Integrating Data, Models, and Reasoning in Critical Care

Thomas Heldt, Bill Long, George C. Verghese, *Fellow, IEEE*, Peter Szolovits,
Roger G. Mark, *Senior Member, IEEE*

Abstract— Modern intensive care units (ICUs) employ an impressive array of technologically sophisticated instrumentation to provide detailed measurements of the pathophysiological state of each patient. Providing life support in the ICU is becoming an increasingly complex task, however, because of the growing volume of relevant data from clinical observations, bedside monitors, mechanical ventilators, and a wide variety of laboratory tests and imaging studies. The enormous amount of ICU data and its poor organization makes its integration and interpretation time-consuming and inefficient. There is a critical need to integrate the disparate clinical information into a single, rational framework and to provide the clinician with hypothesis-driven displays that succinctly summarize a patient's trajectory over time. In this paper, we present our recent efforts towards the development of such an advanced patient monitoring system that aims to improve the efficiency, accuracy, and timeliness of clinical decision making in intensive care.

I. INTRODUCTION

Modern intensive care units (ICUs) employ an impressive array of technologically sophisticated instrumentation to provide detailed measurements of the pathophysiological state of each patient. These advances in sensor technology have been matched by advances in hospital information systems, which now enable the collection and archiving of large amounts of patient data on a routine basis. Yet despite these advances, the paradigm for monitoring patients in the ICU has essentially remained unchanged since the 1970s: the clinical staff gathers large volumes of data from disparate sources, retrospectively screens the data by eye for important events, and relies on clinical experience to synthesize the information and to formulate differential diagnoses and treatment plans. Since the ICU is a data-rich environment, this strategy puts an enormous burden on the clinical staff. The *information overload* that results might actually hinder the diagnostic process, and may even lead to neglect of relevant data, resulting in errors and complications in ICU care. Furthermore, most of the information gathered is only screened for overt clinically important events and – in the absence of such events – is subsequently discarded.
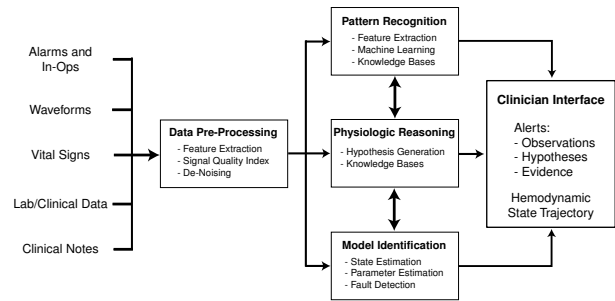
Fig. 1.   The major modules of an advanced patient monitoring system

On the other hand, the richness and detail of the collected data make it possible for a new generation of monitoring systems to track the pathophysiological state of the patient, employing the power of modern signal processing, pattern recognition, computational modeling, and knowledge-based clinical reasoning.

What is needed is a monitoring system that not merely reports all relevant measurements to clinicians, but also synthesizes the information and forms pathophysiologic hypotheses that best explain the observed data. Such systems should reduce the problem of information overload, and provide alarms that are more timely and more specific than today's single-variable limit alarms. The system we envision will integrate and analyze multi-dimensional input streams of continuous and intermittant data as shown diagramatically in Figure 1.

By helping to focus the practitioner's attention on the most significant events and changes in the patient's state, and by suggesting likely physiological interpretations of that state, the use of such systems should permit early detection of complex problems, provide useful guidance on therapeutic interventions, and ultimately lead to improved patient outcomes.

In this paper, we present our recent efforts towards the development of an advanced patient monitoring system that aims to improve the efficiency, accuracy, and timeliness of clinical decision making in intensive care.

II. INTENSIVE CARE DATABASE

The successful development and evaluation of an advanced patient monitoring system for intensive care relies critically on large amounts of well-characterized and rich clinical test data. It is crucial to develop and test algorithms with real data, and to be able to perform such tests repeatedly and reproducibly as algorithm refinements evolve. A central

aim of our research is therefore to create and disseminate a major new database from ICU patients. Previous such efforts include the IMPROVE database (with about 60 patient-days) [1], the Multi-parameter Intelligent Monitoring for Intensive Care (MIMIC) database (with approximately 200 patient-days) [2], and the commercially available APACHE [3] and the PPROJECT IMPACT [4] databases (with several thousands of patients each).

In collaboration with a Boston-area hospital, we are currently building the MIMIC-II database, a comprehensive set of data from ICU patients, including up to four continuous waveform signals, up to 30 multi-parameter physiologic trends, bedside monitor alarms, clinical laboratory data, medication records, fluid balance, nursing progress notes, discharge summaries, imaging reports, to name only the most important data sources. Data are collected from multiple surgical and medical intensive care units. We currently have complete sets of data, including high-resolution ECG and blood pressure waveforms, from approximately 2,200 patients (age range: 19–100 years; mean: 66.3 years, median: 68), totaling approximately 16,000 patient-days (mean stay in the ICU: 7.2 days, median: 3 days, maximum: 150 days). In addition, we have 15,000 patient records (including 4,000 neonatal patient records) that are lacking only the high resolution waveform data. In its current state, the entire database occupies over one terabyte of storage space.

A major focus of our attention is the annotation of a carefully selected subset of the total database. The annotations will characterize and document the physiological state of these patients as a function of time, primarily on the basis of progress and discharge notes, documented therapeutic interventions, and significant changes in the pattern of vital signs, and laboratory tests. Such physician-validated annotations are invaluable for algorithm development, especially with a database as heterogeneous as the MIMIC-II database.

Once thoroughly de-identified, the annotated and unannotated portions of the database will be made available for the research community at large via the NCRR-funded Research Resource for Complex Physiological Signals (http://www.physionet.org).

III. ALGORITHM DEVELOPMENT

Many opportunities exist for algorithm development in the context of intensive care, ranging from improving ECG and blood pressure alarm specificity to automated detection of various disease states.

An important area of research in advanced patient monitoring rests with the automated assessment of a measure of signal quality.

Zong recently demonstrated the importance of signal quality assessment for the reduction of false arterial blood pressure alarms in the ICU [5]. Currently, the arterial blood pressure data are processed by bedside monitors that generate alarms if measures of arterial blood pressure (systolic, mean, or diastolic) exceed or fall below pre-set thresholds. This strategy leads to a large proportion of false positive alarms and it contributes in no small measure to the noise pollution

and desensitization of the clinical staff in the ICU. With a signal quality metric in place, and taking into account information from both the ECG and the arterial blood pressure waveform, Zong showed that 98.2% (159 of 162) of false arterial blood pressure alarms could be rejected, reducing the false ABP alarm rate from 26.8% to 0.5%, while accepting 99.8% (441 of 442) of true ABP alarms.

Similarly, signal qualities on multiple data channels must be assessed and noise and artifacts must be attenuated without loss of physiological information before such information is passed to alarm detectors or higher reasoning or identification modules.

IV. MODEL-BASED PATIENT MONITORING

Computational models, when coupled to model identification schemes, have the potential to integrate available data streams into a single, rational framework that succinctly synthesizes our *a priori* understanding of basic physiology and pathophysiological processes. Such models thus relate properties of the available data streams to the physical properties of the underlying physiological system. Ideally, these models will allow for prediction of impending physiological crises (such as hemodynamic instabilities) and for testing and optimization of pharmacological or surgical interventions before such they are actually made on the sickest of patients.

Our group has begun explorations along the lines of model-based diagnoses and tracking of patient state. Here, we highlight three application areas of computational models in the clinical decision making process.

A. Lumped-parameter Models and Model Identification

In some areas of clinical medicine, such as hemodynamic monitoring, enough quantitative knowledge about the underlying physiological system is available to formulate computational models capable of simulating normal and abnormal human physiology. Through parameter tuning, such models can be used to track the patient's state automatically. Frequently, however, the models are built at a resolution too high to allow for accurate estimation of *all* the model parameters from the available data streams. This problem leads to an *ill-conditioned* estimation problem, in which the parameter estimates become very sensitive to the noise in the data. To overcome this ill-conditioning, we have explored subset selection-based optimization algorithms [6], a methodology that improves the conditioning of the estimation problem by reducing its dimensionality. Subset selection identifies those *well-conditioned* parameters that have strong independent effects on the model output, and thus can be estimated robustly.

To evaluate subset selection-based parameter estimation for a commonly-used lumped-parameter model of the cardiovascular system (shown in Figure 2, Samar [7] conducted a simulation study during both steady-state and transient conditions and used single-cycle waveforms of arterial blood pressure, central venous pressure, and pulmonary artery pressure as observable data during cases of simulated hemorrhage. The subset selection algorithm returned distending
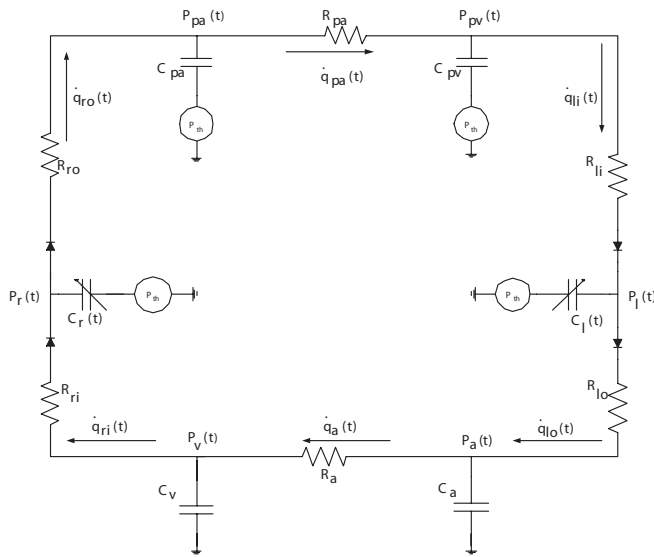
Fig. 2.   Circuit analog of the cardiovascular model. The model represents the lumped effects of the right, and left ventricles, the systemic arteries and veins, and the pulmonary artery and veins.

blood volume, DBV, systemic peripheral resistance, R_a, and the end-diastolic compliance of the right ventricle, C_r^{ed}, as the well-conditioned or active parameters. Figure 3 shows the correlations between estimated and real values for the three active parameters. The correlation coefficients for the three active parameters are 0.98, 0.97, and 0.54 for ΔDBV, R_a, and C_r^{ed}, respectively. More importantly, though, by focussing the estimation problem on the most identifiable of parameters, the mean relative estimation error could be reduced by a factor of two as compared to the full-order estimation problem.

Our simulation results indicate that the use of subset selection to identify and estimate only the well-conditioned parameters significantly improves the reliability of the estimates.

B. Model-based Expert System for Interpretation of Hemodynamic Data

Using the same lumped-parameter model of the cardiovascular system to generate data for a variety of simulated disease states, Zhao used a rule-based expert system approach to identifying all 23 parameters of the cardiovascular model shown in Figure 2 [8]. A clinical domain expert provided the rules for updating the parameter values given a discrepancy between the "measured" and simulated variables. This knowledge base was coded into an inference engine that subsequently guided the parameter search for common disease states such as hemorrhage, septic shock, heart failure, hypertension, and myocardial infarction to name but a few. Of the fifteen disease states considered, model identification in ten cases was deemed excellent (mean relative error of 10.2%) by the domain expert, two were considered acceptable (mean relative error of 54%), one was
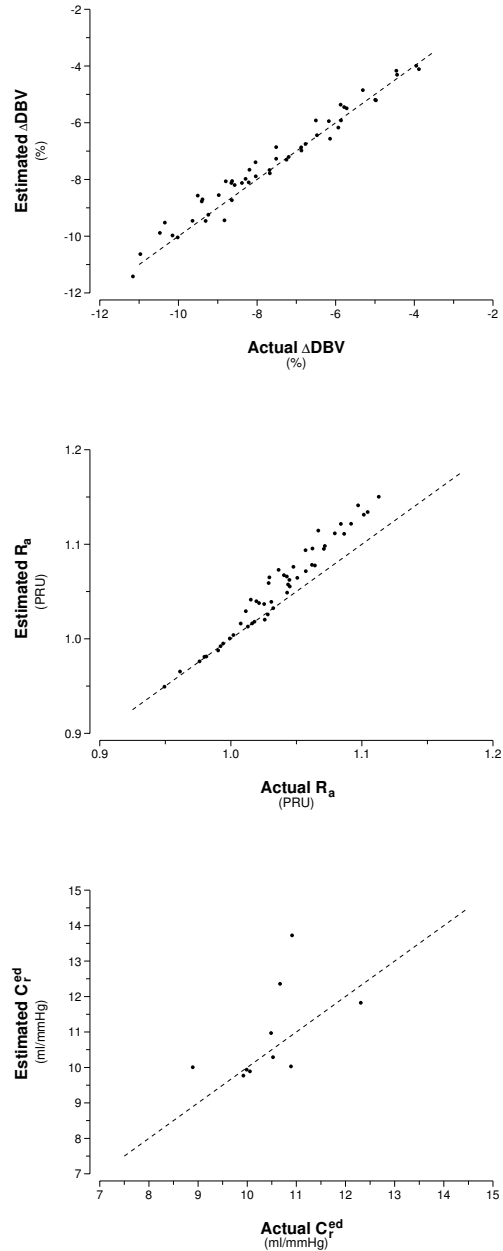

Fig. 3.   Estimated vs. actual parameter values.

deemed unsatisfactory (mean relative error of 120%); the parameter search was stuck in local minima in two cases.

The problem addressed by Zhao was very ambitious as one is rarely interested in the identification of *all* parameters of a cardiovascular model. Rather it is a few important parameters that clinicians aim to track over time. Samar's mathematical approach and Zhao's expert system represent two polar opposite approaches to the problem of identifying model parameters from observable data streams. It is quite conceivable that a combination of the two approaches will yield the best possible amount of information about a par-

ticular patient's state.

C. Bayesian Network Models of Cardiovascular Function

Lumped-parameter models capture the dynamic behavior of hemodynamic variables (pressure, flows, volumes) and their control parameters (resistance, heart rate, contractility, venous tone) quite well. Yet, in the ICU a lot of data is available that is not of the hemodynamic type, and consequently, lumped-parameter models generally fail to account for changes in such data as lab values, levels of medications, and qualitative assessments such as "the patient looks pale". Furthermore, lumped-parameter models generally fail to account for the uncertainty ascribed to each particular measurement. This type of information can best be coded in a probabilistic framework, in which each measured variable is considered random. The appropriate probability densities can chosen to be discrete or continues depending on the particular nature of the variable; they can, and generally will, depend on values of other variables in the system, thus taking account of the physiological relationships that exist between different measured variables.

Our efforts in building probabilistic network models for diagnosis of cardiovascular diseases and prediction of therapy effects date back to the late eighties [9], [10]. In one line of research, we used a Bayesian network model to establish a hierarchy of differential diagnoses given the findings of a physical exam (this program can be accessed via http://groups.csail.mit.edu/medg/). In a second line of research, we used a similar model to predict the effects of therapeutic interventions. In both cases, the performance of the model was satisfactory and en par with the expectation of trained clinical staff (for example, 36 out of 42 correct differential diagnoses [10]).

More recently, we have begun adapting a Bayesian network model of the cardiovascular system to allow for sequential identification of unmeasurable variables, taking into account both general population distributions and the recent history of a particular patient [11].

Ultimately, lumped-parameter models, probabilistic models, and expert systems will need to communicate with each other and allow for information to be passed from one type of model to the other in order to arrive at a holistic picture of the physiological state of each patient.

V. CONCLUSIONS

Providing life support in the ICU is becoming an increasingly complex task because of the growing volume of relevant data from clinical observations, bedside monitors, mechanical ventilators, and a wide variety of laboratory tests and imaging studies. The enormous amount of ICU data and its poor organization makes its integration and interpretation time-consuming and inefficient. There is a critical need to integrate the disparate clinical information into a single, rational framework and to provide the clinician with hypothesis-driven displays that succinctly summarize a patient's trajectory over time. Such a framework will inevitably be based on a hierarchy of computational models, whose parameters will be identified on-line as more and more data becomes available for a particular patient.

VI. ACKNOWLEDGMENTS

REFERENCES

[1] K. Nieminen, R. Langford, C. Morgan, L. Takala, and A. Kari. A Clinical Description of the IMPROVE Data Library. *IEEE Eng Med Biol Mag*, vol. 16, 1997, p 21–24.

[2] G.B. Moody and R.G. Mark, A Database to Support Development and Evaluation of Intelligent Intensive Care Monitoring, *Comput. Cardiol.*, vol. 23, 1996, pp 657–660.

[3] M. Suistomaa, M. Niskanen, A. Kari, M. Hynynen, and J. Takala. Customized prediction models based on APACHE II and SAPS II scores in patients with prolonged length of stay in the ICU. *Int Care Med*, vol. 28, 2002, pp 479–485.

[4] http://www.piccm.com.

[5] W. Zong, G.B. Moody, R.G. Mark. Reduction of false arterial blood pressure alarms using signal quality assessment and relationships between the electrocardiogram and arterial blood pressure. *Med Biol Eng Comput*, vol. 42(5), 2004, pp 698–706.

[6] M. Burth, G.C. Verghese, M. Vélez-Reyes. Subset selection for improved parameter estimation in on-line identification of a synchronous generator. *IEEE Trans Power Systems*, vol. 14(1), 1999, pp 218–225.

[7] Z. Samar, T. Heldt, G.C. Verghese, R.G. Mark. Model-based cardiovascular parameter estimation in the Intensive Care Unit. *Comput Cardiol*, vol. 32, 2005, pp 635–638.

[8] R. Zhao. A model-based expert system for interpretation of hemodynamic data from ICU patients. Technical Report TR-718, Massachusetts Institute of Technology, Laboratory of Computer Science, May 1997.

[9] W.J. Long, S. Naimi, M.G. Criscitiello, R. Jayes. Using a physiological model for prediction of therapy effects in heart disease. *Computers in Cardiology Proceedings* 1986, pp 15–20.

[10] W.J. Long, S. Naimi, M.G. Criscitiello, G. Larsen. Differential Diagnosis generation from a causal network with probabilities. *Computers in Cardiology Proceedings* 1988, pp 185–189.

[11] J.M. Roberts, T.A. Parlikar, T. Heldt, G.C. Verghese. Bayesian Networks for Cardiovascular Monitoring. *Proceedings of the IEEE Engineering in Medcine and Biology Conference, New York, 2006.*