

A Database to Support Development and Evaluation of Intelligent Intensive Care Monitoring

George B Moody and Roger G Mark

Harvard-MIT Division of Health Sciences and Technology, Cambridge, MA, USA
Cardiology Division, Beth Israel Hospital, Boston, MA, USA

Abstract

Development and evaluation of automated decision support systems requires large amounts of well-characterized, reproducible test data. The MIMIC (Multi-parameter Intelligent Monitoring for Intensive Care) Database is intended to meet these needs. The database, currently nearing completion, will include 100 patient records, each typically containing between 24 and 48 hours of continuous data recorded from patient monitors in the medical, surgical, and cardiac intensive care units of Boston's Beth Israel Hospital. Each record will be accompanied by detailed clinical data derived from the patient's medical record and from the hospital's on-line medical information systems. We select patients to record from those likely to be hemodynamically unstable during the planned recording period. We expect to complete the selection of the recordings to be included in the database by the end of 1996, and to make the database available to other researchers shortly thereafter.

1 Introduction

As the sophistication of diagnostic and monitoring devices grows, care of acutely ill patients requires rapid assessment and accurate interpretation of a large and ever-growing volume of medical data. Automated decision support systems capable of assimilating these numerous data streams, tracking and anticipating clinically significant changes, and presenting pathophysiologic hypotheses, would be of great value to clinicians facing "information overload" in the intensive care unit (ICU). We have begun to explore the design of "intelligent patient monitors" to accomplish these goals. To support this effort, we are creating the MIMIC Database.

Reference databases [1, 2, 3] are essential resources for developers and evaluators of algorithms and systems for analysis of physiologic data. The complexity of such data renders hopeless efforts to develop analytic techniques by pure reasoning. It is essential to

test algorithms with realistic data, and to be able to perform these tests repeatedly and reproducibly as algorithm refinements are proposed. The requirement of reproducibility stems from the complexity of all but the most trivial analysis algorithms: without a strictly reproducible test, one can never establish with certainty if an observed difference in the results of analysis is attributable to a difference in the analytic method, or to a difference in the input data. It follows from this reasoning that a fair evaluation of an algorithm, whether against a standard or in comparison with another algorithm, also requires a reproducible test. Digital recordings are ideal test data for these purposes.

Experience with our earlier databases has also demonstrated their potential for stimulating basic research. By making well-characterized clinical data available to researchers, these databases make it possible to formulate and answer numerous physiologic questions, without the necessity of developing a new set of reference data at great cost in each case. Even non-specialists, who may lack other means of access to clinical data, can readily participate in such efforts. Since the databases are generally available, research results based on them can be readily replicated by others, fostering both healthy competition and wide-ranging collaborations with minimal barriers to entry.

These databases have also had value in medical education, by providing well-documented case studies of both common and rare but clinically significant pathologies. In clinical practice, however, it is common to have much more information to support diagnoses than has generally been available in existing reference databases.

In the current project, we have taken several steps beyond our previous reference databases. First, the recordings themselves are far longer than in our previous work. We sought to make the recordings long enough to be able to illustrate both the development of critical events and the results of interventions, in

many cases several times in a single recording. Second, the MIMIC database includes comprehensive data from each subject's medical record, providing context for each recording. Although existing ICU monitors do not often make use of these data, it is clear that any process intended to provide competent support for medical decisions must be aware of all of what is known about the patient. Thus we regard this component of the database as essential for the task of developing intelligent ICU monitoring in the future. Third, although we have previously published a database of polysomnographic recordings containing blood pressure and other signals, the MIMIC database project is our first extensive and systematic effort to record hemodynamic variables together with multi-lead ECGs and other real-time signals.

We select patients from those considered likely to be hemodynamically unstable during the planned recording period. We chose to focus on this group of patients because their care in the ICU presents a particularly demanding challenge to physicians and nurses. The cardiovascular control system functions primarily as a regulator of blood pressure. Sudden changes in blood pressure hence signify a seriously compromised patient, whose control system is unable to react appropriately (whether by increasing cardiac output or by vasoconstriction) to a challenge such as hemorrhage, sepsis, trauma, medication, or cardiac arrhythmia (among other possibilities). These events require rapid intervention from the ICU staff, but the appropriate response depends on the cause. Thus, a timely and appropriate intervention requires rapid assimilation of a large amount of data from a diverse set of sources, including not only the current signals and recent trends from a bedside monitor, but also (for example) the patient's symptoms, medications, fluid balance, and response to previous therapeutic interventions.

2 Methods

With the cooperation and assistance of the manufacturer, we developed special-purpose software (available on request from the authors) for logging the signals acquired by the Hewlett Packard CMS ("Merlin") bedside monitors in the medical, surgical, and cardiac ICUs of Boston's Beth Israel Hospital. The process requires the use of a pair of optional interface cards in the bedside monitor, communicating over two serial links at 38400 baud with a standard PC equipped with a Digiboard PC/4e "smart" serial interface, running our software. This configuration allows us to record three ECG signals each sampled at 500 Hz, and four or five other signals each sampled at 125 Hz (or two ECG signals and six other signals), in addition to the

monitor's periodic measurements, alarms, and monitor status messages. These limits are determined by the amount of data that can be passed through the two serial links at the maximum speed; in almost all cases, they permit us to log all of the signals available to the monitor, and all of its outputs. (We are not limited to the subset of signals displayed on the monitor's screen; if additional signals are available, they can be logged even if they are not displayed.) Although we initially planned to record for 24 hour periods, the data logging PCs proved to be so reliable that we have been able to record for 48 hours routinely, and in some cases as much as 80 hours per record. (We limit the total amount of data logged to about 600 Mb per record, so that each record can be stored on a single CD-ROM.)

To limit possible loss of data from power interruptions, the continuous recording is written in ten-minute segments, which are stored during logging on disk and copied to tape automatically at the conclusion of the logging process. Our software permits these segments to be read later as individual records or as a single continuous record without gaps between segments. In principle, this facility could permit us to produce seamless records that span multiple CD-ROMs.

At the beginning of each logging session, we record a sample of the monitor's calibration pulses. When reviewing the recordings, we use these pulses to correct for inter-signal sampling skew between the ECG signals. (Uncorrected skews of 16 or 32 ms may remain between the other signals, due to limitations in the monitor design). The signals are calibrated, and the record is copied to CD-R media for archival storage.

We attempt to capture all of the patient data available to the ICU staff during the recording period, so that the recording can serve as a high-fidelity patient simulation for development and evaluation of intelligent monitors. Our protocol is designed to be minimally disruptive to the normal monitoring procedure, so we do not require that specific signals be recorded; hence the MIMIC Database contains a variety of signal sets representative of standard practice in the ICUs at the Beth Israel Hospital. The records contain two or three ECG signals, and any other available signals, usually including arterial blood pressure, respiration, and pulse oximeter signal (a pulsatile signal used by the bedside monitor to derive measurements of oxygen saturation). The monitor's measurements of slowly-changing variables are recorded at intervals of 1.024 seconds (though some of these measurements are updated less frequently). These typically include heart and respiration rates, oxygen saturation, and systolic, diastolic, and mean arterial and pulmonary arterial blood pressures; frequently other measurements, such




Figure 1: A sample of MIMIC Database record 248, as it appears when viewed using WAVE. Note the annotation at the center, indicating that the monitor issued an alarm for multiform PVCs (the beats that triggered the alarm occurred several seconds earlier; the marker indicates the time when the alarm was issued).

as blood temperature, inspired minimum and end-tidal CO_2 , fractional inspired O_2 , and cardiac output, are also available.

For each record, clinical data derived from the patient's medical records (including symptoms, diagnoses, progress notes, and medications administered) and from the hospital's clinical computing systems (primarily laboratory results) are also included in the MIMIC Database. The available on-line data are downloaded into a relational database, and additional data abstracted from the paper medical record are hand-entered into another relational database. These databases are provided on the MIMIC Database CD-ROMs, and are also used to create an annotation file for each record that provides a timeline for simulations or case studies. Thus, for example, it is possible using our WAVE software [1] (see figure 1) to review the record in a fashion that allows "discovery" of events, laboratory results, and other findings in chronologic order.

For certain studies, the time series of the monitor's *measurements* may be sufficient. In addition to a presentation of these measurements together with the digitized signals, we have also prepared collections of these time series presented as signals for analysis by standard signal-processing software (see figure 2). Since these data require approximately two orders of magni-




Figure 2: A three-hour excerpt of the time series of monitor measurements for MIMIC Database record 216. Note the monitor's alarms (recorded as annotations, shown above the associated time series) surrounding the blood pressure changes that occur beginning at about 11:00:00.

tude less storage than the original signals, a complete collection of them for the entire MIMIC database, representing about 6000 hours of data, can fit on a single CD-ROM. This format is particularly convenient for studies that involve post-processing of the monitor measurements, as illustrated in figures 3 and 4.

Study of the simultaneous changes in several physiological variables will often illuminate the underlying physiology, and may be helpful in detecting artifactual data. For example, the scatter plot in the upper left panel of figure 3 demonstrates that changes in mean ABP and diastolic PAP are often correlated. This is to be expected in the context of normal myocardial function, where preload and afterload tend to move in the same direction. When these pressures are anti-correlated, one suspects changes in myocardial contractility as the cause (see fig.4). In the complete absence of correlations between the two pressures, one suspects the possibility of artifactual data. For example, the spike in PAP near 20 hours in the lower panels of figure 3 is not accompanied by changes in ABP, and examination of the raw signals confirms that the monitor alarms for that event were erroneous due to artifact.

3 Discussion and Conclusions

Our decision to focus on hemodynamically unstable patients means that, although the MIMIC Database should be representative of the full range of pathophysiologies that result in sudden blood pressure changes, it does not represent the entire ICU population. As au-

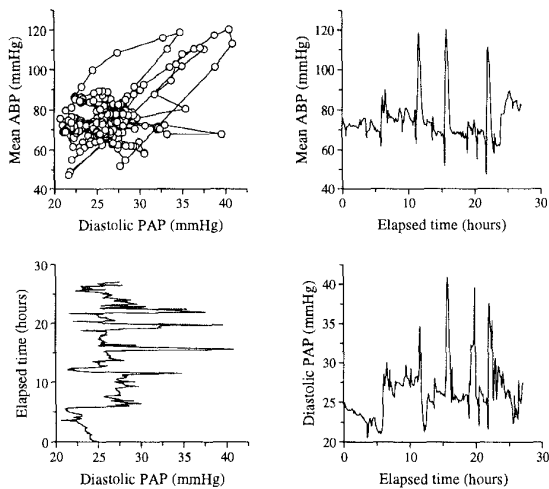


Figure 3: Mean ABP and diastolic PAP changes are usually correlated in record 216.

tomated decision support systems for the ICU develop and mature, there will be a need for recordings that represent other significant groups of patients whose problems also require rapid assessment and appropriate intervention.

The very long time series now available for study in the MIMIC Database will be of particular interest in investigations of heart rate, blood pressure, and respiratory dynamics and their interactions. We also anticipate that the clinical information accompanying each record will make the MIMIC Database a valuable source of well-characterized case studies for medical education, particularly in the intensive care setting.

Data collection began in late 1994 and is now nearly complete. As of August 1996, we had recorded roughly 250 patient-days. We will finish the selection of the recordings to be included in the MIMIC Database by the end of 1996, and the database will subsequently be made available to other researchers.

Acknowledgments

Development of the MIMIC Database has been made possible with the support of the Hewlett Packard Foundation and of Nihon Kohden America, Inc. Technical support from Hewlett Packard was essential to the project, and was freely and generously given to us; we especially wish to thank Dick Myreck and John Ames of HP, who arranged for the long-term loan of a bedside monitor for bench testing. At Boston's Beth Israel Hospital, we received the cooperation and assistance of many individuals to whom we are

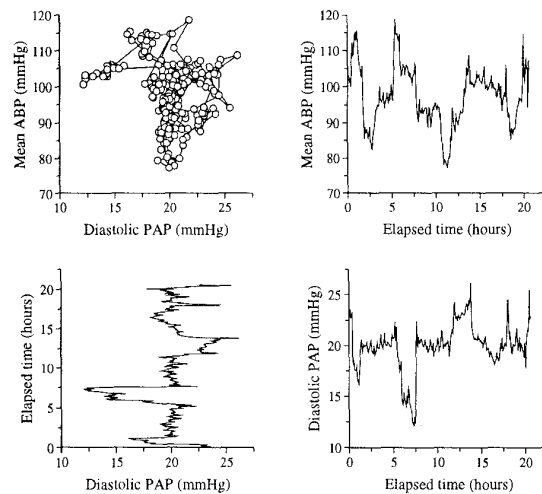


Figure 4: In MIMIC Database record 415, there is an anticorrelation between mean ABP and diastolic PAP.

deeply grateful. We wish especially to thank the patients who participated in the project and their families; the staffs of the Medical, Surgical, and Cardiac Intensive Care Units; the Biomedical Engineering Research and Arrhythmia Laboratories; and the hospital's electrical safety engineers. We also thank our students, Bryant Lin, Adam Hoytya, and Ruilin Zhao, all of whom contributed to the success of this project. Finally, GBM wishes to thank Edna Moody of the BIH Center for Clinical Computing for her inestimable help in obtaining data from the hospital's on-line medical information systems, and for everything else.

References

- [1] Moody GB, Mark RG. The MIT-BIH arrhythmia database on CD-ROM and software for use with it. *Computers in Cardiology 1990*; 185-188.
- [2] Taddei A, Distante G, Emdin M, Pisani P, Moody GB, Zeelenberg C, Marchesi C. The European ST-T database: standard for evaluating systems for the analysis of ST-T changes in ambulatory electrocardiography. *Eur. Heart J.* (1992) 13:1164-1172.
- [3] Moody GB, Feldman CL, and Bailey JJ. Standards and applicable databases for long-term ECG monitoring. *J. Electrocardiology* (1993) 26 Suppl:151-155.

Address for correspondence:

George B. Moody
 MIT Room 20A-113
 Cambridge, MA 02139 USA
george@hstbme.mit.edu
<http://scheherazade.mit.edu>