User Guide and Documentation for the MIMIC II Database

Gari D. Clifford[1,2], Daniel J. Scott[1,2] and Mauricio Villarroel[1,2]

[1]Harvard-MIT Division of Health Sciences & Technology, Rm E25-505,
45 Carleton St., Cambridge MA 02142, USA
[2]Massachusetts Institute of Technology Technology, Rm E25-505,
77 Massachusetts Av., Cambridge MA 02139, USA

MIMIC-II database version 2, release 1.

Rev: 239 LastChangedDate: 2010-02-25 11:21:40 -0500 (Thu, 25 Feb 2010)

Preface

This user guide is intended for clinicians with some knowledge of programming and/or graduate-level researchers with knowledge of biomedical signal processing. The user is expected to have a working knowledge of `SQL`. Basic knowledge of a statistical (signal processing) package such as `Matlab` or **R** is useful. Knowledge of `C/C++` and/or `Java` may also help, but is not essential.

Many of the signal processing algorithms and data sets described in this guide are available from, or described in papers posted at the following URLs:

`http://www.physionet.org/`
`http://mimic.mit.edu/`

Most of the algorithms posted at the above URLs have been written either in C or `Matlab`. Libraries for reading these databases are also freely available. We hope that through these URLs this database will continue to evolve and add to the growing body of open (repeatable) biomedical research.

Gari Clifford, Li-wei Lehman, Roger Mark, Mohammed Saeed,
George Moody, Daniel Scott and Mauricio Villarroel

The Laboratory for Computational Physiology,
Cambridge, MA, USA,
January 2009

Acknowledgments

This work was supported by the National Institute of Biomedical Imaging and Bioengineering (NIBIB) of the National Institutes of Health (NIH) (grant number R01 EB001659). We also acknowledge contributions by Philips Healthcare and the National Library of Medicine. The content of this document is solely the responsibility of the authors and does not necessarily represent the official views of the NLM, the NIBIB, the NIH or Philips Medical Systems.

Of course, this database could never had come into existence without the hard work of all the students, staff and faculty. The MIMIC II database is a collective effort driven by several individuals guided by Professor Roger Mark from the Laboratory for Computational Physiology at the Massachusetts Institute of Technology. Many research collaborators were and are directly involved in the constant evolution of the database, including: Omar Abdala, Anton Aboukhalil, Tiffany Chen, Gari Clifford, Anagha Deshame, Margaret Douglass, Thomas Heldt, Isaac Henry, Caleb Hug, Brian Janz, Sherman Jia, Tin Kyaw, Li-Wei Lehman, Bill Long, Qiao Li, Christine Lieu, Atul Malhotra, Benjamin Moody, George Moody, Ishna Neamatullah, Tushar Parlikar, Andrew Reisner, Ali Saeed, Mohammed Saeed, Daniel Scott, Dewang Shavdia, James Sun, Peter Szolovits, Danny Talmor, George Verghese, Mauricio Villarroel and Wei Zong.

Several individuals made possible the data collection infrastructure, including Brian Gross, KP Lee, Larry Nielsen and Greg Raber from Philips Healthcare (Andover, MA) and Philips Research of North America, and John Halamka, Larry Markson, Larry Nathanson and Lu Shen from Harvard Medical School and the Beth Israel Deaconess Medical Center. We also would like to thank numerous other collaborators at MIT, Harvard, Beth Israel Deaconess Medical Center, Philips Research and our advisory committee: James B. Bassingthwaighte, Reed Gardner, Clement McDonald and Michael Shabot.

IRB Approval

This study was approved by the Institutional Review Boards of Beth Israel Deaconess Medical Center (Boston, MA) and the Massachusetts Institute of Technology (Cambridge, MA). Requirement for individual patient consent was waived as the study did not impact clinical care and all data were de-identified.

Document License

This document and the data model described are licensed under the terms of the GNU General Public License version 2, as published by the Free Software Foundation.

This document is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MER-

CHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this guide; if not, you can ask for a copy from:
Free Software Foundation, Inc.
675 Massachusetts Avenue
Cambridge, MA 02139
USA.

Comments and questions

We have revised the document to the best of our ability, but you may find that some features have changed since this document was published, that we made some mistakes, or you may simply need more information for a particular section. If so, please notify us by writing to:

MIMIC II Database
Laboratory for Computational Physiology
E25-505
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
USA

Contents

1 Introduction **1**
- 1.1 Overview . . . 1
- 1.2 Background . . . 1
- 1.3 Overview of data collection . . . 3
- 1.4 Data Organization . . . 3
  - 1.4.1 Types of data . . . 3
  - 1.4.2 What is a patient record? . . . 6
  - 1.4.3 Subject ID - Case ID matching . . . 8
  - 1.4.4 De-identification of patients' data . . . 9
- 1.5 Clinical overview of patients in the MIMIC II database . . . 13
- 1.6 Noise, artifacts and missing data . . . 17
- 1.7 Summary and further reading . . . 18

2 Database Description **21**
- 2.1 Overview . . . 21
- 2.2 Clinical database . . . 21
  - 2.2.1 Patient . . . 22
  - 2.2.2 Care Giver . . . 24
  - 2.2.3 Care Unit . . . 25
  - 2.2.4 Patient timeline . . . 27
  - 2.2.5 Patient data . . . 29
  - 2.2.6 Summary . . . 35
- 2.3 High resolution waveforms and associated trends . . . 35
  - 2.3.1 Overview . . . 35
  - 2.3.2 The WFDB software package . . . 37
  - 2.3.3 MIMIC II waveform records . . . 38
  - 2.3.4 Alarms and Inops . . . 40
  - 2.3.5 Signal Quality . . . 42

3 Database Access **47**
- 3.1 Introduction . . . 47

4 Examples of data analysis **49**
4.1 Introduction . . . 49
4.2 Clinical Examples . . . 49
4.2.1 Patient population age statistics . . . 49
4.2.2 SAPS I on first day of admission . . . 53
4.2.3 Calculate mortality rate . . . 53
4.2.4 Statistics for patients admitted with Subarachnoid Hemorrhage on admission . . . 56
4.2.5 Patient BMI calculation . . . 58
4.2.6 Resolving discrepancies between multiple itemIDs for one parameter . . . 58

5 Quick-Start, Frequently Asked Questions and Known problems/issues **61**
5.1 Quick Start . . . 61
5.2 FAQs about the MIMIC II database . . . 62
5.3 FAQs about data access . . . 62
5.4 Known issues/problems . . . 63
5.5 Abbreviations . . . 65

6 Appendix **69**
6.1 Database Schema . . . 69
6.1.1 ADDITIVES . . . 69
6.1.2 ADMISSIONS . . . 70
6.1.3 A_CHARTDURATIONS . . . 71
6.1.4 A_IODURATIONS . . . 72
6.1.5 A_MEDDURATIONS . . . 73
6.1.6 CENSUSEVENTS . . . 74
6.1.7 CHARTEVENTS . . . 75
6.1.8 DB_SCHEMA . . . 76
6.1.9 DELIVERIES . . . 77
6.1.10 D_CAREGIVERS . . . 78
6.1.11 D_CAREUNITS . . . 78
6.1.12 D_CHARTITEMS . . . 79
6.1.13 D_IOITEMS . . . 79
6.1.14 D_MEDITEMS . . . 80
6.1.15 D_PARAMMAP_ITEMS . . . 80
6.1.16 D_PATIENTS . . . 81
6.1.17 ICD9 . . . 81
6.1.18 ICU_STAY . . . 82
6.1.19 IOEVENTS . . . 83
6.1.20 MEDEVENTS . . . 84
6.1.21 NOTEEVENTS . . . 85
6.1.22 PARAMETER_MAPPING . . . 86
6.1.23 POE_MED . . . 87
6.1.24 POE_ORDER . . . 88

6.1.25 REPORTEVENTS . . . 89
6.1.26 TOTALBALEVENTS . . . 89
6.2 Multiple ID mappings . . . 90
6.2.1 Bicarbonate (HCO3) . . . 90
6.2.2 Bilirubin (highest) . . . 91
6.2.3 Blood Pressure . . . 91
6.2.4 Blood Transfusions . . . 91
6.2.5 Cardiac Output (CO) . . . 92
6.2.6 Carbon Dioxide (CO2) . . . 92
6.2.7 Creatinine (highest) . . . 92
6.2.8 Central Venous Pressure (CVP) . . . 92
6.2.9 Glucose Levels . . . 92
6.2.10 Intra-aortic balloon (IABP) pump rate . . . 93
6.2.11 Intra-cranial Pressure (ICP) . . . 93
6.2.12 IV Fluids . . . 93
6.2.13 Lactate . . . 93
6.2.14 Oxygen Saturation (SpO2/SaO2) . . . 93
6.2.15 pH . . . 94
6.2.16 Potassium . . . 94
6.2.17 Pressor Medications . . . 94
6.2.18 Pulmonary Arterial Pressure (PAP) . . . 94
6.2.19 Respiration Rate . . . 95
6.2.20 Sodium . . . 95
6.2.21 Temperature . . . 95
6.2.22 Urine Output . . . 95
6.2.23 Ventilators . . . 96
6.2.24 White Blood Cell Count (WBC) . . . 96
6.3 Commonly used parameters . . . 96
6.4 Frequency of all ICD-9 codes for adult ICU-related hospital admissions . . . 97

Chapter 1

Introduction

1.1 Overview

This *User Guide* is intended to describe the MIMIC II (**M**ultiparameter **I**ntelligent **M**onitoring in **I**ntensive **C**are) database, an Intensive Care Unit (ICU) database which is freely available, together with this guide, from:
http://www.physionet.org/physiobank/database/mimic2db.
This document can be viewed as HTML or downloaded as a PDF from the same location. The MIMIC II database was collected as part of a Bioengineering Research Partnership (BRP) grant from the National Institute of Biomedical Imaging and Bioengineering entitled, "Integrating Data, Models and Reasoning in Intensive Care" (RO1-EB001659). The project was established in October 2003 and included an interdisciplinary team from academia (MIT), industry (Philips Medical Systems) and clinical medicine (Beth Israel Deaconess Medical Center). The objective of the BRP is to develop and evaluate advanced Intensive Care Unit (ICU) patient monitoring systems that will substantially improve the efficiency, accuracy and timeliness of clinical decision making in intensive care.

1.2 Background

ICU patients are typically the most physiologically fragile patients in the hospital and may experience prolonged hospital stays with significant morbidity and mortality. The modern ICU employs an impressive array of technologies that results in the generation of a rich–yet disparate–set of clinical and physiologic data used to guide patient care. ICU clinicians are challenged to interpret all the available ICU data to not only improve patient outcomes, but also to contain costs and adopt evidence-based practices. The enormous amount of ICU data and its poor organization make its integration and interpretation time-consuming and inefficient. The data overload that results may actually hinder the diagnostic process, and may even lead to neglect of relevant data, resulting

in errors and complications in ICU care [1]. In the long term, automated or semi-automated monitoring and clinical decision support systems (CDSS) are needed. These systems must be capable of not only presenting ICU data to human users but also of forming pathophysiological hypotheses that best explain the rich and complex volume of relevant data from clinical observations, bedside monitors, mechanical ventilators and the wide variety of available laboratory tests and imaging studies. Such systems should reduce the ever-growing problem of information overload, and provide much more clinically relevant and timely alarms than today's disparate limit-based alarms. While there have been decades of research in utilizing artificial intelligence and expert-systems for medical data processing and forecasting [2], little research has found its way into widely deployed ICU monitoring and information systems. The development of such systems requires access to large volumes of real-world ICU data that can serve as a testing platform to refine and evaluate such algorithms.

The role of a rich and comprehensive database in this context is twofold. First, through data mining, such a database allows for extensive epidemiological studies that link patient data to clinical practice and outcomes. Such insight can in turn motivate the development of alarms, alerts, or algorithms to improve clinical practice and thus improve patient outcomes. Second, it is essential to develop and test algorithms with real data, and to be able to perform such tests repeatedly and reproducibly as algorithm refinements evolve. Within the critical care community, well-known databases including APACHE [3] and Project IMPACT [4] have resulted in the acquisition of hundreds of thousands of ICU patient cases from dozens of hospitals throughout the United States of America. The purpose of such databases is mostly to assess and compare the severity of ICU patient conditions and outcomes, and the costs of treatment across all participating intensive care units on the basis of very few, highly aggregate pieces of information. Such data abstractions often do not include detailed information regarding temporal relationships between therapeutic interventions and corresponding diagnostic data, and thus, would be insufficient to characterize clinically significant transient events such as hemodynamic instability, or acute organ injury.

The detail and volume of data necessary to support such research as described above has been difficult to gather in the past due to limitations on computational processing power, networking bandwidth, digital storage capacities, proprietary vendor data formats, and concerns related to patient privacy [5]. Through a collaborative effort between academia, industry, and clinical medicine, we have attempted to address these aforementioned challenges, and established a major new, publicly available ICU database, MIMIC II.

For more information on the rationale for assembling this database, the original research proposal can be found here:
http://www.physionet.org/physiobank/database/mimic2db/docs/.

This document provides a detailed overview of the formation and contents of the MIMIC II database. The methodology used in data post-processing and the organization of MIMIC II is also described. In section 2 we provide a characterization of MIMIC II with respect to quantitative data specifications as

well as clinical characterizations using standard metrics such as patient acuities, problem lists, demographics, and mortality rates. In section 3 access modalities are described.

1.3 Overview of data collection

The data were collected over a six year period, beginning in 2001, from the Boston's Beth Israel Deaconess Medical Center (BIDMC). Any patient who was admitted to the ICU on more than one occasion may be represented by multiple patient visits. The adult ICUs (for patients aged 15 years and over) include medical (MICU), surgical (SICU), coronary (CCU), trauma (T-SICU), and cardiac surgery (CSRU) care units. Data were also collected from the neonatal ICU (NICU).

Figure 1.1 illustrates the data acquisition process, which did not interfere with the clinical care of patients, since databases were dumped off-line and bedside waveform data and derived trends were collected by an archiving agent over TCP/IP. Source data for the MIMIC II database consists of a) bedside monitor waveforms and associated numeric trends derived from the raw signals, b) clinical data derived from Philips' CareVue system, and c) data from hospital electronic archives. These data are assembled in a protected and encrypted database (both flat files for the waveforms and trends, and in the form of a relational database for all other data). Once the data have been assembled in a central repository and time aligned, the waveforms and trends for each individual are linked to the corresponding individual's data in the relational database. (See section 1.4.3 for more information.) The data are then de-identified to produce a final set of data for public consumption. (See section 1.4.4 and [6] for more information on this detailed process.)

The resulting records contain realistic patient measures with all the associated challenges (such as noise or missing data gaps) that advanced monitoring and clinical decision support systems (CDSS) algorithms would receive as input data. Noise and artifact examples in the database, together with methods for dealing with these problems are described in sections 1.6 and 5.4.

1.4 Data Organization

1.4.1 Types of data

There are essentially two basic types of data in the MIMIC II database; clinical data stored in a relational database, and bedside monitor waveforms and their associated derived parameters and events stored in flat binary files (with ASCII header descriptors), and sorted with one directory per patient. Only a fraction of the total records in the relational database have associated waveform data. (At the time of the initial release of MIMIC II there were 26,655 total admitted patients, 19,075 of which were adults at the time of admission. Of the 2,769 of the waveform records currently available, 2,430 are linked to adult patients.

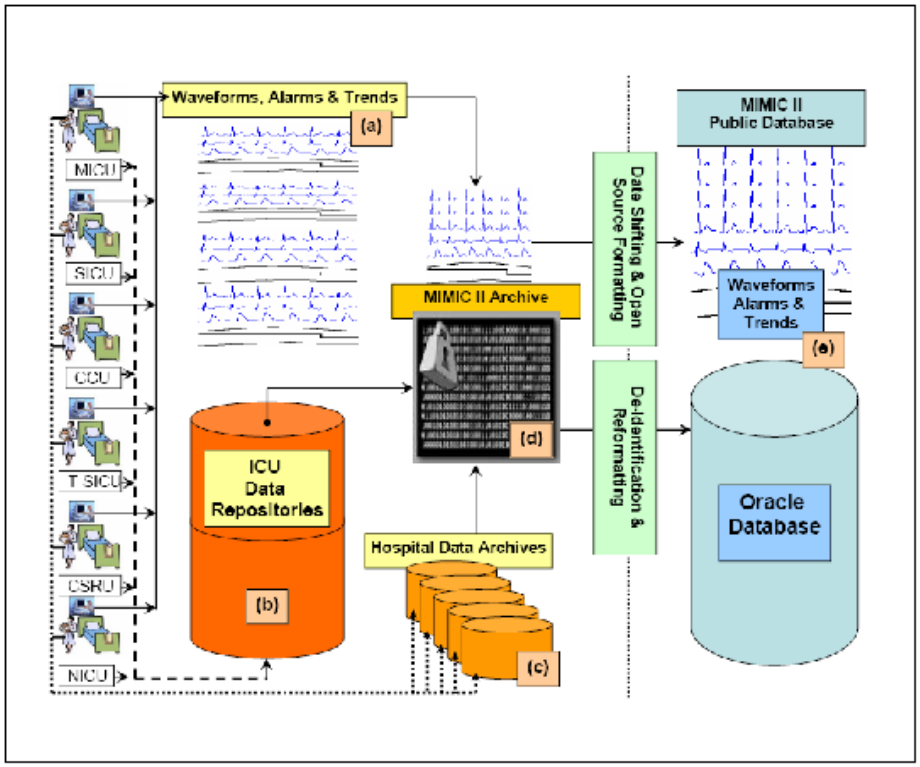


Figure 1.1: Schematic of data collection and database construction. Source data consists of (a) bedside monitor waveforms and trends, (b) the ICU clinical databases and (c) the hospital archives. These data are assembled in a protected and encrypted database (d) which is then de-identified (e) to provide one relational database plus associated flat file bedside waveforms and trends.

Note that we define *waveforms* to be rapidly sampled (125Hz) signals recorded by the bedside monitors such as electrocardiograms (ECG) and arterial blood pressure (ABP) waveforms, illustrated in Figure 1.2. We define *trends* to be a time series of parameters derived from the waveforms by the bedside monitors, such as heart rate, systolic blood pressure, cardiac output and relative oxygen saturation. Of course, time series of repeated clinical measurements are also found in the relational database, such as pH levels, laboratory values and administered medications. Figure 1.3 illustrates a typical set of time series (or 'trends'). The first two channels of data are HR (heart rate) and IBP (invasive blood pressure) which are taken form the flat file trend data. The third trace (NBP: non-invasive blood pressure) is recorded by a nurse from oscillometric cuff inflations and so is sampled much more sparsely. *Events* are automatically generated markers triggered by the bedside monitor algorithms.



Figure 1.2: Typical clean waveform data in the MIMIC II database. From top to bottom: Two leads of ECG (II and MCLI), arterial blood pressure (ABP) and pulmonary arterial pressure (PAP).

These include arrhythmia alarms, error messages (such as cable disconnections) and beat labels. These data are therefore unevenly sampled. Numeric trends are generally produced by the bedside monitors once per second, although after transmission to the central ICU database, they are often stored only once every 5 to 60 minutes. See section 2.3 for more details on these data types. A list of all the possible alarms can be found in table 2.6, ranked by their frequency in the database, together with associated statistics concerning the mean, minimum and maximum values at which the thresholds are set by the clinical staff.

Clinical data are recorded far less frequently than bedside monitor data and come from a variety of databases. These include the laboratory results, pharmacy provider order entry (POE) records, admission and death records, discharge summaries, ICD-9 codes, imaging and ECG reports and the ICU central database (which includes some subset of the bedside monitor trends, drip rates, free text nursing notes and nurse-verified down-sampled trends, amongst other information). A selection of these parameters can be seen in Figure 1.3.

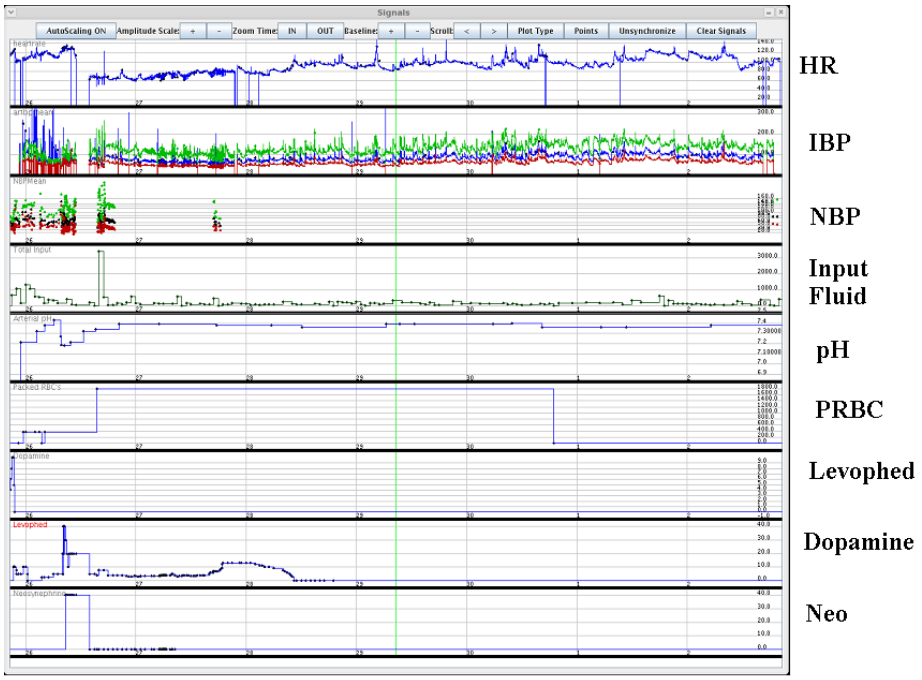


Figure 1.3: Trend data associated with a particular patient stay. Parameters are HR: heart rate, IBP: invasive blood pressure (systolic, mean and diastolic in green, blue and red respectively), NBP: non-invasive blood pressure (with the same color coding), Input Fluid: total fluids given to the patient per hour, pH: acidity/alkalinity of patient, PRBC: packed red blood cell administration, Levophed: Levophed administration, Dopamin: dopamine levels, and Neo: neosynephrine.

A more detailed description of the content of these databases can be found in section 2.1.

1.4.2 What is a patient record?

Since a patient may have been admitted several times during the period in which our data were collected, it is important to understand exactly how to identify patients and their individual patient stays.

There are essentially four identifiers for data associated with any given patient:

- Subject ID (*Subject_ID*) - an integer number identifying a particular patient. This can be thought of as a substitute for a unique medical record number. In the flat file data posted on PhysioNet, the number represent-

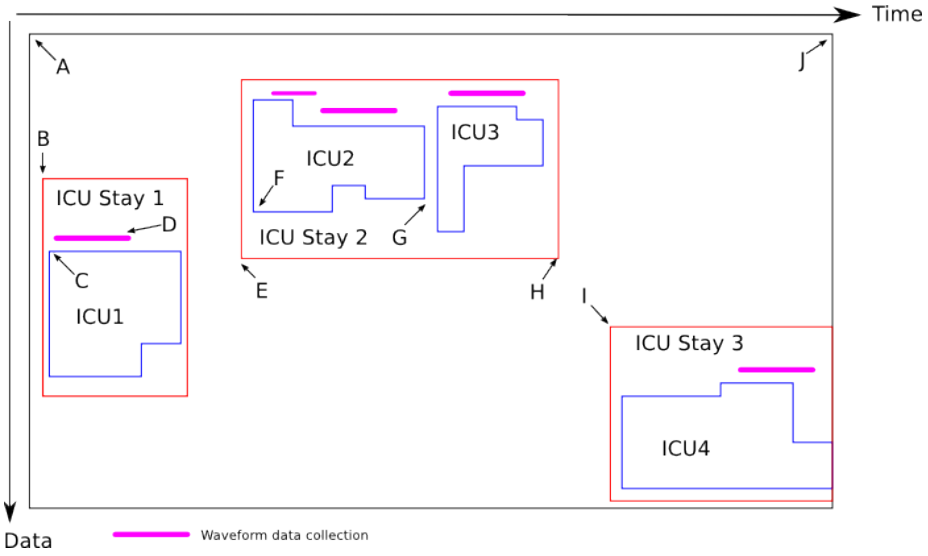


Figure 1.4: Schematic of a patient record. Note that the patient may experience several ICU stays, for which differing amounts of data are available. See section 1.4.2 for details.

ing the Subject_ID is left padded with zeros to five digits and preceded by the letter *s*. In the relational database, the Subject_ID is an integer and therefore has no preceding letter.

- Hospital admission ID (*Hadm_ID*) - an integer number identifying a particular admission to the hospital. Each patient may have many *Hadm_ID*'s associated with their unique *Subject_ID*.

- ICU stay ID (*ICUstay_ID*) - an integer number identifying an ICU stay. An ICU stay, refers to the period of time when the patient is cared for continuosly in an Intensive Care Unit. Each patient may have one or more ICU stays associated. An ICU stay is considered to be continuous if any set of ICU events (such as bed transfers or changes in type of service) belonging to one Subject_ID are less than 24 hours apart. Longer breaks in the patient's stay automatically cause a new *ICUstay_ID* to be assigned.

- Case ID (*Case_ID*) - This is a five digit number preceded by the letter *a* (for adults) or *n* (for neonates). This ID indicates a set of waveforms associated with a given patient. For various reasons (described in section: 1.4.3 below), there may be multiple case IDs associated with a given patient.

Figure 1.4 illustrates the connection between the above labels for a given individual, with each hospital stay surrounded by a red box. Time progresses from left to right, and the type/quantity of information is represented by the

vertical axis. The first piece of information concerning a patient is recorded at time **A** and the patient is then admitted to the first ICU (ICU1) at time **B**. The first ICU data point is collected about the patient at time **C**. Notice also that the amount of information reduces part way through the ICU stay (indicated by the contraction of the vertical width of the blue box). This is common in many patients, where less signals are monitored as a patient improves. Note that at time **D** the waveforms and trends end. This may be because the patient was disconnected from the bedside monitors, prior to discharge, or because there was some issue that led to an interruption in data collection. At time **E** (which is more than 24 hours after the discharge from ICU1, the patient is readmitted to the ICU (although this time it is a different ICU (ICU2). Note that this is not always the case, a patient may be discharged and re-admitted to the same or a different ICU. If the time is less than 24 hours, then the patient is considered to be still in the ICU and no new ICU stay is created, even if they are transferred to a new ICU (e.g. ICU3). After discharge (at time **H**) and readmittance (at time **I**) longer than 24 hours, a new ICU stay (3) is created.

Note that a patient may move between ICUs during any given admission. If the move is longer than 24 hours, we define it to be a new admission. Note also that the amount of data varies during and between ICU stays and that data are often missing - see section 1.6.

Note also that for the preliminary release, any patients that have stays which overlap their 90^{th} birthday have been removed from the database. Patients used in the test set for the PhysioNet / Computers in Cardiology Competition 2009 have also been removed in this first release.

1.4.3 Subject ID - Case ID matching

Given that the MIMIC II data are collected from different sources, they must be matched to a unique patient and temporally aligned. The bedside monitor-generated data included a unique identifier (the *Case_ID*), assigned automatically by the monitor, and fields for patient name (first and last name) and medical record number (MRN). The name and MRN fields were manually entered by nurses into the networked central station when a patient was admitted. Unfortunately in approximately 30% of cases one or more identifier fields were not completed for admitted patients. Moreover, human errors are likely to exist in the manually recorded name and MRNs. The CareVue clinical information system also included a unique patient identifier (that maps to our *ICUstay_ID*) for each ICU stay of a patient. The subject's CareVue data also includes identifying information such as a patient's name and MRN which was automatically input from the hospital-wide information system when a patient is admitted to a unit.

When waveform files included the patient's identifying information (name, and MRN), the physiologic data records (indexed by a *Case_ID*) were matched to the corresponding clinical information records from CareVue. There were two stages to the merging process. The first stage included matching names and medical record numbers (when available and accurately recorded) from

the monitor-generated data records to those of the clinical data records from CareVue. The second stage included comparing the similarity of the physiologic trends from the higher resolution monitoring data (approximately 1 sample per minute) with the nurse-validated vital sign trends in the clinical information system sampled on an hourly basis.

Briefly, determination of trend similarity included four stages:

- Determining a temporal overlap between the available trends that were present in the physiologic and clinical data.
- Identifying if unusual parameters had been recorded at the same time (e.g. cardiac output).
- Correlating median filtered, down-sampled heart rate and blood pressure numeric trends with heart rate and blood pressure trends in the clinical data.
- Visual inspection of a subset of files to verify the results.

However, it is possible that some of the matches may be incorrect. After manual review we believe we have caught most of the inconsistencies, but anomalies may still be present.

1.4.4 De-identification of patients' data

The process for the removal of protected health information (PHI) in the the MIMIC II database is fully described in our publication [6] which can be freely accessed at the following URL: http://www.biomedcentral.com/1472-6947/8/32 A labeled subset of the data, together with a public version of the code can be found on PhysioNet at: http://www.physionet.org/physiotools/deid/.

Figure 1.5 illustrates the de-identification process. Briefly, the salient points for the user of our database are:

- All dates were shifted 10 ± 3.12 years into the future. The date shift for each patient was independently assigned by sampling from a uniform distribution.
- All ICU dates for a given patient were shifted by the same amount to preserve inter-admission time gaps.
- The day of the week and season of the year were preserved.
- Patients who turned 90 during one of their admissions have been removed form the database. They may be included at a later date.
- Patients older than 89 years at the date of first admission have had their dates of birth shifted so that they appear to be 200 years old at the time of their first admission. They will therefore show up as extreme outliers. Their inter-admission timings are still preserved.

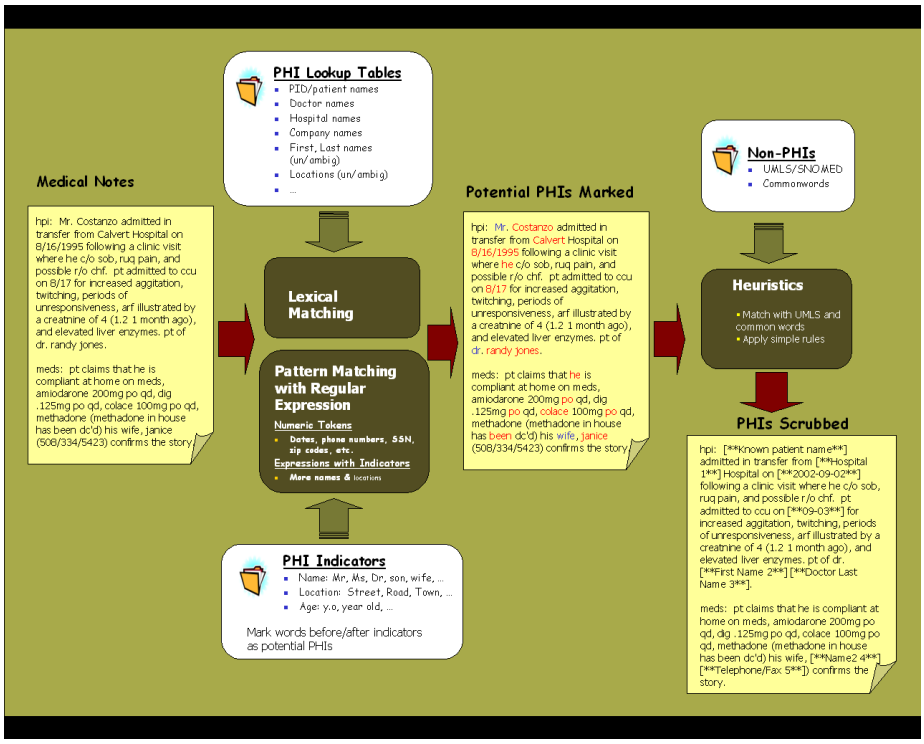


Figure 1.5: De-identification process

- Since date shifts were randomly assigned, longitudinal studies that involve changes in patient care practices over time cannot be supported by the fully de-identified data. Support for studies that require the year of admission will be considered on an individual basis by special request.

- All HIPAA-defined types of PHI were removed, plus care-giver and hospital-specific identifiers.

- The algorithm achieved an overall recall of 0.967 and precision of 0.749 on a 'gold standard' test corpus, which out-performs a single human de-identifier and performs at least as well as a consensus of two human de-identifiers [6].

Examples of a de-identified nursing progress note and discharge summary can be found in figures 1.6 and 1.7 respectively. Note that a few of the de-identified sections of the nursing note are false positives, and a small fraction of the clinical information may have been lost. However, all dates and names (the only PHI in this document) were caught by our algorithm. Note also the the high prevalence of abbreviations such as S/0 (sign out), D/C'd (discontinued,

2013-10-18 22:41:00
NPN 7AM-11PM:
S/O: Pt has had a very eventful day. At~6:45 AM he was noted to have SBP 40's by NBP, with HR 60's. Initially responsive, but rapidly decreasing responsiveness followed by respiratory arrest. Pt was ambued with 100% FiO2, then [**Month/Day **]. An A-line was placed; we have consistently been able to easily draw blood from the line, but it appears dampened and reads quite a bit lower than the NBP, so we have been using the NBP all day. He soon required pressors for SBP 70's. He was started initially on Neo, which was titrated up to a max of 120 mcg/min with little if any effect. he was then started on Levo. Over several hours, with some difficulty, the Neo was weaned to off with the Levo as high as 40 mcg/min. He was transiently on Dopa, as high as 10 mcg/kg/min, but it was soon D/C'd d/t HR into the 140's. Around 1PM his BP again began to fall, into the 50's. His extremities were cold, and HR dropped into the 60's again. He was given 250cc fluid bolus, and Dopa was again attempted, at a lower dose. This time, however, he began to have lots of ventricular ectopy, including short runs of VT. Dopa was again D/C'd, Levo increased more, and he again stabilized for a few hours. About 7:45 he suddenly went into sustained VT. A-line tracing was flat (though is has never been reliable). In the interest of saving time, a cuff pressure was not checked. He was unresponsive, and was defibrillated once with 200J. He converted initially to ST with lots of ectopy, then settled down into NSR after a few minutes. He has remained in NSR since. BP is borderline on high-dose Levo. EKG shows ST depressions, but not much changed from yesterday. CK's, Troponin added to earlier labs.
F/E: Pt is dialysis-dependant. He has had >2.5L fluid since MN, and will be dialyzed tomorrow. Lytes have been followed closely; Mg repleted after episode of VT, and he has been given 15gm Kaexolate for borderline hyperkalemia.
NEURO: Pt initially unresponsive this AM. Over the day he has been agitated with ANY intervention. Initially well-sedated on [**Month/Day 15**], but he was changed to Fentanyl gtt with prn Ativan to try to avoid hypotension from the [**Month/Day 15**]. Fentanyl has been increased a couple of times. He is OK when left alone, but easily agitated.
[**Month/Day **]: Hct 30-32, stable. Coags greatly elevated with INR 5.1 this AM. He was given 2mg Vit K SQ, but coags worse afterwards. No further intervention at present.
GI: Vomitted brown OB+ material both before and after intubation. Belly soft, obese, obviously tender. Too unstable to go to CT. Plan was for U/S, but he was hypotensive to 50's when they came, so it was deferred. Medium loose brown, foul-smelling stool this AM (sent for C-diff). On Protonix.
ID: Temp rising to max of 101.7 this evening. He has been fully cultured and is on multiple abx. Ampho dose which was up when he arrested this AM was stopped with ~half of it infused. He did not recieve the rest.....HO aware. WBC 30-40K, Lactate has risen to 7.9. He has a worsening metabolic acidosis, with bicarb now down to 12.
RESP: Intubated, vented. Current settings A/C .5/750/24/PEEP 5. ABG's show adequate oxygenation, compensated metabolic acidosis. LS diminished. He has minimal secretions, but he was found to have green beans in the back of his throat on intubation, and we have suctioned a few pieces out...none since this AM.
SKIN: He has 2 small decubs on buttocks, covered with Duoderm. Also has open area in left groin.
ACCESS: A-line as described above. He has a right femoral tunneled [**Male First Name (un) 139**] catheter. A clotted left EJ line was removed this AM. Multiple attempts at other access have been made by many people without success.
SOCIAL: pt has a sister [**Last Name (un) 140**] who was in. He also has a very involved home care nurse named [**First Name8 141**] [**Last Name 142**] who was extremely upset about his condition. She was in to visit this evening, and was here for the VT episode. The pt's lawyer also came in briefly. He does not have a proxy; SW notified by case manager of his admission, serious condition, and need for proxy determination.
A: septic shock with multiple potential sources.
P: continue abx, follow cx results. Support BP and resp as needed. Follow labs closely. Anticipate possible need for CVVHD is does not tolerate HD. SW consult for proxy.

Figure 1.6: Example of a de-identified progress note. Sub-headings have been capitalized in bold face type for easier reading. Removed text is denoted by square brackets. True positives are colored green, false positives are colored red.

==========

DISCHARGE SUMMARY

Name: [**Known patient lastname**], [**Known patient firstname**]

[**Unit Number 626**]

Admission Date: [**2016-11-07**]

Discharge Date: [**2016-11-22**]

Date of Birth: [**1972-09-20**]

Sex: F

HISTORY OF PRESENT ILLNESS: Patient is a 44-year-old lady status post living related kidney transplant on [**2016-10-19**], who presented at [**Hospital 36**] for end-stage renal disease secondary to type 1 diabetes mellitus.

She presented to [**Hospital1 **] on [**2016-11-07**] with increased drainage from her surgical wound and JP, increased abdominal pain, and anuria x4 days. The patient reported constipation for a week. She denies flatus. She was complaining of nausea and vomiting. Her abdominal pain had become progressively worse left lower quadrant most notable. There is no radiation to the back or elsewhere. She denied any fevers, chills. She noted decreased p.o. intake recently. Her drainage from her wound incision and JP was notable for yellowish clear urine smelling fluid.

==========

Figure 1.7: Example of a section of a de-identified discharge summary. All de-identified elements are denoted by square brackets. No false positives exist in this example.

or discharged), Neo (neosynephrine), NSR (normal sinus rhythm), F/E (fluid and electrolytes), GI (gastrointestinal), HEME (hematology), ID (infectious disease), A (assessment), P (plan), etc. Note also the low degree of structure in the nursing note, broken into a few categories; **S/O, F/E, NEURO, GI, HEME, ID, RESP, SKIN, ACCESS, SOCIAL, A,** and **P**. The boldface type has been added to this figure to highlight these categories, but is not available in the notes.

1.5 Clinical overview of patients in the MIMIC II database

In the MIMIC II database there are 26,655 patients (with unique *Subject_ID*'s) of which 19,075 were adults (defined to be $\geq$ 15 years old at time of last admission) and 6,538 were neonates ($\leq$ 1 month old at the time of first admission[1]), These patients experienced a total of 29,545 hospital admissions (22,916 adults and 6,585 neonates). However, there were only 29,505 ICU-related hospital admissions (with 22,880 for adults and 6,582 for neonates). That is, a small fraction of the patients was admitted but never made it to the ICU for various reasons.

[1]Note that there are 4 other babies in our database admitted to the NICU that were aged between 2 and 5 months inclusively. Note also that 1200 other patients were aged between 6 months and 14 years inclusively at the time of admission, or did not have a date of birth recorded.

Table 1.1: Distribution of major categories of primary ICD-9 codes for adult ICU-related hospital admissions (N=22,880).

Category	**Code Range**	**Number of Admissions**	**%**
Ischemic heart disease	410 - 414	4084	17.85%
Trauma	800 - 959	2337	10.21%
Digestive disease	520 - 579	2228	9.74%
Pulmonary disease	460 - 519	2070	9.05%
Infectious diseases	001 - 139	1603	7.01%
Neoplasms	140 - 239	1562	6.83%
Cerebrovascular disease	430 - 438	1426	6.23%
Other forms of heart disease	420 - 429	1367	5.97%
Arteries and veins	440 - 459	783	3.42%
Complications peculiar to certain specified procedures	996	674	2.95%
Heart failure	428	656	2.87%
Metabolic disorder	240 - 279	634	2.77%
Renal insufficiency	580 - 629	448	1.96%
Poisoning	960 - 989	356	1.56%
Other complications of procedures, NEC	998	349	1.53%
Chronic rheumatic heart disease	393 - 398	300	1.31%
Symptoms, signs, and ill-defined conditions	780 - 799	271	1.18%
Neurologic disease	320 - 389	255	1.11%
Diseases of the musculoskeletal system & connective tissue	710 - 739	247	1.08%
Diseases of pulmonary circulation	415 - 417	175	0.76%
Mental disorders	290 - 319	153	0.67%
Congenital anomalies	740 - 759	153	0.67%
Complications affecting specified body systems, not elsewhere classified	997	150	0.66%
Complications of pregnancy, childbirth, & puerperium	630 - 677	145	0.63%
Hypertensive disease	401 - 405	142	0.62%
Supplementary classification of factors influencing health status and contact with health services	V01 - V86	92	0.40%

continued on next page

continued from previous page			
Category	**Code Range**	**Number of Admissions**	**%**
Diseases of the blood and blood-forming organs	280 - 289	88	0.38%
Diseases of the skin and subcutaneous tissue	680 - 709	64	0.28%
Other and unspecified effects of external causes	990 - 995	54	0.24%
Without primary ICD9 code		6	0.03%
Complications of medical care, not elsewhere classified	999	6	0.03%
Acute Rheumatic fever	390 - 392	2	0.01%
Total		**22,880**	**100.00%**

The total number of ICU stays (number of *ICUstay_ID*'s) was 33,361 (25,852 for adults and for 7,302 neonates), giving an average of 1.25 hospital stays per patient; 1.36 for adults and 1.12 for neonates. Details on how to calculate these numbers are given in chapter 2. Of the adult population, 38% was from the MICU, 26% from the SICU, 20% from the CSRU and 15% from the CCU. 2,430 *Case_ID*'s (with associated waveforms and trends) have been associated with unique patients (*Subject_ID*'s) Although we have collected neonatal waveforms, these will likely be released at a later date when the associated clinical data has been matched and verified. The rest of the description of the data in this section is therefore limited to the adult population.

Table 1.1 lists the thirty-one major primary[2] ICD-9 code categories for hospital admissions and the relative frequency for the adult patient population. Note that of the 22,916 adult hospital admissions (*Hadm_ID*), there was a total of 22,880 that were admitted to the ICU and therefore received ICD-9 codes. Note also that there are a range of codes associated with each condition. Furthermore, many non-primary codes are assigned to each ICU-related hospital admission with an average of 9 codes per admission. Table 6.1 in appendix 6.4 lists the frequency of all the ICD-9 codes (including the primary codes).

Figure 1.8 illustrates the age distribution of adult patients at the time of each ICU stay (N=25,852). Therefore, a patient will be represented more than once if they were admitted to the ICU more than once. The mean age at time of admission was 63.44 years and the median age was 65.33 years. Note small drops in admission rates for patients during their mid 20's to mid 30's, and again around 65 years of age. The distribution of length of stay for each ICU admission is illustrated in figure 1.9. Note that the distribution is heavy tailed, with a mean of 4.57 days and a median of only 2.15 days.

Figure 1.10 illustrates the distribution of patients' Simplified Acute Physiology Score (SAPS I) at each admission (lower plot) and the associated *hospital*

[2]The primary ICD-9 code was assigned by the hospital and should be the main reason the patient was admitted.

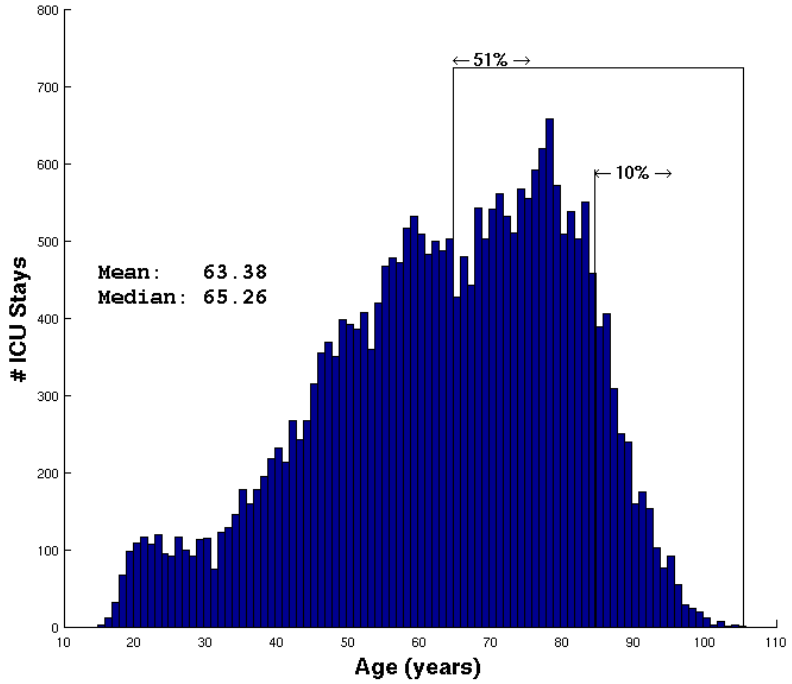


Figure 1.8: Age distribution of adult patients at time of ICU admission in MIMIC II database. (N=25,852)

mortality rate[3] (upper plot). Note that the mortality rate begins to climb steeply around a SAPS I value of 23. Although the mortality rate begins to fluctuate wildly above this value, this is likely an artifact of low numbers. The mean SAPS I was 13.68 and the median was 13. Figure 1.11 illustrates how this mortality varies during a hospital stay by showing the percentage of patients in the MIMIC II database that die in the ICU, and after discharge to the floor (HOSP) for all admissions (upper left) and for the subset with CMO (Comfort Measures Only) (lower right). Notice that mortality curve as a function of SAPS 1 values is not perfectly smooth, particularly above a SAPS 1 value of 20. Above a SAPS 1 value of 30, the mortality rate fluctuates wildly, but this may be an artifact of the small numbers of individuals admitted with such a value.

[3]We only have statistics for who died in the ICU and the hospital, but not after discharge.

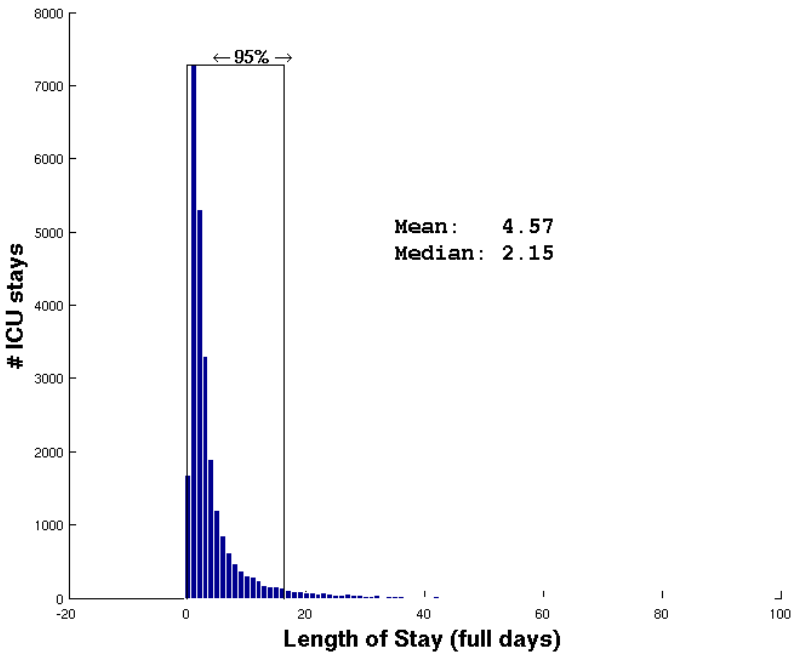


Figure 1.9: Length of stay of patients in MIMIC II database.(N=25,852)

1.6 Noise, artifacts and missing data

The data we have collected is highly representative of that which can be found in the ICU, and therefore is replete with noise and artifacts due to patient movement, sensor degradation, transmission errors, electromagnetic interference and human error.

We have begun to develop signal quality indices to label useful and noisy sections of the data. Currently we have signal quality indices for the ECG and blood pressure waveforms, and we are close to completion for the pulse oximeter and the respiratory waveforms. A more detailed description of the algorithms for signal quality can be found in section 2.3.5 and in Zong *et al.* [7], Sun *et al.* [8] and in Li *et al.* [9, 10].

Signal quality indices for trend data are almost impossible to generate, except by using thresholds on gradients and absolute values that are physiologically impossible. Generally it is better to refer back to the original underlying waveform to derive a signal quality metric.

Data are also missing due to machine or patient disconnections, transmission and recording errors, or human omissions. Some data are also not requested very frequently, and so, although not technically missing, important events may go

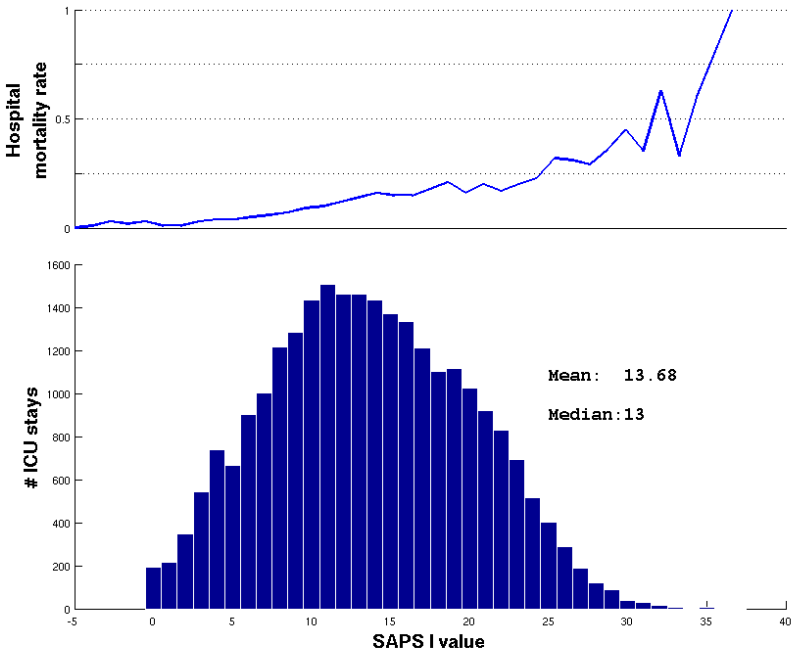


Figure 1.10: Distribution of patients' Simplified Acute Physiology Score (SAPS) I in the MIMIC II database at time of admission (lower plot) and associated hospital mortality rate (upper plot) at end of admission.(N=25,852)

unobserved.

Although short-term missing data can be mitigated somewhat through interpolation, much of our non-waveform and trend data is sparsely sampled. Moreover, the data is not missing at random, since it can be due to changes in shifts or staff-to-patient ratios, or simply because a clinician or nurse did not think that the data were important. Interpolation, or imputation is therefore impossible, unless a model of how the data are missing can be constructed.

Apart form these problems, there may also be errors in the data matching and alignment. Section 5.4 details these and other known issues with the data.

1.7 Summary and further reading

Details of the issues surrounding data collection and annotation can be found in [5] [11] and [12], and details related to de-identification can be found in [6]. Although the database described in this document is large and detailed,

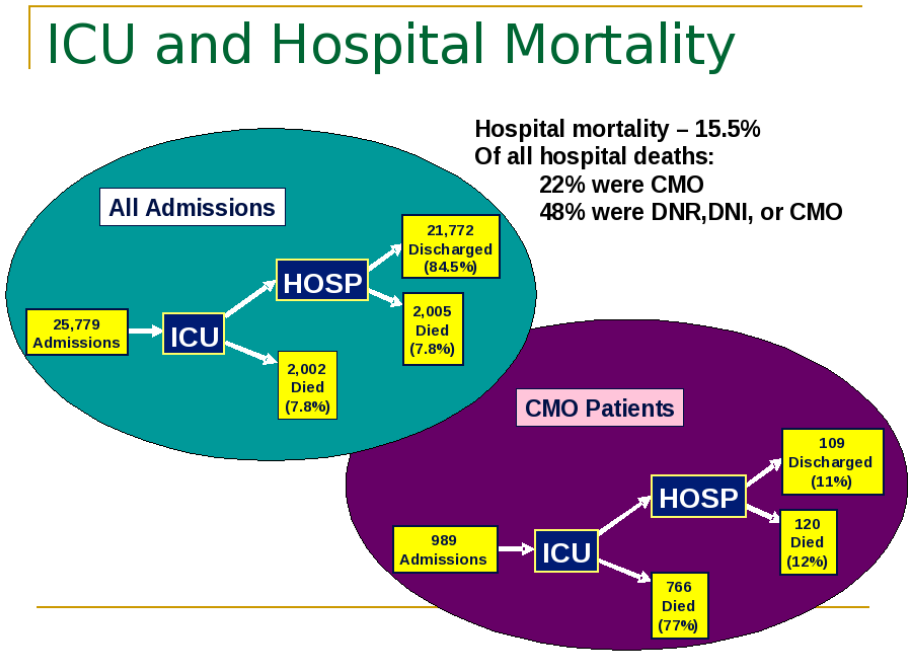


Figure 1.11: Percentage of patients in the MIMIC II database that die in the ICU, and after discharge to the floor (HOSP) for all admissions (upper left) and for the subset with CMO (lower right). CMO = Comfort Measures Only, DNR = Do Not Resuscitate, DNI = Do Not Intubate. (N=25,852)

caution should be taken when analyzing the data, since the clinical systems for monitoring and archiving the data are not perfect and some data may be incorrect.

Moreover, we were unable to capture all of the data associated with a patient's stay, since we did not have data collection facilities in the OR, ER, or step-down wards. Neither were we able to obtain data from other hospitals for any given individual, so visits to outpatient care, or to other hospitals will not be reflected in a patient's stay. Finally, no patient data before the start of the project is available, other than the pre-existing history recorded at a particular stay. Despite these limitations, there is a wealth of data to be found in MIMIC II database.

Note that this document does not provide detailed information on how to extract data from the WFDB flat-file versions of the data found in the MIMIC II database. More information on this can be found here:

http://www.physionet.org/physiobank/tutorials/using-mimic2/

and here:

http://www.physionet.org/physiotools/getting-started.shtml.

Chapter 2

Database Description

2.1 Overview

The MIMIC II database is composed of two distinctive groups of data. The first group, the clinical database, consists of data integrated from different information systems in the hospital and contains diverse information such as: patient demographics, medications, results of lab tests and more. The second group, contains high resolution waveforms recorded from the bedside monitors in the intensive care units. Chapter 3 describes the libraries and methods to access the database.

2.2 Clinical database

The MIMIC II clinical database, is a relational database. Although not strictly required, some familiarity with the structured query language (SQL) will help to understand how the data is stored and obtained. The purpose of SQL is to provide an easy interface to a given database. In relational databases, a typical SQL sentence has the following structure:

```
SELECT column_list
FROM datasources
WHERE constraints
```

Where:

- The SELECT statement, specifies which columns should be returned.
- The FROM statement, denotes the name of the tables where the data is stored.
- The WHERE clause, filters the data retrieved by some criteria.

The query in listing 2.1, shows a real case example: "Extract basic information for a subset of patients". The result of the query is shown in Table 2.1. Although this query is relatively simple, most of the queries throughout this document follow the same structure. The queries presented in the following sections were written with the goal to be illustrative and are not optimized in any way. Detailed explanation of SQL language and query optimizations are beyond the scope of this document, but can be easily be found in SQL [13].

Listing 2.1: Extracting basic information about patients.

```
SELECT subject_id, sex, dob
  FROM d_patients
 WHERE subject_id in (7049, 7060, 7072,
       7078, 9181, 9185, 9195)
```

Figure 2.1 summarizes the major database components, their corresponding attributes and how they relate with a particular patient. The full database schema is much more complex than the one displayed in Figure 2.1, a full description of each table, along with column data types, naming convention, indices, relationships and constraints are provided in the Appendix (Section 6.1).

In the following sections, we provide a general overview, the list of database tables involved, and some sample data for each major group.

2.2.1 Patient

A patient in the MIMIC II database is uniquely identified by an integer number called *Subject_ID*, it can be thought of as a medical record number (MRN) normally found in hospital information systems. The basic information for any given patient is stored in the table D_PATIENTS, normally referred as the patient table throughout this document. As the database went through a careful de-identification process, the patient table only stores the patient identifier (*Subject_ID*), gender (sex) and date of birth (dob, shifted). Table 2.1 shows a sample content of the patient table resulted from the query in listing 2.1.

As shown in Figure 2.1, the patient identifier (*Subject_ID*) is widely used by most of the tables throughout the database to specify to which patient a given measurement or recording refers to. Figure 2.2 shows an example relating which diagnosis codes (ICD-9) were assigned to a given patient, the *Subject_ID*

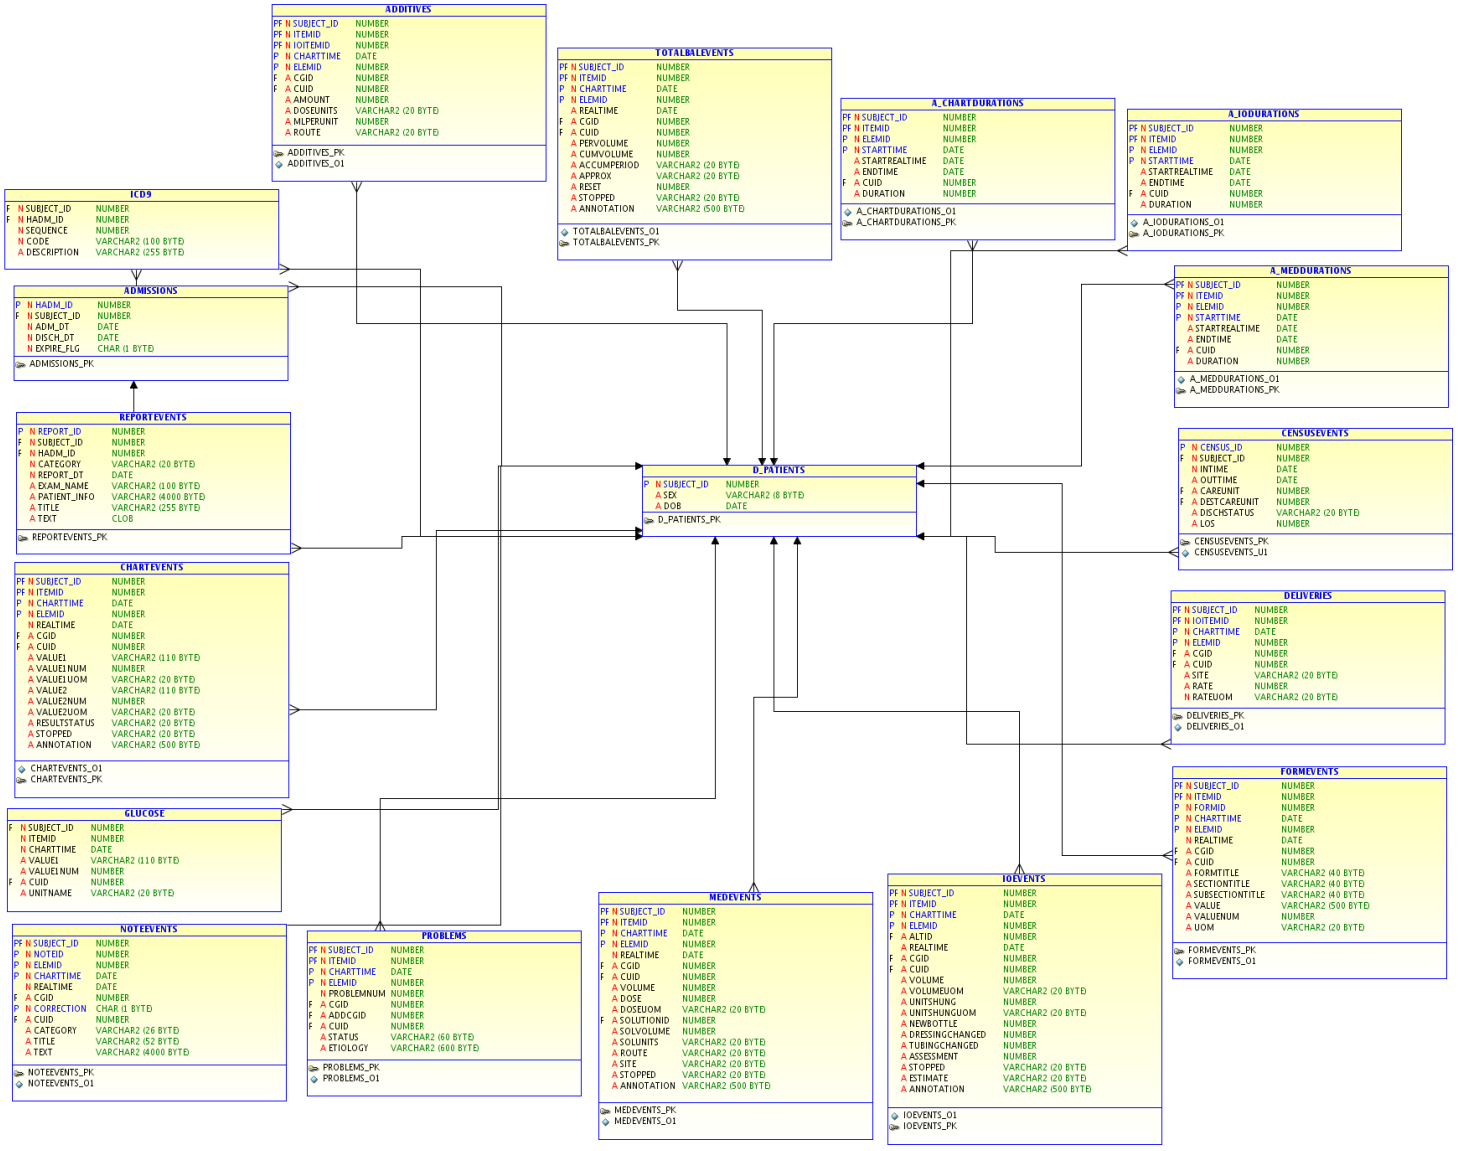


Figure 2.1: Major MIMIC II clinical database components. The patient table D_PATIENTS is central to the database model, other information such as admission or medications can be readily accessed once a particular patient is identified. The tables above show the relationship between the major components in the database.

SUBJECT_ID	SEX	DOB
7049	M	04/10/1952
7060	F	08/01/1932
7072	M	02/22/1928
7078	F	11/11/1967
9181	F	03/11/1960
9185	M	02/28/1927
9195	F	12/19/1974

Table 2.1: Sample content of the patient table. Each patient has a *Subject_ID*, gender and date of birth.

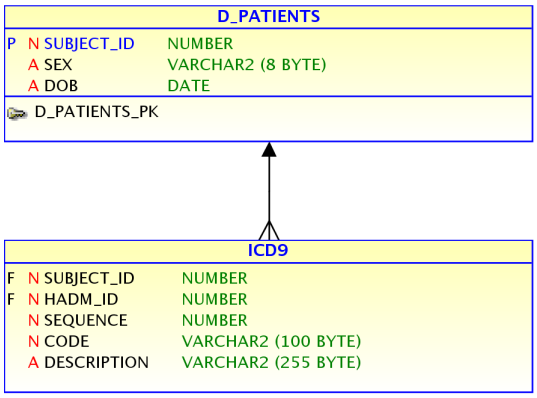


Figure 2.2: Patient to ICD-9 relationship. The patient's ICD-9 codes are related to the patient record in the patient table via a foreign key on the *Subject_ID* field.

field links the ICD-9 and the patient tables. The ICD-9 table records the ICD-9 codes applied to a particular patient during a specific hospitalization period.

2.2.2 Care Giver

Caregivers are the medical staff who are responsible for patients during their hospital stay, such as: nurses, residents or other clinicians. The caregivers are stored in the D_CAREGIVER table, and are uniquely identified by their care giver id (cgid). The caregivers table is related to many other tables such as medevents, noteevents and chartevents and is used to record the care giver who performed a particular operation, procedure or event. Figure 2.3 shows the inter

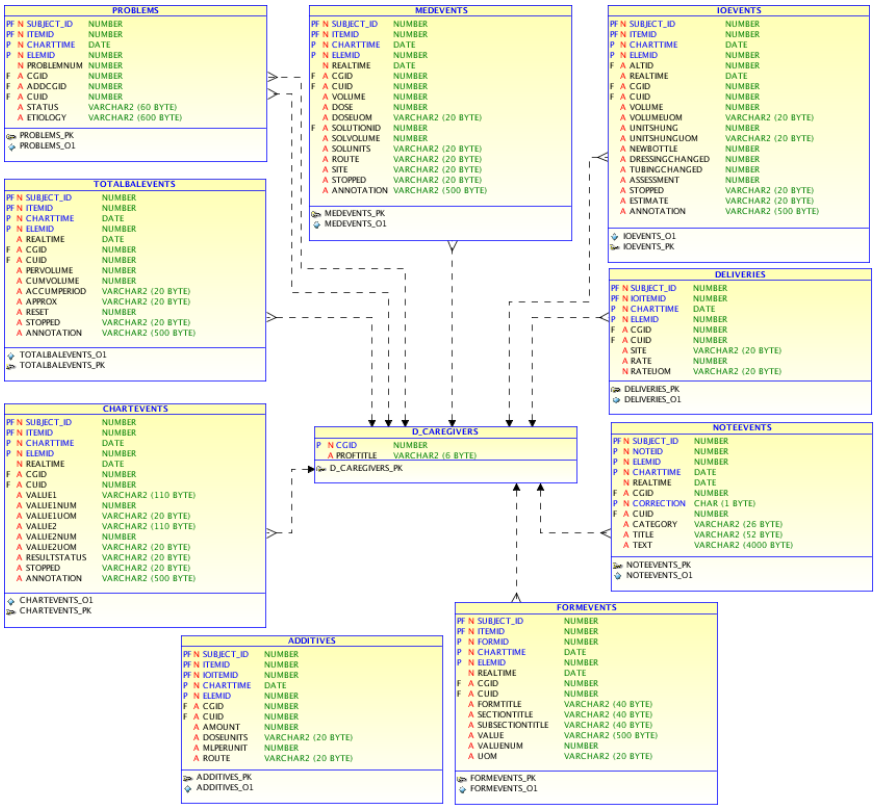


Figure 2.3: Caregivers table relationships. Although a simple table with only 2 columns, the caregivers table is related to many other tables in the database. Caregivers are assigned to *inter alia* medical events, problems and notes.

table relationships of the caregivers table.

2.2.3 Care Unit

The Careunits table, stores information pertaining to the different ICU rooms in the hospital. Figure 2.4 shows the careunit table and its relationship with other event tables in the database. When a note is filed, a problem occurs or a chart event is entered, the particular care unit in which the event took place is recorded. The careunit table only contains a cuid (care unit id) and the unit name, but is used in many places to record the location in which events took place. Table 2.2 shows some sample content.

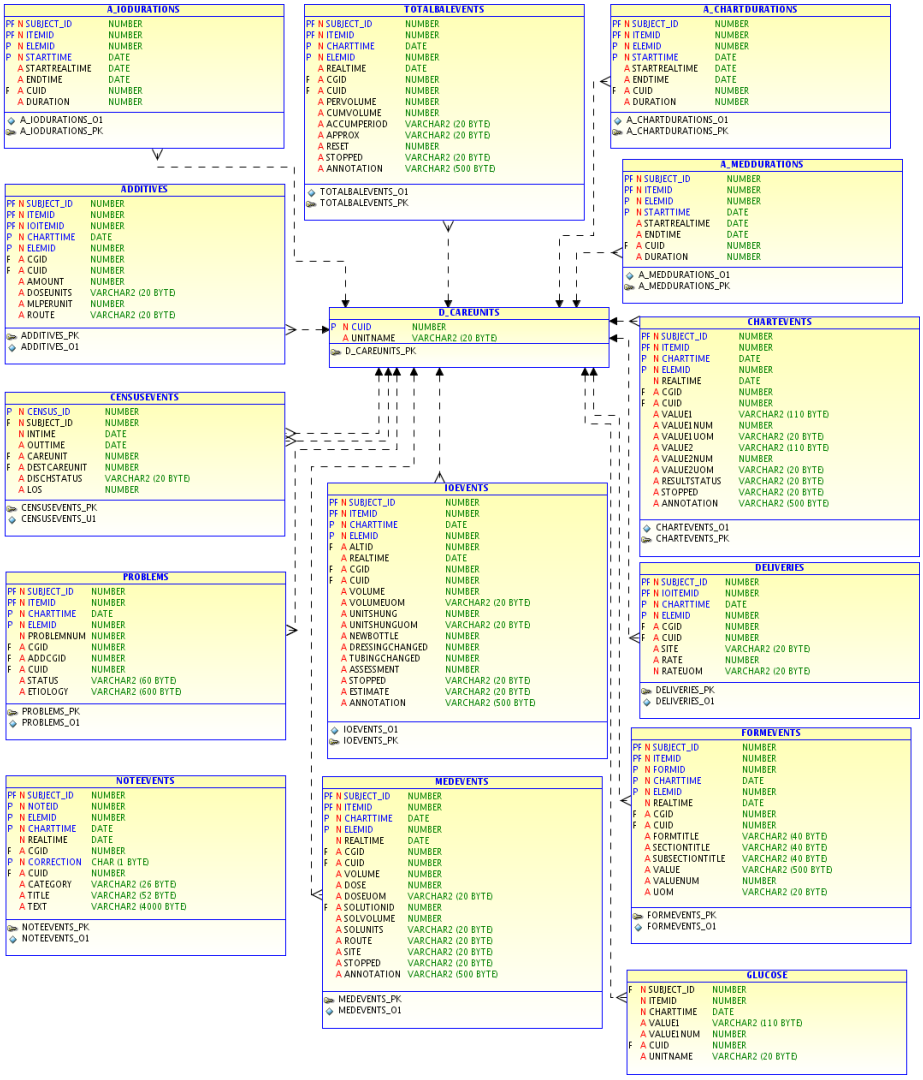


Figure 2.4: Careunits table and relationships. The careunits table is related to many other tables and contains information about the particular care unit where an event occurred.

UNIT NAME
CCU
CSRU
Nursery
Regular ward
Labor and Delivery
Outside hospital Cardiac ICU

Table 2.2: Sample content of the care units table.

2.2.4 Patient timeline

When patients arrive at the hospital, and during their course of their stay, they can be transferred between different units, sent to the operating room for surgery, sent to the floor for recovering or undergo other procedures. To better describe these events, Figure 2.5 shows an excerpt from a discharge summary for a typical patient.

Figure 2.6 is the visual representation of the events presented in Figure 2.5. We can identify the following events:

Hospital admission :

A hospital admission, covers the period from the patient's admission to the hospital, until the patient's discharge from the hospital. It includes any visits to different wards (such emergency room, regular floor, and even different stays in an ICU room).

Patient admissions are recorded in the admissions table. As well as recording the *Subject_ID* of the admitted patient, each admission has a unique identifier (Hadm_ID) an admission and discharge time, and a flag (expire_flg) noting if the patient died (Y) or not (N) in the hospital.

Figure 2.7) shows the relationship between the admission table and other tables in the database.

ICU Stay :

An ICU stay is a combination of one or more ICU census events that are separated by 24 hours or less.

ICU census event :

Each time a patient enters or leaves a particular care unit, an event is recorded into the database in the table censusevents. Each of these events (identified by the column census_id), contain the time and date of entrance and exit of the care unit, the current unit the patient was hosted, the destination care unit the patient was transferred to, and the length of stay in the ICU room for that particular event.

...Chief Complaint:
74 year old female admit to [**Hospital1 80**] [**Hospital Unit Name 26**] [**2018-05-16**] in resp distress, pna, UTI, mild CHF initially on NRB, but then intubated on [**2018-05-18**] (extubated [**2018-05-23**]). Hosp course noted for bradycardia (AV block) during swan placement, CHF. PMH: recent MI, CHF, a fib, CVA, GERD, gastritis, TIAs, Bell's palsy, lower GI polyps...

* * *

...transfered from outside hospital status post embolecotomy for R brachial emboli with history of severe aortic stenosis and anemia for cardiac work-up. Patient of Dr. [**Last Name (STitle) **], found to have colonic polyps on colonoscopy for anemia work-up at OSH. Pt admitted to receive medical clearance for future procedure. Found to have a urinary tract infection with signs of sepsis severe respiratory difficulty and severe aortic stenosis...

* * *

...On [**2018-06-27**], Ms. [**lastname 6384**] was taken to the operating room where she underwent an aortic valve replacement utilizing a 21mm [**Last Name (un) **] [**Doctor Last Name **] pericardial bioprosthesis. Postoperatively she was taken to the cardiac surgical intensive care unit for monitoring....

* * *

...She developed atrial fibrillation and underwent cardioversion on [**2018-06-30**]. Ms. [**lastname 6384**] was only able to hold a normal sinus rhythm for less then two minutes and amiodarone was started. Heparin and coumadin were started for anticoagulation with the plan for a repeat cardioversion in a month. Tube feeds were started for nutritional support and calorie counts were started. On postoperative day seven, Ms. [**lastname 6384**] was transferred to the cardiac

Figure 2.5: An excerpt from a patient's discharge summary, describing typical events during the patient stay in the hospital.

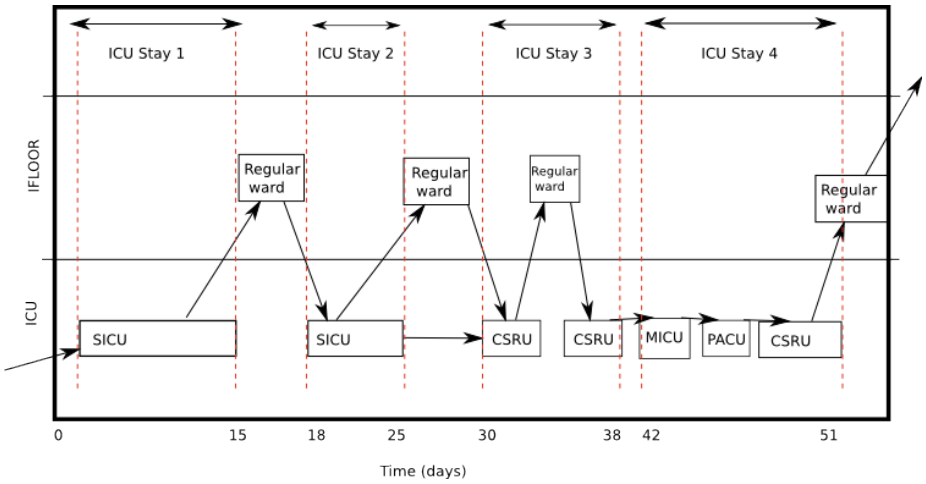


Figure 2.6: Typical events during the patient hospitalization period.

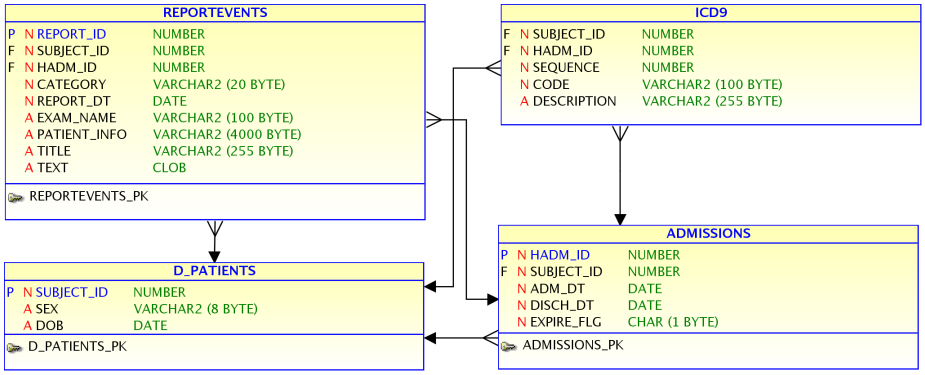


Figure 2.7: Relationship between a hospital admission and other database tables.

2.2.5 Patient data

During a patient's hospital stay, various information is collected about a patient. Demographics, vital signs, laboratory tests, medications, fluid balance, nursing notes, imaging reports, etc. can all be recorded in the database.

- Medication records: Medications prescribed and administered either via computer controlled iv or oral. Computer controlled administration via iv is automatically administered. Although manual prescription is recorded for pharmacy orders, there is no guarantee that the medication was actu-

ally administered to the patient.

- Fluid records: Fluids withdrawn/administered from/to a patient are also recorded. This provides physicians with an up-to-date and accurate measure of a patient's fluid levels.
- Notes: Notes, including discharge summaries are recorded in free-text fields in the database. Nurses can enter any information here.
- Reports

Events

There are numerous event tables which record events occurring throughout the hospital stay.

- MedEvents
- ChartEvents
- CensusEvents
- IOEvents
- NoteEvents
- ReportEvents

Durations

The durations tables record the duration of events. As well as recording the patient, event type and element recorded, the duration tables record the start and end times.

- A_MedDurations
- A_ChartDurations
- A_IODurations

Items

The items tables record the items which can be recorded for a particular event. As such, they are related to both the corresponding event and duration table. Each item has a unique itemid, as well as a label and a category.

- D_MedItems
- D_ChartItems
- D_IOItems
- D_ParamMap_Items
- D_ProblemItems

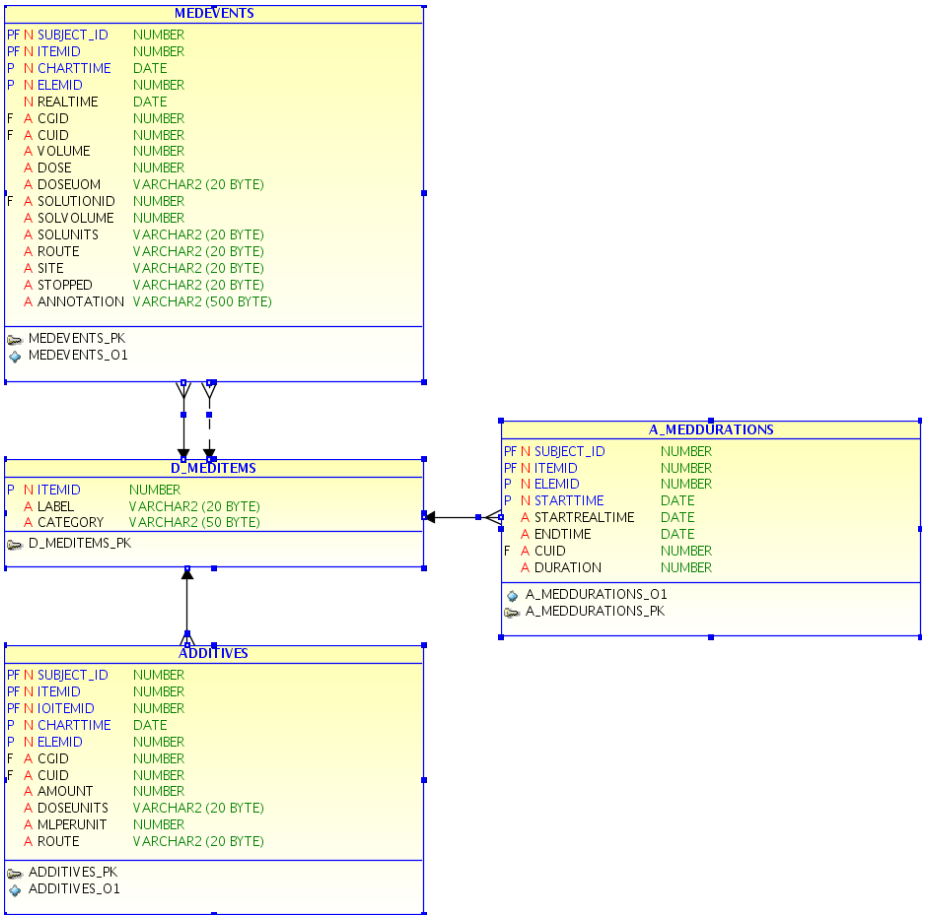


Figure 2.8: Patient medication is stored in 4 tables. The medevents, d_meditems, a_meddurations and additives tables record all data related to patient medication.

Medications

Medication(s) given to a patient are recorded in the medevents, d_meditems, a_meddurations and additives tables. The relationships are shown in Figure 2.8.

The tables contain details of the available drugs and information related to the particular administration. The med_events table contains information pertaining to the patient, dosage and annotation. The additives table contains information related to the additives which are included with the drug administration. The a_meddurations table contains information about the time and duration of the medication and the d_meditem table links the other tables together.

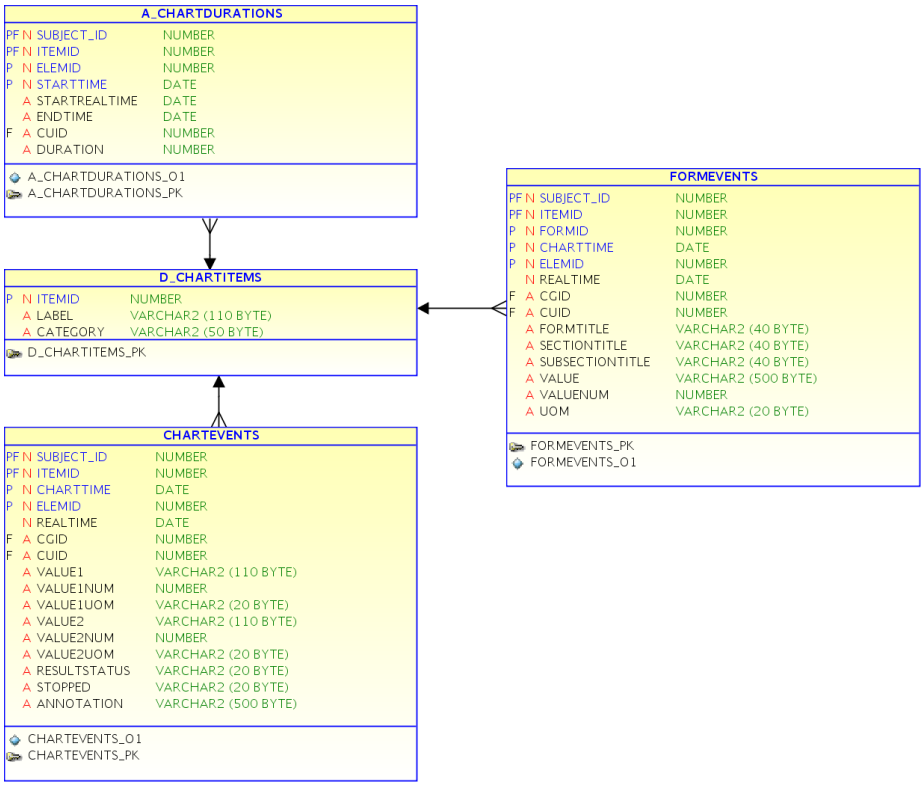


Figure 2.9: Patient chart data is stored in 4 tables. The chartevents, d_chartitems, a_chartdurations and form events tables record all data related to patient charts.

Charts

Patient medical chart data is recorded in the chartevents, d_chartitems, a_chartdurations and formevents tables. The relationships are shown in Figure 2.9.

Fluids

Patient input/output (IO) data is recorded in the ioevents, d_ioitems, a_iodurations, deliveries, totalbalevents and additives tables. The relationships are shown in Figure 2.10.

Notes

Patient notes are recorded in the noteevents table. The relationships are shown in Figure 2.11.

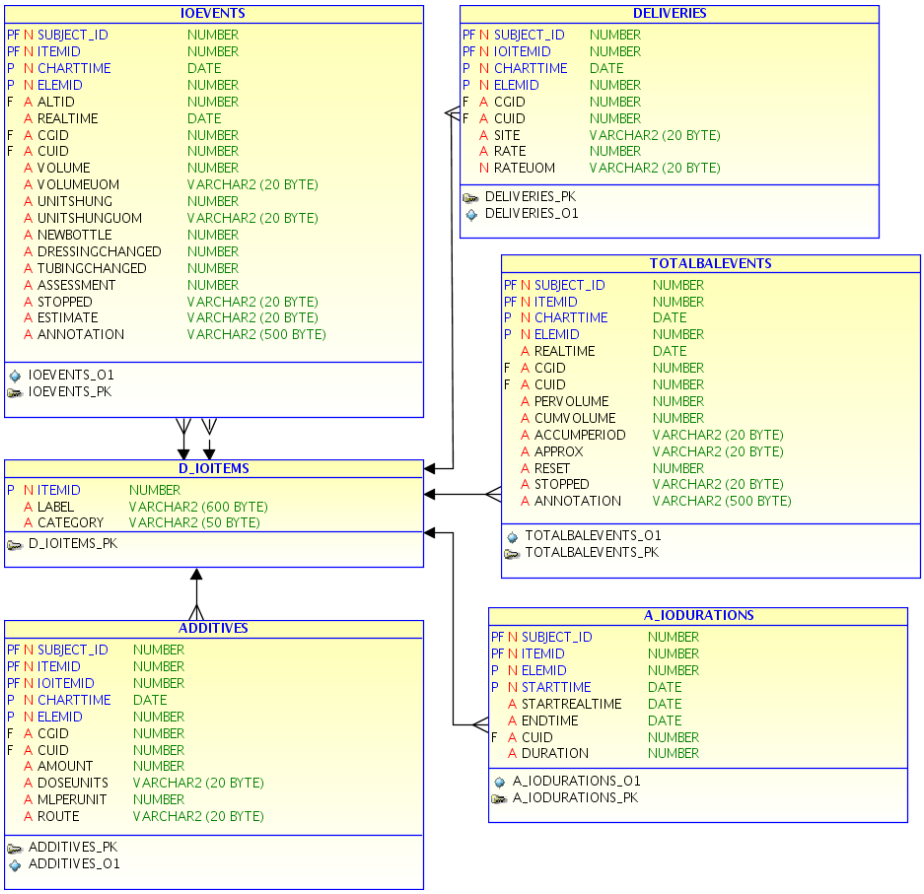


Figure 2.10: Patient IO data is stored in 6 tables. The ioevents, d_ioitems, a_iodurations, deliveries, totalbalevents and additives tables record all data related to patient charts.

During the course of the patient stay in the ICU, free text "notes" are produced by the hospital staff. Nurses typically write "nursing notes", a summary of events which occurred during their shift period. When the patient is discharged from the hospital, the responsible physician dictates a summary of the entire hospitalization period, known as the "discharge summary". These reports are recorded in the MIMIC II database.

Progress or Nursing notes and discharge summaries are stored in the noteevents table and are linked to patients through the *Subject_ID*. Notes are linked to care_givers and care_units through their unique identifiers. The noteevent table also contains relevant meta-data such as the data and time of entry and various other categories and IDs. A sample discharge summary has been

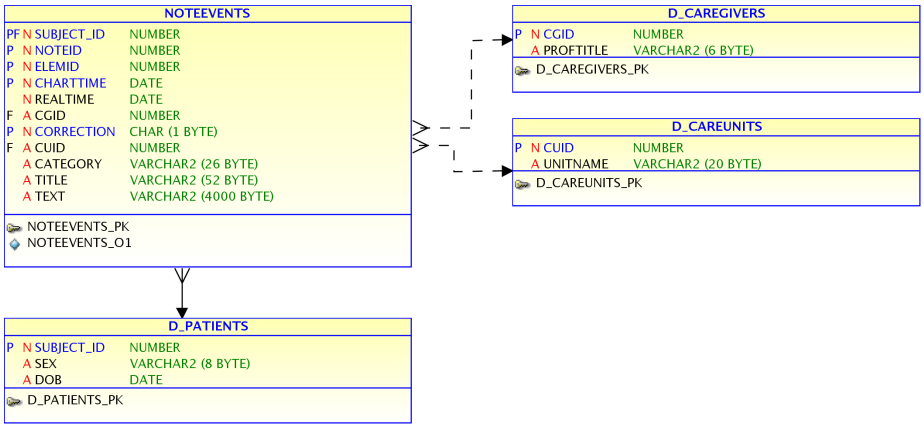


Figure 2.11: Patient notes are stored in the noteevents table.

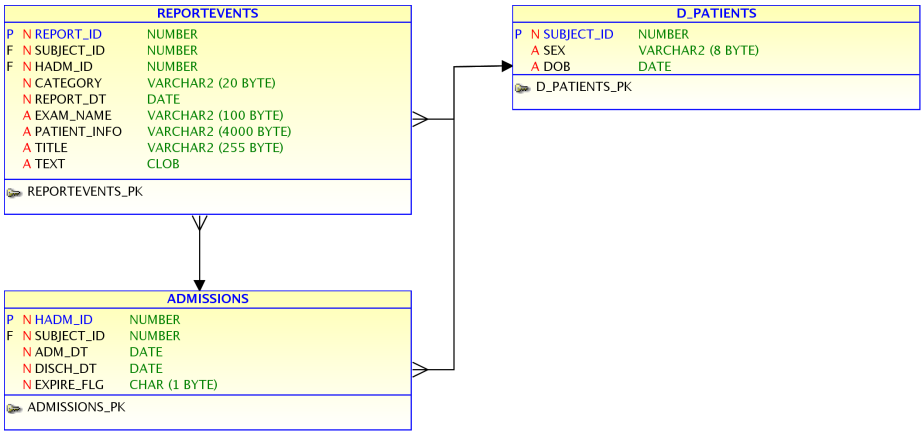


Figure 2.12: Report table schema and relationships.

shown previously, in Section 1.7. The data is free text and can contain anything entered by the nurse.

Reports

Reports are the output of various diagnostic equipment. Reports from X-rays, echos and ECGs are stored in dedicated tables. The reportevent table also contains relevant meta-data such as the data and time of entry and various other categories and IDs. Figure 2.12 shows the reportevents table schema and its relationships.

Figure 2.13 shows a sample radiology report. The data is free text and contains information obtained from radiology.

2.2.6 Summary

All of the information described so far can be thought of "discrete" patient data. It is generally recorded manually and only requires infrequent updates. For example, admission/discharge only occurs once during a patient stay. ICU transfer will only occur a few times. Medication will only occur a few times a day and reports will be added when particular diagnostics are performed.

In contrast, the high resolution waveforms which are discussed next are recorded constantly. Measurements are recorded automatically by computer 125 times per second. The relational database described above is a poor device for recording data of this type and a separate system is used to store these waveforms.

2.3 High resolution waveforms and associated trends

2.3.1 Overview

High resolution waveforms were converted from a proprietary format into MIT's WFDB format (explained later in this chapter). Our waveform collection efforts were essentially broken into two groups, which reflect different versions of the Philips waveform archiving software. The first group includes approximately 2,800 records, each consisting of:

- Up to 4 simultaneous channels sampled at 125 Hz, with an amplitude resolution of 8 bits, normally containing 2 leads of ECG, ABP and a pulmonary artery pressure (PAP) if available. If the PAP was not available, a respiratory signal (impedance pneumogram) an oxygen saturation waveform (photoplethysmogram) was recorded.
- Up to 30 parameters recorded once a minute with an amplitude resolution of 8 bits . Parameters include heart rate (HR), systolic/mean/diastolic blood pressure (SBP/MBP/DBP), peripheral oxygen saturation (SpO2), and cardiac output (CO).
- All bedside monitor-generated alarms (Such as arrhythmias) and in-ops (such as sensor disconnects).

Upgrades in hardware and software allowed us to improve data collection and gather a second group of records to include:

- Up to 7 simultaneous channels sampled at 125 Hz.
- Analog to digital resolution of 10 bits.

```
Reason: CHECK ETT TUBE PLACEMENT, ?PNA, CHF
[**Signature 1**]
UNDERLYING MEDICAL CONDITION:
 85 y/o male s/p acute mi and catherization now in ccu with
cardiogenic shock
REASON FOR THIS EXAMINATION:
CHECK ETT TUBE PLACEMENT
?PNA
CHF
[**Signature 1**]
 FINAL REPORT
CLINICAL INDICATION:  Assess endotracheal tube placement in
patient with congestive heart failure.

Comparison is made to previous study of one day earlier.
An endotracheal tube is present, in satisfactory position.
A Swan-Ganz catheter terminates in the proximal left
pulmonary artery and has been withdrawn in the interval.
An intraaortic balloon pump terminates about 3.3 cm below
the superior aspect of the aortic knob, and a nasogastric
tube terminates in the region of the gastroduodenal
junction.

Cardiac and mediastinal contours are stable in the interval
and pulmonary vascularity is within normal limits for
technique.  There has been improvement in the left
retrocardiac opacity and there remains a patchy right
basilar opacification which is slightly increased.  A small
amount of fluid is seen in the minor fissure.

IMPRESSION:

1) Lines and tubes in satisfactory position, as detailed
above, with no evidence of pneumothorax.

2) Improved left retrocardiac opacity and worsened right
lower lobe opacity likely due to atelectasis.

JPE

   DR. [**First Name11 25**] [**Initials 5**] [**Last Name
26**]Approved: SAT [**13-09-01**] 7:27 PM
```

Figure 2.13: Sample radiology report. The text field of the radiology report table contains information obtained from radiology.

Component	Description
WFDB library	This is a set of functions for reading and writing files in the formats used by PhysioBank databases. The library supports reading directly from remote servers allowing applications linked with the WFDB library to view or analyze data without the need to download entire records and to store them locally.The WFDB library is implemented in C but provides interfaces for software written using Perl, Python, C# (and other .NET languages), Java, Matlab, PHP, Ruby, TCL, and several versions of Lisp.
WFDB applications	A large set of well-tested, interoperable command-line tools for signal processing and automated analysis.These applications are described in detail in the WFDB Applications Guide [14].
Visualization tools: WAVE	Extensible interactive graphical environment for manipulating sets of digitized signals with optional annotations

Table 2.3: WFDB software package major components

- Trend parameters with a temporal resolution of once a second.

This second, improved, data collection effort is ongoing, and will be added to the public database as we process them.

2.3.2 The WFDB software package

Over the past twenty years, the team at PhysioNet has developed a large collection of open source software to store, analyze and manipulate physiological measurements [14].The WFDB software package is written in highly portable C and can be used on all popular platforms, including GNU/Linux, Mac OSX, MS-Windows, and all versions of Unix. A set of wrappers allow the integration of the WFDB library with other programming languages and interfaces so that the tools can also be run from within visualization tools or other programming environments. Table 2.3 summarizes the major components of the WFDB Software Package; a more detailed description is available at the PhysioNet web site (`http://www.physionet.org/`).

File type	Extension	Description
Header	.hea	Contain signal file names and attributes in plain text format.
Signal	.dat	Contain signals binary data.
Annotations	-dependent, e.g. ".al" for alarms, ".wqrs" for ECG beat annotations	Contain signal custom annotations in binary format
Calibration	.cal	Contain signal calibration specifications.

Table 2.4: WFDB file types

2.3.3 MIMIC II waveform records

All MIMIC II waveforms are stored in WFDB format. Table 2.3 summarizes the types of files you will find for the MIMIC II waveform database.

The records vary in length; some are several weeks in duration. It is common for the signal sources to be interrupted or changed occasionally during recordings of such a long duration. In a typical waveform database, you will find a directory layout including several record names or "cases". All files associated with each record are gathered in a sub-directory named after the record. For example, the files associated with record a40001 are all located within the directory named a40001.

Table 2.4 presents some useful WFDB commands to navigate through the waveform records. In a typical WFDB database record, a header file specifies the names of the associated signal files and their attributes, briefly:

- Record name: a string of characters that identifies the record.
- Number of channels: an integer greater than zero.
- Sampling rate: an integer or floating-point number, interpreted as samples per second.

Figure 2.14, shows the output of the "wfdbdesc" command, which outputs a human-readable description of a waveform record.

If you want to display the signal contents of a particular record, you can use the "rdsamp" command. Figure 2.15 shows the output of this program.

Another option is to use a visualization tool like WAVE, or ATM (http://www.physionet.org/cgi-bin/ATM), to display the contents of a particular waveform record. Figure 2.16 shows the display of a particular waveform record. Note that the signal processing algorithms can be run from this viewer.

Command	Description
wfdbdesc	Reads specifications for the signals described in the header file for record.
rdsamp	Reads signal files for the specified record and writes the samples as decimal numbers on the standard output. Each line of output contains the sample number and samples from each signal, beginning with channel 0, separated by tabs.
wave	Can be used to view the specified WFDB record or records on any display controlled by an X11 server. It includes facilities for interactive annotation editing. The keyboard and mouse are used to control the display interactively

Table 2.5: Useful WFDB commands

```
Starting time: [16:24:28.848 30/03/2011]
Length: 1:28:00.000 (660000 sample intervals)
Sampling frequency: 125 Hz
4 signals
Group 0, Signal 0:
  File: a42174\_000006.dat
  Description: II
  Gain: 55 adu/mV
  Initial value: 1
  Storage format: 80
  I/O: can be unbuffered
  ADC resolution: 8 bits
  ADC zero: 0
  Baseline: 0
  Checksum: 19538
Group 0, Signal 1:
  File: a42174\_000006.dat
  Description: V
  Gain: 39 adu/mV
  Initial value: 4
  Storage format: 80
  I/O: can be unbuffered
  ADC resolution: 8 bits
  ADC zero: 0
  Baseline: 0
  Checksum: -9315
Group 0, Signal 2:
  File: a42174\_000006.dat
  Description: ABP
  Gain: 1.25 adu/mmHg
  Initial value: 184
  Storage format: 80
  I/O: can be unbuffered
  ADC resolution: 8 bits
  ADC zero: 0
  Baseline: -100
  Checksum: -3865
Group 0, Signal 3:
  File: a42174\_000006.dat
  Description: PAP
  Gain: 2.5 adu/mmHg
  Initial value: 125
  Storage format: 80
  I/O: can be unbuffered
  ADC resolution: 8 bits
  ADC zero: 0
  Baseline: -100
  Checksum: -4501
```

Figure 2.14: Sample output for wfdbdesc

```
time     II     V      ABP     PAP
(sec)   (mV)   (mV)   (mmHg)  (mmHg)
----------------------------------------
0.000   0.018 0.103 22.400 90.000
0.008   0.018 0.077 22.400 90.000
0.016   0.036 0.051 22.400 90.000
0.024   0.018 0.051 22.400 90.000
0.032 -0.018 0.051 60.800 90.000
0.040   0.000 0.026 60.800 90.000
0.048 -0.018 0.026 60.800 90.000
0.056 -0.036 0.000 60.800 90.000
0.064 -0.036 0.026 22.400 90.000
0.072 -0.018 0.000 22.400 90.000
0.080 -0.055 0.000 22.400 90.000
0.088 -0.073 0.000 22.400 90.000
0.096 -0.055 0.000 60.800 90.000
0.104 -0.055 0.000 60.800 90.000
0.112 -0.036 0.000 60.800 90.000
0.120 -0.036 0.000 60.800 90.000
0.128 -0.055 0.000 22.400 90.000
0.136 -0.073 0.000 22.400 90.000
0.144 -0.073 0.000 22.400 90.000
0.152 -0.073 0.000 22.400 90.000
0.160 -0.073 0.000 60.800 90.000
0.168 -0.073 0.000 60.800 90.000
0.176 -0.073 0.000 60.800 90.000
```

Figure 2.15: Sample output of rdsamp

2.3.4 Alarms and Inops

The following description about alarms and inops corresponds to the first group of waveforms collected (approximately 2,800 patient records). All are adults where the *Case_ID*'s are numbered less than a44000.

Simultaneously with approximately 10,000 patient-days of waveforms and trends, we have collected over 450,000 alarms and inops. This amounts to a frequency of one alarm or alert every thirty minutes for each patient in the hospital, although most of these are not life-threatening. Tables 2.6 and 2.7 list the types of alarms and *inops* (non-physiological alerts such as machine disconnections) gathered by our data collection system. A three-star alarm is a potentially life-threatening condition that requires immediate attention. Two-star alarms require less immediate attention, although may provide warning of an increased risk of adverse problems in a patient over time. Note that most of these alarms are not verified as correct. In particular, the *.alarm* are all unverified. A subset of alarms (with the *.alM* extension) are verified by humans, and are described in the next section.

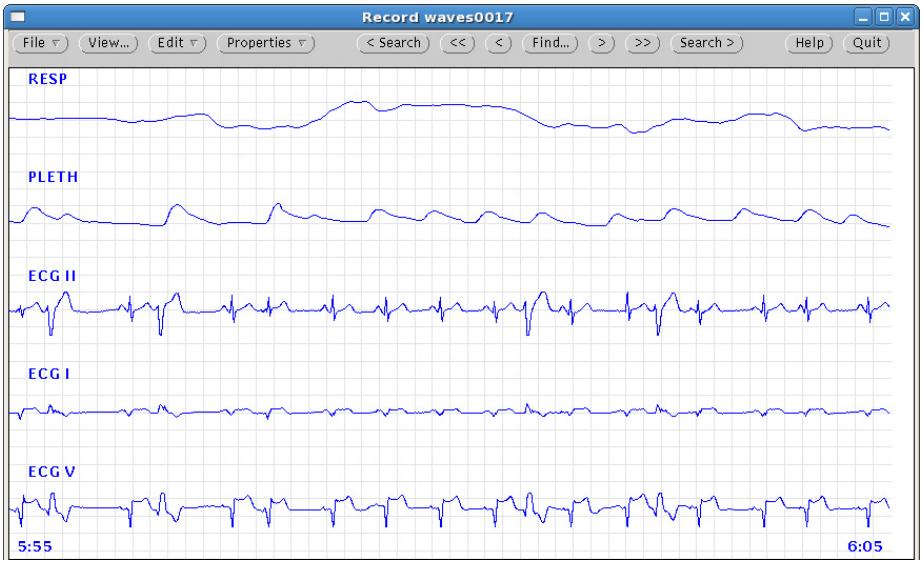


Figure 2.16: Sample waveform record

Annotated alarms

Since no large annotated dataset of alarms is publicly available, a set of gold standard alarms to support the development and testing of a false alarm suppression algorithm was generated from the above alarms. Initially we have concentrated on life-threatening arrhythmia alarms, namely; Asystole, Extreme Bradycardia, Extreme Tachycardia, Ventricular Tachycardia and Ventricular Fibrillation. In order to assemble such a database we first searched for patient records in the MIMIC II database that met the following two criteria:

1. Record contains at least one of the 5 above listed critical alarm categories.
2. At least one of the alarms is associated with simultaneous ABP and ECG waveforms.

Our initial search yielded 496 patient records with a total of 45,370 hours of simultaneous ECG & ABP waveforms containing 8,636 alarms. Each alarm was manually reviewed by two independent experts, and discrepancies were adjudicated by a third expert.

Alarm repetitions referring to the same event, were removed. Furthermore, all 48 patients that possessed active intra-aortic balloon pumps (IABP) were removed, since their ABP waveforms do not appear as "physiologically normal". The final set comprises 448 patients with 5,386 alarms with simultaneous ABP & ECG waveforms. These annotations have been posted on PhysioNet with the file extension *.alM*.

Alarm Label	Definition
***ASYSTOLE	No QRS for 4s
***BRADY m < n	Bradycardia < 40BPM
** HR m < n	Low Heart Rate
** HR m > n	High Heart Rate
** IRREGULAR HR	Irregular Heart Rate
** MISSED BEATS	Missed Heart Beats
** MULTIFORM VPBs	Multiform Ventricular Premature Beats
** PACER NOT CAPT	Pacemaker not capturing
** PAIR VPBs	Pair of Ventricular Premature Beats
** R-ON-T VPBs	R on T type beats
** RUN VPBs 3 - 9	Run of 3-9 Ventricular Premature Beats
** RUN VPBs > 9	As above, but >9
** STi m.m<n.n	ST depression in mV
** STi m.m>n.n	ST elevation of in mV
***TACHY m > n	Tachycardia, HR > n BPM
** VENT BIGEMINY	Ventricular Bigeminy
***VENT FIB/TACH	Fibrillatory Waveform for 4s or more
** VENT RHYTHM	Ventricular Rhythm
***VENT TACHY	Run of >=5 Ventricular Beats with HR>100
** VENT TRIGEMINY	Ventricular Trigeminy

Table 2.6: ECG alarms in the MIMIC II database together with their definitions.

Full details of how these alarms were annotated is available in Aboukhalil *et al* [11], together with an evaluation of their statistics.

2.3.5 Signal Quality

Although our search engines allow a researcher to determine what signals exist for which patients, this is no guarantee of quality, and sometimes the data can be so noisy that no useful clinical information can be extracted from the data. To avoid requesting noisy data, and using this data for further processing, we have developed a set of signal quality indices (SQI's) for both electrocardiogram and blood pressure data. We have also used these indices and a multi-channel weighting algorithm, to generate our best estimates of heart rate and blood pressure for every 10 second window of waveform data.

Alarm Label	Definition
***ABP m < n	Hypotension (Extremely Low Blood Pressure)
** ABP m < n	Low Blood Pressure
** ABP m > n	High Blood Pressure
***APNEA	No Respiratory Effort Detected
** ART m < n	Low Blood Pressure (secondary line)
** ART m > n	High Blood Pressure (secondary line)
** CVP m < n	Low Central Venous Pressure
***Desat m < n	Desaturation (of SPO2) (#)
** ICP m > n	Low Intra-Cranial Pressure
** LAP m > n	High Left Arterial Pressure
** NBP m < n	Low Non-Invasive Blood Pressure
** NBP m > n	High Non-Invasive Blood Pressure
** PAP m < n	Low Pulmonary Arterial Pressure
** PAP m > n	High Pulmonary Arterial Pressure
** P1 m < n	Low (Generic) Pressure
** RESP m > n	Hyperventialaation
**SpO2 m < n	Low Oxygen Saturation
**Tblood m.m<n.n	Low Blood Temperature (in deg C)
** UAP m < n	Low Umbilical Arterial Pressure
** UAP m > n	Low Umbilical Arterial Pressure
** UVP m > n	High Umbilical Venous Pressure
***ABP DISCONNECT	Invasive Arterial Line Disconnect
***PAP DISCONNECT	Pulmonary Arterial Catheter Disconnect
***UAP DISCONNECT	Uterine Pressure Line Disconnect

Table 2.7: Non-ECG alarms and Inops in the MIMIC II database. (#) indicates that this alarm is available only on the latest version (revision F) of the bedside monitor software.

The ECG signal quality metrics are a combination of statistical measures, in both the time and frequency domains, multi-channel QRS detector performance, and correlation to past data. A more in-depth description of the formation of these annotations, together with an evaluation of a robust HR and ABP tracking algorithm that utilizes this information can be found in [9] and [10]. The blood pressure signal quality metric is based upon two earlier developed metrics by Zong [7] and Sun [8].

There are several annotations associated with each beat in the ECG and ABP signals. Table 2.8 describes the available SQI annotations for waveform records. Full descriptions of how to interpret the SQI output for ECG and ABP can be found in in [9] and [10] respectively. Note that although many of the artifact types for each of these signals have been incorporated, and known errors

in heart rate and blood pressure calibrated to the SQI output, some artifacts are not well represented. In particular, the tricky problem of blood pressure damping is not yet fully solved in our ABP SQI metric. Any analysis of blood pressure should therefore be tempered by the fact that damping may lead to an error, and in particular, an under-estimation of the SBP and pulse pressure.

Table 2.8: Per beat SQI annotations files. Where "file" is the waveform record name such as "a41000". SDR = Spectral Distribution Ratio. DF = digital filter. LT = length transform. EPLTD and WQRS are the beat detectors that use the DF and LT methods respectively. See [9] and [10] details.

Annotation	**Description**
file.ecgsqid	A combined beat-by-beat ECG SQI created by selecting the best ECG SQI between different ECG leads.
file.ecgsqid$n$	The ECGSQI annotation of the n^{th} ($n = 0, 1, 2, ...$) ECG lead, used to create a combined annotation file'.multid'
file.epltd$n$	The EPLTD (DF) annotation of the n^{th} ($n = 0, 1, 2, ...$) ECG lead, used to create EPSQI and ICHSQI
file.epsqi$n$	The EPSQI (bSQI) annotation of the n^{th} ($n = 0, 1, 2, ...$) ECG lead, used to create Kurtosis (kSQI) and SDR (sSQI)
file.ichsqi$n$	The ICHSQI (iSQI) annotation of the n^{th} ($n = 0, 1, 2, ...$) ECG lead, used to create ECGSQI
file.kurtd$n$	The Kurtosis (kSQI) and SDR (sSQI) annotation of the n^{th} ($n = 0, 1, 2, ...$) ECG lead, used to create ECGSQI and the sample-and-hold HR (HRsh1) of the first ECG lead
file.multid	An all-in-one annotation include all leads of ECG beats with ECGSQI and ABP beat with ABPSQI, used to create the HR and ABP
fileTd	A WFDB trend data file that includes different calculations of HR, HRSQI and ABP sampled at 0.1Hz
fileTd.hea	A header file of fileTd
file.wqrs$n$	The WQRS (LT) annotation of the n^{th} ($n = 0, 1, 2, ...$) ECG lead, used to create EPSQI (bSQI)
fileT	A WFDB trend data file similar to fileTd, but without baseline wander filter to calculate kurtosis.

continued on next page

continued from previous page	
Annotation	**Description**
fileT.hea	Header file for fileT
file.wsqi	The WSQI annotation of the ABP lead, used to create ABPSQI
file.jsqi	The JSQI annotation of the ABP lead, used to create ABPSQI
file.abpsqi	The ABPSQI annotation of the ABP lead, used to create '.multid'

Chapter 3

Database Access

3.1 Introduction

This chapter provides instructions for connecting to and extracting data from the MIMIC II database. There are two methods for accessing the data:

- Flat file download from PhysioNet http://physionet.org
- Web-based "MIMIC Explorer" http://mimic.physionet.org/

In order to gain access to the database, you must follow the instructions posted on PhysioNet (http://www.physionet.org/physiobank/database/mimic2cdb/restricted/)

When you have completed and signed the data use agreement, your application will be processed and your access details emailed to your account.

Once you have completed this process, an account will be created for you to access the restricted areas of the MIMIC II website and the MIMIC Explorer.

Chapter 4

Examples of data analysis

4.1 Introduction

This chapter provides some introductory examples for obtaining data/statistics from the database. We hope that these examples will enable users to become familiar with the tables and the data they contain. These examples will also help to illustrate the complexity of the database and the difficulties encountered in obtaining certain data. Please note that the examples shown here have been tested on database version MIMIC2V24. Although direct copy-and-paste of the examples is possible, some examples may result in "Invalid Character" errors. For example, the asterisk (*) character may not be copied-and-pasted correctly. In this case, simply type the correct character in its place.

4.2 Clinical Examples

4.2.1 Patient population age statistics

We begin by providing some simple queries for extracting age/ICU admission data. Query 4.1 simply counts the number of unique subject ids in the database.

Listing 4.1: SQL to obtain the number of subject ids in the database

```
-- Subject IDs: Total          26655
  select count(*)
    from d_patients
```

The next two queries provide the ages of adults (4.2) and neonates (4.3) at the time of their most recent hospital admission.

Listing 4.2: SQL to obtain the age of adult patients at the time of their most recent hospital admission

```
-- Subject IDs: Adults:      19075  (Age at last
    hospital admission)
  select count(*)
   from (
        select distinct p.subject_id,
               p.sex, p.dob,
               first_value(a.adm_dt)
                 over ( partition by p.subject_id
                        order by p.subject_id, a.adm_dt
                            DESC) as adm_dt,
               first_value(a.disch_dt)
                 over ( partition by p.subject_id
                        order by p.subject_id, a.adm_dt
                            DESC) as disch_dt,
               first_value(a.expire_flg)
                 over ( partition by p.subject_id
                        order by p.subject_id, a.adm_dt
                            DESC) as expire_flg
          from d_patients p,
               admissions a
         where a.subject_id = p.subject_id
           and round(months_between(a.adm_dt, dob) /12,
               2) >= 15
         order by p.subject_id
         )
```

Listing 4.3: SQL to obtain the age of neonate patients at the time of their most recent hospital admission

```
-- Subject IDs: Neonates: 6538   (Age at first hospital
     admission)
  select count(*)
   from (
          select distinct p.subject_id,
                 p.sex, p.dob,
                 first_value(a.adm_dt)
                   over ( partition by p.subject_id
                          order by p.subject_id, a.
                              adm_dt ASC) as adm_dt,
                 first_value(a.disch_dt)
                   over ( partition by p.subject_id
                          order by p.subject_id, a.
                              adm_dt ASC) as disch_dt,
                 first_value(a.expire_flg)
                   over ( partition by p.subject_id
                          order by p.subject_id, a.
                              adm_dt ASC) as expire_flg
            from d_patients p,
                 admissions a
           where a.subject_id = p.subject_id
             and months_between(a.adm_dt, dob) <= 1
           order by p.subject_id
     )
```

Queries 4.4 and 4.5 return the number of admissions for adults and neonates respectively. Since any particular patient can be admitted to the hospital multiple times, there are more admissions than patients.

Listing 4.4: SQL to obtain the number of adult hospital admissions

```
-- Hospital admissions: Adults: 22916
  select count(*)
   from (
        select p.subject_id, p.sex, p.dob,
               round(months_between(a.adm_dt, dob) /12,
                     2) age,
               a.adm_dt, a.disch_dt,
               a.expire_flg
          from d_patients p,
               admissions a
         where a.subject_id = p.subject_id
           and round(months_between(a.adm_dt, dob) /12,
                2) >= 15
         order by p.subject_id, a.adm_dt
     )
```

Listing 4.5: SQL to obtain the number of neonate hospital admissions

```
-- Hospital admissions:       Neonates: 6585
  select count(*)
   from (
        select p.subject_id, p.sex, p.dob,
               round(months_between(a.adm_dt, dob) /12,
                     2) age,
               a.adm_dt, a.disch_dt,
               a.expire_flg
          from d_patients p,
               admissions a
         where a.subject_id = p.subject_id
           and months_between(a.adm_dt, dob) <= 1
         order by p.subject_id, a.adm_dt
     )
```

In addition to hospital admissions, ICU stays are recorded in the database. A patient can have multiple ICU stays during a hospitalization period and so there are more ICU stays than there are hospital admissions. Query 4.6, shows the "Stay" query. "Stay" is defined by obtaining records from the d_patients, icu_stay and censusevents tables, grouping the records which pertain to the same patient and providing the minimum "intime" and maximum "outtime" of the ICU stays of a patient. Since hospital admission/discharge records are not necessarily available, we assume that the earliest "intime" represents the hospital admission and the latest "outtime" represents discharge and so the

"Stay" query provides the hospital admission/discharge summary for a patient

Listing 4.6: SQL to obtain the number of ICU admissions where the stay length is three days or longer

```
-- ICU stays:  total: 33361
  select count(*)
   from (
        -- Returns all the adult ICU stays where the
            patient
        -- was least 3 days in the ICU
        select s.icustay_id, s.subject_id, s.
            icustay_intime, s.icustay_outtime
        from icustay_detail s
        where s.icustay_los/(24*60) > 3 -- LOS is in
            minutes
        )
```

Since the distinction between adult and neonate patients has been provided in the admissions queries, we simply provide the code to obtain ICU stays of three days or longer. The code to separate this data into adults and neonates is left as an exercise for the reader. The code is also provided in the accompanying source code distribution.

4.2.2 SAPS I on first day of admission

The calculation of SAPS I scores is complex, involving 14 different parameters. In the MIMIC-II database, the SAPS I scores have been precalculated and stored in itemid 20001. Therefore, to obtain a SAPS I score for a patient, query 4.7 is required.

Listing 4.7: SQL to obtain the SAPS I scores of patients

```
select * from chartevents where itemid = 20001;
```

4.2.3 Calculate mortality rate

The SQL query below is designed to obtain mortality rates for the patients in the MIMIC II database. The results can be categorized by age group and type of clinical service (e.g. MICU, SICU, CCU, etc.).

The query for obtaining ICU stays has been provided previously (Query 4.6) and is simply replaced by the "Stay" subquery in the following queries.

Query 4.8 displays the "StayWUnit" query. This query provides a list of ICU stays along with information about whether or not the patient is a child,

or adult and the room of the last ICU stay. Initially, the ages of the patients are calculated and each person is classified into either "children" for people under 15 years old or "adult" for patients 15 and older. Then the ICUs are grouped, depending on their specific notation. For example, "C-SICU" and "T-SICU" are classified as "SICU"; "CSRU-A" and "CSRU-B" are both "CSRU". Finally, the results are combined with the "Stay" query to obtain the full result-set.

Listing 4.8: SQL query for determining ICU stays along with adult/infant record and the location of the last ICU stay.

```
SELECT DISTINCT subject_id, icustay_id, dob,
    sex, intime, outtime,
    round(months_between(intime, dob) /12, 2)
    age,
  CASE
    WHEN round(months_between(intime,
        dob)/12, 2) >= 15
      THEN 'adult'
    ELSE 'children'
  END AS age_group,
  icustay_first_careunit
FROM icustay_detail icu,
     d_careunits cu
WHERE cu.cuid (+)= icu.icustay_first_careunit
ORDER BY subject_id, icustay_id
```

Query 4.9 shows the "StayWHospAdm" query. This query obtains ICU stays which fall inside a hospitalization period. The query uses the admissions table to obtain the admit and discharge times and compares them with the in and out times of ICU visits from the "StayWUnit" query.

Listing 4.9: SQL query for determining ICU stays which fall outside a hospitalization period. The query uses the results from Query 4.8 to obtain the hospitalization period.

```
SELECT icu.subject_id, icu.icustay_id, icu.dob,
       icu.sex, icu.intime,
       nvl(icu.outtime,had.disch_dt) outtime,
       icu.age, icu.age_group, location,
       had.hadm_id, had.adm_dt AS hadm_dt,
       had.disch_dt AS hdisch_dt, had.expire_flg
FROM StayWUnit icu,
     mimic2v2.admissions had
WHERE icu.subject_id = had.subject_id
  AND icu.intime >= had.adm_dt
  AND icu.intime <= had.disch_dt
```

Query 4.10 shows the "StayWService" query. This query obtains the service type of the ICU stays within a hospitalization period. If the service type is "null", then it will be taken as the location of the previous ICU stay.

Listing 4.10: SQL query for determining the service type of ICU stays.

```
SELECT DISTINCT icu.subject_id, icu.icustay_id,
    icu.dob, icu.sex, icu.intime, icu.outtime,
    icu.age, icu.age_group, icu.hadm_id,
    icu.hadm_dt, icu.hdisch_dt, icu.expire_flg,
    first_value(CASE
      WHEN itemid IS NOT NULL THEN
        DECODE(nvl(ce.value1, 'Unknown'),
          'CSICU', 'SICU',
          'NSICU', 'SICU',
          'MSICU', 'MICU',
          nvl(ce.value1, 'Unknown'))
      ELSE location
    END)
  OVER (partition BY icu.subject_id, icu.hadm_id,
    icu.icustay_id
  ORDER BY icu.subject_id, icu.hadm_id,
    icu.icustay_id, ce.charttime DESC)
  AS service_type
FROM StayWHospAdm icu,
     mimic2v2.chartevents ce
WHERE ce.subject_id (+)= icu.subject_id
AND   ce.itemid (+)= 1125
AND   ce.charttime (+)>= icu.intime
AND   ce.charttime (+)<= icu.outtime
ORDER BY icu.subject_id, icu.hadm_id, icu.icustay_id
```

Other related queries can be found in the `mortality_rate` file which is included in the source code package for this documentation. Other queries available include:

The "LastStayWService" query obtains the last ICU stay for each patient. i.e. It groups the "StayWService" query by the subject_id. The "HospAdmWStayByAge" query performs the "LastStayWService" query and groups the result set by age, to obtain age statistics for ICU stays. The "DeathByAge" query obtains hospital mortality statistics by performing the "HospAdmWStayByAge" and selecting only patients who died. "HospMortByAge" obtains the percentage mortality rates for each age group. "AdultHospAdmByService" obtains the number of adult hospital admissions grouped by service type. "AdultDeathByService" obtains adult hospital mortality numbers grouped by service type while "HospMortByService" provides adult hospital mortality percentages by service type. "DeathInICUByAge" obtains age grouped ICU mortality numbers (in the ICU) while "InICUMortByAge" provides age grouped ICU mortal-

ity percentages.

"AdultDeathInICUByService" obtains service grouped ICU mortality numbers, while "InICUMortByService" provides service grouped ICU mortality percentages.

"DeathOutICUByAge" obtains mortality numbers outside the ICU and "OutICUMortByAge" provides the same results grouped by age. "AdultDeathOutICUByService" provides service-grouped numbers while "OutICUMortByService" provides percentages.

4.2.4 Statistics for patients admitted with Subarachnoid Hemorrhage on admission

This query obtains statistics for patients who were admitted with subarachnoid hemorrhage. The query is built in three parts: The first obtains the patients who stayed in the ICU for at least three days. The second obtains the last ICU stay for adult patients with subarachnoid hemorrhage. The third and final part calculates the statistics.

As in Section 4.2.3, the "Stay" query has been defined previously (Query 4.6) and is simply replaced by "Stay" in the following queries.

Query 4.11 performs the second part. The ICU stay query (Query 4.6) is sub-selected to obtain the last ICU stay and then further reduced by selecting only those patients with subarachnoid hemorrhage (ICD-9 code 430).

Listing 4.11: SQL query for obtaining only the last ICU stay for each patient and further restricting the query to patients with ICD-9 code 430 (Subarachnoid hemorrhage).

```
SELECT DISTINCT subject_id,
first_value(intime) OVER ( partition BY subject_id
    ORDER BY subject_id, intime DESC) AS intime,
first_value(outtime) OVER ( partition BY subject_id
    ORDER BY subject_id, intime DESC) AS outtime,
first_value(expire_flg) OVER ( partition BY subject_id
    ORDER BY subject_id, intime DESC) AS expire_flg,
first_value(age) OVER ( partition BY subject_id
    ORDER BY subject_id, intime DESC) AS age,
first_value(sex) OVER ( partition BY subject_id
    ORDER BY subject_id, intime DESC) AS sex
FROM (
  -- Returns any ICU stay that falls between the
  -- hospitalization period
  SELECT had.subject_id, icu.icustay_id, icu.intime,
    icu.outtime, had.expire_flg, icu.age, icu.sex
    FROM (
   -- Returns the last hospital admissions for patients
   -- with the required ICD9 code
   SELECT DISTINCT subject_id,
     first_value(adm_dt)
       OVER (partition BY subject_id
         ORDER BY subject_id, adm_dt DESC) AS adm_dt,
     first_value(disch_dt)
       OVER (partition BY subject_id
         ORDER BY subject_id, adm_dt DESC) AS disch_dt,
     first_value(expire_flg)
       OVER (partition BY subject_id
         ORDER BY subject_id, adm_dt DESC) AS
             expire_flg
     FROM (
       -- Returns all the hospital admissions for
           people
       -- who had the ICD9 code: Subarachnoid Hemorrage
       SELECT a.hadm_id, a.subject_id, a.adm_dt,
              a.disch_dt, a.expire_flg
         FROM mimic2v2.admissions a,
              mimic2v2.icd9 icd
         WHERE icd.hadm_id = a.hadm_id
           AND icd.code IN ('584.9')
   )
  ORDER BY subject_id
) had, (
SELECT * from ThreeDayStay
  ) icu WHERE icu.subject_id = had.subject_id
    AND icu.intime >= had.adm_dt
    AND icu.intime <= had.disch_dt
) ORDER BY subject_id
```

The final part of the query calculates the statistics for these patients. This query is very long, owing to the large number of columns required and so is not provided here. The query can be found in the accompanying source code package.

4.2.5 Patient BMI calculation

Patient height and weight is required for Body Mass Index (BMI) calculation. First we must find the itemids which contain patient height measurement. The following query obtains all itemids which reference patient height.

Listing 4.12: SQL query for obtaining the itemids which contain patient height.

```
-- Query to obtain ItemIds containing patient height.
SELECT *
FROM d_chartitems
WHERE itemid in (920, 214, 215, 1394, 3721)
```

Next, we require patient weights. The following query provides all itemids which reference patient weight, including, among others, admit, pre-op and ideal body weight.

Listing 4.13: SQL query for obtaining the itemids which contain patient weight.

```
-- Query to obtain ItemIds containing patient weight.
SELECT *
FROM d_chartitems
WHERE itemid in (7000, 1300, 580, 762, 765, 767, 762,
  3580, 3581, 3582, 4020, 3693, 1393, 1395, 3723, 4183,
  3897, 4032, 765, 767, 733, 3692, 7000, 3718)
```

To calculate BMI, we need to select the actual values of height and weight from the database, 'pivot' the data so that we have a suitable dataset format and then perform the calculation:

Listing 4.14: SQL query for obtaining patient BMI.

4.2.6 Resolving discrepancies between multiple itemIDs for one parameter

Since the database contains many "free text" itemIDs, there is no unique method for representing certain parameters. For example, CPR is mentioned in 9 dif-

ferent itemIDs. Each itemID can contain a variety of values such as “yes”, “no”, “done”, “performed”, etc. In addition multiple different capitalization and spelling errors are found. This makes it extremely difficult to obtain accurate information on whether or not CPR was performed for a particular patient.

As and when certain data is extracted from these “free text” fields, it can be moved into more meaningful fields which will permit simpler data extraction. For example, the CPR itemIDs mentioned above could be translated into a binary field which simply states whether or not CPR was performed.

Examples of multiple mappings, plus code to merge them, can be found in appendix 6.2.

Chapter 5

Quick-Start, Frequently Asked Questions and Known problems/issues

5.1 Quick Start

There really is no quick start to exploring the data in MIMIC II DB. First you need to decide which type of data you want to process:

1. 'waveform' data
2. high (temporal) resolution trend data (numerics)
3. 'clinical' data - low temporal resolution data and/or categorical data.

You may also want to decide if you want to use all of the available patients, or just narrow your focus on a subset cohort. For example, you may wish to concentrate on patients who were admitted with a particular diagnosis, or who were administered with a particular drug. In order to find patient cohorts, you may wish to use our prototype search engine, which can be found at:
http://mimic.physionet.org.
Currently, with this interface you can search in the following ways:

- Locate free text strings the nursing progress notes and discharge summaries.
- Locate patients with particular combinations of waveforms.
- Locate patients with particular demographics (age and gender).

Each of these queries can be 'ANDed' together to construct more restrictive subsets.

Once you have found your subset of patient on which you want to work, you will want to download the data you wish to work on, or perform more detailed queries. Methods for doing this are described in chapter 3. Sample queries can be found in chapter 4. We strongly suggest you work through these before attempting to construct patient cohorts and analyze data. We also strongly advise that you take note of the different definitions of patient identifiers, as described in section 1.4.2. Note that you can analyze patients by individual (subject) ID, hospital admission, ICU stay, or by waveform Case ID. We also suggest that you read the frequently asked questions (FAQs) below.

5.2 FAQs about the MIMIC II database

1. **How can I get an answer to my question?** Please ensure that your question has not already been answered by this document; either in the main body of text or in these FAQs. If you are sure that your question is not addressed by this document, then please email the authors with your question.

2. **What is the MIMIC II database?** The Multiparameter Intelligent Monitoring in Intensive Care II database [15] contains detailed clinical data from patients hospitalized in Intensive Care Units (ICUs).

3. **Where is the MIMIC II database physically located?** The MIMIC II database is hosted by MIT: it is located on the Cambridge Campus.

4. **Who do I contact for more information?** Current contact information is always available on the MIMIC II website (`http://mimic.physionet.org/`)

5.3 FAQs about data access

1. **What methods can I use to access the MIMIC II database?** There are 3 access methods for the MIMIC II Database: Oracle's SQL Developer, JDBC [16] and WFDB. In addition, complete database dumps are available for import into your local systems; such data can then be accessed in any desired method. Please see Chapter 3.1 for more details on data access methods.

2. **How do I gain access to the MIMIC II database?** You will need to agree to and sign our data use agreement in order to gain access to the database. You will also have to fill out a short form detailing your plans for the data. Simply visit our website (`http://mimic.physionet.org`) to apply for access and to obtain further information.

3. **Who can access the MIMIC II database?** Anyone who wishes to perform research on the MIMIC II data will be permitted access to the database. Simply visit our website to apply.

5.4 Known issues/problems

1. **Waveform-trend misalignment.** Although the trends should match up with the parameters derived from the waveforms, this is not always true. This can be due to filter delays, network timing errors or data server timing errors.

2. **Inter-waveform alignment problems.** The method used for MIMIC waveform data extraction was not designed for inter-waveform analysis. The waveform data contain unspecified/unknown filtering delays and/or unknown inter-channel delays, which may not be constant in a given record. Therefore, although the ECGs are time-aligned, there may be a (changing) delay of up to 500ms between any of the other waveforms in the data. Therefore, no pulse transit times can be accepted to be true (absolute or relative).

3. **Missing waveform and trend data.** Every patient will have some level of data missing between the admit and discharge time. This can be for many reasons:

   - The patient was disconnected from the monitor for some period (perhaps for a scan or to replace electrodes).
   - The data collection system or the network over which the data was transmitted crashed.
   - The data that was collected was corrupted and conversion to WFDB was not possible.

4. **The clinical data change dimensionality over time and between patients, and are irregularly and sparsely sampled.** The amount and type of data that are recorded concerning a patient, and the frequency at which it is sampled is a function of both the settings on the monitoring equipment, and the activities of the clinical staff. This in turn is reflective of the clinical team's understanding of the patient's changing condition(s). Many tests are not routine and therefore are only ordered when the clinical team suspects a given condition based on the presenting observations. Therefore, the dimensionality of the data for a given patient may fluctuate over time and no signal is guaranteed at any given point in time. When a patient's condition becomes more acute, data are often sampled more frequently, and the number of sampled parameters increases. This leads to the question of what to do with missing data. Interpolation and imputation schemes generally perform poorly because there are no models of *how* the data are missing[17]. It should also be noted that prediction

or classification algorithms can be 'fooled' by the *presence* or *absence* of a data stream. That is, it may not be the result of a test that causes an algorithm to give a particular output, but rather just the fact that a clinician thought the particular test was needed. Caution should be taken in the interpretation of such results.

5. **Contradictory data.** Some data derived from the waveforms or trends may be incommensurate with each other, or with the data in the relational database. This can be due to noise in the data, the use of different windows and filters to process the data, time alignment errors, or the fact that humans can override the machine transmitted data (in the relational database) with values that they think more correctly reflect the patient's physiology. It should be noted, that these cannot always be trusted [18].

6. **Multiple data streams / itemIDs for a single parameter.** Each parameter may be recorded in a variety of ways by both humans and machines. For example, the heart rate (HR) can be derived from the ECG, ABP and PPG (pulse oximeter). You should not expect these parameters to give the same exact values. They will also respond to artifacts in different manners, and sometimes be affected at different points in time by the artifacts.

   In the relational database, each signal or parameter may be recorded under a variety of different names. For example, Lactic acid values are found in chartevents-818 and chartevents-1531. A current list of the known mappings can be found in Appendix 6.2, although we encourage users to send us other mappings that they discover.

7. **Possible mistakes in the subject-case ID mapping.** Linking data from the bedside monitors and the other hospital databases was not a trivial process. Although names and medical record numbers are sometimes manually entered into the bedside monitor, often they are not, or are done so incorrectly. Furthermore, even when a patient is discharged from the ICU, they are sometimes not 'discharged' from the bedside monitor, and so the next patient may inadvertently inherit the name and MRN of the old patient. Therefore, one should be attentive to this possibility. For the patients with no MRN or name identifiers in the waveform and trend data, we attempted to match the patients based on admit/discharge times, available trends, and numerics of the data. This form of matching is obviously more error prone than MRN or name matching. See section 1.4.3 for more information.

   Although every effort has been made to map the waveform and trend data to the associated clinical data, mistakes will be present. If you think you have discovered such a mistake, please email us with the evidence and we will do our best to answer your query or correct the data.

8. **Possible mistakes in calibration or conversion units** Care should be taken to identify data that appears to be out of range or exhibiting abnor-

mal offsets. For example, temperature may be measured in degrees Centigrade and recorded in degrees Fahrenheit for part of a patient's record. More fundamentally, conversion factors may have become corrupted, and so representations of parameters may not always be correct.

9. **Possible mistakes in waveform labels** We have noticed that in converting to an open format, the data, which was written to disk in a proprietary format using Microsoft .Net, errors have crept into the waveform labeling. Sometimes channels labelled as V (ECG) are actually respiratory waveforms. At other times, labels are "UNKNOWN" and although they are often PPGs, this is not always true.

10. **My drug is having the opposite effect of what I expected** Drugs effects are variable, depending on interactions with other drugs, dosage levels and cardiovascular state. See section 6.2.17

11. **The nursing note does not make complete sense or contradicts the data**. Nursing notes are 'free-text' notes that can contain typos, errors or hard to understand short-hand. While we have tried to provide a list of useful abbreviations in section 5.5, this is not complete and errors may still exist. Note also that the numerical data may be in error.

We are always striving to improve our database, and so if you notice any anomalies, and/or have any suggestions on how to fix them, please do contact us.

5.5 Abbreviations

- **A** - Assessment
- **ABP** - Arterial Blood Pressure
- **ABPSQI** - Arterial Blood Pressure Signal Quality Index
- **BIDMC** - Beth Israel Deaconess Medical Center
- **BP** - Blood Pressure
- **BPM** - Beats (or Breaths) per Minute
- **BUN** - Blood Urea Nitrogen (also known as Urea or Urea nitrogen)
- **CAREVUE** - The Philips bedside ICU workstation for clinicians
- **CCU** - Cardiac Care Unit
- **CDSS** - Clinical Decision Support System
- **CMO** - Comfort Measures Only
- **CO** - Cardiac Output

- **CSRU** - Cardiac Surgery Recovery Unit
- **CVP** - Central Venous Pressure
- **DBP** - Diastolic Blood Pressure
- **DF** - Digital Filter
- **DNI** - Do Not Intubate
- **DNR** - Do Not Recusistate
- **D/C'd** - Discontinued, or Discharged
- **ECG** - Electrocardiogram
- **ECGSQI** - Electrocardiogram Signal Quality Index
- **ECO** - Estimated Cardiac Output
- **EKG** - Electrocardiogram
- **F/E** - Fluid and Electrolytes
- **GCS** - Glasgow Coma Scale (or sometimes *Score*
- **GI** - Gastrointestinal
- **HEME** - Hematology
- **HIPAA** - Health Insurance Portability and Accountability Act
- **HR** - Heart Rate
- **IBP** - Invasive Blood Pressure
- **IABP** - Intra-Aortic Balloon Pump or Invasive Arterial Pressure
- **ICD** - Implantable Cardioverter Defibrillator
- **ICD-9** - International Statistical Classification of Diseases and Related Health Problems (version 9)
- **ICU** - Intensive Care Unit
- **ID** - Identifier or Infectious Disease
- **ISM** - Information Support Mart
- **LT** - Length Transform
- **MBP** - Mean Blood Pressure
- **MICU** - Medical Intensive Care Unit

- **MRN** - Medical Record Number
- **NBP** - Non-invasive Blood Pressure
- **Neo** - Neosynephrine
- **NIBP** - Non-invasive Blood Pressure
- **NICU** - Neonatal Intensive Care Unit
- **NSR** - Normal Sinus Rhythm
- **P** - Plan
- **PAP** - Pulmonary Arterial Pressure
- **PCWP** - Pulmonary Capillary Wedge Pressure (or *Wedge Pressure*)
- **PPG** - Photoplethysmogram
- **PRBC** - Packed Red Blood Cells
- **RR** - Respiration Rate
- **SAPS** - Simplified Acuity Score
- **S/0** - Sign Out
- **SaO2** - Arterial Oxygen Saturation
- **SBP** - Systolic Blood Pressure
- **SICU** - Surgical Intensive Care Unit
- **SQI** - Signal Quality Index
- **SPO2** - Peripheral Oxygen Saturation
- **TCO** - Thermodilution Cardiac Output
- **T-SICU** - Trauma Surgical Intensive Care Unit
- **WFDB** - Waveform Database software package – See http://www.physionet.org/

Chapter 6

Appendix

6.1 Database Schema

This section describes each base table for the MIMIC II database, in alphabetical order. A list of foreign key references, primary keys, and a description of the columns is provided for each table.

The MIMIC II database schema is composed of several additional tables that deal with user authentication, database security, event logging and other application specific storage requirements. The description of these tables is outside the scope of this document.

6.1.1 ADDITIVES

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		
ITEMID	NUMBER	N		
IOITEMID	NUMBER	N		
CHARTTIME	DATE	N		
ELEMID	NUMBER	N		
CGID	NUMBER	Y		
CUID	NUMBER	Y		
AMOUNT	NUMBER	Y		
DOSEUNITS	VARCHAR2	Y		
MLPERUNIT	NUMBER	Y		
ROUTE	VARCHAR2	Y		

Primary key(s)

ADDITIVES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ITEMID			

ADDITIVES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	IOITEMID			

ADDITIVES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	SUBJECT_ID			

ADDITIVES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CHARTTIME			

ADDITIVES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ELEMID			

Foreign key(s)

6.1.2 ADMISSIONS

Columns

Column Name	**Data Type**	**Nullable**	**Default value**	**Comments**
HADM_ID	NUMBER	N		Table record unique identifier, the hospital admission ID
SUBJECT_ID	NUMBER	N		The patient for the admission.
ADM_DT	DATE	N		Date and time the patient was admitted to the ICU
DISCH_DT	DATE	N		Date and time the patient was discharged from the ICU
EXPIRE_FLG	CHAR	N	'N'	

Primary key(s)

ADMISSIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	HADM_ID			

Foreign key(s)

6.1.3 A_CHARTDURATIONS

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The patient for the chart event.
ITEMID	NUMBER	N		The chart item for the duration record
ELEMID	NUMBER	N		
STARTTIME	DATE	N		The start time of the chart item
STARTREALTIME	DATE	Y		
ENDTIME	DATE	Y		The end time of the chart item
CUID	NUMBER	Y		The care unit where the chart event took place
DURATION	NUMBER	Y		The duration of the chart event

Primary key(s)

Foreign key(s)

6.1.4 A_IODURATIONS

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The patient for the IO event.
ITEMID	NUMBER	N		The IO item for the duration record
ELEMID	NUMBER	N		
STARTTIME	DATE	N		The start time of the IO event
STARTREALTIME	DATE	Y		
ENDTIME	DATE	Y		The end time of the IO event
CUID	NUMBER	Y		The care unit where the IO event took place.
DURATION	NUMBER	Y		The duration of the IO event

Primary key(s)

A_IODURATIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	SUBJECT_ID			

A_IODURATIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ITEMID			

A_IODURATIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ELEMID			

A_IODURATIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	STARTTIME			

Foreign key(s)

6.1.5 A_MEDDURATIONS

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The patient for the medication
ITEMID	NUMBER	N		The medication item for the duration record
ELEMID	NUMBER	N		
STARTTIME	DATE	N		The start time of the medication event
STARTREALTIME	DATE	Y		
ENDTIME	DATE	Y		The end time of the medication event
CUID	NUMBER	Y		The care unit where the medication event took place
DURATION	NUMBER	Y		The duration of the medication event.

Primary key(s)

A_MEDDURATIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	SUBJECT_ID			

A_MEDDURATIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ITEMID			

A_MEDDURATIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ELEMID			

A_MEDDURATIONS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	STARTTIME			

Foreign key(s)

6.1.6 CENSUSEVENTS

Columns

Column Name	Data Type	Nullable	Default value	Comments
CENSUS_ID	NUMBER	N		Table record unique identifier, the census event id.
SUBJECT_ID	NUMBER	N		The partient for the census event.
INTIME	DATE	N		The time when the patient entered the ICU.
OUTTIME	DATE	Y		The time when the patient left the ICU.
CAREUNIT	NUMBER	Y		The originating care unit for the event.
DESTCAREUNIT	NUMBER	Y		The destination ICU
DISCHSTATUS	VARCHAR2	Y		The discharge status of the patient.
LOS	NUMBER	Y		

Primary key(s)

CENSUSEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CENSUS_ID			

Foreign key(s)

6.1.7 CHARTEVENTS

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The patient for the chart event.
ITEMID	NUMBER	N		The chart item for the event record.
CHARTTIME	DATE	N		The time of the chart event.
ELEMID	NUMBER	N		
REALTIME	DATE	N		
CGID	NUMBER	Y		The care giver for the chart event.
CUID	NUMBER	Y		The care unit where the chart event took place.
VALUE1	VARCHAR2	Y		The first chart event value.
VALUE1NUM	NUMBER	Y		The first chart event value (cast to numeric value).
VALUE1UOM	VARCHAR2	Y		The units of measurement for the first chart event value.
VALUE2	VARCHAR2	Y		The second chart event value.
VALUE2NUM	NUMBER	Y		The second chart event value (cast to numeric value).
VALUE2UOM	VARCHAR2	Y		The units of measurement for the second chart event value.
RESULTSTATUS	VARCHAR2	Y		
STOPPED	VARCHAR2	Y		
ANNOTATION	VARCHAR2	Y		

Primary key(s)

CHARTEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	SUBJECT_ID			

CHARTEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ITEMID			

CHARTEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CHARTTIME			

CHARTEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ELEMID			

Foreign key(s)

6.1.8 DB_SCHEMA

Columns

Column Name	**Data Type**	**Nullable**	**Default value**	**Comments**
CREATED_DT	DATE	N	SYSDATE	The schema creation date.
CREATED_BY	VARCHAR2	N	USER	The user who created the database schema.
UPDATED_DT	DATE	N	SYSDATE	The date that the database schema was updated.
UPDATED_BY	VARCHAR2	N	USER	The user who updated the database schema.
SCHEMA_DT	DATE	N	SYSDATE	The schema date.
VERSION	VARCHAR2	N		The schema version number.
COMMENTS	VARCHAR2	Y		Schema comments.

Primary key(s)

Foreign key(s)

6.1.9 DELIVERIES

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The patient for the particular delivery.
IOITEMID	NUMBER	N		The IO item id for the delivery.
CHARTTIME	DATE	N		The chart time of the IO delivery.
ELEMID	NUMBER	N		
CGID	NUMBER	Y		The care giver for the delivery.
CUID	NUMBER	Y		The care unit for the delivery.
SITE	VARCHAR2	Y		The site of the delivery.
RATE	NUMBER	Y		The delivery rate.
RATEUOM	VARCHAR2	N		The units of measurement of the rate of delivery.

Primary key(s)

DELIVERIES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	SUBJECT_ID			

DELIVERIES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	IOITEMID			

DELIVERIES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CHARTTIME			

DELIVERIES_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ELEMID			

Foreign key(s)

6.1.10 D_CAREGIVERS

Columns

Column Name	Data Type	Nullable	Default value	Comments
CGID	NUMBER	N		Table record unique identifier, the care giver id.
PROFTITLE	VARCHAR2	Y		Title for care giver, like "Nurse", "MD".

Primary key(s)

D_CAREGIVERS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CGID			

Foreign key(s)

6.1.11 D_CAREUNITS

Columns

Column Name	Data Type	Nullable	Default value	Comments
CUID	NUMBER	N		Table record unique identifier, the care unit id.
UNITNAME	VARCHAR2	Y		Name of the unit, like "MICU", "CCU".

Primary key(s)

D_CAREUNITS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CUID			

Foreign key(s)

6.1.12 D_CHARTITEMS

Columns

Column Name	Data Type	Nullable	Default value	Comments
ITEMID	NUMBER	N		Table unique identifier, the chart item id.
LABEL	VARCHAR2	Y		
CATEGORY	VARCHAR2	Y		The chart item category.

Primary key(s)

D_CHARTITEMS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ITEMID			

Foreign key(s)

6.1.13 D_IOITEMS

Columns

Column Name	Data Type	Nullable	Default value	Comments
ITEMID	NUMBER	N		Table record unique identifier, the IO item ID.
LABEL	VARCHAR2	Y		
CATEGORY	VARCHAR2	Y		The IO item category.

Primary key(s)

D_IOITEMS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ITEMID			

Foreign key(s)

6.1.14 D_MEDITEMS

Columns

Column Name	Data Type	Nullable	Default value	Comments
ITEMID	NUMBER	N		Table record unique identifier, the medication item id.
LABEL	VARCHAR2	Y		
CATEGORY	VARCHAR2	Y		The medication item category.

Primary key(s)

D_MEDITEMS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ITEMID			

Foreign key(s)

6.1.15 D_PARAMMAP_ITEMS

Columns

Column Name	Data Type	Nullable	Default value	Comments
CATEGORY	VARCHAR2	N		
DESCRIPTION	VARCHAR2	Y		

Primary key(s)

D_PMAPITEMS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CATEGORY			

Foreign key(s)

6.1.16 D_PATIENTS

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		Table record unique identifier, the subject id.
SEX	VARCHAR2	Y		The patient gender
DOB	DATE	N		The date of birth of the patient

Primary key(s)

D_PATIENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	SUBJECT_ID			

Foreign key(s)

6.1.17 ICD9

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The patient for the ICD9 code record.
HADM_ID	NUMBER	N		The hospital admission for the ICD9 code record.
SEQUENCE	NUMBER	N		
CODE	VARCHAR2	N		The ICD9 code.
DESCRIPTION	VARCHAR2	Y		A description of the ICD9 code.

Primary key(s)

Foreign key(s)

6.1.18 ICU_STAY

Columns

Column Name	Data Type	Nullable	Default value	Comments
ICUSTAY_ID	NUMBER	N		Table record unique identifier, the ICU stay id.
CENSUS_ID	NUMBER	N		The census event ID for this ICU stay.

Primary key(s)

ICU_STAY_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ICUSTAY_ID			

ICU_STAY_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CENSUS_ID			

Foreign key(s)

6.1.19 IOEVENTS

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The parient for the IO event.
ITEMID	NUMBER	N		The IO item id for the event.
CHARTTIME	DATE	N		The chart time of the IO event.
ELEMID	NUMBER	N		
ALTID	NUMBER	Y		
REALTIME	DATE	Y		
CGID	NUMBER	Y		The care giver for the IO event.
CUID	NUMBER	Y		The care unit for the IO event.
VOLUME	NUMBER	Y		The volume of the IO event.
VOLUMEUOM	VARCHAR2	Y		The units of measurement of the volume of the IO event.
UNITSHUNG	NUMBER	Y		
UNITSHUNGUOM	VARCHAR2	Y		
NEWBOTTLE	NUMBER	Y		
DRESSINGCHANGED	NUMBER	Y		
TUBINGCHANGED	NUMBER	Y		
ASSESSMENT	NUMBER	Y		
STOPPED	VARCHAR2	Y		
ESTIMATE	VARCHAR2	Y		
ANNOTATION	VARCHAR2	Y		

Primary key(s)

Foreign key(s)

6.1.20 MEDEVENTS

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The patient for the medication event.
ITEMID	NUMBER	N		The medication item id for the event.
CHARTTIME	DATE	N		The chart time of the medication event.
ELEMID	NUMBER	N		
REALTIME	DATE	N		
CGID	NUMBER	Y		The care giver for the medication event.
CUID	NUMBER	Y		The care unit for the medication event.
VOLUME	NUMBER	Y		The volume of the medication event.
DOSE	NUMBER	Y		The doseage of the medication event.
DOSEUOM	VARCHAR2	Y		The units of measurement of the doseage of the medication event.
SOLUTIONID	NUMBER	Y		The solution used in the medication event.
SOLVOLUME	NUMBER	Y		The volume of solution used in the medication event.
SOLUNITS	VARCHAR2	Y		The units of the volume of solution used in the medication event.
ROUTE	VARCHAR2	Y		
SITE	VARCHAR2	Y		The site of the medication event.
STOPPED	VARCHAR2	Y		
ANNOTATION	VARCHAR2	Y		

Primary key(s)

MEDEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	SUBJECT_ID			

MEDEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ITEMID			

MEDEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CHARTTIME			

MEDEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ELEMID			

Foreign key(s)

6.1.21 NOTEEVENTS

Columns

Column Name	Data Type	Nullable	Default value	Comments
SUBJECT_ID	NUMBER	N		The patient for the note event.
NOTEID	NUMBER	N		The note item for the event.
ELEMID	NUMBER	N		
CHARTTIME	DATE	N		The chart time of the note event.
REALTIME	DATE	N		
CGID	NUMBER	Y		The care giver for the note event.
CORRECTION	CHAR	N		
CUID	NUMBER	Y		The care unit for the note event.
CATEGORY	VARCHAR2	Y		
TITLE	VARCHAR2	Y		
TEXT	VARCHAR2	Y		The text of the note event.

Primary key(s)

NOTEEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	SUBJECT_ID			

NOTEEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	NOTEID			

NOTEEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	ELEMID			

NOTEEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CHARTTIME			

NOTEEVENTS_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	CORRECTION			

Foreign key(s)

6.1.22 PARAMETER_MAPPING

Columns

Column Name	Data Type	Nullable	Default value	Comments
PARAM1_STR	VARCHAR2	Y		
PARAM1_NUM	NUMBER	Y		
CATEGORY	VARCHAR2	N		
PARAM2_STR	VARCHAR2	Y		
PARAM2_NUM	NUMBER	Y		
ORDER	NUMBER	Y		
VALID_FLG	CHAR	N	'Y'	
COMMENTS	VARCHAR2	Y		

Primary key(s)

Foreign key(s)

6.1.23 POE_MED

Columns

Column Name	Data Type	Nullable	Default value	Comments
POE_ID	NUMBER	N		
DRUG_TYPE	VARCHAR2	N		
DRUG_NAME	VARCHAR2	N		
DRUG_NAME_GENERIC	VARCHAR2	Y		
PROD_STRENGTH	VARCHAR2	Y		
FORM_RX	VARCHAR2	Y		
DOSE_VAL_RX	VARCHAR2	Y		
DOSE_UNIT_RX	VARCHAR2	Y		
FORM_VAL_DISP	VARCHAR2	Y		
FORM_UNIT_DISP	VARCHAR2	Y		
DOSE_VAL_DISP	NUMBER	Y		
DOSE_UNIT_DISP	VARCHAR2	Y		
DOSE_RANGE_OVERRIDE	VARCHAR2	Y		

Primary key(s)

Foreign key(s)

6.1.24 POE_ORDER

Columns

Column Name	Data Type	Nullable	Default value	Comments
POE_ID	NUMBER	N		
SUBJECT_ID	NUMBER	N		
HADM_ID	NUMBER	N		
START_DT	DATE	Y		
STOP_DT	DATE	Y		
ENTER_DT	DATE	N		
MEDICATION	VARCHAR2	Y		
PROCEDURE_TYPE	VARCHAR2	Y		
STATUS	VARCHAR2	Y		
ROUTE	VARCHAR2	Y		
FREQUENCY	VARCHAR2	Y		
DISPENSE_SCHED	VARCHAR2	Y		
IV_FLUID	VARCHAR2	Y		
IV_RATE	VARCHAR2	Y		
INFUSION_TYPE	VARCHAR2	Y		
SLIDING_SCALE	CHAR	Y		
DOSES_PER_24HRS	NUMBER	Y		
DURATION	NUMBER	Y		
DURATION_INTVL	VARCHAR2	Y		
EXPIRATION_VAL	NUMBER	Y		
EXPIRATION_UNIT	VARCHAR2	Y		
EXPIRATION_DT	DATE	Y		
LABEL_INSTR	VARCHAR2	Y		
ADDITIONAL_INSTR	VARCHAR2	Y		
MD_ADD_INSTR	VARCHAR2	Y		
RNURSE_ADD_INSTR	VARCHAR2	Y		

Primary key(s)

POE_ORDER_PK	**Type:**	**Primary Key**	**Unique:**	Y
Description:	Primary key for the table			
Columns:	POE_ID			

Foreign key(s)

6.1.25 REPORTEVENTS

Columns

Column Name	Data Type	Nullable	Default value	Comments
REPORT_ID	NUMBER	N		Table record unique identifier, the report event id.
SUBJECT_ID	NUMBER	N		
HADM_ID	NUMBER	N		
CATEGORY	VARCHAR2	N		
REPORT_DT	DATE	N		
EXAM_NAME	VARCHAR2	Y		
PATIENT_INFO	VARCHAR2	Y		
TITLE	VARCHAR2	Y		
TEXT	CLOB	Y		

Primary key(s)

Foreign key(s)

6.1.26 TOTALBALEVENTS

Columns

Column Name	Data Type	Nullable	Default value	Comments
RESET	NUMBER	Y		
STOPPED	VARCHAR2	Y		
ANNOTATION	VARCHAR2	Y		
SUBJECT_ID	NUMBER	N		
ITEMID	NUMBER	N		
CHARTTIME	DATE	N		
ELEMID	NUMBER	N		
REALTIME	DATE	Y		
CGID	NUMBER	Y		
CUID	NUMBER	Y		
PERVOLUME	NUMBER	Y		
CUMVOLUME	NUMBER	Y		
ACCUMPERIOD	VARCHAR2	Y		
APPROX	VARCHAR2	Y		

Primary key(s)

Foreign key(s)

ADDITIVES	
ADMISSIONS	The list of patient admissions to the ICU or other hospital room
A_CHARTDURATIONS	Chart event durations
A_IODURATIONS	Duration of IO events
A_MEDDURATIONS	Duration of medication events
CENSUSEVENTS	Events which record ICU transfers
CHARTEVENTS	Events which occur on a patient chart.
DB_SCHEMA	Database schema information.
DELIVERIES	
D_CAREGIVERS	Care giver records
D_CAREUNITS	Care unit records
D_CHARTITEMS	Chart items
D_IOITEMS	IO items.
D_MEDITEMS	Medication items.
D_PARAMMAP_ITEMS	Mappings between parameters.
D_PATIENTS	Patient records
ICD9	ICD9 code records.
ICU_STAY	ICU stay records.
IOEVENTS	IO events.
MEDEVENTS	Medication events.
NOTEEVENTS	Note events.
PARAMETER_MAPPING	Mappings between parameters.
POE_MED	
POE_ORDER	
REPORTEVENTS	Report events.
TOTALBALEVENTS	Total balance events.

6.2 Multiple ID mappings

Except for “invasive diastolic arterial blood pressure”, all parameter sample values are in the ”value1num” column in the CHARTEVENTS table. The systolic and diastolic invasive arterial blood pressure are stored under the same ITEMID; with the systolic value in “value1num”, and diastolic in “value2num”.

6.2.1 Bicarbonate (HCO3)

```
ITEMID  LABEL      CATEGORY
-------------------
787     Carbon Dioxide    Chemistry
812 HCO3              [value1num]
3810 Total CO2
```

6.2.2 Bilirubin (highest)

```
ITEMID  LABEL      CATEGORY
--------------------
848     Total Bili (0-1.5)    Chemistry
1538    Total Bili            Chemistry
```

6.2.3 Blood Pressure

Invasive (Arterial) Blood Pressure (IABP/IBP)

```
ITEMID  LABEL      CATEGORY
--------------------
52 Arterial Blood Pressure (Mean)
51 Arterial Blood Pressure (Systolic)  [value1num]
51 Arterial Blood Pressure (Diastolic) [value2num]
```

Invasive blood pressures are generally more accurate than non-invasive blood pressures.

Non Invasive Blood Pressure (NIBP)

```
ITEMID  LABEL      CATEGORY
--------------------
455     nabpsys  value1num
455     nabpdias  value2num
456     nabpmean  value1num
1149    NBP:
751     zzzNBP
```

6.2.4 Blood Transfusions

```
ITEMID  LABEL      CATEGORY
--------------------
  734 Packed RBC's 350.0ml
  31 RBC'S
  144 Packed RBC's
  172 OR Packed RBC's
  397 Washed PRBC's
  980 Packed RBC's 150.0ml
  1011 Packed RBC's 750.0ml
  1106 Packed RBC's 75.0ml
  1141 Packed RBC's 372.0ml
  1585 Packed RBC's 100.0ml
  1737 Packed RBC's 125.0ml
  1738 Packed RBC's 400.0ml
  2909 Packed RBC's 325.0ml
```

```
  3735 Packed RBC's 95.0ml
  3992 Packed RBC's 500.0ml
  4245 Packed RBC's 3.0ml
  4258 Packed RBC's 450.0ml
  4422 Packed RBC's 300.0ml
```

6.2.5 Cardiac Output (CO)

```
ITEMID  LABEL      CATEGORY
-------------------
  90 C.O.(thermodilution)
  89 C.O. (fick)
  1601 C.C.O
  2112 continuous C.O
```

Note that thermodilution CO calculations are generally more accurate than those calculated through the Fick method.

6.2.6 Carbon Dioxide (CO2)

```
ITEMID  LABEL      CATEGORY
-------------------
777     ArtCO2Calc    [value1num]
4199    ArtCO2Calc    [value1num]
787     CarbonDioxide [value1num]
3808    CarbonDioxide [value1num]
3810    CarbonDioxide [value1num]
```

6.2.7 Creatinine (highest)

```
ITEMID  LABEL      CATEGORY
-------------------
791 Creatinine (0-1.3)
3750 Creatinine   (0-0.7)
1525 Creatinine
```

6.2.8 Central Venous Pressure (CVP)

```
ITEMID  LABEL      CATEGORY
-------------------
  1103 cvp
  113 CVP
```

6.2.9 Glucose Levels

```
ITEMID  LABEL      CATEGORY
-------------------
```

```
811 Glucose (70-105)    [value1num]
1529 Glucose             [value1num]
```

6.2.10 Intra-aortic balloon (IABP) pump rate

```
ITEMID  LABEL      CATEGORY
--------------------
  224 IABP Mean
  225 IABP setting
  2162 IABP TIMING ADJUSTED
  2391 IABP
  2515 IABP-BP
  2865 iabp-bp
  6424 IABP BP
```

6.2.11 Intra-cranial Pressure (ICP)

```
ITEMID  LABEL      CATEGORY
--------------------
  1374 ICP Right
  226 ICP
  2045 icp left
  2745 ICP LEFT
  5856 icp
```

6.2.12 IV Fluids

```
ITEMID  LABEL      CATEGORY
--------------------
1       HourlyIn           pervolume
2       HourlyOut          pervolume
2       DailyOut           cumvolume
18      HourlyIV           pervolume
29      NetHourlyBalance   pervolume
```

6.2.13 Lactate

```
ITEMID  LABEL      CATEGORY
--------------------
818 Lactic Acid(0.5-2.0)
1531 Lactic Acid
```

6.2.14 Oxygen Saturation (SpO2/SaO2)

```
ITEMID  LABEL      CATEGORY
--------------------
```

```
  1148 SpO2:
  646 SpO2
  834 SaO2
```

6.2.15 pH

```
ITEMID  LABEL      CATEGORY
-------------------
865     pH value1num
1126    pH value1num
780     pH value1num
4202    pH value1num
4753    pH value1num
```

6.2.16 Potassium

```
ITEMID  LABEL      CATEGORY
-------------------
829 Potassium (3.5-5.3) [value1num]
1535 Potassium           [value1num]
3792 Potassium  (3.5-5.3) [value1num]
```

6.2.17 Pressor Medications

```
ITEMID  LABEL      CATEGORY
-------------------
 46     Isuprel
 47 Levophed
 120 Levophed-k
 43 Dopamine
 307 Dopamine Drip
 44 Epinephrine
 119 Epinephrine-k
 309 Epinephrine Drip
 51 Vasopressin
 127 Neosynephrine
 128 Neosynephrine-k
```

Note: Drugs interact, and can often have the oppositcve effect you might expect when the drug dose is low or high. For example, the vasodepressor action of Isuprel is reversed to vasopressor action by small doses of ergotamine or ergotoxine. This reversal is associated with a marked increase in the amplitude of ventricular contraction, in pulse pressure and in rate [19].

6.2.18 Pulmonary Arterial Pressure (PAP)

```
ITEMID  LABEL      CATEGORY
```

```
---------------------------
  491 PAP Mean
  492 PAP S/D
```

6.2.19 Respiration Rate

```
ITEMID  LABEL      CATEGORY
---------------------------
  614 Resp Rate (Spont)
  615 Resp Rate (Total)
  618 Respiratory Rate  (*)
  653 Spont. Resp. Rate
  1151 Respiratory Rate:
  1635 HIGH Resp Rate
  1884 Spont Resp rate
  2117 Low resp rate
  3603 Resp Rate
  3337 Breath Rate
```

(*) indicates preferred ITEMID.

6.2.20 Sodium

```
ITEMID  LABEL      CATEGORY
---------------------------
837 Sodium (135-148)   [value1num]
1536 Sodium             [value1num]
3803 Sodium  (135-148)  [value1num]
```

6.2.21 Temperature

```
ITEMID  LABEL      CATEGORY
---------------------------
676 Temperature C          [value1num]
677 Temperature C (calc)   [value1num]
678 Temperature F          [value1num]
679 Temperature F (calc)   [value1num]
```

6.2.22 Urine Output

```
ITEMID  LABEL      CATEGORY
---------------------------
26 Urine Out Total
3053 URINE OUT
3165 Urine Output Total
3175 Urine
```

```
3462 urine
3519 urine amnt
```

6.2.23 Ventilators

```
ITEMID  LABEL      CATEGORY
--------------------------
505     peep              [value1num]
506     peep              [value1num]
535     PeakInspPressure  [value1num]
543     PlateauPressure   [value1num]
544     Plateau Time (7200)
545     Plateau-Off
619     Respiratory Rate Set
 39     Airway Size
535     Peak Insp. Pressure
683     Tidal Volume (Set)
720     Ventilator Mode
721     Ventilator No.
722     Ventilator Type
732     Waveform-Vent
```

6.2.24 White Blood Cell Count (WBC)

```
ITEMID  LABEL      CATEGORY
--------------------------
1542 WBC                 [value1num]
1127 WBC   (4-11,000)    [value1num]
861 WBC (4-11,000)      [value1num]
4200 WBC 4.0-11.0        [value1num]
```

6.3 Commonly used parameters

Together with the above parameters in section 6.2, the following list may be helpful.

```
ITEMID  LABEL      CATEGORY
--------------------------
770     AST         [value1num]
781 BUN (6-20)  [value1num]
198 GCS         [Total]
828     Platelets   [value1num]
211     heartrate   [value1num]
813     Hematocrit  [value1num]
20001   SAPS1       [value1num]
504  PCWP        [value1num]
```

GCS = Glasgow Coma Scale. BUN = Blood Urea Nitrogen (also known as Urea or Urea nitrogen). SAPS1 indicates Simplified Acuity Score (version 1). PCWP = Pulmonary Capillary Wedge Pressure (or simply 'Wedge Pressure').

6.4 Frequency of all ICD-9 codes for adult ICU-related hospital admissions

Table 6.1 lists the most frequent ICD-9 codes (including the primary ICD-9 codes) for the same population for the thirty-two major ICD-9 categories. Note that there is one extra category in the non-primary code categories (see table 1.1) "Supplementary classification of external causes of injury and poisoning" (code range E800 to E999). These ICD-9 codes are not used for primary classification.

Table 6.1: Distribution of major categories of ICD-9 codes for adult ICU-related hospital admissions (n = 211,416)

Category	**Code Range**	**Number of Codes**	**%**
Metabolic disorder	240-279	23927	11.32%
Pulmonary disease	460-519	18694	8.84%
Other forms of heart disease	420-429	13203	6.25%
Digestive disease	520-579	13181	6.23%
Supplementary classification of factors influencing health status and contact with health services	V01-V86	12801	6.05%
Ischemic heart disease	410 - 414	12794	6.05%
Renal insufficiency	580-629	11262	5.33%
Hypertensive disease	401-405	10970	5.19%
Symptoms, signs, and ill-defined conditions	780-799	8777	4.15%
Diseases of the blood and blood-forming organs	280-289	8744	4.14%
Trauma	800-959	7975	3.77%
Heart failure	428	7470	3.53%
Infectious diseases	001-139	7388	3.49%
Mental disorders	290-319	7337	3.47%
Supplementary classification of external causes of injury and poisoning	E800-E999	6632	3.14%

continued on next page

continued from previous page			
Category	**Code Range**	**Number of Codes**	**%**
Arteries and veins	440-459	5828	2.76%
Neoplasms	140-239	5403	2.56%
Neurologic disease	320-389	5140	2.43%
Diseases of the musculoskeletal system & connective tissue	710-739	3632	1.72%
Other complications of procedures, NEC	998	2936	1.39%
Cerebrovascular disease	430-438	2849	1.35%
Diseases of the skin and subcutaneous tissue	680-709	2799	1.32%
Complications affecting specified body systems, not elsewhere classified	997	2551	1.21%
Complications peculiar to certain specified procedures	996	2403	1.14%
Other and unspecified effects of external causes	990-995	1941	0.92%
Chronic rheumatic heart disease	393-398	1438	0.68%
Diseases of pulmonary circulation	415-417	1233	0.58%
Congenital anomalies	740-759	673	0.32%
Complications of pregnancy, childbirth, and the puerperium	630-677	627	0.30%
Poisoning	960-989	590	0.28%
Complications of medical care, not elsewhere classified	999	213	0.10%
Acute Rheumatic fever	390-392	5	0.00%
Total		**211416**	**100.00%**

Bibliography

[1] Tsien CL, Fackler JC. An annotated data collection system to support intelligent analysis of intensive care unit data. In IDA '97: Proceedings of the Second International Symposium on Advances in Intelligent Data Analysis, Reasoning about Data. London, UK: Springer-Verlag. ISBN 3-540-63346-4, 1997; 111–121.

[2] Norris PR, Dawant BM. Closing the loop in ICU decision support: Physiologic event detection, alerts, and documentation. J Am Med Inform Assoc 2002;9(90061):S102–107. URL http://www.jamia.org/cgi/content/abstract/9/6_suppl_1/S102.

[3] Abbott RR, Setter M, Chan S, Choi K. APACHE II: prediction of outcome of 451 ICU oncology admissions in a community hospital. Ann Oncol Sep 1991;2(8):571–574.

[4] Glance LG, Osler TM, Dick AW. Identifying quality outliers in a large, multiple-institution database by using customized versions of the Simplified Acute Physiology Score II and the Mortality Probability Model II0. Crit Care Med Sep 2002;30(9):1995–2002. URL http://dx.doi.org/10.1097/01.CCM.0000026324.64324.59.

[5] Clifford GD, Long WJ, Moody GB, Szolovits P. Robust parameter extraction for decision support using multimodal intensive care data. Phil Trans Royal Soc A January 2009;367(1877):411–429. URL http://www.ncbi.nlm.nih.gov/pubmed/18936019. Special issue on Signal Processing in Vital Rhythms and Signs.

[6] Neamatullah I, Douglass M, Lehman LH, Reisner A, Villarroel M, Long WJ, Szolovits P, Moody GB, Mark RG, Clifford GD. Automated de-identification of free-text medical records. BMC Med Inform Decis Mak 2008;8(32). URL http://www.biomedcentral.com/1472--6947/8/32/.

[7] Zong W, Moody GB, Mark RG. Reduction of false arterial blood pressure alarms using signal quality assessment and relationships between the electrocardiogram and arterial blood pressure. Med Biol Eng Comput 2004; 42:698–706.

[8] Sun JX, Reisner AT, Mark RG. A signal abnormality index for arterial blood pressure waveforms. Comput Cardiol September 2006;33:13–16.

[9] Li Q, Mark RG, Clifford GD. Robust heart rate estimation from multiple asynchronous noisy sources using signal quality indices and a Kalman filter. IOP Physiol Meas January 2008;29(1):15–32. (Awarded the Martin Black Prize for Best Paper in Physiological Measurement in 2008).

[10] Li Q, Mark RG, Clifford GD. Artificial arterial blood pressure artifact models and an evaluation of a robust blood pressure and heart rate estimator. BMC Biomed Eng Online 2009;8(13). URL http://www.biomedical-engineering-online.com/content/8/1/13.

[11] Aboukhalil A, Nielsen L, Saeed M, Mark RG, Clifford GD. Reducing false alarm rates for critical arrhythmias using the arterial blood pressure waveform. J Biomed Inform June 2008;41(3):442–451. ISSN 1532-0464. URL http://www.ncbi.nlm.nih.gov/pubmed/18440873. Available online March 21st 2008.

[12] Saeed M, Clifford GD, Villarroel M, Moody GB, Lehman L, Abdala O, Douglass MM, Frassica J, Heldt T, Hug C, Janz BA, Kyaw TH, Lieu C, Long WJ, Moody B, Nielsen L, Neamatullah I, Raber G, Reisner A, Szolovits P, Verghese G, Mark RG. MIMIC-II: A major new database to support research in ICU patient monitoring and clinical decision support systems. Critical Care 2009;In Submission.

[13] Beaulieu A. Learning SQL. O'Reilly, August 2005. ISBN 0-596-00727-2.

[14] Goldberger AL, Amaral LAN, Glass L, Hausdorff JM, Ivanov PC, Mark RG, Mietus JE, Moody GB, Peng CK, Stanley HE. Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. Circulations 2000;101(23):e215–e220. URL http://www.physionet.org.

[15] Moody G, Mark R. A database to support development and evaluation of intelligent intensive care monitoring. Computers in Cardiology 1996 Sep 1996;657–660.

[16] Java Database Connectivity (JDBC). URL http://java.sun.com/products/jdbc/overview.html.

[17] Abdala OT, Saeed M. Estimation of missing values in clinical laboratory measurements of ICU patients using a weighted K-nearest neighbors algorithm. Comput Cardiol 2004;31:693–696.

[18] Hug C, Clifford GD. An analysis of the errors in recorded heart rate and blood pressure in the ICU using a complex set of signal quality metrics. Comput Cardiol September 2007;34:641–645.

[19] Lands AM, Luduena FP, Grant JI, Ananenko E, Tainter ML. Reversal of the depressor action of n-isopropylarterenol (isuprel) by ergotamine and ergotoxine. J Pharmacol Exp Ther 1950;100(3):284–297. URL http://jpet.aspetjournals.org/cgi/content/abstract/100/3/284.