



HHS Public Access

Author manuscript

J Biomed Inform. Author manuscript; available in PMC 2024 March 01.

Published in final edited form as:

J Biomed Inform. 2023 March ; 139: 104269. doi:10.1016/j.jbi.2022.104269.

Mining for equitable health: Assessing the impact of missing data in electronic health records

Emily Getzen^{a,*}, Lyle Ungar^b, Danielle Mowery^a, Xiaoqian Jiang^c, Qi Long^{a,*}

^aDepartment of Biostatistics, Epidemiology, and Informatics, University of Pennsylvania, Philadelphia, PA, United States

^bDepartment of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, United States

^cUTHealth School of Biomedical Informatics, Houston, TX, United States

Abstract

Electronic health records (EHR) are collected as a routine part of healthcare delivery, and have great potential to be utilized to improve patient health outcomes. They contain multiple years of health information to be leveraged for risk prediction, disease detection, and treatment evaluation. However, they do not have a consistent, standardized format across institutions, particularly in the United States, and can present significant analytical challenges— they contain multi-scale data from heterogeneous domains and include both structured and unstructured data. Data for individual patients are collected at irregular time intervals and with varying frequencies. In addition to the analytical challenges, EHR can reflect inequity— patients belonging to different groups will have differing amounts of data in their health records. Many of these issues can contribute to biased data collection. The consequence is that the data for under-served groups may be less informative partly due to more fragmented care, which can be viewed as a type of missing data problem. For EHR data in this complex form, there is currently no framework for introducing realistic missing values. There has also been little to no work in assessing the impact of missing data in EHR. In this work, we first introduce a terminology to define three levels of EHR data and then propose a novel framework for simulating realistic missing data scenarios in EHR to adequately assess their impact on predictive modeling. We incorporate the use of a medical knowledge graph to capture dependencies between medical events to create a more realistic missing data framework. In an intensive care unit setting, we found that missing data have greater negative impact on the performance of disease prediction models in groups that tend to have less access to healthcare, or seek less healthcare. We also found that the impact of missing data on disease prediction models

*Corresponding authors. getzene@penmedicine.upenn.edu (E. Getzen), qlong@upenn.edu (Q. Long).

CRediT authorship contribution statement

Emily Getzen: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Lyle Ungar:** Methodology, Writing – review & editing. **Danielle Mowery:** Methodology, Writing – review & editing. **Xiaoqian Jiang:** Methodology, Writing – review & editing. **Qi Long:** Conceptualization, Methodology, Resources, Writing – review & editing, Supervision, Funding acquisition.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

is stronger when using the knowledge graph framework to introduce realistic missing values as opposed to random event removal.

Keywords

Missing data; Electronic health records; Knowledge graph; Fairness; Health disparities

1. Introduction

In 2009, the Health Information Technology for Economic and Clinical Health (HITECH) act was passed to promote the adoption of health information technology in the United States [12]. For more than a decade, the use of Electronic Health Records systems has increased tremendously, enabling the availability of patient health histories for analysis. EHR are the digital version of a patient's paper medical chart. The patient record contains detailed medical history, diagnoses, medications, treatment plans, laboratory test results, among other health information necessary for understanding a patient's health status as well as their engagement with care providers within the inpatient, outpatient, and ambulatory care settings. Therefore, EHR data have been leveraged to address a variety of use cases including clinical research, quality improvement, clinical decision support, and population health. To support these use cases, computational methods are applied to lower-level tasks e.g., extracting medical events that can inform higher-level tasks, e.g., disease prediction [24]. As of December 2019, there are slightly under 2.1 million papers published on EHR for biomedical research [23]. By harnessing the power of EHR, we could improve early-stage disease detection, risk prediction, and treatment evaluation; thus leading to a significant reduction in patient costs and improvement in patient outcomes. Table 1.

In the context of applying machine learning to solve problems for risk prediction, disease detection, and treatment evaluation, EHR pose many challenges— they do not have a consistent, standardized format across institutions particularly in US, can contain human errors and introduce collection biases. In addition, some institutions or geographic regions do not have access to the technology or financial resources necessary to implement EHR, thus resulting in vulnerable and disadvantaged communities not being electronically visible [23].

1.1. Defining levels of information complexity for EHR data

EHR data come in a variety of forms; however, to our knowledge, there is no terminology to delineate the levels of complexity. For example, raw data in EHR is very different from EHR data that have been heavily pre-processed and curated to generate a matrix form with a priori selected and well-defined features. To describe different levels of EHR data, we suggest the following terminology: **Level 0** data refers to the raw data that resides in EHR systems without any pre-processing steps. **Level 1** data refers to data after limited data pre-processing steps including harmonization, integration and curation. **Level 2** data refers to EHR data in matrix form that includes variables/features extracted through chart reviews, or other mechanisms (see Fig. 1 for an example). Level 0 data lack structure or standardization (e.g., narrative text, non-codified fields). Level 1 data are complex – they

oftentimes come in the form of a sequence of events with heterogeneous structure (e.g., templated text, codified prescriptions and diagnoses, lab tests, vitals mapped to standard terminologies and vocabularies). In addition, data for individual patients are collected at irregular time intervals and frequencies. These challenges prevent researchers from being able to directly apply classical statistical and machine learning methods for analysis of Level 1 data. While Level 2 data is amenable to such analyses, there is significant time spent pre-processing such as feature extraction/engineering through chart review or other means, and there may be substantial information loss from Level 1 to Level 2 EHR data. This information loss can be a result of data that are non-conformant or non-computable becoming “missing” or lost. Thus, for machine learning models to be adopted in a clinical setting, it is highly advantageous to build models that can use Level 1 data directly. Note that in this setting, our time series windows are denoted by medical visits (or encounters) where measurements may be summarized for a single visit in order to predict a future disease (See Fig. 1). However, there are other scenarios in which time windows can be defined differently (for example, hourly windows to predict complications within a single hospital admission).

Individual patients typically have varying amounts of data in their EHR for a variety of reasons such as lack of collection / documentation, less medical visits [26]. Many issues can contribute to biased data collection, particularly with regard to trends in health-decision making of different patients. For example, studies have found that individuals from vulnerable populations (e.g., immigrants, low socioeconomic status, psycho-social issues) are more likely to visit multiple healthcare institutions to receive care— this leads to data fragmentation, which appears as “missingness” to end users. In addition, low socioeconomic status patients may receive fewer diagnostic tests and medications for chronic disease [8]. Racial and ethnic minority groups have sub-optimal access to healthcare, with Black Americans in particular facing disparities in terms of health status, mortality, and morbidity [10]. It has also been found that male patients tend to seek healthcare less than female patients [15]. Age also plays a role in health-seeking behaviors – studies found that adults 65 years or older tend to have more consultations with family physicians and undergo an annual health check [3]. The barriers that exist for young people seeking healthcare include discontinuities in care, lack of payment for transition support, and lack of preparedness for an adult model of care [16]. The consequence is that less data and information are documented in EHR for patients belonging to certain groups due to fragmented care [18].

Machine learning models trained on data where certain groups are prone to have less data may exhibit unfair performance for these populations [7]. This can be viewed as a missing data problem. There is a growing recognition that ubiquitous missing data in EHR, even when analyzed using powerful statistical and machine learning algorithms, can yield biased findings and unfair treatment decision, further exacerbating existing health disparities [8,18]. Furthermore, when it comes to EHR use for research, oftentimes investigators search for patients whose data are complete enough for analytical purposes [25], more likely excluding patients from under-served populations.

1.2. Assessing impact of missingness on EHR data

Strategies for dealing with missing data tend to rely on assumptions about the nature of the mechanism that causes the missingness. As such, missing data are typically categorized into one of three classes: **Missing Completely at Random (MCAR)**, **Missing at Random (MAR)**, and **Missing Not at Random (MNAR)**. MCAR, a term coined by [20], describes data where the probability of being missing is the same for all cases, thus implying that the cause of missingness is unrelated to the data itself. An example could be a lab technician forgetting to input data points into a patient's health record, regardless of any attributes that describe a patient. The second class, Missing at Random (MAR), describes data in which missingness depends on the observed data. For example, a patient's demographic characteristics which are recorded are associated with seeking less healthcare and therefore having sparser medical records. MAR is more general and realistic than MCAR in practice. When data are MCAR or MAR, the response mechanism is termed ignorable, or a researcher can ignore the reasons for missingness in the analysis and thus simplify model-based methods [11]. If neither MCAR nor MAR holds, the missingness is categorized as Missing Not at Random (MNAR). This means that the probability of being missing may depend on missing values [2]. This type of missingness can also be classified as non-ignorable if a patient's unobserved underlying condition (e.g., undiagnosed depression) prevents them from traveling to see a healthcare provider [22].

There is a growing body of work on handling missing data in health records, and as such there exists current methods for simulating missingness based on the three mechanisms. [1] outlined considerations for dealing with missingness in EHR data based on their own numerical experiments. For their MAR simulations, missingness in a feature depends on another feature value. For MNAR simulations, missingness in a feature depends on its own value being in a selected quartile. [13] conducted simulation studies to compare phenotyping methods to their Bayesian latent class approach. For MAR, missingness was simulated with a Bernoulli distribution based on patient-specific missingness probabilities varying based on age, race, and body mass index (BMI). For MNAR, missingness probabilities depended on Type 2 Diabetes Mellitus status. Both papers simulate missingness in Level 2 EHR data, or EHR data contained in a matrix with a priori defined features. Again, from Level 1 EHR data to Level 2, there is a sizable loss of information in addition to the labor-intensive process of manually and/or automatically converting to Level 2 data. More notably, none of the existing works account for dependencies between EHR events when artificially generating missing data in their experiments. An example of dependency is the relationship between diabetes and insulin. Most of the time, insulin is used to treat diabetes. Thus, if an instance of diabetes is removed, the insulin prescription should also be removed most of the time to reflect a realistic clinical scenario.

1.3. Modeling medical event relatedness using knowledge graphs

To determine the relatedness between medical events, one mechanism might be to train a knowledge graph. Its information could then be leveraged to simulate realistic missingness. Knowledge graphs have previously been very useful in the analysis of electronic health records. Goodwin and Harabagiu [9] automatically constructed a graph of clinically related concepts and presented an algorithm for determining similarity between medical concepts.

Santos et al. [21] developed a knowledge graph that integrates proteomics and clinical data in order to assist with clinical decision-making. Rotmensch et al. [19] also developed knowledge graphs that link symptoms and diseases directly from electronic health records. They developed three types of graphs based on logistic regression, naïve Bayes classifier, and a Bayesian network using noisy OR gates. The constructed knowledge graphs were evaluated against Google's manually-constructed knowledge graph and against physician opinions. They found that the noisy OR model produced a high quality graph that outperforms the other models. None of these methods have been applied for the purpose of accounting for event dependencies when simulating realistic missing data. We use the noisy OR-gate model from Rotmensch et al. [19] in our experiments due to its high performance compared to naïve Bayes and logistic regression approaches, its flexible assumptions, and ease of use.

There is currently no framework for introducing realistic missing values in Level 1 EHR data. There is also a gap in investigating the effects of missing data with regard to these data. As stated previously, Level 1 EHR represent a sequence of events with heterogeneous structure with features that are not well-defined. This also includes repeated measures at different time points for different patients. Events are oftentimes coded multiple times in a single patient medical visit, and frequently there exists causal pathways between medical events in the patient's record. As a result, we cannot rely on traditional methods to simulate realistic missing data in EHR and assess their impact.

In this work, we develop an innovative framework to simulate missingness under the three mechanisms (namely, MCAR, MAR, and MNAR) that mimic real-life situations and health-seeking behaviors of various groups of patients for Level 1 EHR. We account for dependencies between medical events through the use of a medical knowledge graph. We create models for Level 1 EHR for predicting future diagnoses at a patient's next medical visit and quantify the extent to which prediction performance is affected by missing data under MCAR, MAR, and MNAR by the amount of missing data and by health-seeking behaviors that may lead to health disparities.

2. Materials and methods

An overview of our methodology is presented in Fig. 2.

2.1. Data

We choose to simulate missing data in the MIMIC-III (Medical Information Mart for Intensive Care III) dataset [14]. The MIMIC-III dataset consists of de-identified health-related data of patients who stayed in critical care units at the Beth Israel Deaconess Medical Center between 2001 and 2012. It is important to note that the MIMIC-III data is from the intensive care unit (ICU) and thus may be different from EHR from outpatient and primary care providers. Once a patient has been admitted to the ICU, access to care likely would have less of an impact on the quantity and quality of data collected. However, in simulating scenarios where patients have less access to healthcare, a dataset less likely to already reflect issues in access to care makes more sense to work with. This dataset is a good Level 1

candidate for simulating missing data because the numbers of medical events are evenly distributed across most groups of interest (see Table 2).

2.2. Simulating realistic missing data within patient groups

We extract the sequence of structured and codified events in the MIMIC-III database for each patient. Structured and codified events represented include prescriptions (p_(drug names and NDC codes)), lab tests (l_(identifiers associated with lab measurements; LOINC codes)), conditions (c_ICD-9 codes), symptoms (s_ICD-9 codes), and diagnoses (d_ICD-9 codes). For example, d_250 represents diagnosis (d) of diabetes mellitus with an ICD-9 code of 250. Because we are evaluating a prediction tool to assist physicians with future diagnoses, we define the patient's medical history in two ways: for patients with the diagnosis of interest in their records, we modeled the events in the medical visits prior to the current visit that contains the diagnosis of interest. For patients that do not have the diagnosis of interest, events in the visits prior to their most recent visit are used (see Fig. 3). In this data structure, oftentimes structured and codified events are coded multiple times in the same visit.

2.3. Simulating missing data under MCAR, MAR, MNAR

When data are Missing Completely at Random (MCAR), missing values are independent of both observed and unobserved data, and occur at random. In mathematical notation, we write this as:

$$P(M | Y_o, Y_m) = P(M)$$

where M refers to the missing indicator, Y_o refers to the observed data, and Y_m refers to unobserved data. To induce realistic MCAR missingness, we mimic an EHR situation in which a lab tech forgets to input data points. This means that even if a structured event is coded multiple times in the patient record, one instance of that event being removed does not affect the other instances. In our simulations, the number of patients that have clerical errors in their medical records (X_i) follows a binomial distribution with probability 0.75 ($X_i \sim \text{Bin}(N, 0.75)$). For patients selected for clerical errors, we vary the proportion of events removed from each medical record from 0 to 0.75.

When data are Missing at Random (MAR), the events that lead to missingness are dependent on the observed data. In mathematical notation, we write this as:

$$P(M | Y_o, Y_m) = P(M | Y_o)$$

In our experiments, we explore variables such as sex, age, race, and insurance status (to infer socioeconomic status). We use a logistic regression model to simulate these conditional probabilities. In mathematical notation, this is:

$$P(M = 1 | Y_o) = \frac{1}{1 + \exp(-Y_o * \beta)}$$

where Y_o now refers to our observed predictors (sex, age, insurance status, race) and β represents the coefficients for the predictors.

We define younger patients (age 17–54), patients of color (Asian, Hispanic, Black, multi-racial, Middle Eastern, Hawaiian), male patients, and patients on government insurance as “data-lacking” groups as we would expect there to be more missing data in medical records from primary care providers and other types of EHR due to either less access to care or seeking less care. We define older patients (age 55–70 +), female patients, patients on private insurance, and White patients as “data-rich” groups due to such groups either having more access or seeking more care. The racial categories of American Indian and Alaska Natives were not included in the dataset. Note that in the experiments conducted in Section 3.1, we include all non-white patients. In the 2nd set of experiments in Section 3.2, we examine only Black vs White comparisons. The reason for this is because the biggest health disparities exist between these groups. For age, The cut-off point was chosen to ensure that the mean and median number of records would be evenly distributed among groups.

Since it has been shown that male patients are less likely to visit a healthcare provider than females, we select a β coefficient that reflect this difference in probabilities for sex. For insurance status, we separate patients on government insurance (Medicare, Medicaid) from patients with private insurance to represent lower and higher socioeconomic status. For race, we stratify Black and White patients. The β_{sex} coefficient is chosen such that $P(M|Male) = 0.88$ for male patients, and $P(M|Female) = 0.35$. We use the same β coefficient to determine probabilities for insurance status (with patients on government insurance having the higher probability for missing data), race (with non-white having the higher probability for missing data), and age (< 55 , ≥ 55). Our model creates these differences in probabilities for these groups. Since there is no prior data on what these differences should be, we chose values that would demonstrate the potential impact on the data and predictive validity.

When neither MCAR nor MAR hold, the data are Missing Not at Random (MNAR). This means that after accounting for available observed information, the missingness still depends on unobserved data. In mathematical notation, this is:

$$P(M | Y_o, Y_m) = P(M | Y_o)$$

A realistic MNAR scenario in EHR data is a patient’s unobserved underlying condition preventing them from going to health-related appointments (such as chronic pain, depression). For each diagnosis that we try to predict in the next medical visit, we select three other conditions / diagnoses to be “underlying conditions“. To select these “underlying conditions”, we first isolate medical events that have a cosine similarity > 0.7 , and manually investigate whether or not they could realistically be responsible for missing appointments. For example, if we are predicting a future diagnosis of Diabetes, we might find that ‘d_357’ (Inflammatory and toxic neuropathy) has a cosine similarity of 0.843 and ‘c_V5867’ (Long-term insulin use) has a cosine similarity of 0.854. It is more likely that toxic nerve damage would prevent an individual from attending a medical visit due to the fact that it involves pain as opposed to long-term insulin use.

Table 3 displays diagnoses that we will predict in our numerical experiments, their accompanying “underlying conditions”, and corresponding cosine similarity.

We use the same logistic regression model used in the MAR experiments with an added covariate for the “unobserved” underlying condition to determine probabilities for a patient to have missing data, or $P(M | Y_o, D_j)$ where D_j is an indicator for having one of the specified underlying conditions. D_1 corresponds to having at least one underlying condition in the medical history, and D_0 corresponds to having no underlying conditions in the medical history. $P(M | D_1) = 0.88$ and $P(M | D_0) = 0.35$ in our model.

2.4. Applying a knowledge graph to induce realistic missingness

To account for dependencies between EHR events, we employ the use of a knowledge graph based on the noisy OR-gate model [19]. Noisy OR is a conditional probability distribution that describes causal mechanisms by which children nodes are affected by parent nodes. Parent nodes can be defined as diseases; children nodes can be defined as symptoms. In a deterministic noise-free setting, the presence of a disease would always cause symptoms to be observed, but in real life the process is less deterministic. The model deals with inherent noise in the process by introducing failure and leak probabilities. For example, the presence of a disease y_j might fail to turn on its child symptom x_i with probability f_{ij} . The leak probability l_j denotes the probability that a symptom is on even if all parent diseases are off.

The probability of a child being present is defined as

$$P(c_i = 1 | y_1, \dots, y_n) = 1 - (1 - l_i) \prod f_{ij}^{y_j}.$$

Then the importance measure is defined as

$$IMPT_{noisy-or} = 1 - f_{ij}$$

due to the fact that higher importance indicates that a disease is more likely to turn on or generate an observed corresponding symptom. Model parameters are learned using maximum likelihood estimation. The importance measure is derived from the conditional probability distributions, and thus no assumptions about the prior distribution of parent nodes are made. This model was shown in numerical experiments to outperform naïve Bayes and logistic regression models. Unlike other models, such as naïve Bayes and logistic regression, it does not make any assumptions about the prior distribution of outcomes. This is important since patients present with multiple types of medical events in their records.

In our experiments, we have five different event categories: diseases, abnormal lab tests, prescriptions, symptoms, and conditions. We determined that conditions would not be included in the knowledge graph due to the variety of non-clinical events that can be classified as a ‘condition’ in the MIMIC-III data set (for example, organ donor status). Thus, we learn four different noisy OR-gate models with each event category considered to be the parent node, and all other event categories considered to be the children nodes (see Fig. 4). We verified the knowledge graph through a few relationships and what we would expect

from their relatedness. For example, Diabetes and Insulin had an importance measure of 0.65, while Diabetes and ECG test had an importance measure of 0.10.

Because the MCAR experiments are based on clinical errors and do not capture any clinically meaningful situations, we do not use the knowledge graph for the MCAR event removal process. Instead, in the patients that are selected to have missing data, we simply remove events at random up to a certain proportion of data denoted by p , where $p \in [0.0, 0.2, 0.4, 0.6, 0.75]$ of events in a medical record if a patient is selected for missingness. This allows us to see the impact of missing data at various proportions. See Fig. 5 for the MCAR event removal process.

For the MAR experiments, a patient might miss a medical visit or have less diagnose/medication data due to having less access to or seeking less care. We first remove entire medical visits from the patient's record up to a desired proportion. In each of the remaining visits, we remove individual medical events up to a desired proportion. The total amount of data removed from visits and from medical events in the remaining visits should be equal to p as defined above. In this type of situation, if a patient misses a lab test, the subsequent diagnosis and prescription that would have resulted from that abnormal lab test should also be removed. We utilize the knowledge graph to account for these dependencies.

Rather than determining a threshold on the importance measure to learn an edge for a knowledge graph, we take a more probabilistic approach for event-removal. The process for removing events is as follows: If event A is chosen for removal, we use the noisy OR-gate model that treats event A's category as the parent node (for example, event A is a diabetes diagnosis, and thus we choose the model that treats diagnoses as the parent nodes, and treats prescriptions, lab tests, and symptoms as the children nodes). From there, each event in the patient's record has a failure probability associated with it. The importance measure described above is calculated for each event (the likelihood that the diabetes diagnosis will "turn on" the event). We treat each corresponding importance measure as the probability that the corresponding event will also be removed with event A. Of the events removed using this method, the event with the highest importance measure is chosen as the new parent node and its corresponding importance measures for the remaining medical events are calculated and the process repeats. If no additional events are removed, a random medical event is chosen and the process repeats. This occurs until reaching a certain threshold of events removed. See Fig. 6 for the MAR event removal process using the medical knowledge graph.

We also use the knowledge graph for removing events in the MNAR experiments.

2.5. Modeling disease prediction based upon medical histories

In our numerical experiments, we focus on assessing the impact of missing data in EHR on prediction accuracy. To build disease prediction models using EHR data, we applied a two-step approach. In the first step, we generate an embedding for the medical history of each patient by leveraging word embedding algorithms [4,6]. Word embedding algorithms can be extremely useful in the analysis of Level 1 EHR data. These algorithms convert words to vectors of real numbers such that the cosine distance between two vectors reflects the similarity of two words. In using the structured and codified data from Level 1 EHR data

for word embedding, one can think of a “word” as a medical event in the patient record, and a “sentence” as the chronological sequence of all events in a patient record. We use the algorithm Word2Vec [17] to generate embeddings for medical events. Word2Vec has two training mechanisms: continuous bag of words (CBOW) and skip-gram. CBOW utilizes the surrounding context (other words around the target word) to predict a target word, and skip-gram uses a single word to predict a target context. The hidden layer following training contains the vectors that correspond to each word (Fig. 7).

2.6. Developing and evaluating disease prediction models

In the second step, we build three models for predicting a diagnosis of interest in a future medical visit, including diabetes (prevalence = 0.318), chronic kidney disease (prevalence = 0.203), heart failure (prevalence = 0.371), hypertensive chronic disorder (prevalence = 0.180), and lipid metabolism disorder (prevalence = 0.275). The ICD-9 codes are coarsened to the disease category.

Once we generate embeddings for individual events in a patient’s medical history, we derive an embedding for the patient’s entire medical history by multiplying each individual embedding by a temporal factor and summing up the vectors in the patient record. The equation that represents this process is as follows:

$$\frac{\sum V_i \times e^{\lambda \times (i - n + 1)/n}}{\sum e^{\lambda \times (i - n + 1)/n}}$$

where V_i is the embedding of event i in a patient’s medical record, λ is the time decay factor, and n is the total number of events in the record. For our experiments, a decay factor of 5 is used. Fig. 8 is useful for visualizing this process.

We develop and evaluate prediction models using three approaches: patient diagnosis projection similarity (PDPS), lasso (least absolute shrinkage and selection operator) regression, and artificial neural network.

Patient Diagnosis Projection Similarity [4] involves projecting patient sequences into the vector space while accounting for the temporal impact of events, as described previously. The cosine similarity between the patient vector and the diagnosis of interest is used to predict whether or not a patient is at risk for developing that disease. One of the main benefits of the method is its ability to predict risks for multiple diseases simultaneously. We apply PDPS to evaluate the effects of missingness on results (see Fig. 9).

As an alternative approach, we also leverage the patient vectors to build disease-specific prediction models by using lasso (least absolute shrinkage and selection operator) regression and deep learning based on a three-layer artificial neural network with mean-squared error loss. Each coordinate in a patient vector can be viewed as a feature used for the prediction of a medical outcome. Fig. 10 displays an example of this data structure: each row in the matrix corresponds to a patient vector, and each column corresponds to a coordinate feature. The last column indicates whether or not the patient is labeled as positive or negative for having a diagnosis of interest in a future medical visit.

We evaluate model performance by area under the operating curve (AUC) at various missingness proportions for the patients selected to have missing data by their given model. There are two scenarios for this. In the first scenario, we only remove medical events from the training data and leave the testing data complete. This is a useful approach because it provides a consistent benchmark to compare results at different missingness proportions; however, it does not reflect a realistic situation in which the training and testing data come from similar populations. In the second scenario, we remove medical events from patients in both the training and testing data. Evaluating with different testing sets means that the subsequent results are not directly comparable; however, the changes in prediction performance would reflect more of a real-life situation where models are trained and evaluated with missingness incorporated.

In each experiment, we report the results that are averaged over 200 Monte Carlo datasets. We present our experiments evaluated with the neural network in the main results (Figs. 10–17). Since the results for lasso regression were nearly identical, we do not include them. PDPS performed very poorly (AUC averaging at 0.5 for all diagnoses), possibly due to the fact that in this work we remove visits that contain the target diagnosis in the patient history. For this reason, we do not present these results.

3. Results

We outline the key findings of disease prediction accounting for three types of missingness: MCAR, MAR, and MNAR among data-lacking and data-rich patient groups.

3.1. Assessing data missingness across disease types

In the first set of results, we look at five disease prediction models (diabetes, chronic kidney disorder, heart disease, hypertensive chronic disorder, lipid metabolism disorder). The models are evaluated at increasing proportions of missingness in patients selected to have missing data in their medical records. We look at two scenarios: medical events removed only from the training set, and medical events removed from both the training and testing data. The former has the advantages of being able to evaluate on the same test set each time; however, the latter reflects a more realistic scenario in terms of evaluating the impact of missing data on underserved populations.

In the MCAR experiments, we randomly choose 78 percent of patients to have missing data to determine the extent of the impact of missing data on prediction results. If a patient is chosen to have missing data, we remove events at random from their medical history. The model performances when missingness is only induced in the training data does not diminish very much. When missingness is simulated in both the training and testing data, we observe a slight gradual decline in AUC for predicting all diagnoses at increasing levels of missingness (Fig. 11).

In the MAR experiments, we use a logistic regression model with four covariates (age, sex, insurance status, race) to determine whether or not a patient is selected to have missingness simulated in their medical record. ~ 78 percent of patients overall were selected for missingness based on this model. Missingness is induced by removing visits and unique

events at random for selected patients. If a patient is selected to have a visit removed, and they only have one visit in their history, they are removed from the analysis. However, their data is factored into the proportion of missing data. We first assess the impacts of missingness on all patients in our data. We see a stronger decline in performance as the proportion of missing data increases when we remove events from both the training data and test data as opposed to just the training data (see Fig. 12). We also see a slightly stronger decline in performance as we increase the proportion of missing data for MAR experiments compared to MCAR.

In the MNAR experiments, we use the same model as MAR, but include an additional covariate that accounts for an “unobserved underlying condition”. We choose a β coefficient such that $P(M) = 0.88$ for patients with a selected underlying condition, and $P(M) = 0.35$ for patients without. If a patient is selected for having missing data and the conditions are present in the record, the conditions are also subsequently removed after missing data is simulated. The overall effect looks similar to the MAR experiment, with declines in model performance as proportion missing in those selected for missingness increases (see Fig. 13) The declines are stronger than in the MCAR experiments.

For the complete test set experiments, we see that model performance remains more constant across different levels of missingness if missingness is not simulated in the test set, compared to when it is. A possible reason for this could be the fact that multiple similar medical events are used to predict a future diagnosis, and these events would always be present in the test set. For the incomplete test set experiments, we see stronger decreases in model performance for each missingness mechanism.

In the incomplete test data MCAR experiments, we do not observe as notable of a decrease as we do for MAR and MNAR. This is expected because events can be coded multiple times in an EHR record, so to remove an event at random does not necessarily mean it does not exist elsewhere in a patient’s record. Dependencies are also not removed and thus likely help improve prediction performance since similar events probably have similar vector representations. In the MNAR models, the model performances across various levels of missingness look very similar to the MAR model performances. Thus, it appears that the addition of an underlying condition covariate does not contribute to further declines in disease prediction performance. This is surprising due to the fact that the underlying conditions were chosen such that they would be associated with the diagnosis of interest that we were trying to predict.

3.2. Assessing the impact of data missingness for data-lacking groups

Because we are interested in assessing the impact of missingness on populations with less access to healthcare or behaviors leading to less engagement with the healthcare system (data-lacking groups), in the next set of numerical experiments for MAR, we stratify by demographic group by splitting the testing data to evaluate each group separately. All patients begin with an average of 15 percent of their data missing. Patients that belong to a particular data-lacking group will have this average increase in increments. Patients that belong to the data-rich group will have this average stay at 15 percent. Thus, it is important to note that the differences shown in Fig. 18 are a result of only increasing the amount of

missing data in the data-lacking group. We do three different versions of this to look at the effects of insurance status, age, and race independently. In these experiments, we only present experiments where the events are removed from both the training and the test set (see Fig. 18).

First, we stratify our test set by insurance status. In the logistic regression model, we chose a β coefficient such that $P(M) = 0.35$ for patients on private / self insurance and $P(M) = 0.88$ for patients on government insurance to reflect less access to healthcare. We see initial disparities for patients on government insurance in predicting all diagnoses except for chronic kidney disease, noting that this could be due to the fact that end stage renal disease uniquely qualifies for government insurance. As proportion of data missing increases for patients on government insurance, we see a stronger decline in model performance at the various increments. We also stratify our test data by younger patients (< 55) and older patients (≥ 55). We see that initially the AUC is higher for predicting all diagnoses for younger patients as opposed to older patients. As missingness in younger patients increases, we see declines in model performance, but it only leads to worse performance in predicting lipid metabolism disorder at the highest proportion of missing data. Lastly, we stratify by White patients vs Black patients to see how race plays a factor in impact of missing data. We see strong initial disparities in predicting hypertensive chronic disorder for Black patients. We see very strong declines in model performance for predicting Diabetes as the proportion of missing data increases. We do not see these declines when predicting heart disease and lipid metabolism disorder.

When assessing the impact of missing data on data-lacking groups, we observe that for most diagnoses, model performance for such groups are negatively impacted by missing data. Notably, some model performances for certain diagnoses are more affected by missing data than others. Models for diabetes, chronic kidney disease, and hypertensive chronic disorder are all more affected by missingness than models for heart disease. One possible explanation for this phenomenon could be that there are fewer types of events at smaller frequencies that are predictive of diagnoses like diabetes, chronic kidney disorder, and hypertensive chronic disorder. On the other hand, there could be many events at high frequencies that are predictive of heart disease, thus causing the model to not be as affected by missing data. We also see that certain data-lacking groups are more affected by missing data than others (for example, the decline in performance for younger patients as missingness increases is not as strong as the decline in performance for Black patients for certain diagnoses). Due to differences in quality of care, it is possible that Black patients have less predictive events for these diagnoses as well.

3.3. Comparing event removal with and without a knowledge graph

Next, we directly compare scenarios where all patients are selected for missingness and either have events removed at random (like in the MCAR experiments) or events removed using the knowledge graph (removing clusters of related events together). In these experiments, all patients have missing data. We vary the proportion of the amount of missing data in each patients' record to assess the impact as we did before. In these experiments, entire visits are not removed as they are in the MAR experiments. We simply use the

knowledge graph event-removal methodology from section 2.3 to remove random clusters of dependent events from each medical record.

We observe that in both scenarios (removing events from only the training data and removing events from all data), there are notably stronger declines in model performance as the proportion of data missing increases if events are removed using the medical knowledge graph, compared to without using the medical knowledge graph. When events are randomly removed without using the medical knowledge graph, it is likely that among a set of clinically related events (say, diabetes diagnosis and insulin prescription), some of them (say, diabetes diagnosis) are removed but some (say, insulin prescription) are not. Then, the information from the remaining events (say, insulin prescription) can still be used in downstream learning tasks such as prediction.

4. Discussion

In this work, we propose a novel approach for simulating realistic missing data in Level 1 EHR and evaluate the impact of such missing data on downstream leaning tasks through the use of prediction models that use the entire patient history. This general approach can be adapted and refined further. For example, we used the noisy OR-gate model to obtain a knowledge graph for medical events in our numerical experiments. Alternatively, we can use knowledge graphs from other existing databases for medical events to capture dependencies and pathways for simulating missing data or other research purposes. In our work, we observe diminished model performance as a result of all types of missingness, and this diminished performance is more subtle in some scenarios (MCAR) as opposed to others (MAR, MNAR). We also see that disease prediction performance for data-lacking groups is more significantly impacted by missing data. Lastly, we find that for the same amount of missing data, accounting for relatedness of medical events and concepts in EHR leads to greater impact on model performance. Overall, however, we do observe that to detect meaningful differences in model performance, missingness needs to be quite high (70–75 percent missing). Thus, it is positive to see that these models are fairly robust to missingness.

Our results indicate that disproportionate missing data in patients in certain demographic groups does impact disease model performance in a negative way. What this means is that the same disease prediction model may not be as effective for patients that have less data in their medical records. This has real implications for patients in under-served populations, particularly patients of lower socioeconomic status and patients of color – as advances in AI are incorporated into the clinic, these groups could get left behind from progress in early disease detection and risk prediction that these technologies afford. The divide in health outcomes with regard to attributes such as socioeconomic status and race is already wide. If we were to naively implement and integrate these models into the clinic, the gap could widen even more. Thus, it is of utmost importance to address missing data in EHR for under-served populations when applying powerful ML methods. As scientists and engineers, it is our responsibility to ensure that our contributions do not leave people behind.

Dependencies exist between EHR events in the real world. Diagnoses typically happen as a result of abnormal lab tests, and prescriptions typically occur to treat a diagnosis. We observe that removing events via knowledge graph results in more of an impact on prediction performance as opposed to random removal. Although we cannot know the exact causal pathways between these events, we can do a better job in accounting for these correlations when simulating missing data in Level 1 EHR. The difference in results when we simulate missing data via knowledge graph compared to random event removal displays the importance of generating realistic missing data. When our simulations more-so mimic situations that happen in the real world, we observe a stronger decline in model performance. This allows us to see more clearly the impact that missing data might have on underserved populations, whereas without this mechanism in place we might not have seen as much evidence of disparities.

One limitation of this work is the imprecise characterization of the actual causal pathways that occur in medical data. We remove clusters of related events as a proxy for these causal pathways, but in reality this may not accurately capture how missing data truly happens in practice. For example, a diagnosis might never occur without a certain abnormal lab test, and a prescription might never occur without a certain diagnosis. Thus, the causal pathway might be lab test \rightarrow diagnosis \rightarrow prescription. In all of our experiments, however, there is no direction in the pathways. We do not account for the fact that the probability of removing the lab test upstream in the causal pathway should yield a higher probability of removing the prescription downstream. Additionally, using the noisy OR-gate knowledge graph does not allow for the generation of failure probabilities relating events of the same type (such as two diagnosis events). Dependencies between events of the same type are still accounted for through their relationships with other events, but the method would be better if this were more explicit. Another limitation that exists in the work is despite the fact that groups may have similar events in the dataset, the data quality might differ as well (there may be less diversity of medical codes for data-lacking groups even after accounting for differences in healthcare utilization). This is something that we did not account for in this work.

For future work, we might consider enlisting the help of clinicians that can identify general causal pathways that exist between medical events to further refine the knowledge graph and remove events in a more realistic way. We also might consider finding a way to map our EHR events to existing databases that detail such relationships, such as the Google Health Knowledge Graph or the PennTURBO Knowledge Graph [5]. It is also of interest to account for missingness by developing novel imputation methods that can handle Level 1 EHR and incorporate relatedness of medical concepts and events. There is a large body of literature on imputing missing data in EHR, but little work on imputation models that can incorporate relatedness of medical concepts and events. Additionally, we might also explore the relationship between illness severity, access, and healthcare utilization. We would expect to see that healthier patients with less access to healthcare have more missing data than sicker patients with less access. Finally, given the importance of intersectionality in healthcare analysis, there is a need for future work to focus on intersectional differences (for example, white patients on government insurance vs black patients) to learn more about the impact of disproportionate missing data in disease prediction models using EHR data.

Funding

This work is partly supported by NIH grant R01-GM124111. The content is the responsibility of the authors and does not necessarily represent the views of NIH.

References

- [1]. Beaulieu-Jones BK, Lavage DR, Snyder JW, Moore JH, Pendergrass SA, Bauer CR, Characterizing and Managing Missing Structured Data in Electronic HealthRecords: Data Analysis, *JMIR Med Inform* (2018), 10.2196/medinform.8960.
- [2]. Buuren SV, Flexible Imputation of Missing Data, CRC Press (2018), 10.1201/9780429492259.
- [3]. Deeks A, Lombard C, Micheltore J, Teede H, The effects of gender and age on health related behaviors, *BMC Public Health* 9 (2009), 10.1186/1471-2458-9-213.
- [4]. Farhan W, Wang Z, Huang Y, Wang S, Wang F, Jiang X, A predictive model for medical events based on contextual embedding of temporal sequences, *JMIR medical informatics* 4 (4) (2016) e39. [PubMed: 27888170]
- [5]. Freedman HG, Williams H, Miller MA, Birtwell D, Mowery DL, and Stoeckert CJ (2020). A novel tool for standardizing clinical data in a semantically rich model. *Journal of Biomedical Informatics* 112. Articles initially published in *Journal of Biomedical Informatics*: X 5–8, 2020, 100086. ISSN: 1532–0464. DOI: 10.1016/j.yjbinx.2020.100086. URL: <https://www.sciencedirect.com/science/article/pii/S2590177X20300214>.
- [6]. Getzen E, Ruan Y, Ungar L, and Long Q (2022). Mining for Health: A Comparison of Word Embedding Methods for Analysis of EHRs Data. medRxiv. DOI: 10.1101/2022.03.05.22271961.
- [7]. Ghassemi M, Naumann T, Schulam P, Beam AL, Chen IY, Ranganath R, A Review of Challenges and Opportunities in Machine Learning for Health, *AMIA Joint Summits on Translational Science* (2020), 10.48550/arXiv.1806.00388.
- [8]. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G, Potential biases in machine learning algorithms using electronic health record data, *JAMA Internal Medicine* (2018), 10.1001/jamainternmed.2018.3763.38.
- [9]. Goodwin T and Harabagiu SM (2013). “Automatic Generation of a Qualified Medical Knowledge Graph and Its Usage for Retrieving Patient Cohorts from Electronic Medical Records”. In: 2013 IEEE Seventh International Conference on Semantic Computing 363–370. DOI: 10.1109/ICSC.2013.68.
- [10]. Hall WJ, Chapman MV, Lee KM, Merino YM, Thomas TW, Payne BK, Eng E, Day SH, Coyne-Beasley T, Implicit Racial/Ethnic Bias Among Health Care Professionals and Its Influence on Health Care Outcomes: A Systematic Review, *American journal of public health* 105 (12) (2015) e60–e76, 10.2105/AJPH.2015.302903. URL: <https://pubmed.ncbi.nlm.nih.gov/26469668>.
- [11]. Heitjan DF, Basu S, Distinguishing “Missing at Random and “Missing Completely at Random”, *The American Statistician* 50 (1996) 207–213, 10.2307/2684656.
- [12]. Hoerbst A, Ammenwerth E, Electronic health records. A systematic review on quality requirements, *Methods Inf. Med* 49 (2010) 320–336, 10.3414/ME10-01-0038. [PubMed: 20603687]
- [13]. Hubbard RA, Huang J, Harton J, Oganisian A, Choi G, Utidjian EI, Bailey LC, Chen Y, A Bayesian latent class approach for EHR-based phenotyping, *Statistics in Medicine* 38 (2018) 74–87, 10.1002/sim.7953. [PubMed: 30252148]
- [14]. Johnson AEW, Pollard TJ, Shen L, Lehman LH, Feng M, Ghassemi M, Moody B, Szolovits P, Anthony Celi L, Mark RG, MIMIC-III, a freely accessible critical care database, *Scientific Data* 3 (2016), 160035, 10.1038/sdata.2016.35. [PubMed: 27219127]
- [15]. van Loenen T, van den Berg MJ, Faber MJ, Westert GP, Propensity to seek healthcare in different healthcare systems: analysis of patient data in 34 countries, *BMC Health Services Research* 15 (1) (2015) 465, 10.1186/s12913-015-1119-2. [PubMed: 26453459]
- [16]. Medicine, I of and Council, NR (2015). Investing in the Health and Well-Being of Young Adults. Washington, DC: National Academic Press (US). DOI: 10.39.

- [17]. Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J, Distributed Representations of Words and Phrases and their Compositionality, *Advances in Neural Information Processing Systems* 7 (2013) 3111–3119, 10.48550/arXiv.1310.4546.
- [18]. Rajkomar A, Hardt M, Howell M, Corrado G, Chin M, Ensuring Fairness in Machine Learning to Advance Health Equity, *Annals of Internal Medicine* (2018), 10.7326/M18-1990.
- [19]. Rotmensch M, Halpern Y, Tlimat A, Horng S, Sontag D, Learning a Health Knowledge Graph from Electronic Medical Records, *Scientific Reports* 7 (1) (2017) 5994, 10.1038/s41598-017-05778-z. [PubMed: 28729710]
- [20]. Rubin DB, Inference and Missing Data, *Biometrika* 3 (1976) 581–592, 10.2307/2335739.
- [21]. Santos A, Colaco AR, Nielsen AB, Niu L, Strauss M, Geyer PE, Coscia F, Albrechtsen NJW, Mundt F, Jensen LJ, and Mann M (2022). A knowledge graph to interpret clinical proteomics data. *Nature Biotechnology*. DOI: 10.1038/s41587-021-01145-6. URL: Doi: 10.1038/s41587-021-01145-6.
- [22]. Schafer Shafer JL (1997). *The Analysis of Incomplete Multivariate Data*. New York: Chapman and Hall / CRC. DOI: 10.1201/9780367803025.
- [23]. Shinozaki A, *Electronic Medical Records and Machine Learning Approaches to Drug Development, Artificial Intelligence in Oncology Drug Discovery and Development* (2019), 10.5772/intechopen.92613.
- [24]. Solares JRA, Raimondi FED, Zhu Y, Rahimian F, Canoy D, Tran J, Gomes ACP, Payberah AH, Zottoli M, Nazarzadeh M, Conrad N, Rahimi K, Salimi-Khorshidi G, Deep learning for electronic health records: A comparative review of multiple deep neural architectures, *Journal of Biomedical Informatics* 101 (2020), 103337, 10.1016/j.jbi.2019.103337.40. [PubMed: 31916973]
- [25]. Weber GM, Adams WG, Bernstam EV, Bickel JP, Fox KP, Marsolo K, Raghavan VA, Turchin A, Zhou X, Murphy SN, Mandl KD, Biases introduced by filtering electronic health records for patients with "complete data", *Journal of the American Medical Informatics Association* 24 (2018) 1134–1141, 10.1093/jamia/ocx071.
- [26]. Wells B, Chagin KM, Nowacki AS, and Kattan MW (2013). Strategies for handling missing data in electronic health record derived data. *eGEMs*. DOI: 10.13063/2327-9214.1035.41.

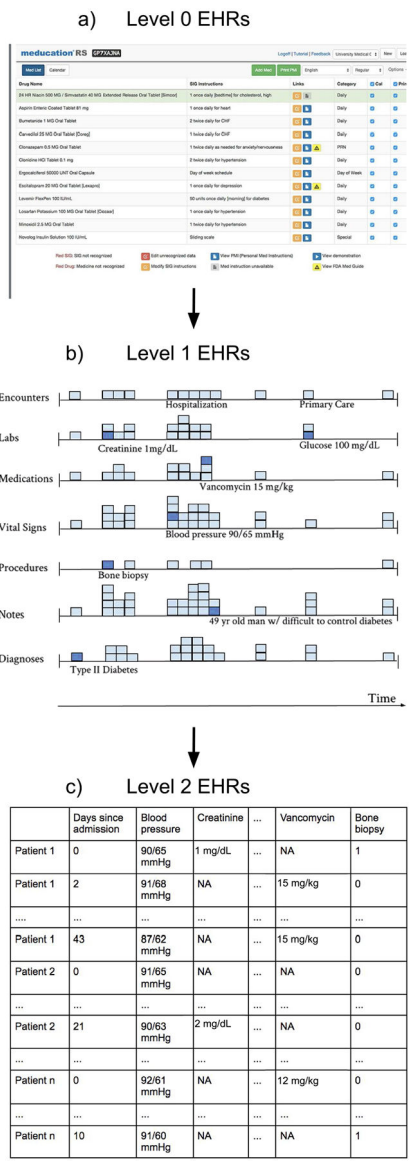


Fig. 1. Different levels of EHR data. a) displays Level 0 EHR data (i.e., data still in EHR system without any pre-processing). b) displays Level 1 EHR data after having been extracted from the EHR system and after minimal processing such as curation and harmonization. Each box in the figure represents a single medical event in the patient history. Note that Level 1 EHR data are in sequential form with structured, codified, and unstructured data and no explicit missing values. c) displays Level 2 EHR data that have been pre-processed into matrix form. Going from Level 1 EHR data to Level 2 EHR data requires significant effort spent pre-processing such as feature extraction/engineering through chart review or other means and may lead to substantial information loss. There are explicit missing values in Level 2 data due to its matrix form.

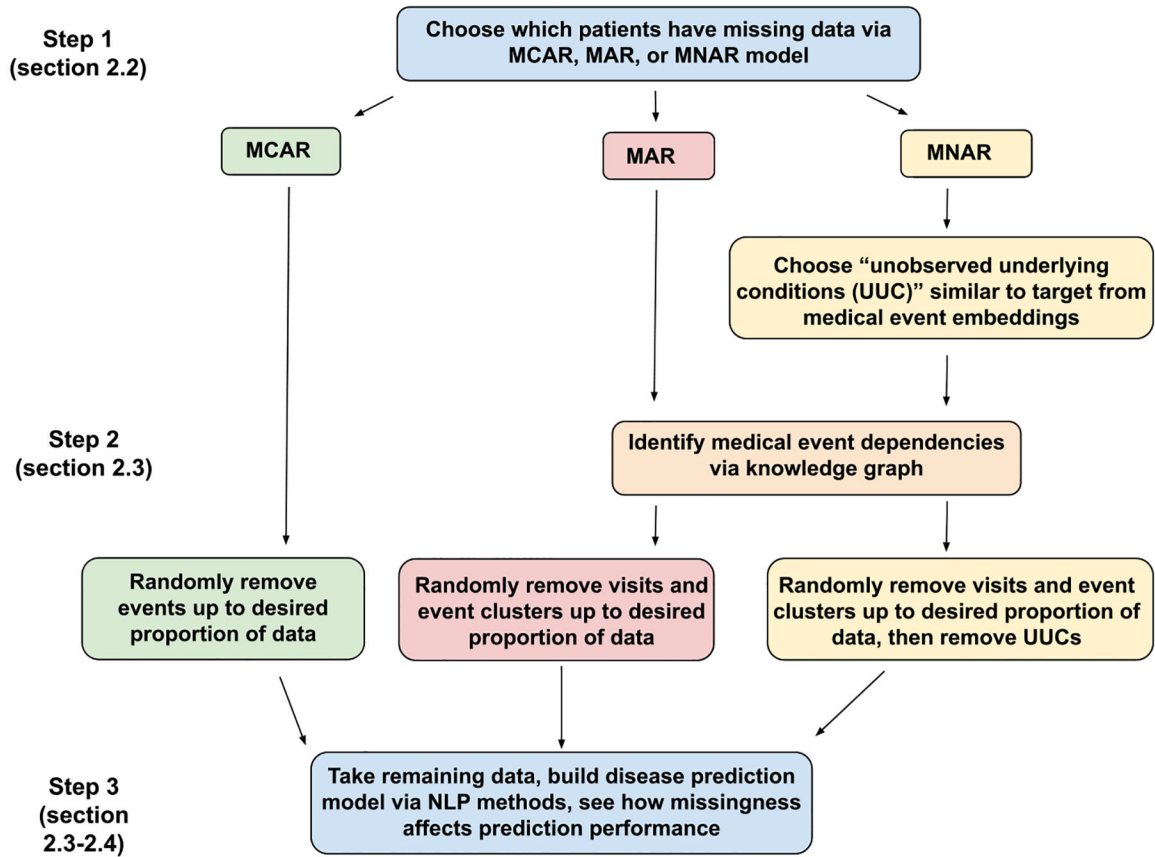


Fig. 2. Overview of the methods section. We begin by using either our MCAR model, MAR model, or MNAR model to choose which patients will have missing data simulated into their medical records. From there, we describe how data are removed under each missingness mechanism, and we take our adjusted input data and build disease prediction models using NLP methods.




Fig. 3. Sequence of structured and codified events used for prediction. Left panel refers to the events in the patient history, right panel refers to events in the next medical visit. Prefixes: c = condition, d = diagnosis, p = prescription, l = laboratory result.




Fig. 4. Noisy OR model example with diagnoses as the parent nodes and all other events as the children nodes.

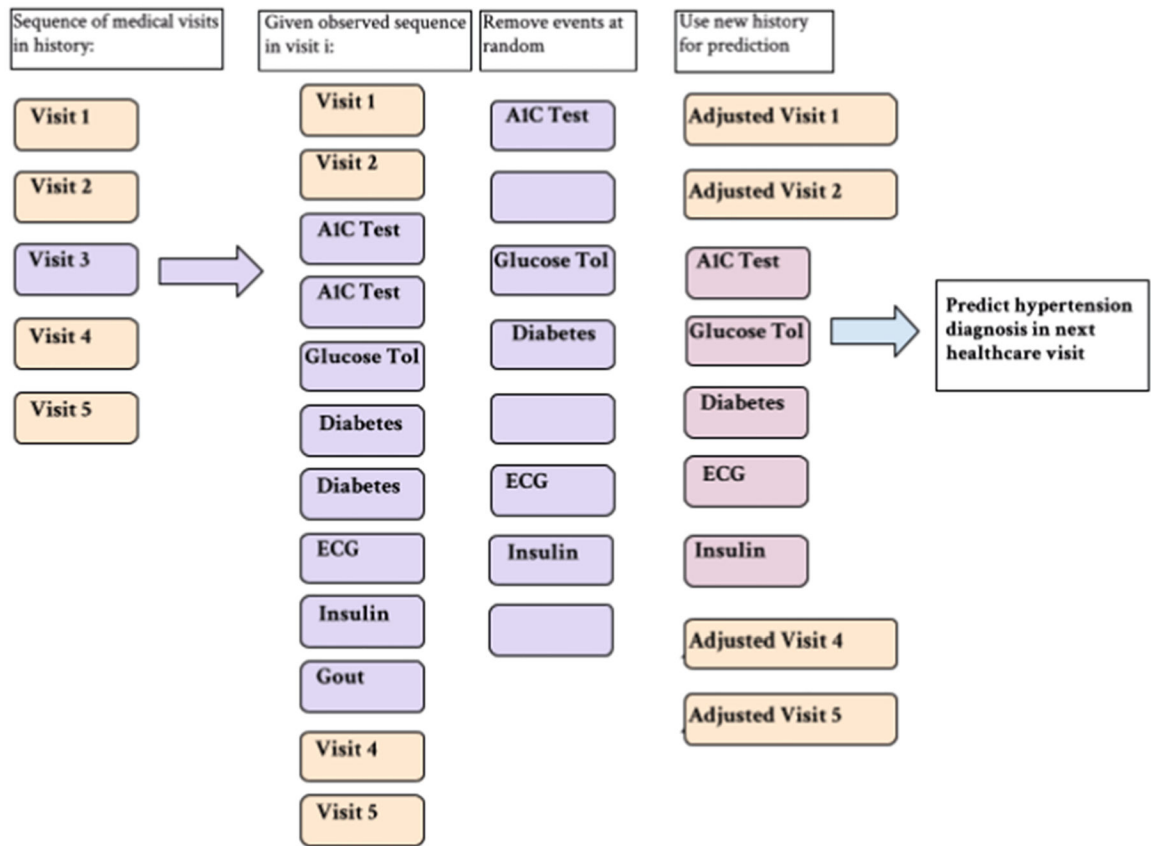


Fig. 5. Clerical error simulation example. Events are removed at random up to desired proportions p , regardless of whether or not they appear multiple times in the record.

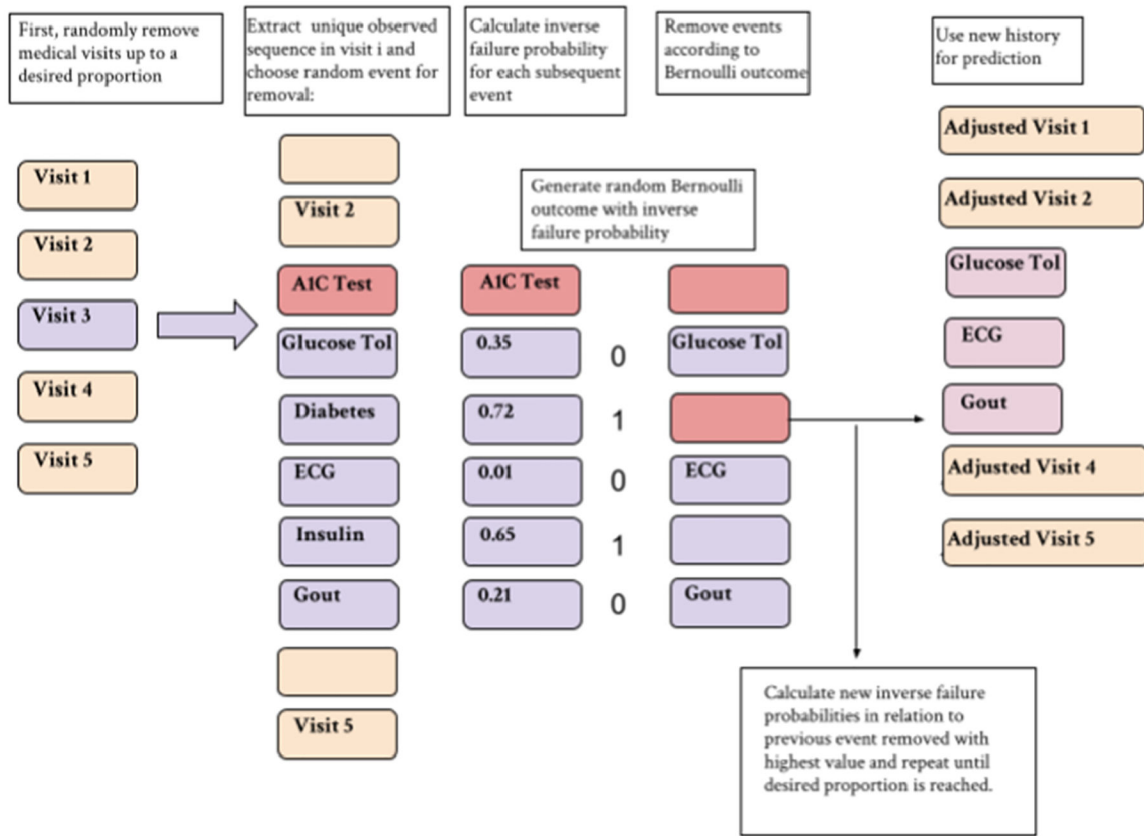


Fig. 6. Event removal process for MAR experiments that incorporate the use of the noisy OR-gate knowledge graph. Visits and clusters are removed such that the total data removed from the record reaches desired proportions p .




Fig. 7. Training mechanisms for FastText – CBOW (left) and Skip-Gram (right). CBOW trains by using a context to accurately predict a target word, Skip-Gram trains by using each word to predict its target context.

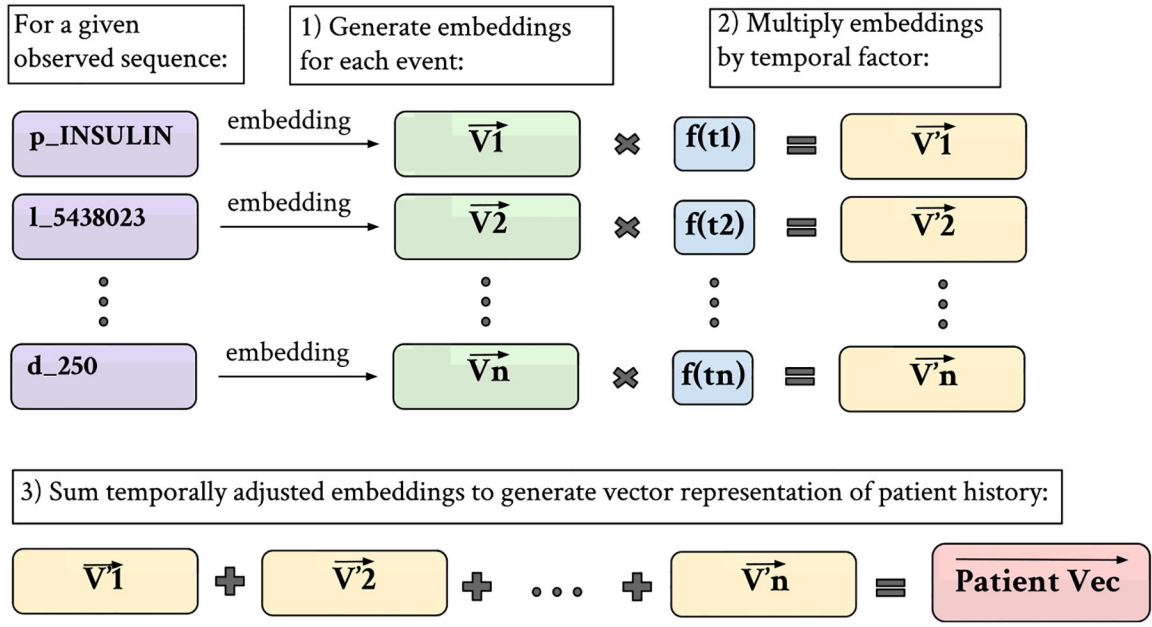


Fig. 8. Patient vector generation process– embeddings are generated for each medical event in a sequence. Each embedding is multiplied by a function of its ordering, such that more recent events have more importance, and the temporally adjusted embeddings are summed to create an overall patient vector.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

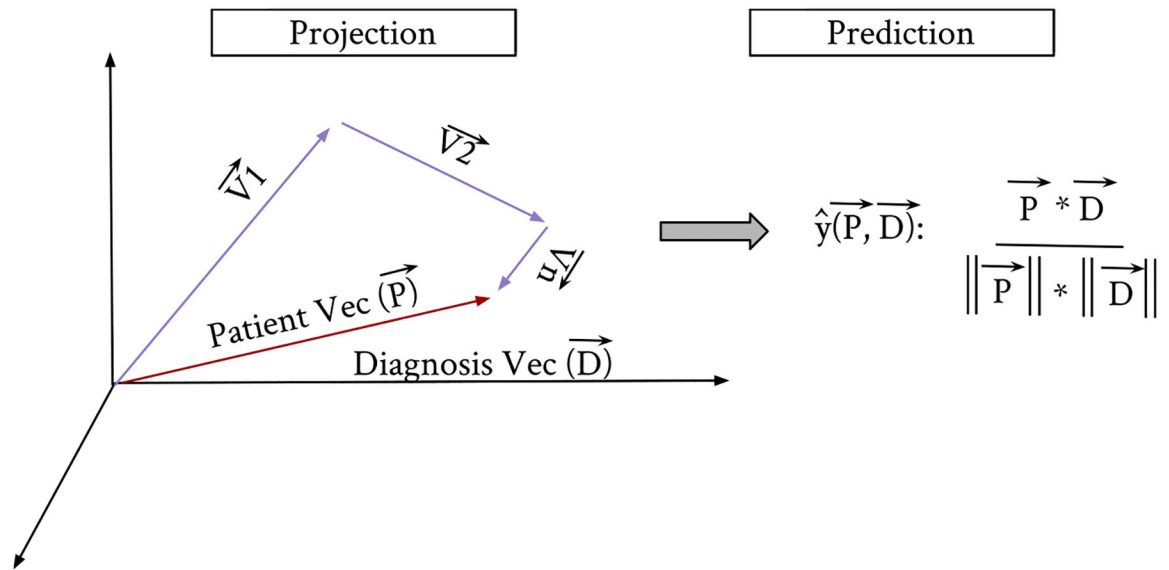


Fig. 9. Patient Diagnosis Projection Similarity by [4]. Patient vectors and target diagnosis vectors are projected into the vector space. Prediction is based on cosine similarity between vectors.

	8	9 ...	91	92	93	94	95	96	97	98	99	d_250
0.063576	-0.043763	...	0.205411	0.086000	-0.079307	-0.094209	0.122335	0.075008	-0.101978	0.171773	-0.086154	0
0.138456	0.015987	...	0.223654	-0.127352	-0.133022	0.098714	0.124876	-0.054426	-0.260688	0.102227	0.034120	0
0.056957	0.078515	...	0.232537	-0.058674	-0.010017	-0.078177	0.137958	0.077087	-0.066291	0.246547	0.089741	0
0.066291	-0.001621	...	0.300620	-0.015476	-0.019463	0.023604	0.064522	-0.072543	-0.104260	0.109070	-0.018547	0
0.138132	-0.023823	...	0.229600	-0.078178	-0.106999	-0.055058	0.061134	0.023506	-0.136760	0.219645	-0.132420	0
...
0.102177	-0.078634	...	0.258427	-0.113346	-0.064184	-0.051446	0.015254	0.051326	-0.135009	0.196235	-0.049943	0
0.095207	0.033359	...	0.407409	-0.093856	-0.198592	-0.122854	0.096119	-0.053041	0.015567	0.146110	-0.266336	0
0.108092	-0.132961	...	0.245729	-0.102928	-0.078311	-0.137644	0.038119	-0.006738	-0.021117	0.330455	-0.162577	0
0.120340	-0.001397	...	0.279110	-0.087301	-0.122910	-0.167694	0.035806	0.042163	-0.091745	0.215741	-0.138780	0
0.187881	0.023371	...	0.323166	-0.065256	-0.182132	0.034015	0.090142	0.077512	-0.127541	0.110238	-0.147341	0

Fig. 10. Sample patient vector data matrix. Each row in the matrix corresponds to a 100-dimensional patient vector. Each coordinate can be used as a feature for prediction of a future diagnosis (Diabetes in this example).

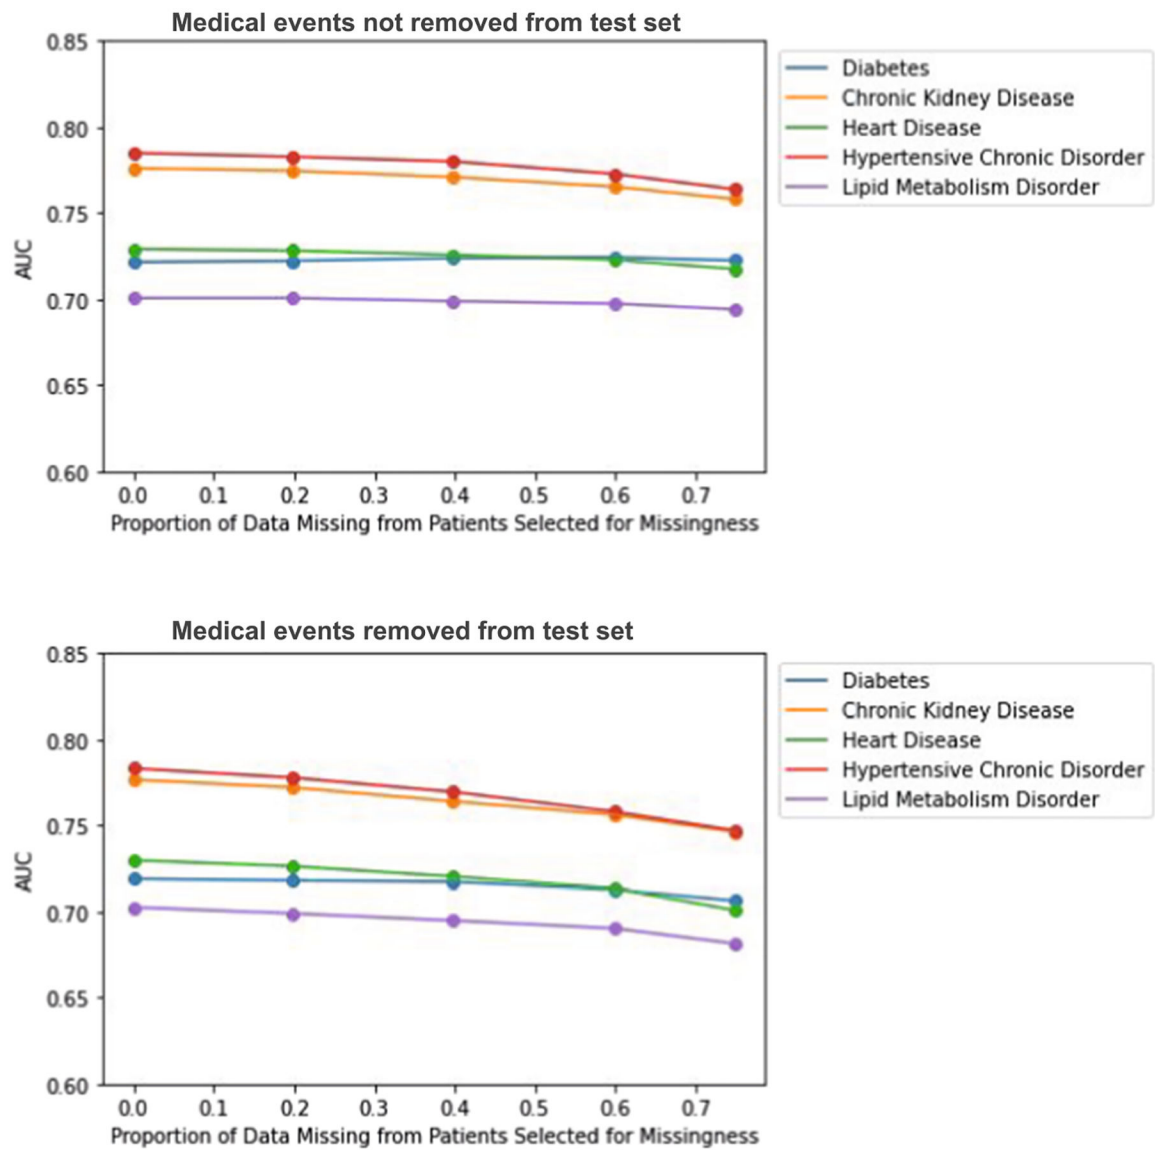


Fig. 11. Impact of MCAR missing data on disease prediction model performance as the proportion of missing data (p) varies. Top figure corresponds to complete test set experiments, bottom figure corresponds to incomplete test set experiments.

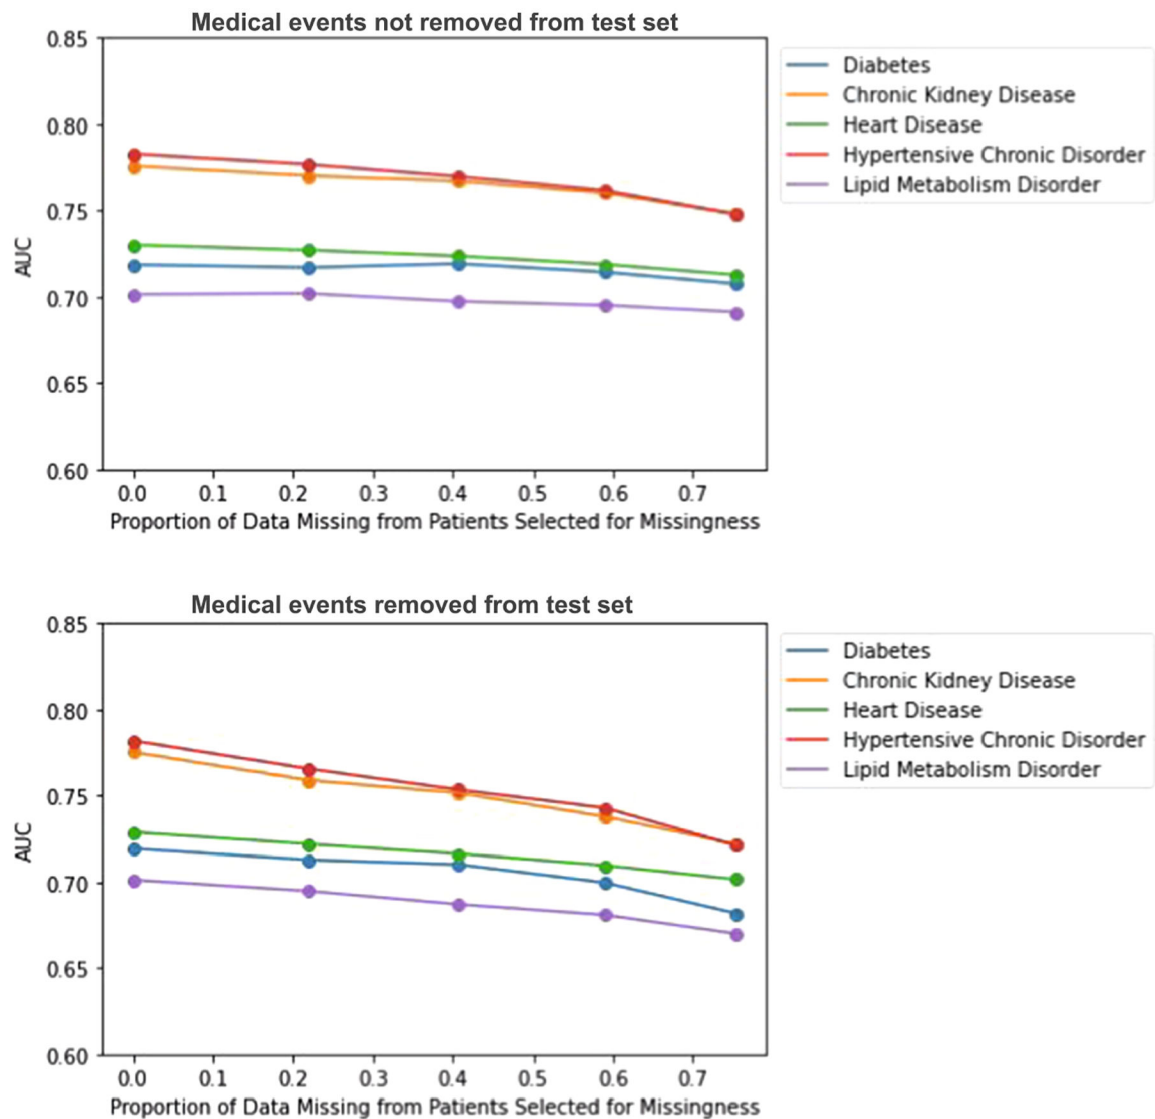


Fig. 12. Impact of MAR missing data on disease prediction model performance as the proportion of missing data (p) varies. Top figure corresponds to complete test set experiments, bottom figure corresponds to incomplete test set experiments.

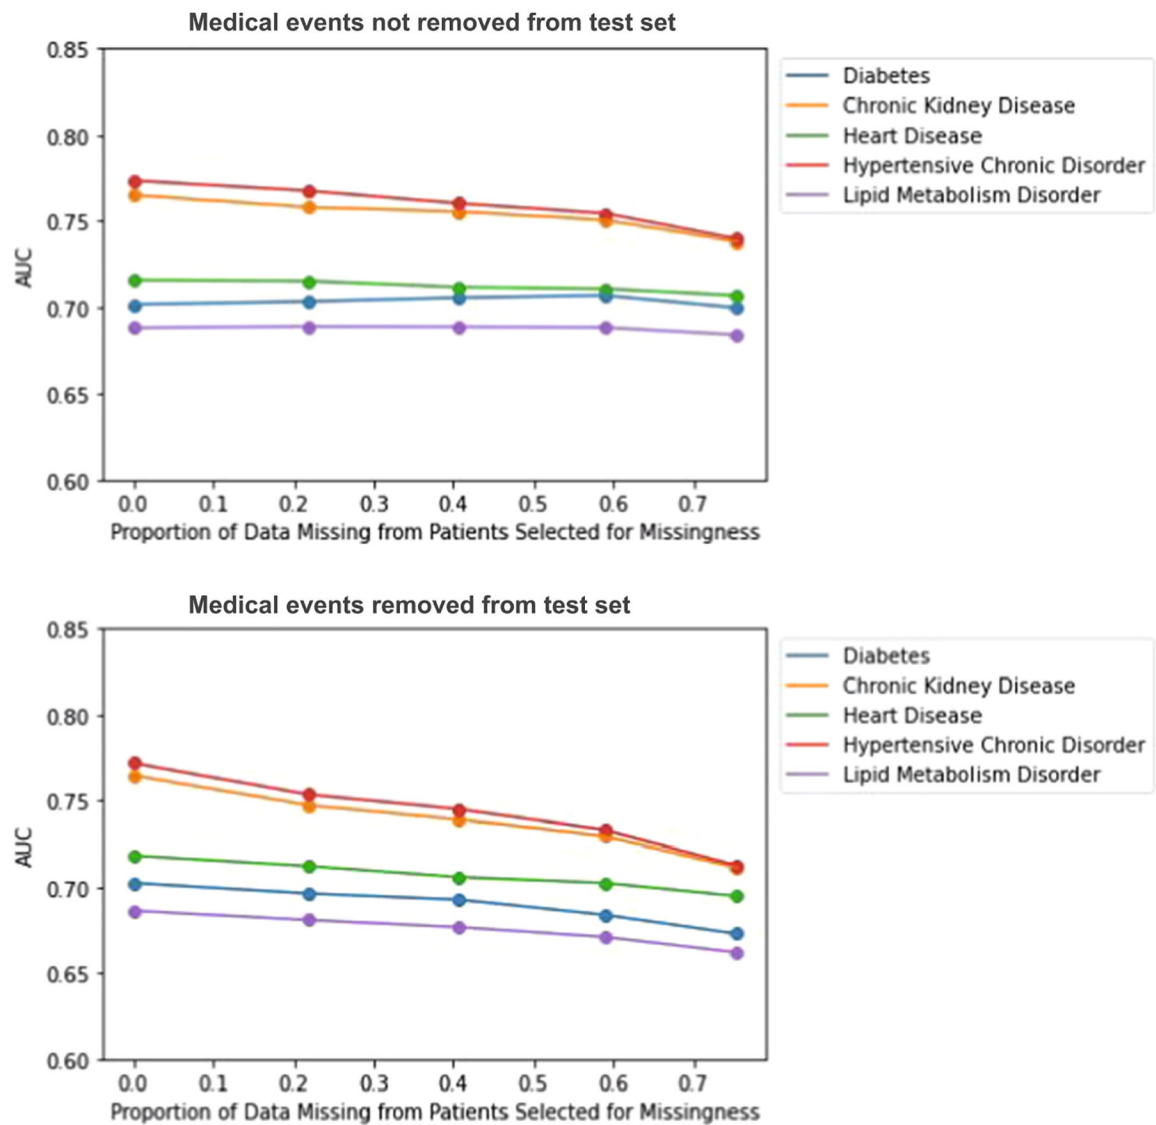


Fig. 13. Impact of MNAR missing data on disease prediction model performance as the proportion of missing data (p) varies. Top figure corresponds to complete test set experiments, bottom figure corresponds to incomplete test set experiments.

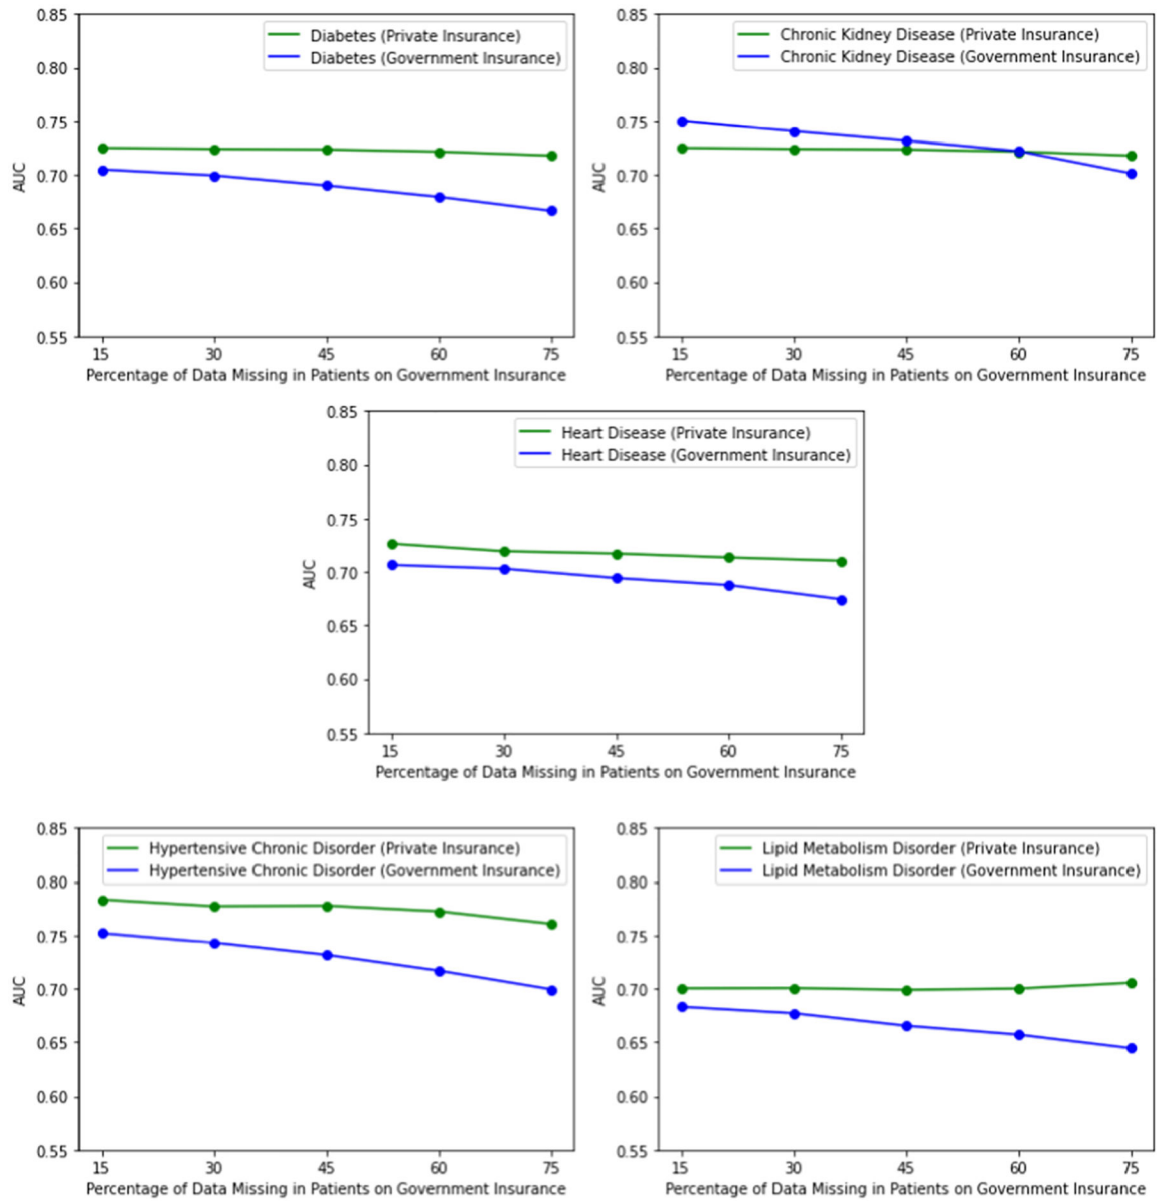


Fig. 14. Assessing the impact of missing data on data-lacking groups. Proportion missing for patients on private insurance is held constant while proportion missing for patients on government insurance is varied.

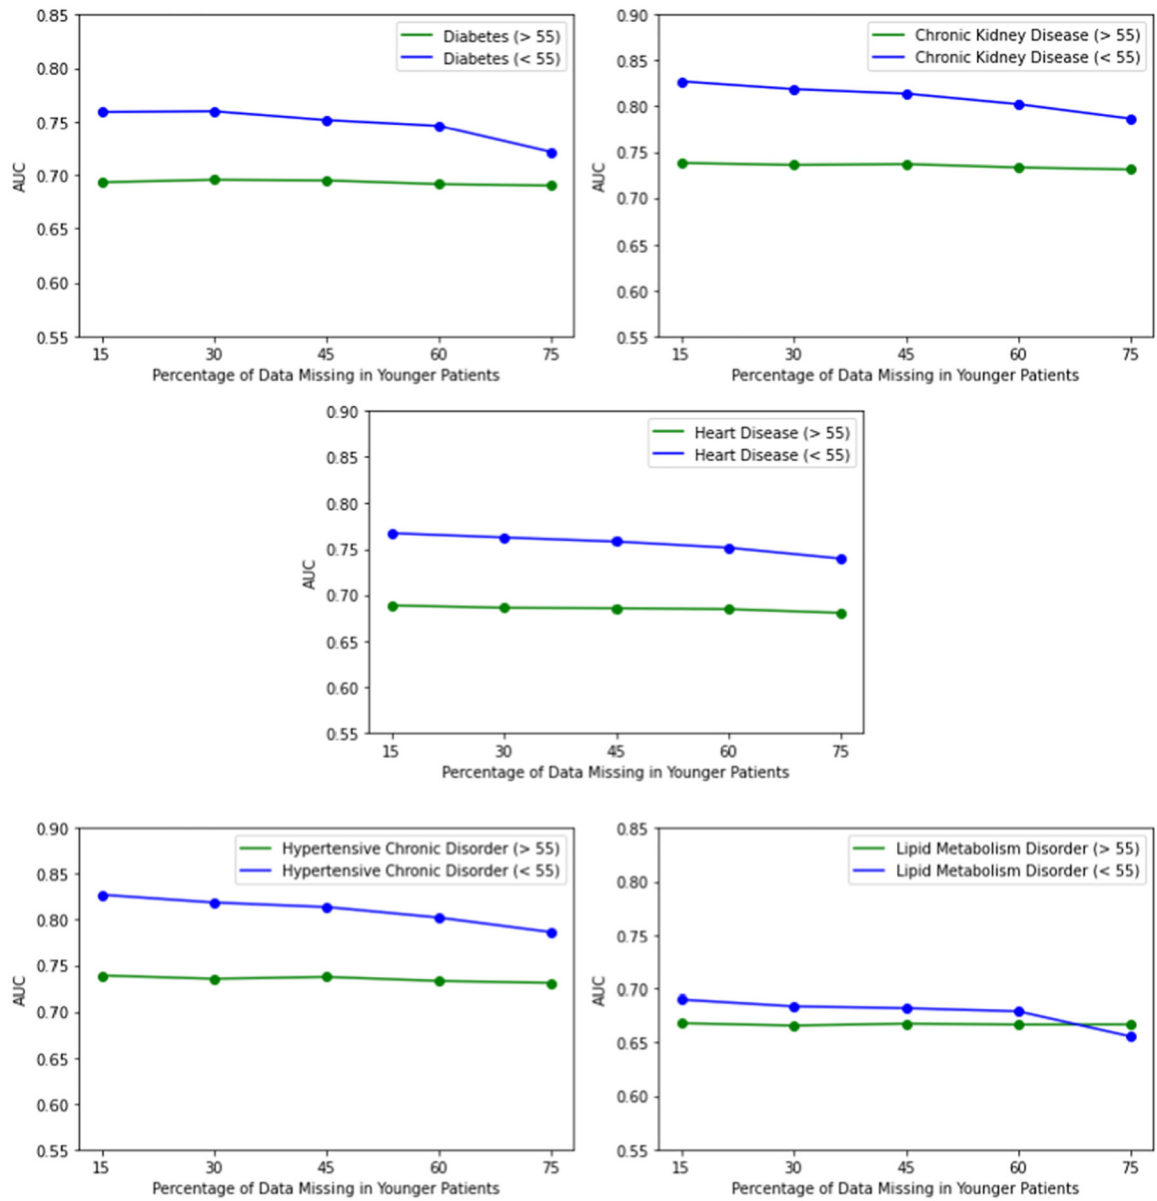


Fig. 15. Assessing the impact of missing data on data-lacking groups. Proportion missing for older patients is held constant while proportion missing for younger patients is varied.

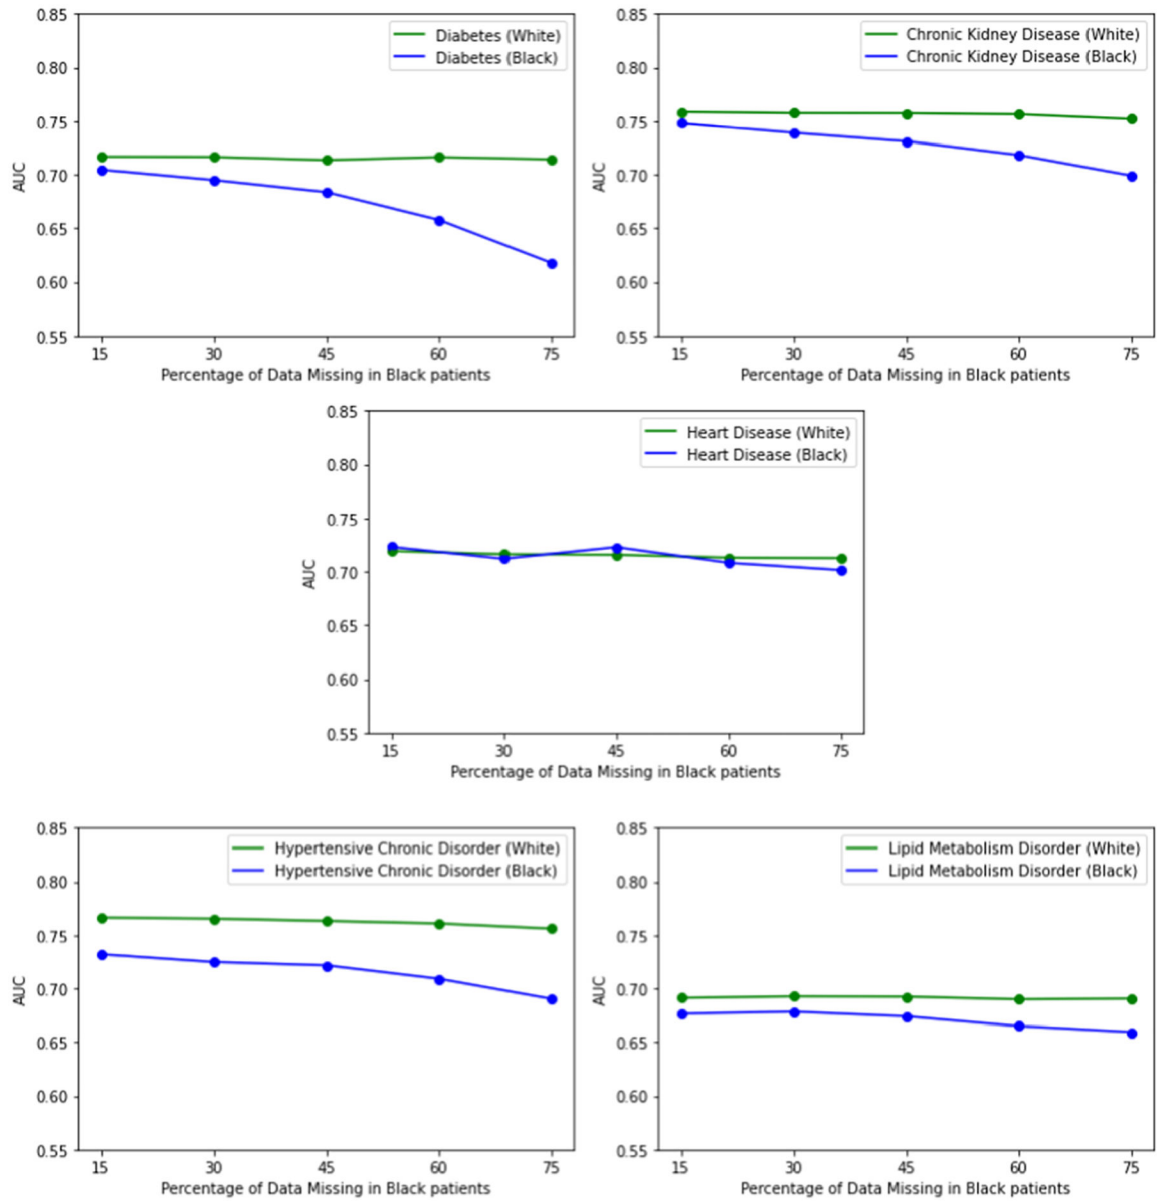


Fig. 16. Assessing the impact of missing data on data-lacking groups. Proportion missing for White patients is held constant while proportion missing for Black patients is varied.




Fig. 17.

Impact of knowledge graph in event removal as the proportion of missing data (p) varies.

Events are only removed from the training data. On the left, we assess the impact of removing medical events at random. On the right, we assess the impact of removing medical events using the knowledge graph.




Fig. 18. Impact of knowledge graph in event removal as the proportion of missing data (p) varies. Events are removed from both the training data and testing data. On the left, we assess the impact of removing medical events at random. On the right, we assess the impact of removing medical events using the knowledge graph.

Table 1

Statement of Significance.

Problem or issue	There is currently no framework for generating realistic missing data in electronic health records (EHR) and quantifying its impact on underserved populations.
What is already known	Although there are existing methods for simulating missing data in EHR, these methods rely on data that have well-defined features. They also do not take medical event dependencies into account.
What this paper adds	We propose a terminology for defining three levels of EHR data and develop a new framework for generating realistic missing data in EHR that can handle structured, sequential, codified data. Our framework incorporates event dependencies into the event-removal process via a medical knowledge graph, thus simulating more realistic missing data. This work also explores the impact of realistic missing data on performance of disease prediction models in groups that tend to have less access to healthcare or seek less healthcare-particularly patients of lower socioeconomic status and patients of color.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Distributions of medical events across various demographic groups. The data in this table correspond to MIMIC-III data that are unaltered (i.e. no missingness has been simulated). There are 3950 patients in the training set and 1692 patients in the test set. The mean number of medical events for each patient is 315, with a minimum of 7 and maximum of 3463. Across all patients, there are a total of 994,320 medical concepts. Mean and median number of medical events per patient record are evenly distributed across age groups. There are more older patients than younger patients in the training data. We also see that mean and median number of medical events per patient record are evenly distributed across sex. Proportions of male-to-female patients are relatively even in the training data.

Age	Mean Events	Median Events	Standard Deviation	Proportion of Patients
17-54	336	252	310	0.254
55-70+	315	261	221	0.716
Sex				
Male	323	261	236	0.541
Female	317	258	258	0.428
Insurance Type				
Government	323	262	251	0.706
Private or Self	313	250	239	0.263
Race				
Black	355	280	325	0.119
White	317	259	228	0.722

Table 3

“Underlying conditions” for MNAR prediction models.

Target diagnosis	“Underlying condition”	Cosine similarity
Diabetes	Toxic Nerve Damage	0.843
	Fluid / Acid-Base Disorder	0.817
	Anemia	0.782
Chronic Kidney Disease	Glomerulosclerosis	0.877
	Nephritis	0.855
	Impetigo	0.771
Heart Failure	Chronic Pulmonary Disease	0.885
	Cardiomyopathy	0.868
	Mitral Valve Disorder	0.843
Hypertensive Chronic Disorder	Toxic Nerve Damage	0.720
	Fluid / Acid-Base Disorder	0.720
	Chronic Pulmonary Disease	0.719
Lipid Metabolism Disorder	Hypothyroidism	0.831
	Gout	0.826
	Fluid / Acid-Base Disorder	0.800

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript