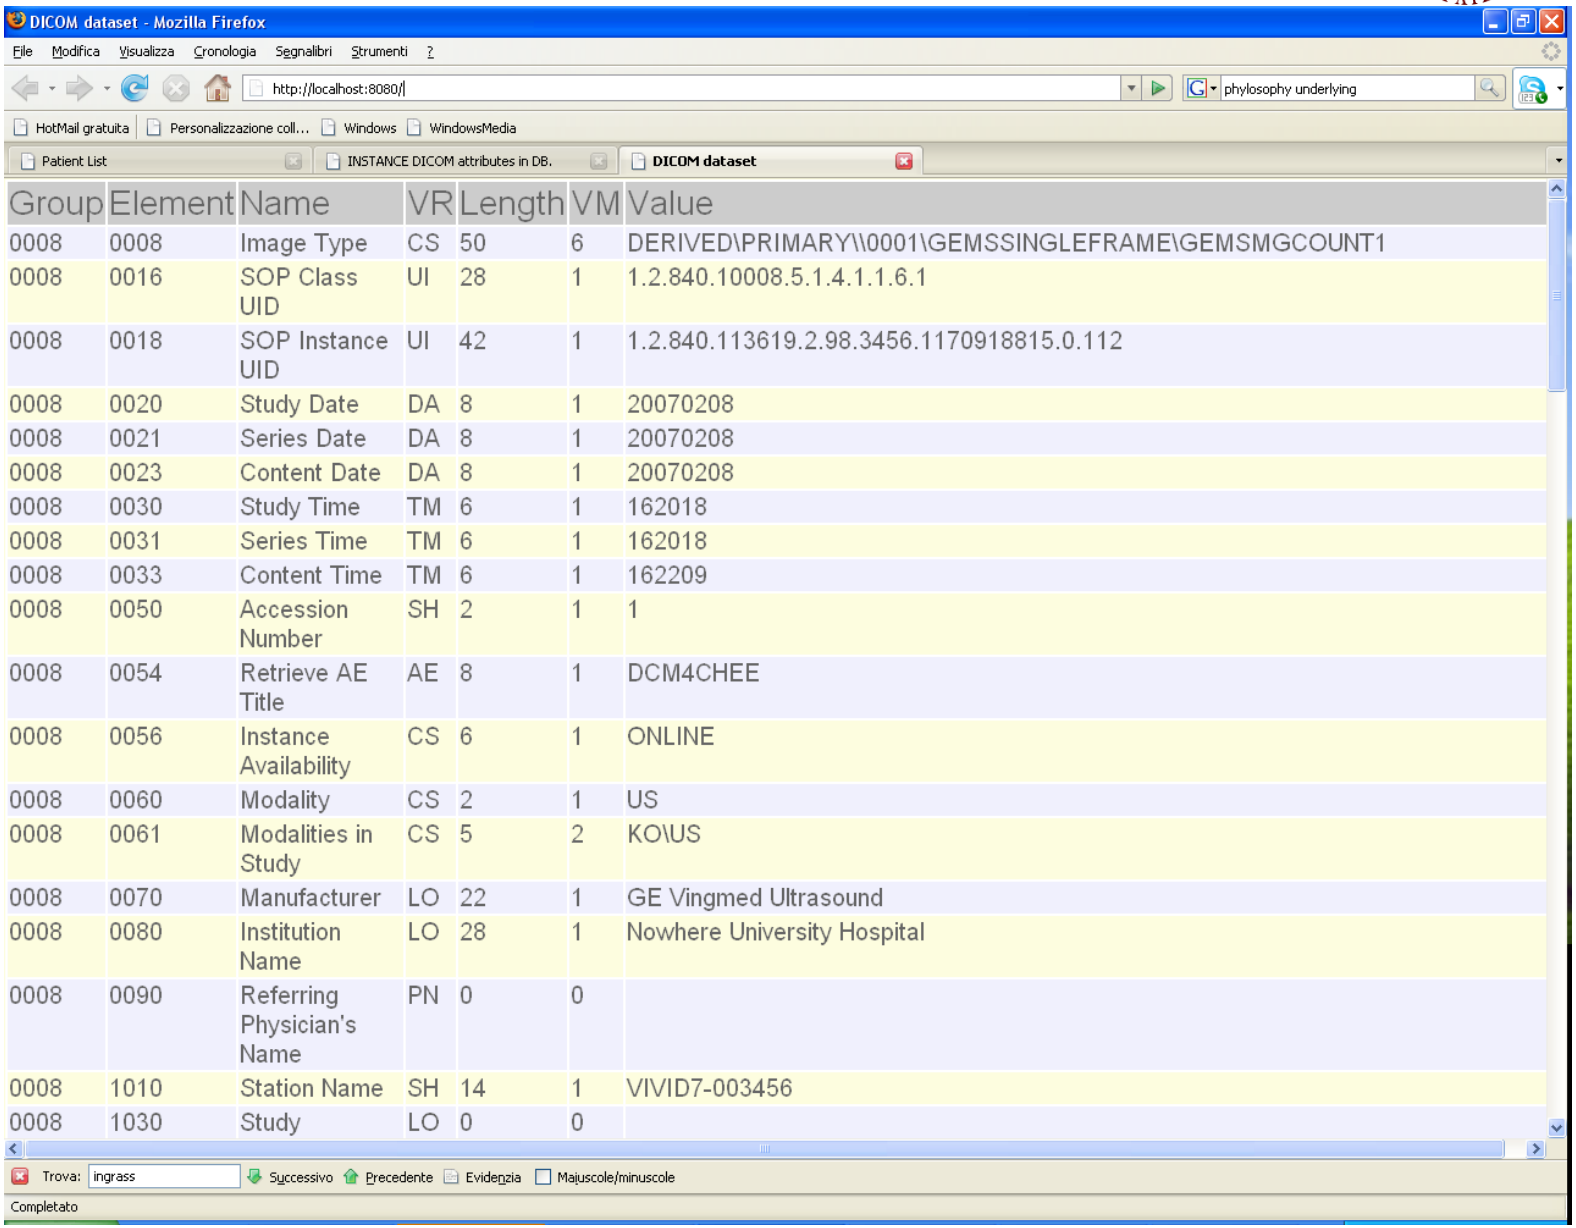

Group	Element	Name	VR	Length	VM	Value
0008	0008	Image Type	CS	50	6	DERIVED\PRIMARY\\0001\GEMSSINGLEFRAME\GEMSMGCOUNT1
0008	0016	SOP Class UID	UI	28	1	1.2.840.10008.5.1.4.1.1.6.1
0008	0018	SOP Instance UID	UI	42	1	1.2.840.113619.2.98.3456.1170918815.0.112
0008	0020	Study Date	DA	8	1	20070208
0008	0021	Series Date	DA	8	1	20070208
0008	0023	Content Date	DA	8	1	20070208
0008	0030	Study Time	TM	6	1	162018
0008	0031	Series Time	TM	6	1	162018
0008	0033	Content Time	TM	6	1	162209
0008	0050	Accession Number	SH	2	1	1
0008	0054	Retrieve AE Title	AE	8	1	DCM4CHEE
0008	0056	Instance Availability	CS	6	1	ONLINE
0008	0060	Modality	CS	2	1	US
0008	0061	Modalities in Study	CS	5	2	KO\US
0008	0070	Manufacturer	LO	22	1	GE Vingmed Ultrasound
0008	0080	Institution Name	LO	28	1	Nowhere University Hospital
0008	0090	Referring Physician's Name	PN	0	0	
0008	1010	Station Name	SH	14	1	VIVID7-003456
0008	1030	Study	LO	0	0	

Figure 52. Visualization of the DICOM header of an image by the web interface

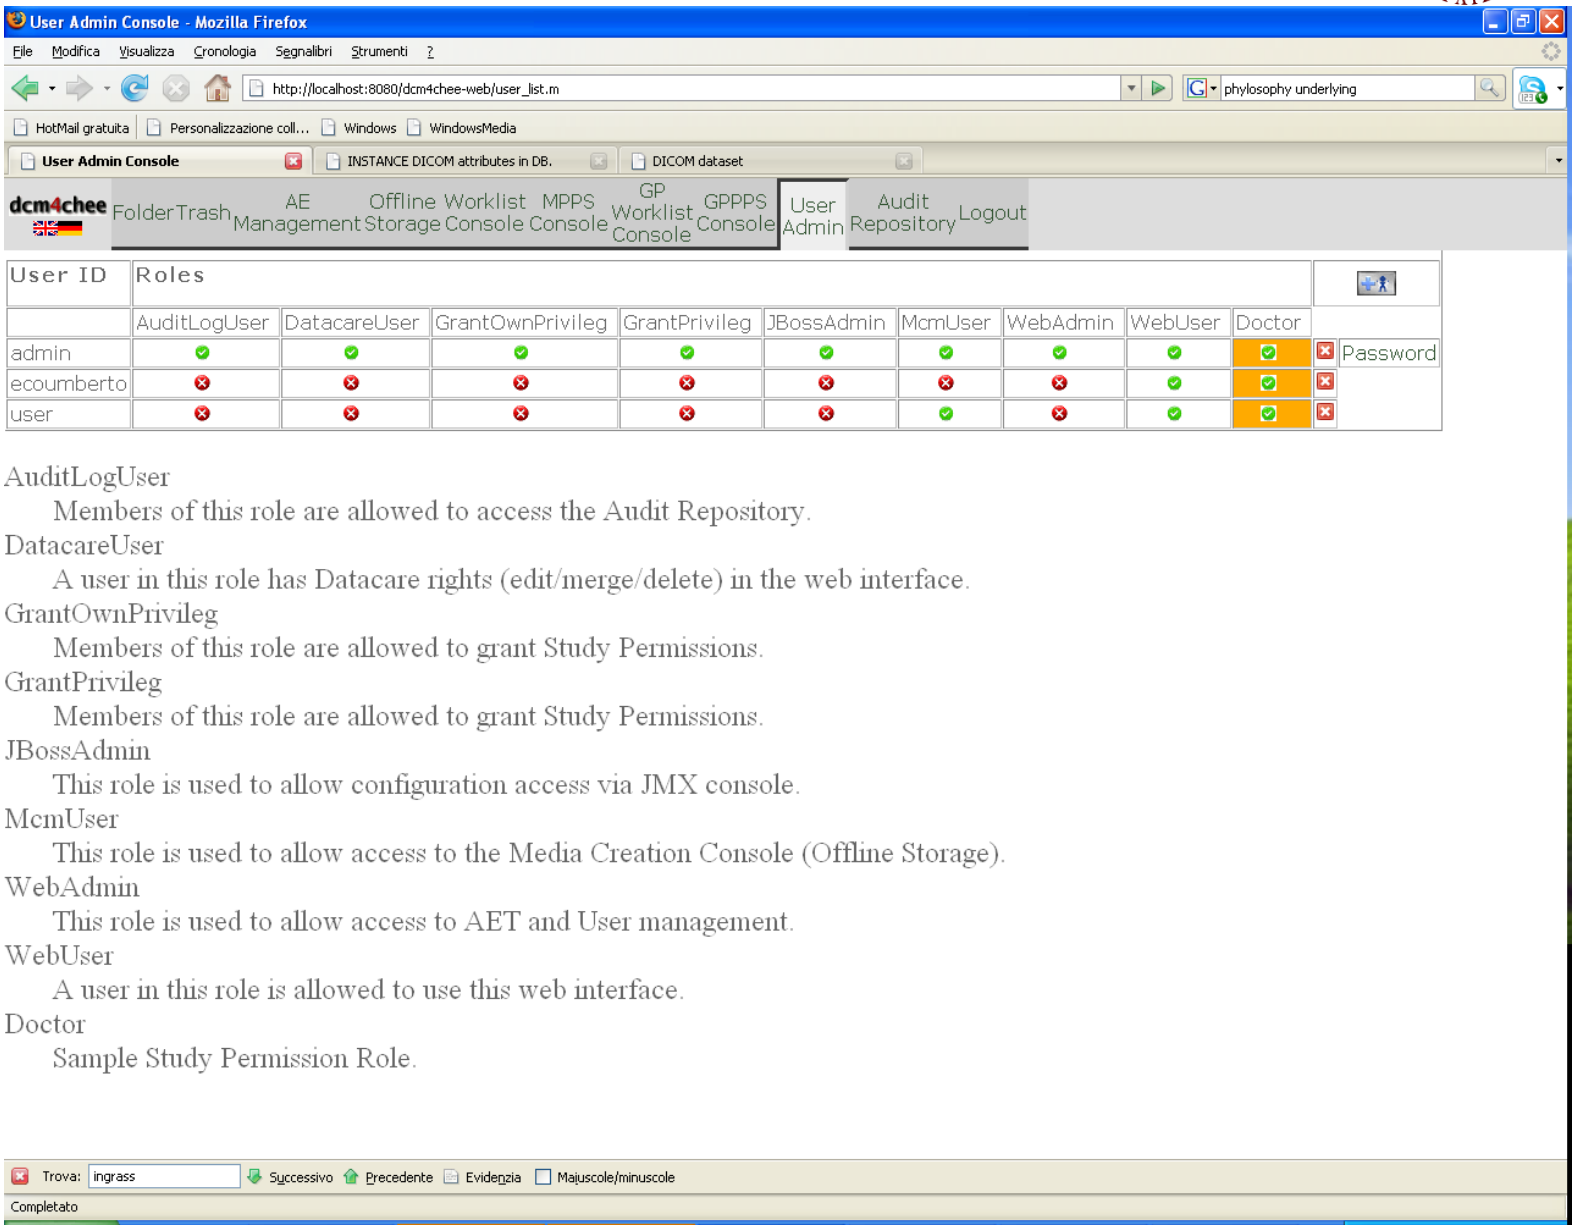

Figure 53. User administration panel. Several roles are available such as Doctor, Datacare User, Web User, Web Administrator, Audit Logger User…

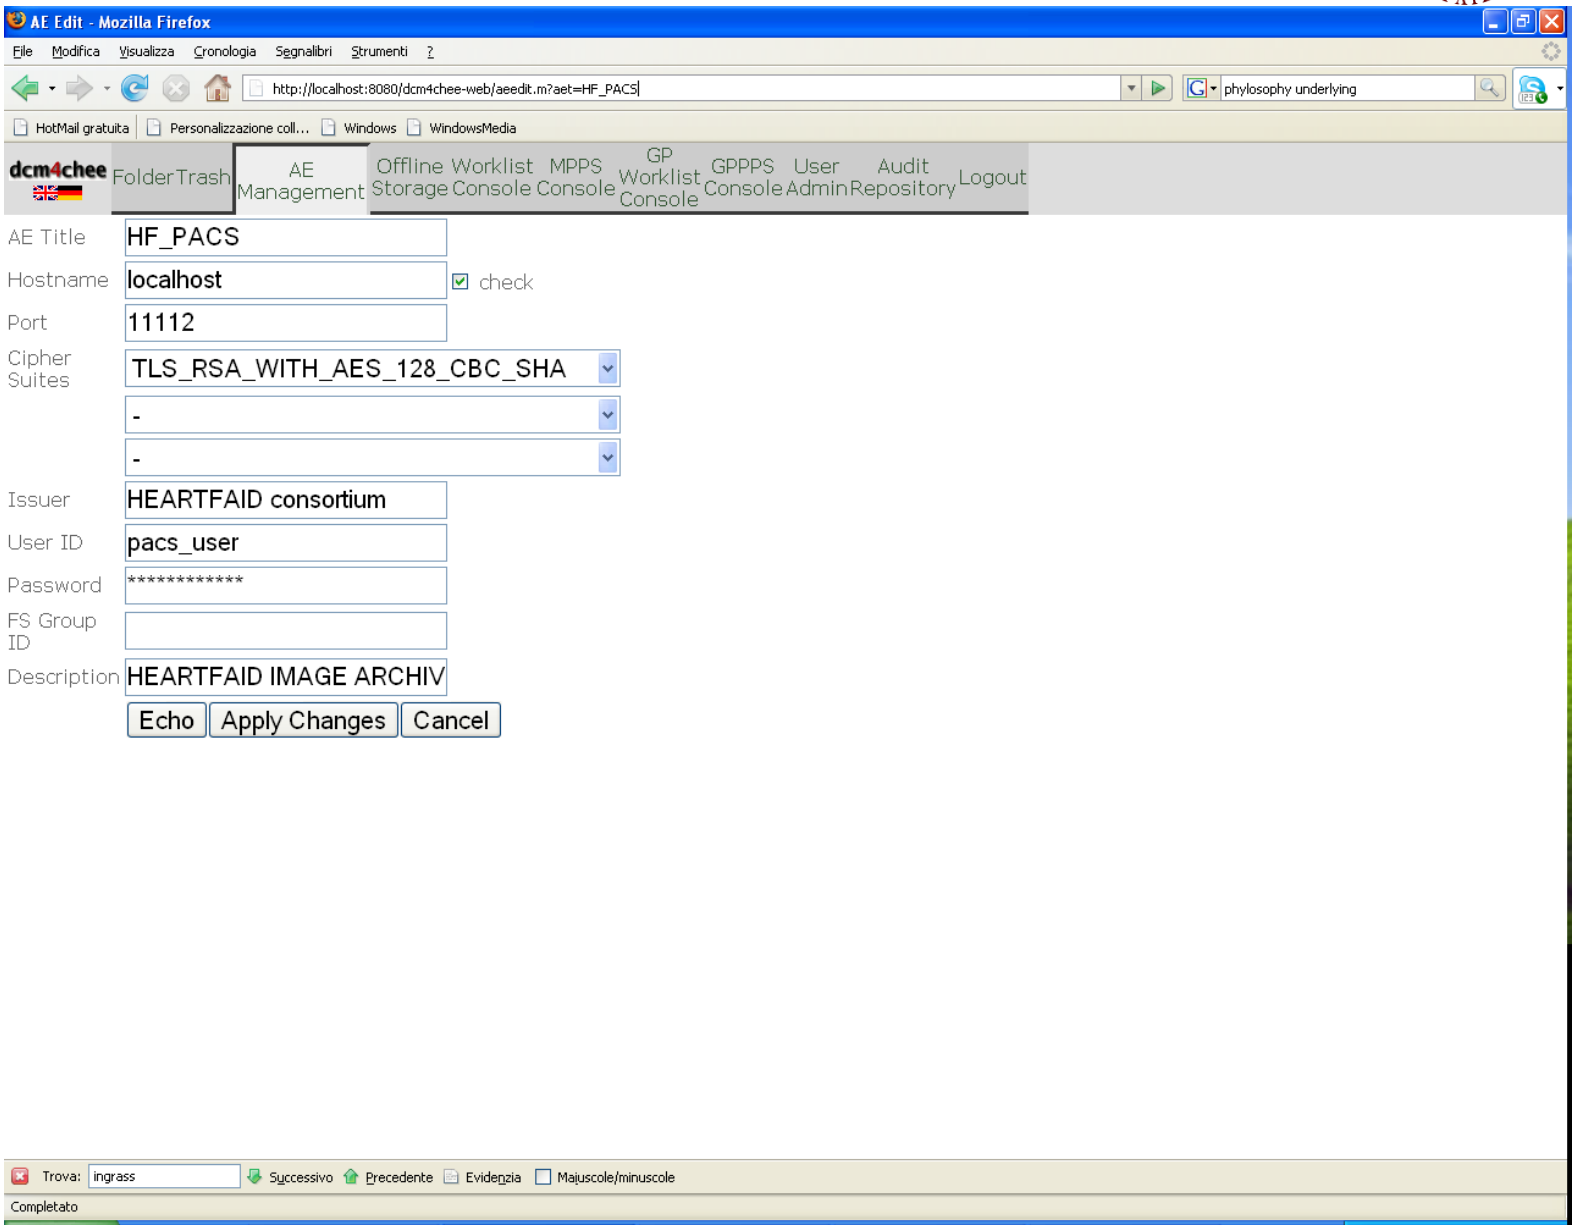

Figure 54. Panel for configuring the DICOM entities allowed to interact with the image archive

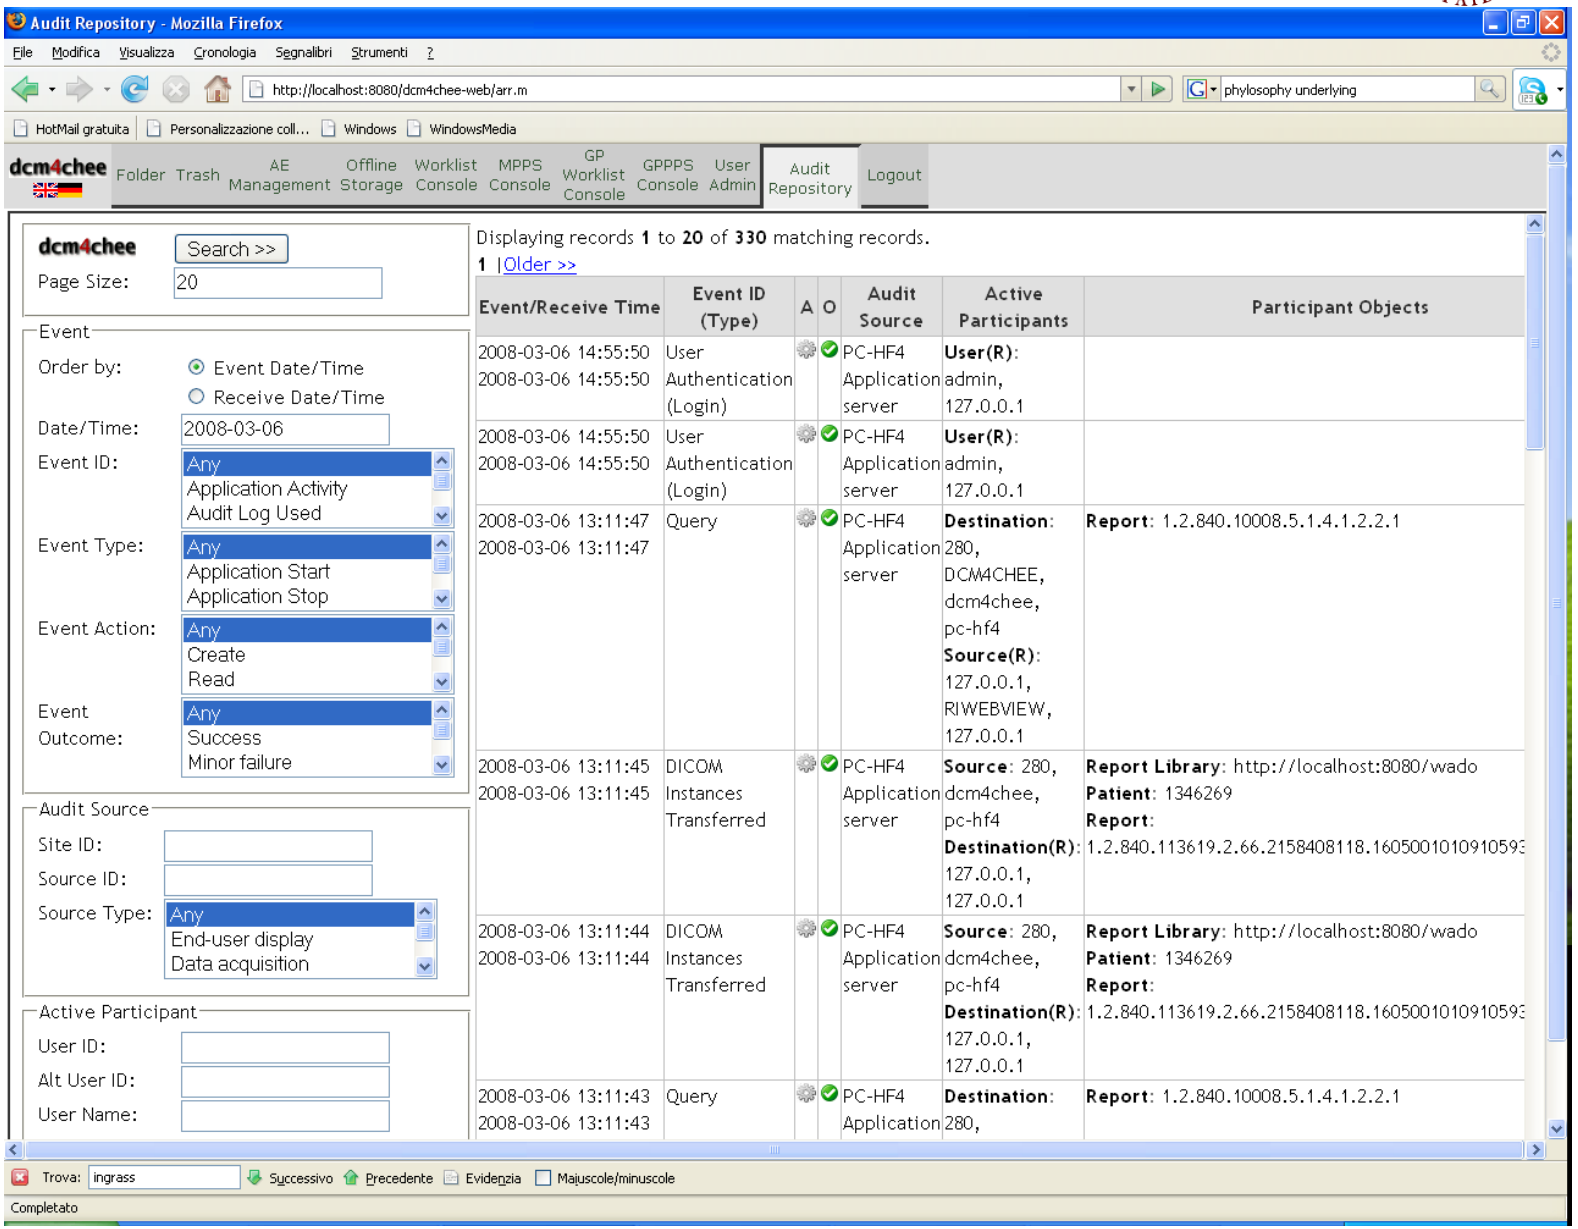

Figure 55. Screenshot of the Audit Repository to monitor accesses and jobs required to the Image Archive

C) Interface for Data Transmission

Clinical partners should be able to upload image to the image archive. If the echo lab is equipped with a PACS connected to the network, connection may be directly established with HEARTFAID Image Archive. However, since this seems not to be feasible in the validation sites, an interface dedicated to image upload should be provided.

The current implementation of the Image Archive has been tested with K-Pacs and Conquest DICOM servers (both acting as clients), DCM4CHE DICOM-send utility and the rather recent ClearCanvas to upload images to the image archive within DICOM standard.

Since K-Pacs is free (although not open source), it could be eventually used for this purposes, setting up a workstation running it in every validation site.

To this end, a so-called DICOM-pack has been prepared and distributed among the interested clinical partners. It contains K-Pacs software, some test images and an internal configuration file. This last file ought to be copied in a suitable directory, making K-Pacs ready to interact with HF Image Archive without any further manual configuration of IPs and DICOM AE Titles. In addition a brief user guide has been prepared. The topics treated in the guide cover K-Pacs installation and configuration, image upload to HF Image Archive, suggested anonymization procedures and basic usage of HF Image Archive web-interface.

The user guide contains text explanations and it is equipped with screenshots annotated with easy-to-understand call-outs. Figure 56, Figure 57 and Figure 58 present some examples taken from the distributed user guide.

Two HEARTFAID validation sites managed to install K-Pacs and upload images. The feedback was positive, also because the use of K-Pacs made easier the internal management of images inside the two validation sites (where no image archive is used for echocardiography but the echocardiography device workstation itself).

Optionally, it is conceivable to develop a web interface calling the standard DICOM C-STORE method (implemented for example in DCM4CHE Toolkit) to upload DICOM images (or more general DICOM media) to the platform from anywhere.

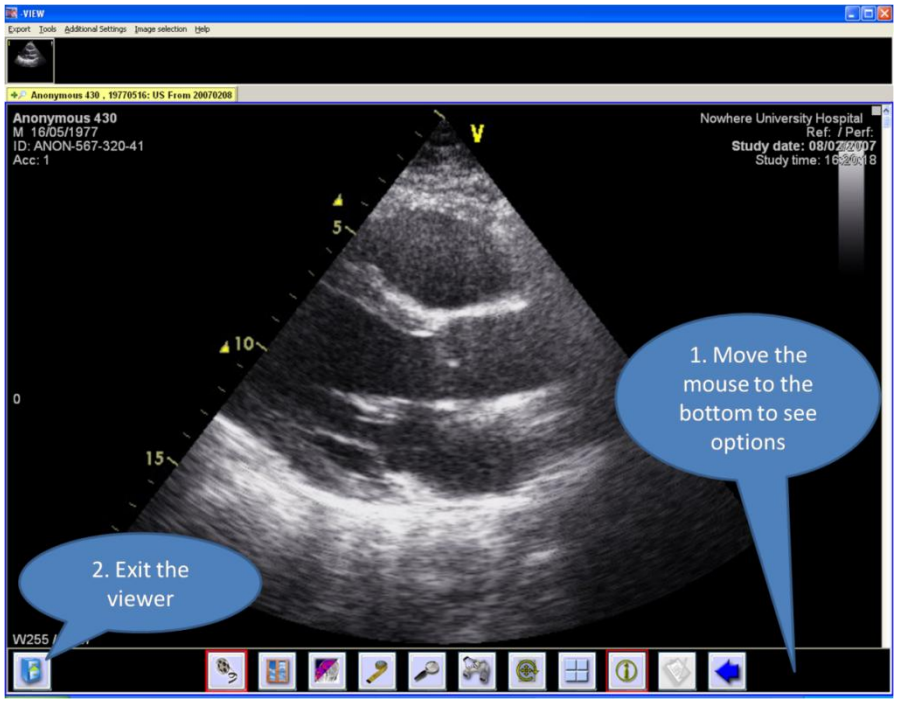

Figure 56. Example of visualization of a test DICOM image contained in the DICOM Pack by using the K-Pacs viewer.

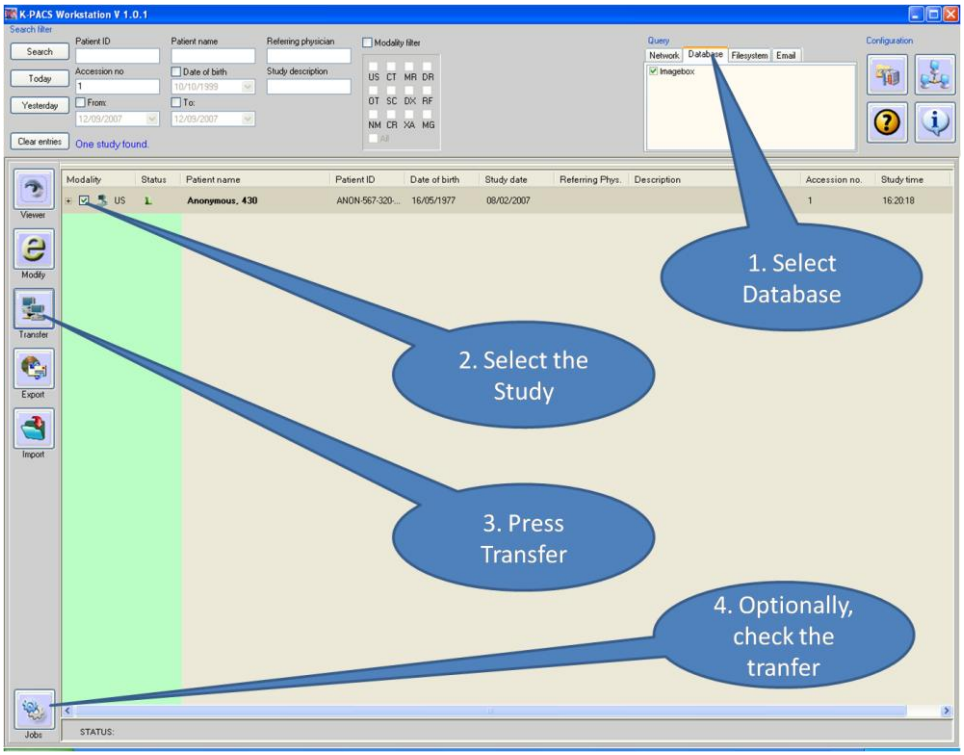

Figure 57. Graphical instructions for sending images to HF Image Archive by using K-Pacs as a DICOM client

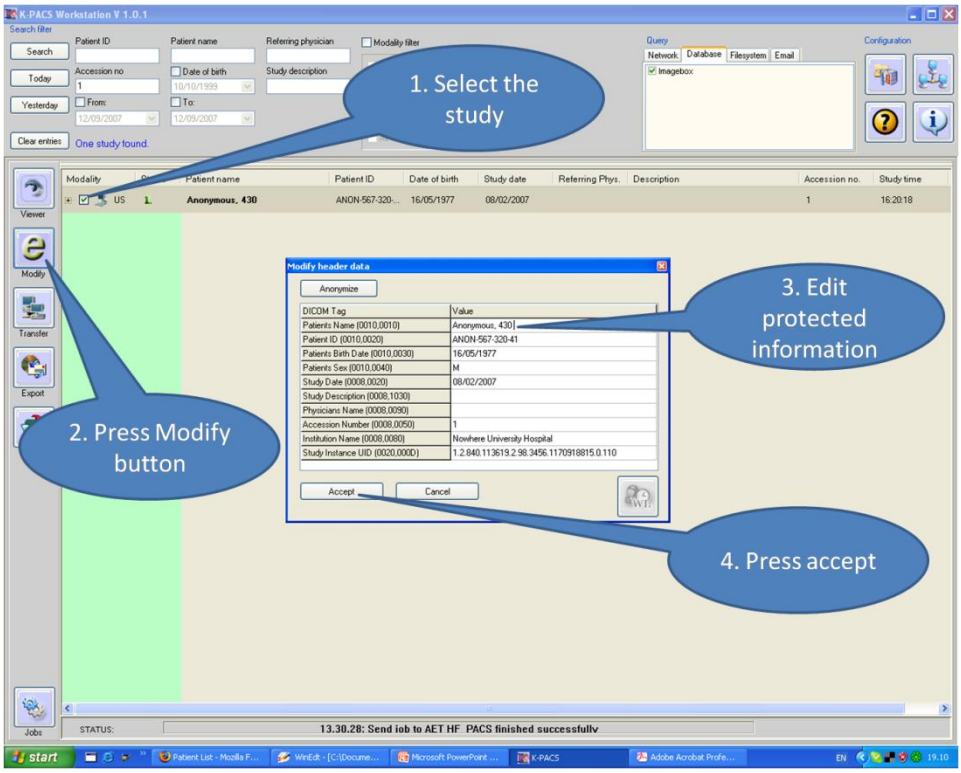

Figure 58. Graphical instructions for anonymizing a study before transferring it to HF Image Archive

D) IHE Retrieve Information for Display (RID)

IHE Retrieve Information for Display (RID) services are implemented in the current installation. In this case the information source is limited to the content of DICOM archive and therefore only allows access to DICOM documents (Structured Reports, Waveforms (RID-ECG) and encapsulated PDF). The enumerated DICOM objects can be queried via RID retrieve information request and are rendered (and cached) as PDF documents (using Apache FOP) for a RID retrieve document request.

DICOM waveform objects are pertinent to HEARTFAID domain and due to the resolution requirements of ECG RID profile rendered as vector graphic (both PDF or xml+svg are supported as MIME types).

The syntax is similar to the one employed for WADO services. For example the URI:

http://heartfaid.isti.cnr.it/IHERetrieveDocument?requestType=DOCUMENT&documentUID=1.3.6.1.5.6.6079.3.989&preferredContentType=application/pdf

allows to retrieve and visualize the PDF document which is rendered in Figure 59.

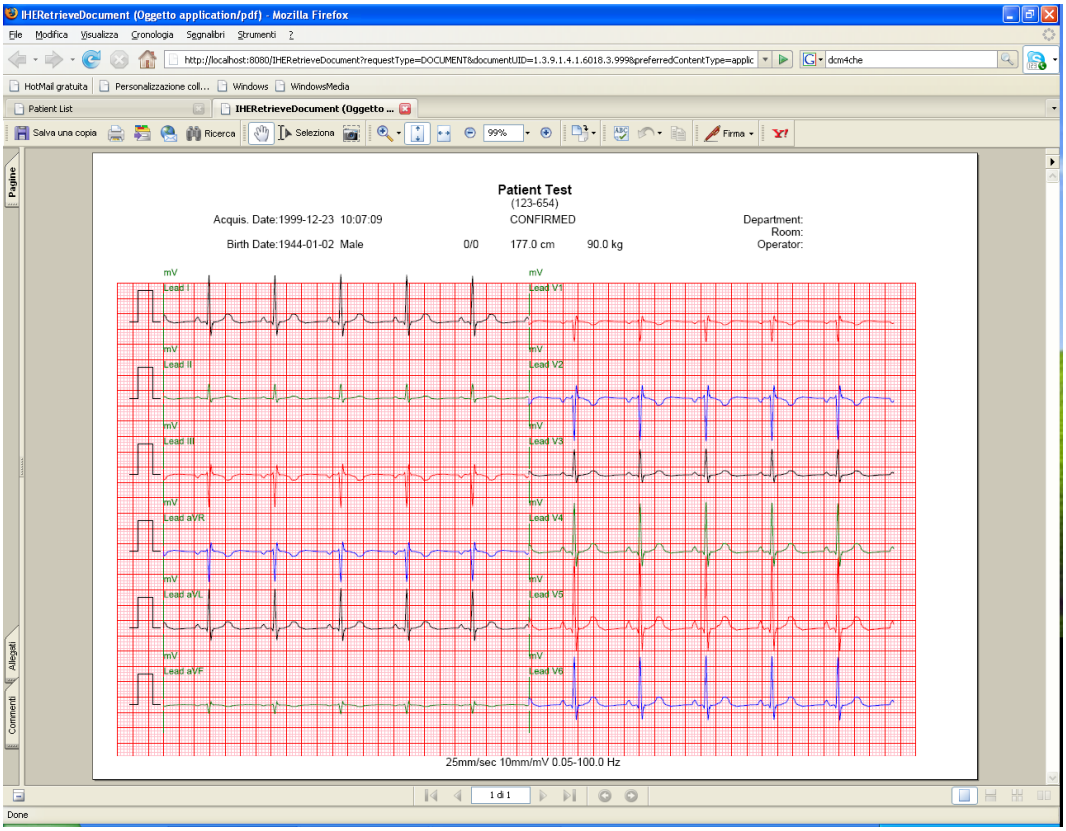

Figure 59. IHE RID of an ECG. The ECG profile is rendered as a vector graphic object and encapsulated in a PDF.

Retrieval of the SVG pictures for rendering inside an integrated web interface is possible by using an URI like:

http://heartfaid.isti.cnr.it/IHERetrieveDocument?requestType=DOCUMENT&documentUID=1.3.6.1.5.6.6079.3.989&preferredContentType=image/svg%2Bxml

Notice that these URIs are used here only for description purposes and do NOT correspond to the actual ones used in the platform.
Finally, by contrast with SCP ECG, it seems that the DICOM waveform format is not suitable for the encoding of the information produced by the implemented signal processing tools.

E) Final remarks

As a final remark, it is worthwhile to point out that DCM4CHE community is pursuing a new child project, called Xero, that provides a new web interface for clinical access to patients and studies (as opposed to the Web-based user interface described above which was mainly intended for administrators). This component is intended for users such as nurses, doctors and perhaps patient's relatives who can't easily install a full radiology

client station. This web-interface should in principle give an answer, among others, to some of the interface problems described in the previous subsections, i.e. user administration, data query/retrieve, data transmission, consistent display of images, display of multiframe images (via CINE), numeric image reports and final reports. Preliminary releases of Xero have been tested at the CNR, but the results (although promising) were not satisfactory and it is not advisable to deploy the current version of Xero inside the platform.

4.2 Integration of image processing results in the Image Archive

The image processing methods described in Chapter 2 may be exploited to obtain, from the images acquired by some modalities, a new set of annotated images together with clinical parameters extracted from them. An exemplifying situation is the computation of left ventricle ejection fraction. While the final result of elaboration is a concise parameter (a percentage), it's computation is obtained on the basis of a set of segmented images (i.e., roughly speaking, a set of images on which the left ventricle borders have been superimposed).
Besides storing the ejection fraction in the patient folder, it is reasonable to keep also track of the intermediate segmented images, since the computed parameter strongly depends on the segmentation accuracy. Thus, saving the final results of image processing in the form of annotated images is important for archiving and reviewing purposes.
The goal of this Section is precisely to explain how the image workflow may be seamlessly integrated in HEARTFAID platform using the available HEARTFAID Image Archive features.
First of all, let's recall that DICOM defines the Secondary Capture (SC) modality. This modality is specially designed to embed the results of image processing (ranging from the application of enhancement filters to more complex image processing procedures) into a DICOM image. The header of the DICOM-SC image may replicate the patient personal information contained in the original DICOM image which is used as input of the image processing algorithms. Further, the header may be used to add a reference to the original DICOM study: in this way the original images and the processed ones are persistently linked together within one DICOM study.
However, when DICOM-SC is used for storing the results of an image processing task such as segmentation, a major limitation is represented by the impossibility to *edit* the segmentation after this has been exported to DICOM-SC. This is because the segmentation (usually represented in a vectorial way inside the software running on a CAD workstation) is burnt into the pixels of the image during the export to DICOM-SC, thus losing the possibility to edit (or even remove) it after exporting.
Although we are aware that DICOM standard will in the future handle segmentation results in the form of meshes and that some DICOM supplements are in an advanced status of preparation (such as DICOM Supplement 132 which aims at defining the so-called Surface Segmentation Storage SOP Class), we believe that the ratio between the benefits for the platform and the effort required to actually implement these new (not

yet standard) DICOM storage classes is too low. Further such kind of work on the implementation of integration standards is far beyond the scopes of WP5, if not out of topic.
Another interesting object specified by DICOM is the Structured Report (SR). which may be used to embed the study report (and in particular the computed parameters). Such representation model is not used in the platform since other already available interfaces and modules replicate partially its functionalities (see Section 5.2).
We refer to (Zhou et al. 2007) for other interesting schemes for the integration of image processing results, also involving IHE integration profiles.
This being the available technical instruments, the proposed solution is depicted in Figure 60. In particular, DICOM network services should be enabled in the image processing workstation; indeed the workstation should be able to retrieve the original images from the Image Archive and, after processing has been accomplished, the resulting images should be stored in the image archive in DICOM-SC format. In general, one may rely on a PACS client application for accomplishing these data transfer procedures (e.g. K-Pacs or ClearCanvas) or, better, DICOM functionalities may be embedded inside the image processing interface.

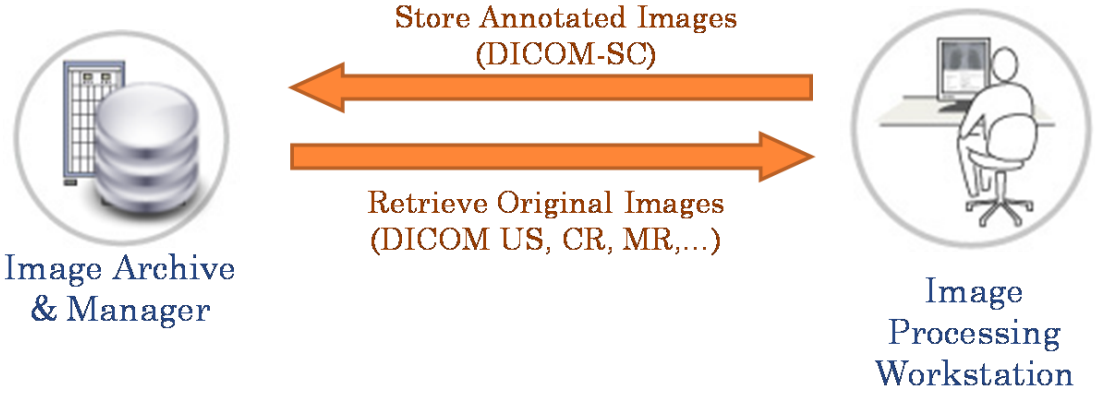

Figure 60. Overview of the integration of the image processing results in the platform by means of the Image Archive

The current prototypical version of the algorithms for the segmentation of the left ventricle and the computation of ejection fraction, developed in MATLAB, is on the way to such an integration. Indeed:

1. The segmented images are saved in DICOM-SC format with a meaningful header. The header may replicate the personal details of the patient contained in the original images and other pieces of information which are not altered during processing. A new series UID is associated to the segmented images, while the study UID (if available in the original images) is kept. Such procedures are handled with routines contained in the MATLAB Image Processing Toolbox.
2. The segmented images are sent to the HEARTFAID Image Archive directly from the image processing application. This is obtained by calling from a MATLAB routine some classes of the JAVA implementation of DICOM provided by the DCM4CHE toolkit.

4.3 Interfaces for Image Visualization and Processing

The purpose of this section is to discuss web and desktop interfaces for image visualization and processing. In particular, we motivate and describe an already implemented interface for web visualization in Section 4.3.1, while in Section 4.3.2 we discuss its possible extension to a lightweight image processing interface for web annotation of images. Since not all image processing algorithms may run in the browser sandbox, possible solutions for this kind of algorithms are briefly presented in Section 4.3.3.

4.3.1 Web viewer interface

A web interface for image visualization should answer to the following needs of HFP:

- Reviewing of images by the referent physicians
- Reviewing of images for second opinion
- Quick access to image data (for example access to data of the previous examination in the same room where the new examination is carried out to appreciate changes in the clinical situation)
- Quick selection of images for post-processing workflow. For example, a physician may want to compute again a) left ventricle ejection fraction or b) end diastolic LV diameter. In case a) very likely he will select an image sequence taken from an apical view, while in case b) he will select a particular M-mode image. In any case, he needs to select quickly and easily a suitable image from the bunch of images in the patient's study.

WADO services, implemented in the HEARTFAID Image Archive, are of course useful, but are not sufficient to deal with image sequences, since no kind of streaming is implemented.

It would be useful to add some features to WADO services, so as to include video streaming in some format, but this depends tightly on the future evolution of DICOM standard.

Motivated by these needs, we have started to work on a new, separate but interoperable web viewer interface.

The goal of this interface is to give to physicians the opportunity to quickly access images and patient data stored in a DICOM Image Archive from anywhere by using a browser. Figure 61 tries to depict the general philosophy underlying the introduction of the web viewer service. In the context of the platform, the web viewer interface mitigates the unnecessary complications of the native interface of HEARTFAID Image Archive and provides more advanced viewing options.

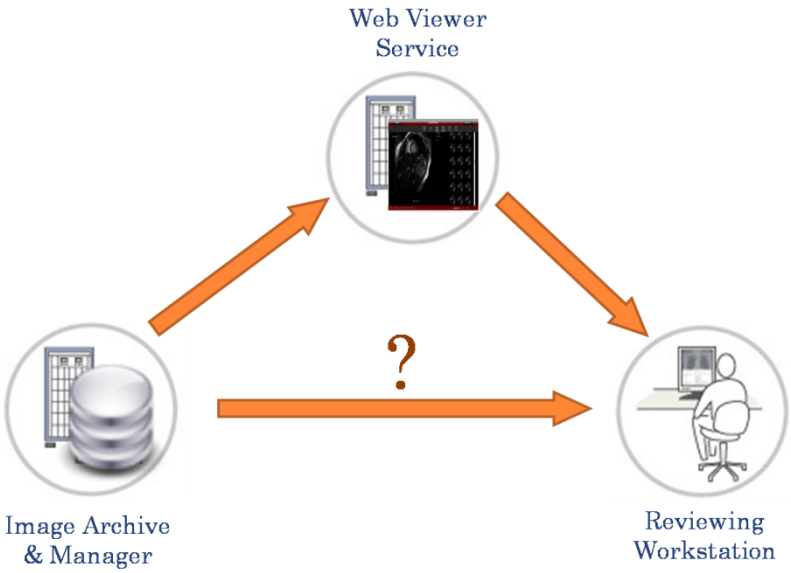

Figure 61. The Web Viewer service provides in an interoperable way a bridge between any DICOM compliant Image Archive and the reviewing workstation

Although up to now, this new web viewer interface is connected to HEARTFAID Image Archive only, in principle it could be connected to any DICOM compatible picture archive equipped with WADO services. Connection to multiple archives is also possible, thus broadening its application field. A simplified version of the flow involving this piece of software is shown in Figure 62.

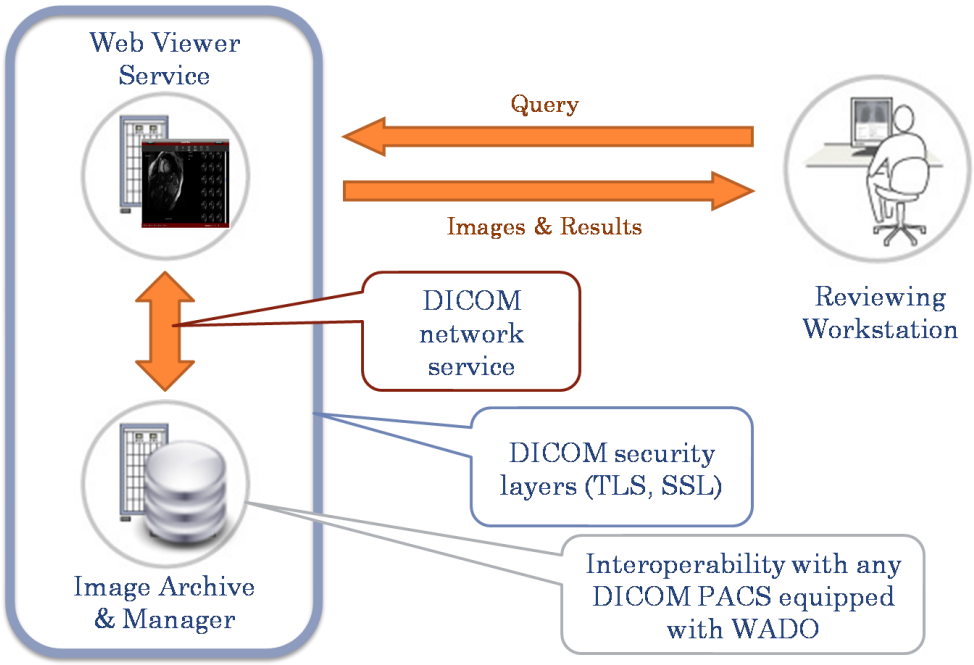

Figure 62. Simplified transactions among the 3 actors of Figure 61

In particular, the transactions between the webviewer and the Image Archive are performed according to DICOM network services and, as such, they can be protected by DICOM compliant security layers (TLS, OpenSSL,…). This interface will be eventually integrated in the HEARTFAID portal. Some of the screenshots of the current version are shown in Figure 63, Figure 64 and Figure 65 (see the following figure pages).

4.3.2 Lightweight image processing web-interface

In principle, it is possible to extend the web viewer interface to allow the possibility to perform common linear and area measurements on an image, directly inside the web browser. Besides the already available pan and zoom tools, this extension requires the development of the tools described in Section 2.3.1 and 2.3.2 and specifically of a "ruler" and a "free hand region selection" tool.
This sort of graphical tools is already employed in several existing applications, showing the feasibility of the solution. Further, conversion from image units to physical units can be achieved through the retrieval of the DICOM header. The retrieval of this header does not pose any further difficulty, since it can be obtained in a way very similar to the retrieval of an image. Actually, a WADO request with XML as content type is sufficient (see Section 4.1.2, *Web services* paragraph).
The performed measurements, suitable associated to tags, can then be stored in the HEARTFAID repository by standard JDBC or, better, via HL7 messaging.
Annotation of ECGs inside the same web interface is also conceivable. Indeed, by the IHE RID service (see Section 4.1.2), an ECG may be visualized via a web browser as a vectorial SVG image. Annotation through a ruler and subsequent result storing requires only minor modifications with respect to image processing.

4.3.3 Desktop image processing interface

Although full web-based diagnostic image analysis is appealing, it is important to notice that some considerations prevent from taking for granted the feasibility of this approach. Actually, image processing algorithms may be very demanding in terms of computational power. Further, since an echo study may be bigger than 100Mb. transferring the original images, the intermediate results of processing (for "nearly real time" user feedback) and the final results may be cumbersome. Thus, computational intensive image processing algorithms are not suited to run inside the browser sand-box. For this reason, it is likely that the prototypical image segmentation algorithms described in Section 2.2 will be difficult to integrate, also in their final form, in a web-based interface. Nevertheless desktop solutions are not difficult to obtain. Our intention is to convert the developed prototypical algorithms from MATLAB to JAVA (exploiting perhaps the JAVA compiler included in the last versions of MATLAB). After this stage, it will not be difficult to integrate them with already available open source interfaces for image analysis, such as ImageJ.

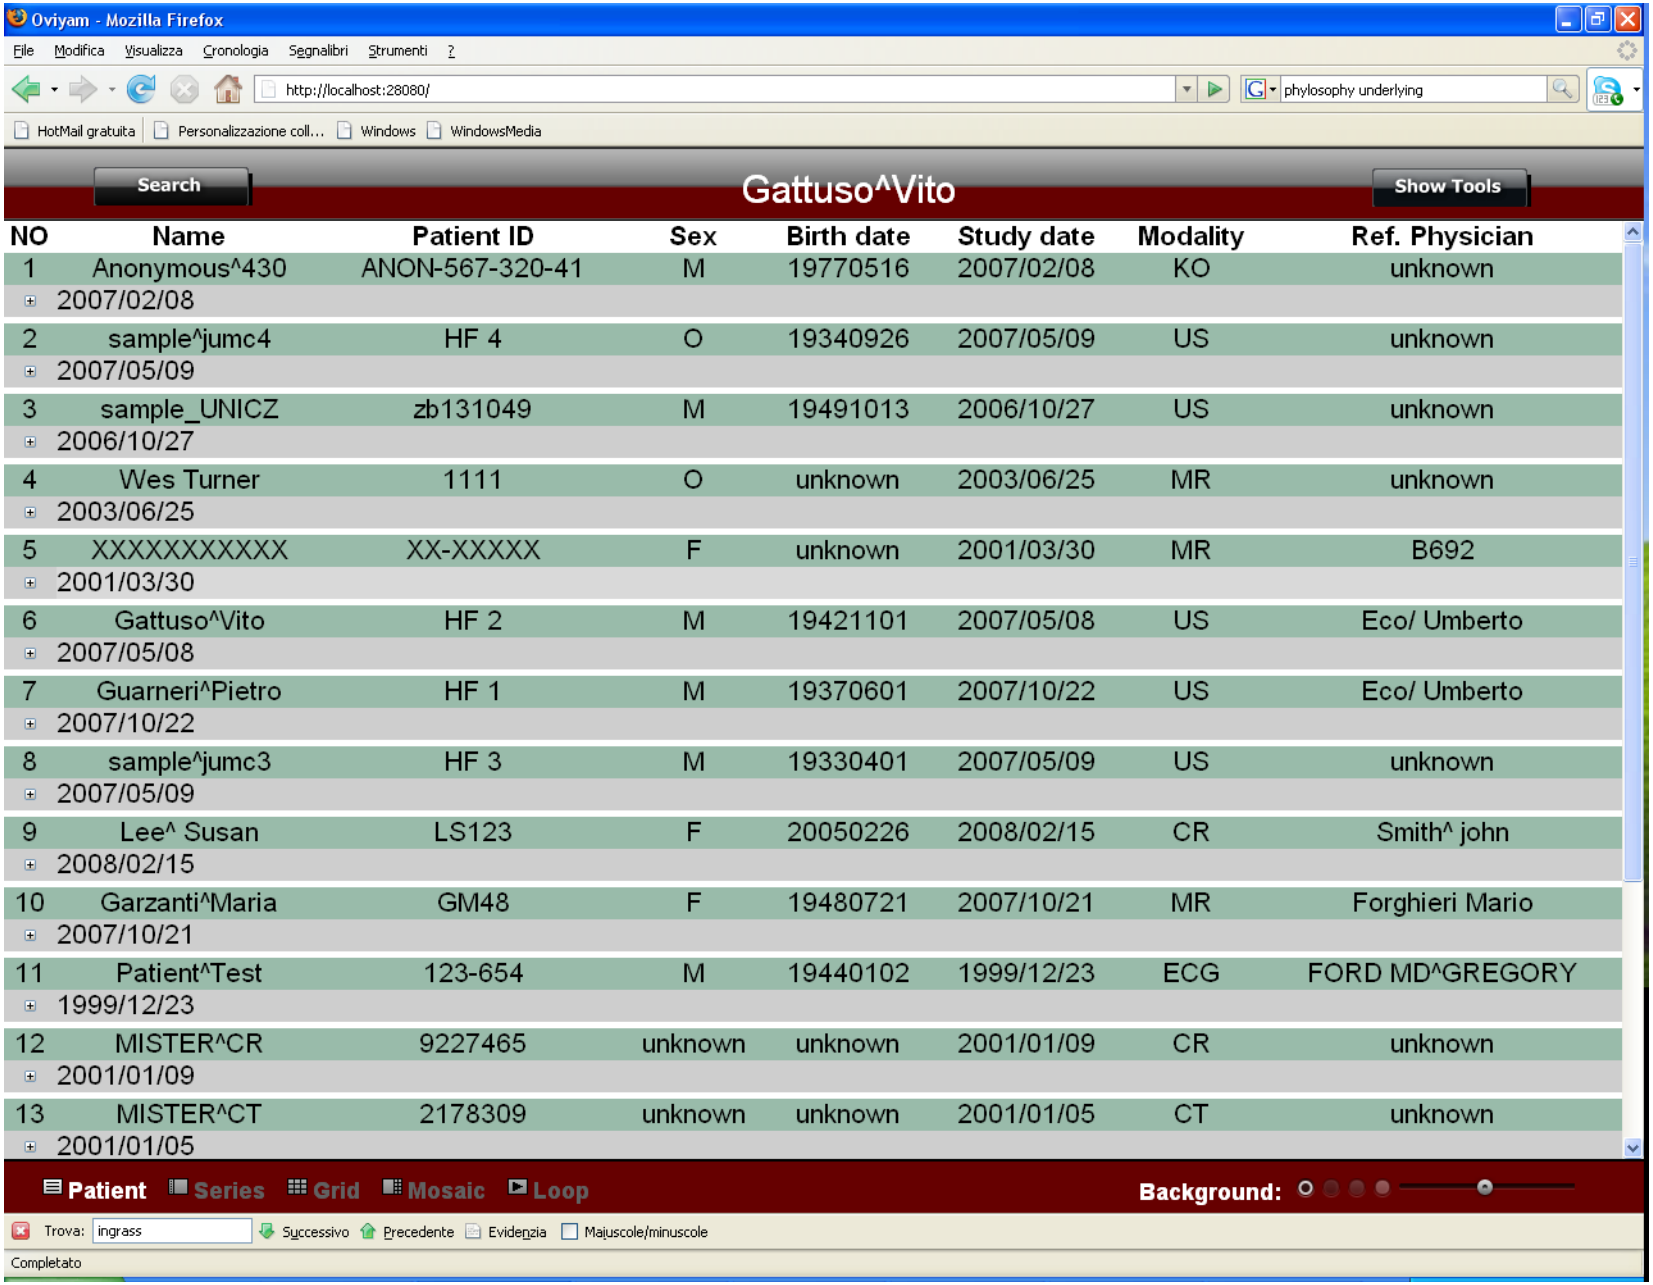

Figure 63. Patient search results in the web viewer interface

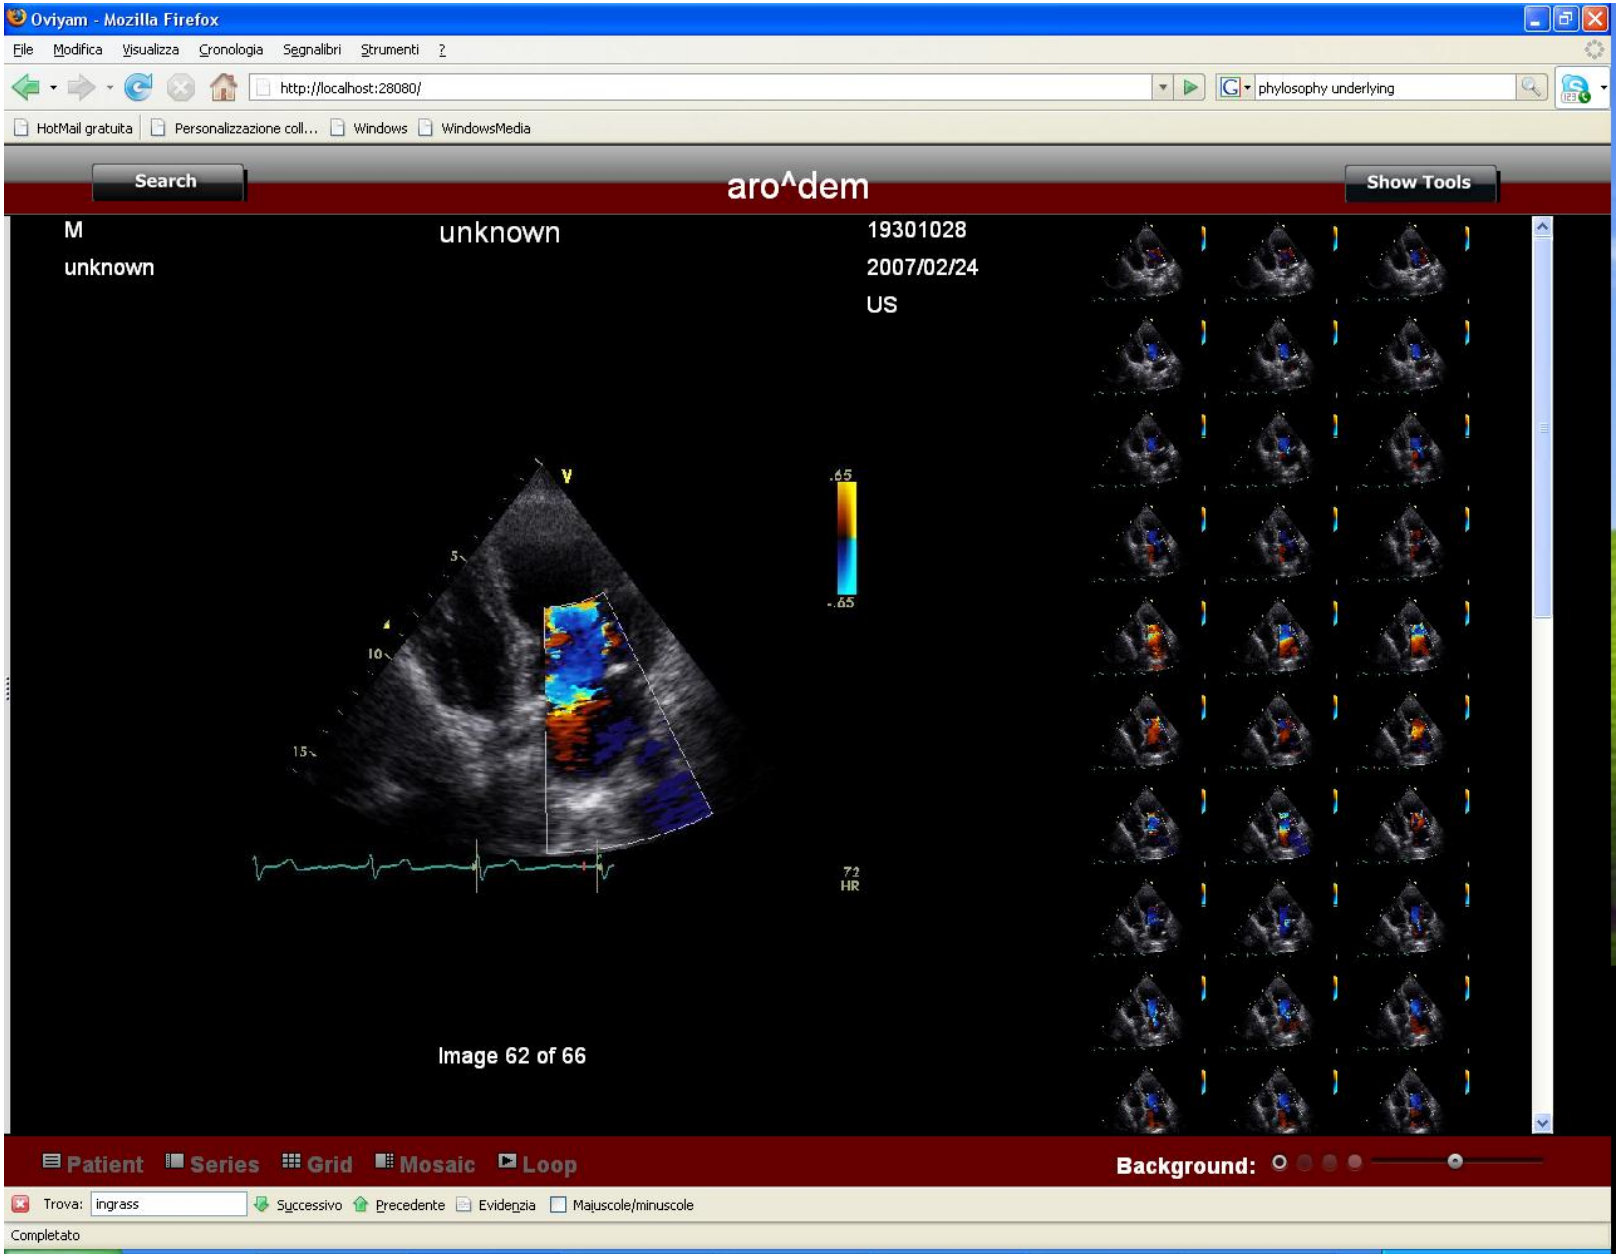

Figure 64. Visualization of a Doppler image sequence. Thumbnails on the right column allow to quickly select images.

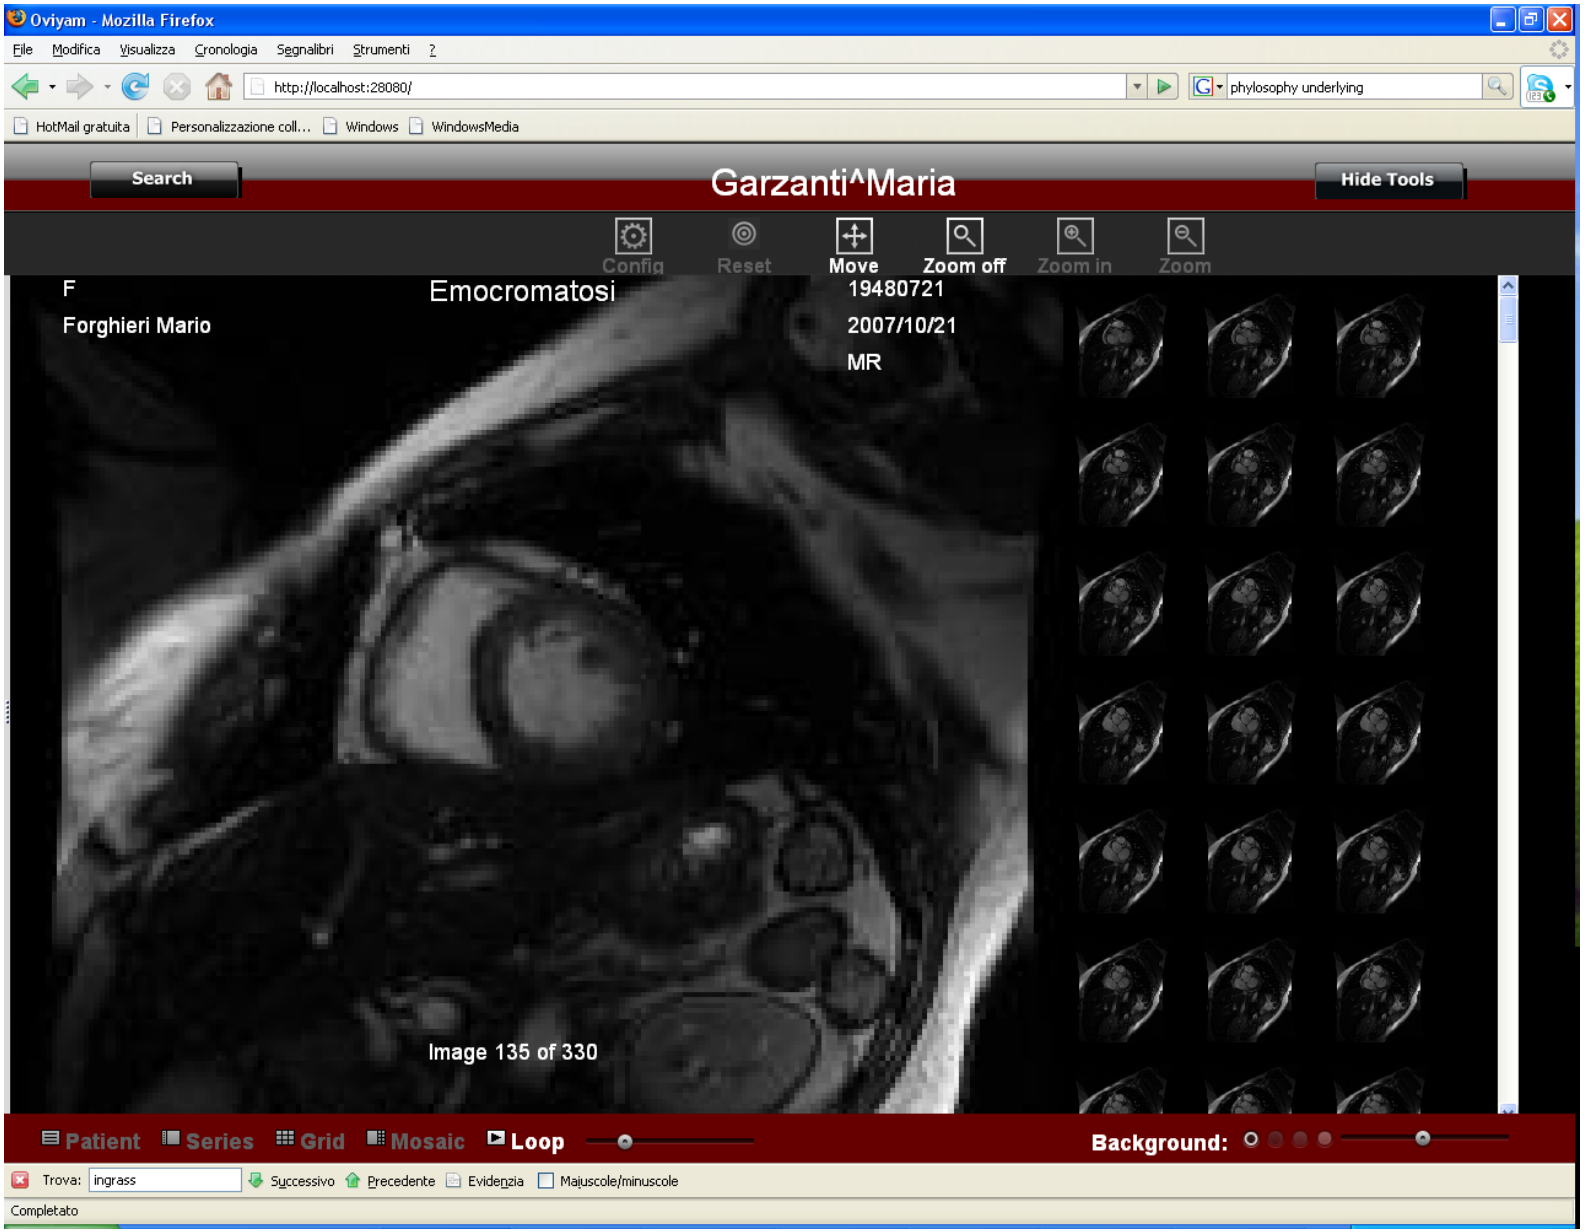

Figure 65. Visualization of an MRI study. The provided pan and zoom tools are used here to focus on the heart region. Finally loop CINE option is activated for visualizing a *movie* of the heart. The frame rate can be adjust by a simple mouse interactions.

4.4 Selected Standards for ECG Storing and Transmission

The ECG files collected by the Archimed 4210 are provided in a dialectal implementation of the SCP-ECG format and they are transferred via floppy disk to the hospital gateway. In the hospital gateway the acquisition program is installed and the routine procedure is to copy the ECGs from the floppy disk in a working directory where the acquisition program will retrieve them. The acquisition program is able to show a list of the ECG files present in the working directory with a summary of their significant associated information.
The file can be viewed and processed with the above described algorithm in order to enrich the information contained in the SCP-ECG file with the averaged dominant beat. Then the ECG file can be encapsulated in an XML container and sent to the XML-based data repository. The ECG viewer is available at any moment for visualizing (and printing when necessary) the ECG in a graphical format. Thus the standard for the ECG storing and transmission is SCP-ECG encapsulated in an XML container.

5. Data Processing for Decision Support

ECG signals, echocardiographic and chest X-ray images processing capability has been considered as a fundamental functionality of HEARTFAID CDSS, in accordance to its relevance in HF patients' management (as extensively discussed in the deliverables D5 and D15 and in the previous sections of this document).
In this chapter, focus is given to the value that signal and image processing adds to the decision support services of HEARTFAID platform; this means how these processes can be integrated within the physicians' workflow and can, then, be employed to assist their important decisional processes. In particular, a specific example, drawn from the *Patient's Management Showcase*, is described in more detail for highlighting just this kind of contributions. Algorithms integration within the CDSS architecture is discussed as well.

5.1 Value added to the CDSS

Signal and imaging investigations are currently a basic step of the diagnostic, prognostic and follow-up processes of heart diseases. Not by chance, in the last decades, the development of CAD schemes has attracted a lot of interest and effort within medical imaging and diagnostic radiology, becoming in some cases a practical clinical approach. The basic concept of CAD is to provide a "second opinion" or a "second reader" that can assist radiologists' by improving the accuracy and consistency of image based diagnoses (Kunio Doi, 2007). As stressed in previous discussions about the usefulness of data processing, the clinical interpretation of diagnostic data and their findings largely depends on the reader's subjective point of view, knowledge and experience. Hence, computer-aided methods, able to make this interpretation reproducible and consistent, are fundamental for reducing subjectivity while increasing the accuracy in diagnosis. As such, they are likely to become an essential component of applications designed to support physicians' decision making in their clinical routine workflow. Other important motivations rely on the limits to reader's ability of data interpretation caused by either the presence of structure noise or the vast amount of data, generated by some devices, which can make the detection of potential disease a burdensome task and may cause oversight errors.
Such considerations drove the choices made while designing the HEARTFAID decision support services. While formalizing the main decisional problems that require the CDSS intervention and, hence, listing up all the pieces of knowledge, data and information relevant for decision making, the importance of considering and interpreting ECG signals and Echo images had come forth.
The first and immediate proof of this relevance is HF diagnosis problem. Indeed, it can be considered as the first stage of HF patients' management which necessarily requires the acquisition and analysis of signal and imagery data. Figure 66 shows the sequence of steps that compose the HF diagnostic workflow: after having assessed the presence of

main signs and symptoms, physicians usually require diagnostic examinations such as ECG, chest X-ray and neuroendocrine evaluations (i.e., Brain Natriuretic Peptides – BNP) in order to check out the diagnosis, confirmed eventually by an echocardiography investigation. Supporting such a decision problem means encoding the workflow into an opportune knowledge base which formalizes, for each step, the set of conditions evaluated by physicians.

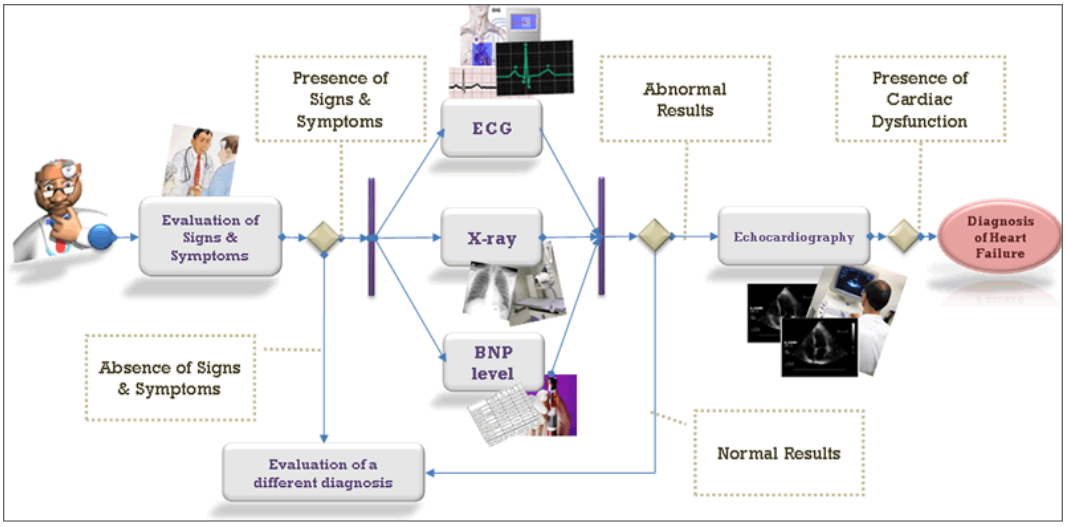

Figure 66. Heart Failure diagnostic process

The first step regards the presence and severity of signs and symptoms such as breathlessness, swelling, fatigue, hepatomegaly, elevated jugular venous pressure, tachycardia, a third heart sound, pulmonary crepitations. Then, ECG signals are acquired for investigating the presence of anterior Q waves and a left bundle branch block, signs of left atrial overload or left ventricular hypertrophy, atrial fibrillation or flutter, ventricular arrhythmia. If ECG abnormalities are present, HF diagnosis is considered carefully possible and further checked out by analyzing chest X-ray. Such an examination is useful for detecting the presence of cardiac enlargement (cardio-thoracic ratio > 0·50) and pulmonary congestion. In parallel, neuroendocrine analysis are performed to test out high levels of natriuretic peptides which suggest the presence of a cardiac disease. Whether the three of these examinations certify the presence of abnormalities, an echocardiographic investigation is performed for documenting a cardiac dysfunction. The most important parameters to be evaluated from such a diagnostic modality is LVEF (n.v. >40%); other relevant data are the fractional shortening, sphericity index, atrioventricular plane displacement myocardial performance index, and left ventricular wall motion index, isovolumic relaxation time, early to atrial left ventricular filling ratio, early left ventricular filling deceleration time, pulmonary venous atrial flow velocity duration and ratio of pulmonary vein systolic and diastolic flow velocities, pulmonary artery pressures. HF diagnosis is finally concluded if symptoms and signs and ECG / X-ray / BNP level / Echocardiographic abnormalities are all present.

The above description underlines the relevance of formalizing diagnostic conditions on features extracted for signal- and image-based examinations. It is likewise evident the

importance of computing such features as objectively as possible by avoiding variability of human interventions in tracing, for instance, linear segments, as required for computing the cardiothoracic ratio, or left ventricle borders, as required for evaluating LVEF.
Some acquisition devices already offer the possibility of automatically computing a set of relevant parameters but are still really expensive and this is the reason why older devices are still very common. In sites with such devices, a great benefit can be obtained from data processing methods applied after the acquisition.
HF diagnosis was a straightforward example of the importance of computer-aided data processing in HF decision making, but other significant contributions can be envisaged.
Overall, among all the profitable applications into decision support workflows, the following can be listed up:

- automatic or semi-automatic computation of parameters relevant in the decisional problems;
- support of physicians' case-based reasoning processes;
- discovery of novel pertinent knowledge.

While the first is typical of routine workflows in relatively simple situations as described in the previous example, the other two can be considered advanced applications that can aid physicians in facing critical cases or critical problems. Actually, not only the parameters extracted from signals and images examinations are significant to physicians for formulating a response but also the data themselves can be useful for having a general overlook of a patient's situation. This means that allowing clinicians to explore data can assure the availability of a lot of other pieces of information hidden in the same data. Moreover, when dealing with a difficult case, comparing the one at hand with assessed responses for other patients' situations can be really helpful. This entails maintaining and making available a database of cases with annotated images and signals which can be retrieved by similarity on a set of computed features (think then to the HEARTFAID Image Archive). Difficult diagnoses and, most of all, prognosis assessment are examples of these situations. Indeed, HF prognosis is a critical and still open problem which has been deeply investigated within HEARTFAID. The clinical partners described the routine currently followed and highlighted its limits. A set of parameters are presently considered relevant; they belong to five main categories (Figure 67): textual, logical and numerical data about the patients (i.e., demographic, clinical and functional); signal data (i.e., electrocardiogram recordings and correspondent parameters); imagery data (i.e., 2D Transthoracic Echocardiography); video data (i.e., multiframe images from Transthoracic Echocardiography). Graphical 3D models could also be considered as part of the relevant clinical data, think for instance to the 3D reconstruction of the left ventricle used for assessing the haemodynamic dysfunctions and impairment of the heart. Unfortunately, at present only some variables are considered as independent and consistent predictors of negative outcome, as reported in Table II borrowed from deliverable D5. Hence, also for prognosis assessment, allowing physicians to explore data that belong to known cases can aid their decision. Plus, for such a critical problem, data processing facilities can have further relevance for the discovery of novel

knowledge by granting the computation of a wide range of parameters which can be explored and put in relations in order to find out new relevant patterns.

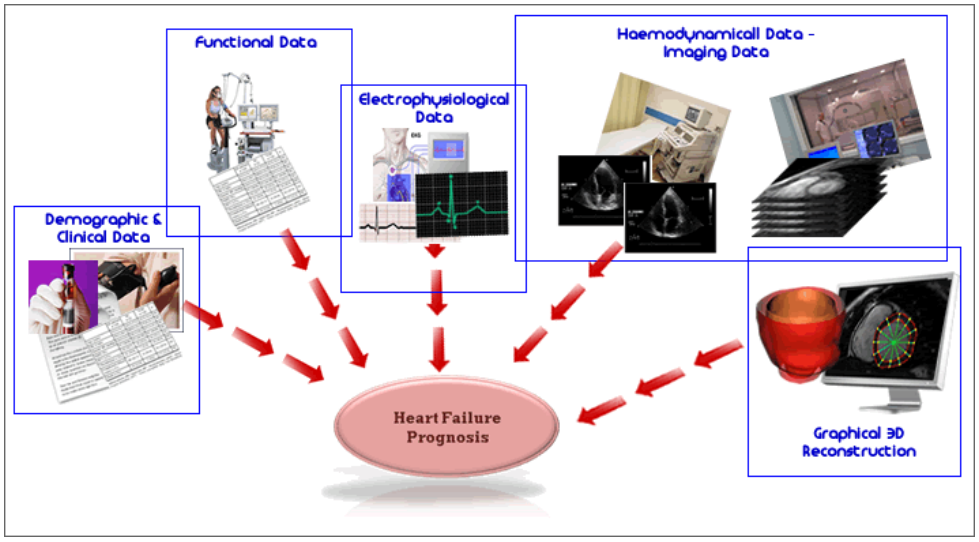

Figure 67. Data relevant for prognosis assessment

Integration within decision support is important not only for the value that signal and image processing can add to physicians' decisional workflows, but also, from the opposite side, for the advantages that opportune knowledge formalization can assure in personalizing diagnostic investigations. This means that adequate conditions could be encoded within the CDSS in order to suggest which kind of parameters could be more usefully evaluated for a given patient during, for instance, a TTE session.

Table II. Independent predictors of a worse prognosis

Demographics and Historical	Clinical	Electrophysiological	Functional	Blood	Hemodynamic
Advanced age	High heart rate	Broad QRS	VO2 max (mL/Kg x min< 10-14)	High serum BNP	Low Left ventricular EF
Coronary aetiology	Persist low Blood Pressure	Low heart rate variability	Low 6 min walking ability	High serum nore-pinephrine	Increased LV volumes
Diabetes	NHYA III-IV	Complex ventricular rhythms	High VE/VCO2 ratio	Low serum sodium	Low cardiac index
Resuscitated sudden death	Involuntary weight loss	T-wave alternants		High serum creatinine	High left ventricular filling pressure
Ethnic group	Ventilatory rhythm and rate disturbance			High serum bilirubine	Restrictive mitral filling pattern
				Anaemia	Impaired right ventricular function
					Cardiothoracic ratio

5.2 Examples of CDSS Functionalities involving Signal and Image Processing

In the previous section, the contribution provided to decision support by data processing facilities has been generally discussed. Herein, a specific example is considered for illustrating a practical application within HEARTFAID platform and, then, detailing how data processing integrates into patients' management. It is drawn from the *Patient's Management Showcase* which pertains to a worsening event during the clinical course of a HF patient and has been selected as a comprehensive showcase of the project.

Patient's Management Showcase

We consider the case of Mr. Pietro Guarneri, a 65 years old patient, already enrolled in the HFP, former smoker, who has being suffering from hypertension for several years. He was enrolled in the HFP six months ago, when also the telemonitoring services offered by the platform were activated. At the baseline visit, the patient referred a slight limitation of his physical activity: he felt comfortable at rest but ordinary activity resulted in fatigue and dyspnoea. Hence, he was assigned to NYHA class II. Anamnesis data were also collected: the patient had an acute myocardial infarction five years before, underwent to aorto-coronary bypass, and 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 the last 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.

- ***Worsening Event***

  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 (Figure 68).

  At the visit, the NYHA class changes from II to III. No variations in the signs are observed by the cardiologist, unless 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 (Figure 68).

  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 fraction 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% (Figure 69).

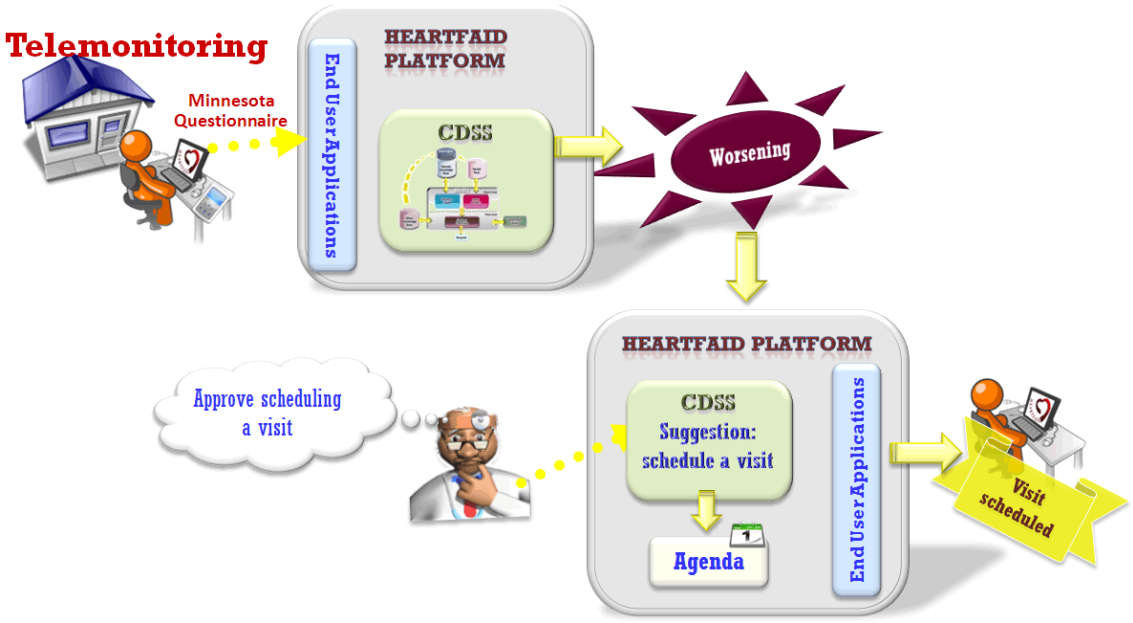

Figure 68. Detection of patient's worsening conditions via telemonitoring

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 sub-costal 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 home***

  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.

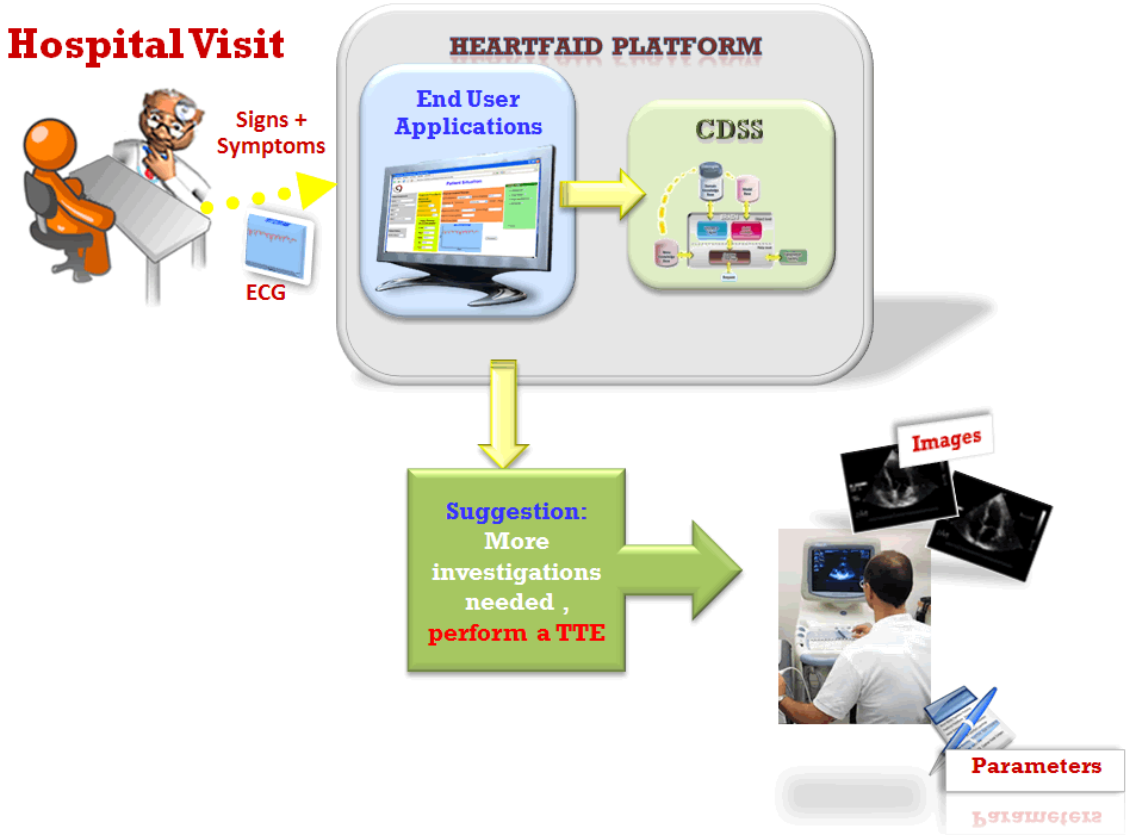

Figure 69. CDSS suggestion: scheduling of a TTE examination

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.

The intervention of data processing methods is clearly evident in the previous description. First of all, the CDSS should suggest the clinician to perform further imaging investigations for assessing the reason of patient's worsening. This can happen because, along with signs and symptoms and some parameters such as heart rate, more pieces of information are needed by the system for suggesting a response in terms of

worsening causes and correspondent therapy suggestion. Once acquired the imagery data, with enclosed a set of parameters computed by the radiographer, image analysis methods developed within HEARTFAID are applied for repeating the measurements in a more objective and consistent fashion. The parameters obtained in this way, i.e., left ventricle volume and ejection fraction, are automatically inputted to the CDSS which can use them for further reasoning on patient's situation.

5.3 Integration within HEARTFAID Clinical Decision Support System

The previous sections have pointed out the relevance of signal and image processing methods and the occurrence of their application as a necessary step in some reasoning processes of HEARTFAID CDSS. That has motivated the specification of such methods as a facility inherent in the CDSS. However, while designing the system, the importance of system modularity and computational complexity has come forth and suggested separating the signal processing facility from the CDSS and integrating it into the eCRF. Actually, the signal processing methods have the duty of assisting the cardiologist in the proper computation of data that have to be filled in the eCRF. When such data are necessary for any CDSS process, they can be obtained by querying the eCRF database.

On the contrary, the methods for processing and analyzing diagnostic images have been kept inside the CDSS module. More precisely, a dedicated component of the system was conceived for collecting and maintaining all the algorithms developed, i.e., the so-called *Model Base* (MB). Originally, a single MB was introduced for both the image processing and the computational reasoning methods. Once again for modularity and efficiency purposes, during the design activity, the two types of methods were divided and considered as belonging to two separate MBs. The *Imaging Model Base* (IMB) contains the image processing algorithms, which are queried for through the *Imaging Model Manager*. A service oriented approach has been selected for integrating IMB and, then, exposing the algorithms as web services.

The approach followed for applying these algorithms can be defined *offline* with respect to CDSS processing workflow, while integrated into the images acquisition process within the platform. For better illustrating the meaning of this, we can consider the case of the Patient's Management scenario described in the previous section. Recall the situation when, during the visit, for better assessing the causes of patient's worsening, the CDSS suggests the clinician to perform an echocardiographic investigation. If s/he approves, the investigation is scheduled and responses are waited for. Then, displaced from the CDSS workstation, in the radiology unit, a sonographer performs the echo exam by recording images and extracting a set of parameters, i.e., the Left Ventricle End Systolic and End Diastolic Volumes (LV ESV and EDV), the Inferior Vena Cava Minimum and Maximum Diameters (IVCDi min and max), the Left Ventricle Ejection Fraction (LVEF), the Inferior Vena Cava Collapsibility Index (IVCI), and the Tricuspid Valve Regurgitation. These can be considered as *subjective parameters* since are computed by applying operator dependent methods. Through the HEARTFAID hospital gateway, data are uploaded into the platform; in particular, images are store into the

DICOM Image Archive, while the parameters are included into the Patients' repository built upon the eCRF. Once images have been uploaded, the image processing tool is launched by a request sent to the Imaging Model Manager; this way the methods for computing objectively the relevant parameters are applied to the recently upload images, thus assuring the correctness of the CDSS. More precisely, the parameters computed are LV -ESV and -EDV, IVCDi min and max, and IVCI. These parameters are stored into the Patients' Repository, together with the already computed ones and a message is sent to the CDSS. Figure 70 illustrates this process.

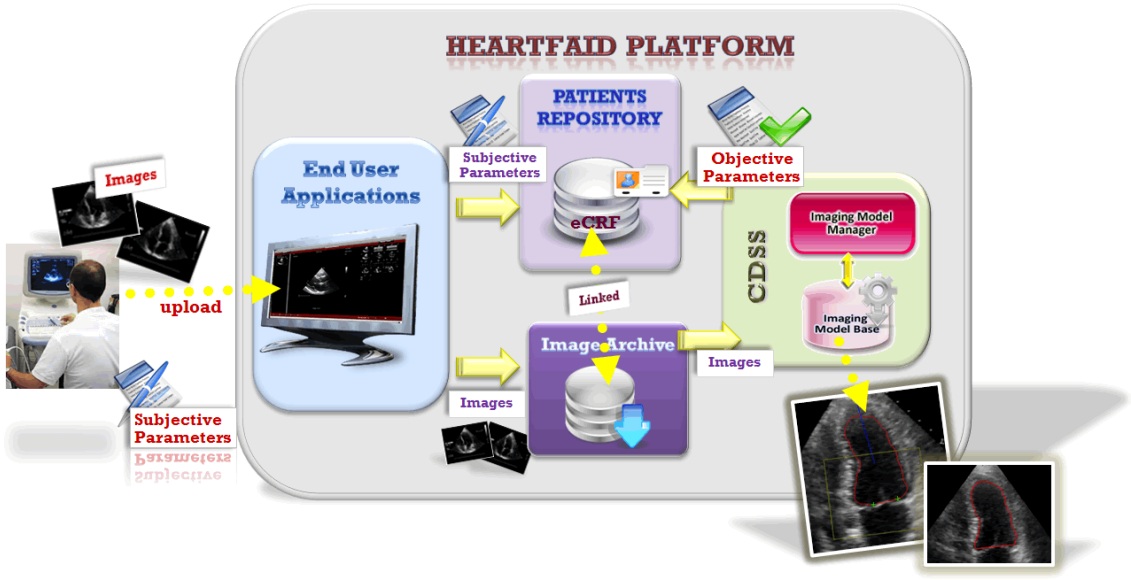

Figure 70. The process of echocardiographic image acquisition and processing

When the clinician logs in again for checking out Mr. Guarneri's situation, the system is aware of newly uploaded parameters and can display them, along with a link to the stored images. Contemporary, the CDSS employs the parameters computed by the image processing tool to estimate a number of additional features, i.e., the Systolic Pulmonary Pressure (SPP), the Pulmonary Hypertension (PH), and the Left Ventricle Ejection Fraction (LVEF). Such parameters are then used for solving patient's worsening; an example of rule formalized for such a task is reported below as it has been elicited in natural language:

```
if a patient has LVEF <= 40 with
or without symphtoms
then the patient is recommended
to take ACE-Inhibitors
```

Conclusions

This deliverable reports the work that has been carried out during the Task 5.2 “Design and development of models and methods for signals and images processing” of HEARTFAID project.

Starting from the preliminary work on the analysis of diagnostic resources and on foundational aspects of biomedical data processing described in deliverable D15, we have described in the current document the activity of study and development of model and methods for image and signal processing relevant to HEARTFAID platform.

In particular, we have motivated the choices made in the selection of the most significant and representative signal and image processing tools for the management of heart failure patients and we have tried to describe the developed algorithms in some technical details, without compromising the overall reading to non-specialist public.

The developed image processing tools include methods for the segmentation of the left ventricle in image sequence obtained by echocardiography (which is recognized as the most important imaging modality in the heart failure realm) and magnetic resonance imaging. Methods for the segmentation and classification of chest X-ray images are also introduced.

It is important to notice that such kind of assisted segmentation methods for echocardiography is not available in the proprietary software running on the diagnostic devices in the HEARTFAID validation sites (and specifically on GE Vivid 7 echocardiography device) and, thus, it represents a value added by the platform.

Further, 3D rendering capabilities have also been included for magnetic resonance imaging. Besides visualization functionalities, this opens the way to the computation of *truly volumetric* heart chamber parameters.

The segmentation results are then deployed by other developed methods to extract clinical significant parameters.
The main example is the estimation of left ventricle ejection fraction from the segmented echocardiographic sequences (by Simpson rule) and from segmented magnetic resonance images (by unbiased volumetric computations). Another interesting example is given by the computation and visualization of wall-thickness on the 3D model of the LV obtained by magnetic resonance imaging.
Other common measurements from echocardiography, magnetic resonance imaging and chest X-ray are included.

The developed signal processing methods are focused on the analysis of ECG data. Besides the clinical relevance of ECG analysis and interpretation, ECG processing is also important for the analysis of data from sensors in research workflows. Actually some home monitoring devices included in the platform are able to acquire a (single-lead) ECG signal.

In particular, an algorithm for QRS detection with an extremely satisfactory performance has been developed. Another algorithm for the construction of the average dominant heart beat has also been developed.
Since the average dominant beat is cleaner from the noise than the original signal, this average beat can be used by the clinicians to perform more accurate measurements using a caliper. This will reduce the inter- and intra- observer variability.

The integration of signal and image processing tools in the platform has been described at two levels. At the technical level, the integration is achieved by using a suitable IT infrastructure, whose main component is a DICOM-compliant Image Archive, in accordance with the general "interoperability" philosophy of the platform.

At a higher level, the CDSS can exploit the clinical parameters more precisely and consistently extracted by the signal and image processing methods for both computing other relevant parameters, which depend on their values, and checking out the validity of formalized conditions on them that are pertinent to patients' diagnosis, prognosis and treatment prescription.
This way, the system can feasibly support clinicians' decision workflows by intervening when needed and providing the required and useful information.

The discussion carried out throughout this document has pointed out several advantages assured by integrating signal and image processing with decision support, and has also highlighted some frontier issues that should be pursued for really and completely enjoying such advantages.

From an operational point of view, the integration within HEARTFAID platform has been detailed by depicting a real situation and explaining data flows and exchanges between the involved platform components. Actually, the application of signal and image processing methods has been designed as *offline* with respect to the CDSS functioning. However, after their application, the CDSS should be able to retrieve the computed and stored parameters so as to reason on them and display them, along with a link to the archived signals or images.

The real scenario, selected for describing the integration of signal and image processing functionalities within the CDSS functioning, has been borrowed from the *Patient's Management Showcase*, which pertains to a worsening event during the clinical course of a HF patient.
Since such a scenario has been selected as a comprehensive showcase of the project, the operational description provided in this document draws the directions of the work needed for the future implementation activities.

List of Abbreviations

AAM	Active Appearance Model
ACC	American College of Cardiology
ACR	American College of Radiology
AHA	American Heart Association
ANN	Artificial Neural Network
AS	Aortic Stenosis
ASE	American Society of Echocardiography
AV	Aortic Valve
BPV	Blood Pressure Variability
CAD	Computer Aided Diagnosis
CD	Chest Diameter
CDSS	Clinical Decision Support System
CNR	Italian National Research Council
CPE	Cardiogenic pulmonary oedema
CT	Computed Tomography
CTR	Cardio-Thoracic Ratio
CW	Continuous Wave Doppler
DB	Data Base
DICOM	Digital Imaging and Communications in Medicine
DICOM-SC	Secondary Capture DICOM Object
DICOM-SR	Structured Report DICOM object
DSS	Decision Support System
ECG	ElectroCardioGram
eCRF	electronic Case Report Form
ED	End diastole
EF	Ejection Fraction
EF	Ejction Fraction
ES	End systole
ESC	European Society of Cardiology
FFT	Fast Fourier Transform
HD	Heart Diameter
HF	Heart Failure
HFP	HEARTFAID platform
HIS	Hospital Information System
HRV	Heart Rate Variability
IDSS	Intelligent DSS

IMB	Image Model Base
IPT	Image Processing Toolbox
IVC	Inferior Vena Cava
KDD	Knowledge Discovery in Databases
LV	Left Ventricle
LV	Left Ventricle
MB	Model Base
MBMS	Model Base Management System
MRI	Magnetic Resonance Imaging
NEMA	National Electrical Manufacturer's Association
NYMC	New York Medical College
PA	Posterior Anterior
PACS	Picture Archiving and Communication System
PW	Pulsed Wave Doppler
RIS	Radiology Information System
SWRL	Semantic Web Rule Language
TDI	Tissue Doppler Imaging
TEE	TransEsophageal Echocardiography
TTE	TransThoracic Echocardiography
URI	Universal Resource Identifiers in WWW
US	UltraSound
WT	WallThickness
XML	eXtensible Markup Language

Bibliography

Akg¨ul et al. 2007a	Akg¨ul, C.B., Sankur, B., Schmitt, F., Yemez, Y.: Multivariate density-based 3d shape descriptors. In: Shape Modeling International. (2007) 3–12
Akg¨ul et al. 2007b	Akg¨ul, C.B., Sankur, B., Yemez, Y., Schmitt, F.: Improving efficiency of density-based shape descriptors for 3d object retrieval. In: MIRAGE. (2007) 330–340
Azhari et al 1990	H. Azhari, S. Sideman, J. L. Weiss, E. P. Shapiro, M. L. Weisfeldt, W. L. Graves, W. J. Rogers, and R. Beyar, "Three-dimensional mapping of acute ischemic regions using MRI: Wall thickening versus motion analysis," Amer. J. Physiol., pt. 2, vol. 259, no. 5, pp. H1492–1503, Nov. 1990.
Barcaro et al. 2007a	Barcaro U., Moroni D., Salvetti O. , Left ventricle segmentation in ultrasound sequences for the automatic computation of ejection fraction. In: *Open German Russian Workshop on Pattern Recognition & Image Understanding* (Ettlingen, Germany, 20-23 August 2007). Proceedings, Forshungsinstitut fur Optronik und Mustererkennung, 2007.
Barcaro et al. 2007b	Barcaro U., Moroni D., Salvetti O. , Automatic Computation of Left Ventricle Ejection Fraction from Dynamic Ultrasound Images. Submitted to *Pattern Recognition and Image Analysis*, Pleiades Publishing House, 2007
Bezdek 1981	Bezdek, L.: Pattern Recognition with Fuzzy Objective Function Algorithm. Plenum Press, New York (1981)
Binder et al. 1999	T. Binder, M. Sussner, D. Moertl, H. Strohmer, T. Baumgarten, G. Maurer, G. Porenta, "Artificial neural networks and spatial temporal contour linking for automated endocardial contour detection on echocardiograms: A novel approach to determine left ventricular contractile function", Ultrasound Med. Biol., 25, pp. 1069, 1076, 1999.
Boyer et al. 2006	Boyer, K.L., Gotardo, P.F.U., Saltz, J.H., Raman, S.V.: On the detection of intra-ventricular dyssynchrony in the left ventricle from routine cardiac mri. In: ISBI. (2006) 169–172
Bregler and Konig 1994	Bregler, C., Konig, Y.: "Eigenlip" for robust speech recognition. In: Acoustics, Speech, and Signal Processing ICASSP-94. Volume II. (1994) 669–672

Bustos et al. 2005	Bustos, B., Keim, D.A., Saupe, D., Schreck, T., Vranic, D.V.: Feature-based similarity search in 3d object databases. ACM Comput. Surv. 37 (2005) 345–387
Caselles et al. 1997	V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours". *International Journal of Computer Vision*, vol. 22, pp. 61-79, 1997.
Chalana et al. 1996	V. Chalana, D.T. Linker, D.R. Haynor, Y. Kim, "A multiple active contour model for cardiac boundary detection on echocardiographic sequences", *IEEE Trans. Med. Imag.*, 15, pp. 290-298, 1996.
Chen et al. 2003	Y. Chen, F. Huang, H.D. Tagare, R. Rao, D. Wilson, E.A. Geiser, "Using prior shape and intensity profile in medical image segmentation", *Proc. IEEE Int. Conf. Comput. Vis.*, Nice, vol. II, pp. 1117-1125, 2003.
Chiarugi et al. 2007a	F. Chiarugi, V. Sakkalis, D. Emmanouilidou, T. Krontiris, M. Varanini, I. Tollis, "Adaptive Threshold QRS Detector with Best Channel Selection Based on a Noise Rating System", Proc. of Computers in Cardiology 2007;34:157-60.
Chiarugi et al. 2007b	F. Chiarugi, M. Varanini, F. Cantini, F. Conforti, G. Vrouchos, "Non-Invasive ECG as a Tools for Predicting Termination of Paroxysmal Atrial Fibrillation", IEEE Transactions on Biomedical Engineering. Aug 2007, Vol. 54, No. 8; 1399-1406.
Christov 2004	I. I. Christov, "Real time electrocardiogram QRS detection using combined adaptive threshold", BioMedical Engineering OnLine 2004:3-28.
Colatonio et al. 2006	Colantonio, S., Moroni, D., Salvetti, O.: Shape comparison and deformation analysis in biomedical applications. In: Eurographics Italian Chapter Conference. (2006) 37–43
Coppini et al. 1995	G. Coppini, R. Poli, and G. Valli, "Recovery of the 3-D shape of the Left Ventricle from Echocardiographic Images", *IEEE Transactions on Medical Imaging*, 14(2), 1995, pp. 301- 317.
Declerk et al. 1997	J. Declerck, J. Feldmar, M.L. Goris, and F. Betting, "Automatic Registration and Alignment on a Template of Cardiac Stress and Rest reoriented SPECT Images", *IEEE Transactions on Medical Imaging*, 16(6), 1997, pp. 727-737.
Di Bona et al. 2003	Di Bona, S., Lutzemberger, L., Salvetti, O.: A simulation model for analyzing brain structures deformations. Physics in Medicine and Biology 48 (2003) 4001–4022

Di Bona et al. 2003b	Di Bona, S., Niemann, H., Pieri, G., Salvetti, O.: Brain volumes characterisation using hierarchical neural networks. Artificial Intelligence in Medicine 28 (2003) 307–322
Di Bono et al. 2004	Di Bono, M., Pieri, G., Salvetti, O.: A tool for system monitoring based on artificial neural networks. WSEAS Transactions on Systems 3 (2004) 746–751
Faber et al. 1991	T.L. Faber, E.M. Stokely, R.M. Peshock and J.R. Corbett, "A Model-Based Four-Dimensional Left Ventricular Surface Detector", *IEEE Transactions on Medical Imaging*, 10(3), 1991, pp. 321-329.
Frangi et al 2001	Frangi, A.F., Niessen, W.J., Viergever, M.A.: Three-dimensional modeling for functional analysis of cardiac images: A review. IEEE Trans. Med. Imaging 20 (2001) 2–5
Grenander and Miller 1998	Grenander, U., Miller, M.I.: Computational anatomy: an emerging discipline. Q. Appl. Math. LVI (1998) 617–694
Hamilton 2002	P. S. Hamilton, "Open Source ECG Analysis Software Documentation", E. P. Limited, Somerville, MA, USA, 2002.
Huang et al 2004	X. Huang, D. Metaxas, and T. Chen, "MetaMorphs: Deformable Shape and Texture Models", in *Proc. of the IEEE Computer Society Conf. on Computer Vision and Pattern Recognition, CVPR'04*, Oral Presentation, Washington, D.C., June, 2004.
Hunt et al. 2005	S.A. Hunt et al., The American College of Cardiology/American Heart Association Task Force on Practice Guidelines, "ACC/AHA 2005 Guideline update for the diagnosis and management of Chronic Heart Failure in the adult", *ACC/AHA*, 82 pages, 2005.
Jekova 2000	I. Jekova, "Comparison of five algorithms for the detection of ventricular fibrillation from the surface ECG", Physiol Meas. 2000; 21:429–39.
Jolly et al. 2001	M.-P. Jolly, N. Duta, G. Funka-Lea, "Segmentation of the Left Ventricle in Cardiac MR Images", *ICCV 2001*, pp. 501-508.
Kass et al. 1988	M. Kass, A. Witkin, and D. Terzopulos, "Snakes: Active contour models", *International Journal of Computer Vision*, 1, 1988, pp. 321-331.
Katsuragawa 2007	S. Katsuragawa, K. Doi, Computer-aided diagnosis in chest radiography. Computerized Medical Imaging and Graphics, Volume 31, Issue 4-5, Pages 212-223, 2007

Keim 1999	Keim, D.: Efficient geometry-based silmilarity search of 3D spacial databases. In: Proceedings of the 1999 ACM SIGMOID, New York, ACM Press (1999) 419–430
Kohler et al. 2002	B.U. Kohler, C. Hennig, R. Orglmeister, "The principles of software QRS detection", Engineering in Medicine and Biology Magazine, IEEE Volume 21, Issue 1, Jan.-Feb. 2002:42 - 57.
Kohonen 1997	Kohonen, T.: Self-Organizing Maps. 2nd edn. Volume 30 of Springer Series in Information Sciences. (1997)
Kunio Doi 2007	Kunio Doi, Computer-aided diagnosis in medical imaging: Historical review, current status and future potential. Computerized Medical Imaging and Graphics 31 (2007) 198–211, Elsevier
Laws 1980	K. I. Laws, "Rapid texture identification," Proc. SPIE, vol. 238, pp. 376-380, 1980.
Li et al. 2005	C. Li, C. Xu, C. Gui, M.D. Fox, "Level set evolution without re-initialization: a new variational formulation". In: *Proceedings of the 2005 Conference on Computer Vision and Pattern Recognition (CVPR'05)*, pp. 430-436, 2005.
Litchell et al. 2002	S. C. Mitchell, J. G. Bosch, B.P. F. Lelieveldt, R. J. van der Geest, J. H. C. Reiber, and M. Sonka, "3-D Active Appearance Models: Segmentation of Cardiac MR and Ultrasound Images", *IEEE Transactions on Medical Imaging*, 21(9), 2002, pp. 1167-1178.
Mignotte and Meunier 2001	M. Mignotte and J. Meunier, "A multiscale optimization approach for the dynamic contour-based boundary detection issue", *Comput. Med. Imag. Graph.*, 25, pp. 265-275, 2001.
Moroni et al. 2006	Moroni, D., Perner, P., Salvetti, O.: A general approach to shape characterization for biomedical problems. In Perner, P., ed.: Industrial Conference on Data Mining ICDM - Workshop on Mass-Data Analysis of Images and Signals. IBAI CD-Report, Leipzig (2006) 56–65
Moroni et al. 2007a	D. Moroni, S. Colantonio, M. Salvetti, O. Salvetti, "Deformable Structures Localization and Reconstruction in 3D images". In: *2nd International Conference on Computer Vision Theory and Applications* (Barcelona, 8-11 March 2007), pp. 215-222, 2007.

Moroni et al. 2007b	Moroni, D., Colantonio, S., Salvetti, M., Salvetti, O.: Deformable structures localization and reconstruction in 3D images. In Ranchordas, A., Araujo, H., Vitria, J., eds.: 2nd International Conference on Computer Vision Theory and Applications. INSTICC, Barcelona (2007) 215–222
Noble and Boukkerroui 2006	J.A. Noble, D. Boukerroui, "Ultrasound image segmentation: a survey", *IEEE Trans. Med. Imag.*, 25, pp. 987-1010, 2006.
Osher and Sethian 1988	S. Osher, J.A. Sethian, "Fronts propagating with curvature dependent speed: algorithms based on Hamilton-Jacobi formulations", *J. Computational Physics*, 79, pp.12-49, 1988.
Papin et al. 2000	Papin, C., Bouthemy, P., M`emin, E., Rochard, G.: Tracking and characterization of highly deformable cloud structures. In: Computer Vision - ECCV 2000. Volume 1843 of LNCS.,Springer Verlag (2000) 428–442
Paragios 2002	N. Paragios, "A Variational Approach for the Segmentation of the Left Ventricle in Cardiac Image Analysis", *International Journal of Computer Vision*, 50(3), 2002, pp. 345-362.
Paragios et al. 2002	N. Paragios, M. Rousson and V. Ramesh, "Knowledge-based registration & segmentation of the left ventricle: a level set approach", in *Proc. Sixth IEEE Workshop on Applications of Computer Vision*, 2002. (WACV 2002), pp. 37 – 42, 3-4 Dec. 2002.
Quekel et al. 1999	L. Quekel, A.Kessels, R. Goei, and J. V. Engelshoven, "Miss rate of lung cancer on the chest radiograph in clinical practice," *Chest*, vol. 115, no. 3, pp. 720–724, 1999.
R.M. Lang et al. 2005	R.M. Lang et al., "Recommendation for Chamber Quantification: A Report from the American Society of Echocardiography's Guidelines and Standards Committee and the Chamber Quantification Writing Group, developed in conjunction with the European Association of Echocardiography, a branch of the European Society of Cardiology", J Am Soc Echocardiogr, 18, pp. 1440-1463, 2005.
Riedmiller and Braun 1993	M. Riedmiller and H. Braun, "A direct adaptive method for faster backpropagation learning: The RPROP algorithm", in *Proceeding of the IEEE International Conference on Neural Networks (ICNN)*, 1993, pp. 586-591.

Sermesant et al. 2003	M. Sermesant, C. Forest, X. Pennec, and N. Ayache, "Deformable biomechanical models: Application to 4D cardiac image analysis", *Medical Image Analysis*, 7, 2003, pp. 475-488.
Sethian 1997	J.A. Sethian, *Level Set Methods and Fast Marching Methods,* Cambridge University Press, 1999.
Stalidis et al. 1999	G. Stalidis, N. Maglaveras, A. Dimitriadis, and C. Pappas, "Using Learning Classification and Four-Dimensional Parametric Modeling for the Analysis of Myocardial Thickening", *Computers in Cardiology*, 26, 1999, pp. 659-662.
Swedberg et al. 2005	K. Swedberg et al., The Task Force for the diagnosis and treatment of CHF of the European Society of Cardiology, "Guidelines for the diagnosis and treatment of Chronic Heart Failure: full text (update 2005)", *European Heart Journal*, 45 pages, 2005.
Talmon 1983	J. L. Talmon, "Pattern recognition of the ECG: a structured analysis", PHD thesis at the Vrije Universitet of Amsterdam, 1983.
van Ginneken et al. 2001	B. van Ginneken, B. M. ter Haar Romeny, and M.A. Viergever, Computer-Aided Diagnosis in Chest Radiography: A Survey, IEEE Trans. on Medical Imaging, vol. 20, no. 12, 2001
Yerushalmy 1969	J. Yerushalmy, "The statistical assessment of the variability in observer perception and description of Roentgenographic pulmonary shadows," *Radiologic Clinics N. Amer.*, vol. 7, no. 3, pp. 381–392, 1969.
Zhou et al. 2007	Z. Zhou et al., CAD–PACS integration tool kit based on DICOM secondary capture, structured report and IHE workflow profiles. *Computerized Medical Imaging and Graphics. Vol.* 31, pp. 346–352, 2007.

Miscellaneous Web Resources

ClearCanvas	http://www.clearcanvas.ca/dnn/
conQuest	http://www.xs4all.nl/~ingenium/dicom.html
Cordonnier 2003	http://medical.nema.org/dicom/workshop-03/pres/cordonnier.ppt, 2003
DCM4CHE	http://www.dcm4che.org/
DCMTK	http://dicom.offis.de/dcmtk.php.en
DICOM	http://medical.nema.org/
IHE	http://www.ihe.net/
IHE Cardiology TF Vol.1	http://www.ihe.net/Technical_Framework/upload/ihe_CARD_tf_vol1_2.pdf
IHE Cardiology TF Vol.2	http://www.ihe.net/Technical_Framework/upload/ihe_CARD_tf_vol2_2.pdf
IHE Echocardiography Workflow	http://www.ihe.net/Cardiology/committees/upload/IHE%20Cardiology%20Echo%20white%20paper%20March%202006%20v.5.pdf
ImageJ	http://rsb.info.nih.gov/ij
K-PACS	http://www.k-pacs.net/
XERO project	http://www.dcm4che.org/confluence/display/ee2/Xero
MIT-BIH Arrhythmia Database (Physionet)	http://www.physionet.org/physiobank/database/mitdb/