FAQs

Q: How do I get access to the MIMIC II Waveform Database or the MIMIC II Clinical Database?

The MIMIC II Waveform Database (bedside monitor trends and waveforms) is freely available. The MIMIC II Clinical Database (all other elements of MIMIC II) is available under terms of a data use agreement (DUA). See **MIMIC II: Getting Access** for details.

Applications may require several days for review; please be patient. If your application is approved, we will contact you with instructions for database access.

If you are having problems with the process, please **contact us**.

Q: What data does the MIMIC II database contain?

A: The MIMIC II database contains clinical data for over 40,000 ICU stays of over 32,000 patients, collected over 7 years. The information includes data pertaining to:

- Patient events such as movement between wards.
- Patient diagnoses using ICD-9 codes.
- Data from bedside monitors such as ECG waveforms, arterial blood pressure and oxygen saturation levels.

- Clinical data such as laboratory test results, medications, patient demographics, nursing progress notes, discharge summaries, etc.

Q: How can I learn how to use the database?

A: The database is accessed with a well known query language, SQL. To learn more about relational databases, and how to use SQL, there are many resources freely available on the web. We recommend this, reasonably short, online book, with an emphasis on the first four chapters: **http://philip.greenspun.com/sql/**.

Q: Can I perform out-of-hospital mortality studies using MIMIC II?

A: Yes, MIMIC II has an in-hospital mortality flag which is recorded in the "d_patients" table, along with the date of death (if available). MIMIC II defines a subject to have died in the ICU if the outtime from the ICU is within 24 hours of the date of death. The database has also been augmented with data from the **Social Security Death Index** (SSDI) to determine out-of-hospital mortality. MIMIC II V2.5 includes SSDI data from 21 months after the last admission, so one can determine 1.75 year mortality. In MIMIC II V2.6, the inclusion of additional subjects narrowed the interval between the last admission and the date of the SSDI data, so one can only calculate 0.75 year mortality.

The SSDI is not perfect and may contain incorrect and/or missing data. Furthermore, updated releases of SSDI can contain corrections/additions to prior information, so it is possible that a patient's date of death may change between releases of MIMIC II. The limitations of the mortality information should always be considered when performing mortality studies.

Neonates often do not have social security numbers before they are discharged from the hospital, and so we are unable to match their dates of death.

Q: What if I want to perform a study that requires knowing in what year the patients were admitted?

A: Since all dates have been shifted randomly into the future, it is impossible to perform any study that requires knowing which patient was admitted before another patient. For some research studies, it is essential to know the approximate date during which care was administered. We make this information available to researchers who require it, subject to an extended data use agreement

Q: Is this all the data concerning a patient?

A: No, not all data falls into a hospital or ICU stay. Patients may be transferred to other hospitals. they are likely to have a medical history before and after the six-year time period of our data collection.

Q: Did you remove any patients?

A: Yes. We removed all VIPs. We also removed all patients who turned 90 during a stay. Patients in the test set for the PhysioNet Computers in Cardiology Competition 2009 were removed while the challenge was in progress, they have since been restored

Q: What is the difference between charttime and realtime?

A: When two times are recorded for the same observation, the earlier time generally indicates when the observation was made, and the later time indicates when information about the event was entered into the electronic medical record. For example, a patient might have had blood drawn at 7:45 for a lab test that was not completed and logged into the record until 10:05. The earlier time is important in understanding the state of the patient at that time; the later time is significant since it may represent the first time at which the observation might have been able to influence the patient's care. In MIMIC II the observation time is 'charttime' and the record time is 'realtime'. In general, you should not use the 'realtime' column, unless you are particularly interested in the time when the data was entered (when studying, for example, practice variation among the staff).

A further complication is that some of these timestamps are generated automatically while others are entered manually, and there is no guarantee that all of the clocks used are precisely synchronized -- so that if two events occur within a few minutes, their timestamps do not unambiguously indicate which event occurred first.

Q: How do I configure pgadmin3 with the MIMIC II VM?

A: The PostgreSQL database in the VM doesn't use password authentication it uses ident authentication, so you won't need a password for MIMIC2. The Host should be blank. The Username should be 'mimic2' . the "Maintenance DB" should be 'MIMIC2'.

Q: How was the SAPS-I score calculated?

A: The SAPS-I score was calculated using the method outlined in this publication: **Le Gall J-R, Loirat P, et al. A simplified acute physiology score for ICU patients. Crit Care Med. 1984; 12: 975-977**.

Q: Why do I keep getting 'certificate' errors?

A: All of our secure web pages are signed using SSL certificates. Install our **certificate authority** to remove these warnings.

Q: Does the database contain codes for the procedures performed on a patient?

A: Yes, in version 2.6, patient procedures can be found in the PROCEDUREEVENTS table. The table contains a 'sequence' number, and a date. The ordering of the dates do not necessarily match the ordering of the sequence. The reason for this is unknown, but we suggest that users trust the dates provided, rather than the sequence.

Q: How do I export entire tables from the database?

A: If you wish to export complete tables (e.g. SELECT * FROM medevents), use the flat files at:
https://physionet.org/works/MIMICIIClinicalDatabase/files/

Use your PhysioNetWorks user name and password to get access to these files; your MIMIC user name and password will not work on PhysioNetWorks.

Q: How can I copy the entire MIMIC II Clinical Database?

A: We recommend using the MIMIC Importer to load the flat files (see the previous question) into a PostgreSQL database that can run on your own Linux, Mac OS X, or Windows computer. See "The MIMIC Importer" in the Downloads section of: **https://physionet.org/works/MIMICIIClinicalDatabase/files/**

Use your PhysioNetWorks user name and password to get access to these files.

Q: How can I get access to the MIMIC II Explorer/Query Builder?

A: First, create a PhysioNetWorks account at

https://physionet.org/users/

and log into it. On your PhysioNetWorks home page, click on the link titled "MIMIC II Clinical Database", and follow the instructions to apply for access. All applications are reviewed, and if yours is approved you will receive separate emails containing instructions for downloading the database from PhysioNetWorks, and for exploring the database on-line via the MIMIC II Explorer/Query Builder web site. Note that approval usually requires one or two working days, and will be delayed if your request is missing any required information.

Q: I am having problems logging in to the system. What should I do?

A: **IMPORTANT**: You have two separate sets of login credentials, which are not interchangeable. Your email address is your PhysioNetWorks user name, and the first part of your email address (up to but not including "@") is your MIMIC user name. The password that you chose when creating your PhysioNetWorks account is your PhysioNetWorks password, and the password that you were assigned in a separate email from the MIMIC project is your MIMIC password.

To login to the MIMIC II Explorer/Query Builder, go to the MIMIC login page:

https://mimic2app.csail.mit.edu/querybuilder/

enter your MIMIC user name and your MIMIC password, and click on "Login". If you have forgotten your MIMIC password, or if you wish to change it for any other reason, request a new one using the MIMIC Contact form at

http://mimic.physionet.org/contact.html

Requests for new MIMIC passwords are reviewed by the MIMIC project, and they are normally answered within one or two business days (usually not on weekends or US holidays).

To login to PhysioNetWorks, go to the PhysioNetWorks login page:

https://physionet.org/users/

click on "Login", enter your PhysioNetWorks user name and your PhysioNetWorks password, and click on "OK". If you have forgotten your PhysioNetWorks password, or if you wish to change it for any other reason, go to the PhysioNetWorks login page:

https://physionet.org/users/

click on "Reset password", and follow the instructions that will be sent to your email address by the autoresponder within a minute or two.

Q: How can I export more than 50 rows from QueryBuilder?

A: QueryBuilder is intended for initial exploration and basic analysis of the data available in MIMIC II. If you need to export large quantities of data, or run large queries which take a long time, you should use an alternative access method. The recommended method is the **MIMIC II Virtual Machine** (PhysioNetWorks username and password, and access to the MIMIC II flat files required) which provides a downloadable image containing a system to build your own local copy of the database.

Q: My query is taking too long to run using QueryBuilder, how can I run longer queries?

A: QueryBuilder is intended for initial exploration and basic analysis of the data available in MIMIC II. If you need to export large quantities of data, or run large queries which take a long time, you should use an alternative access method. The recommended method is the **MIMIC II Virtual Machine** (PhysioNetWorks username and password, and access to the MIMIC II flat files required) which provides a downloadable image containing a system to build your own local copy of the database.

Q: I need to filter CMO patients, how could I do that?

A: MIMIC uses the label name "code status" to label a patient's status. The label is also associated with an itemid. Since the CHARTEVENTS table uses itemid, the user must know the itemid that is associated with label "code status." This can simply be found by running this query:

```
-- find itemid for label "code status"
select itemid
  from mimic2v26.d_chartitems
where lower(label) like '%code%';
```

The result shows "code status" has the itemid "128" value. One can then query rows that matches with this itemid in the CHARTEVENTS table.

```
-- find all chartevents where itemid = 128
select *
from mimic2v26.chartevents
where itemid = '128';
```

The result shows all rows that has itemid equaling to "128." However, the question specifically asks for patients there are in CMO. There is one more filter left to do and this again uses the "like" condition.

```
-- find all chartevents where itemid = 128 and value1 = comfort measures (CMO)
select *
from mimic2v26.chartevents
where itemid = '128' and lower(value1) like '%comfort%';
```

The query goes through CHARTEVENTS table and filters out results that has itemid equaling "128" (which the label is "code status") and value1 equaling "comfort." The results are all rows in CHARTEVENTS table that have the value "Comfort Measure."

It is important to note that INSERT INTO `lcp_faq` VALUES for "code status" can change overtime for a patient's stay. For instance, a patient can change from DNR to CMO, so one needs to look at the charttime carefully to determine if or when a status changes.

Q: I have tried to extract from the clinical database the INSERT INTO `lcp_faq` VALUES for blood pressure, but I wasn't able to find the right ids?

A: MIMIC uses the label name "Arterial BP" for blood pressure. The label is also associated with an itemid. Since the CHARTEVENTS table uses itemid, the user must know the itemid that is associated with label "Arterial BP." This can simply be found by running this query:

```
-- find itemid for label "arterial bp"
select *
from mimic2v26.d_chartitems
where lower(label) like '%arterial bp%';
```

Q: How do you find out when I patients is on mechanical ventilation?

A: MIMIC does not specifically keep timestamps of ventilation time. However, it does keep track of measurements that could point toward the patient being on a ventillatior. It is suggested that if the measurements stop appearing for a 24 hour period, then the patient was extubated. The measurements that we look for are "FiO2" and "FiO2 Set" with itemid 3420 and 190 respectively. The itemid for each can be found by using this query:

```
select *
from mimic2v26.d_chartitems_detail
where label = 'FiO2' or label = 'FiO2 Set';
```

Q: Why do some patients have ages > 200 years?

A: The HIPAA rules for protected health information require that we obscure real ages for patients who are aged over 90 years old (and still alive). We have shifted the DOB for these patients so that their age on first admission is 200 years. The median age for the patients whose age was obscured is 91.4 (Version 2.6) - this value can be used as a surrogate age for these patients. Real ages are shown for patients who have died.

Q: My query is taking far too long. What is the problem?

A: You should check the explain plan to see how your query will be executed. It can help you locate problematic joins and see where the bottlenecks are. The plan only shows what will be done, it does not actually execute the query.

Obtain the explain plan by adding 'explain plan for' in front of your query:

```
explain plan for select * from mimic2v26.d_patients dp, mimic2v26.chartevents ce where dp.subject_id = ce.subject_id;
```

You can then obtain the plan by running:

```
SELECT PLAN_TABLE_OUTPUT FROM TABLE(DBMS_XPLAN.DISPLAY());
```

The plan should show a 'hash join' on d_patients and chartevents. If you make a mistake, and forget the 'where' clause, the plan shows terabytes of memory and incredibly long time predictions.

Q: What is the difference between first_careunit and first_service in the icustay_details view?

A: Basically, careunit is physical location and service refers to who took care of the patient. For example, a SICU patient may have been cared for in the MICU due to a shortage of beds. In general, the careunit data label may not be reliable because the hospital can change the careunit labels several times during the MIMIC data collection time. So service should be more accurate in describing who care to the patient, unless you are interested in study the actual physical locations of patients.

Q: How is vasopressor dosing expressed?

Vasopressor dosing (both maximal and cumulative) may be expressed as “norepinephrine equivalents.” Standard dosing ranges for commonly used vasopressors are as follows: norepinephrine 0–20 µcg/min, epinephrine 0–10 µcg/min, phenylephrine 0–300 µcg/min, and dopamine 0–20 µcg/kg/min. The highest standard range dosage of each medication as listed is assigned a norepinephrine equivalence of 20 and that percentage less than the maximum standard range dosing will proportionally reduce the norepinephrine equivalency. For example, a phenylephrine infusion of 150 µcg/min will be interpreted as a norepinephrine equivalent of 10. The non-titrated, non-catecholamine vasopressin will be assigned a norepinephrine equivalence of 5.

Q: There is an ICU named “FICU”. What does FICU stand for?

A: FICU is a Medical ICU in a building whose name begins with “F”. So you can consider it to be equivalent to MICU.

Q: I notice that there are multiple ICD-9 entries per ICU stay. Why is that?

Each patient typically has many codes. The first code in the list represents the primary diagnosis for the admission.