

Evaluating the Impact of Social Determinants of Health on Prediction of Clinical Outcomes in the Intensive Care Unit

by

Ming Ying Yang

S.B., Computer Science and Engineering and Chemical Engineering
Massachusetts Institute of Technology (2022)

Submitted to the Department of Electrical Engineering and Computer Science
In Partial Fulfillment of the Requirements for the Degree of

MASTER OF ENGINEERING IN
ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

© 2023 Ming Ying Yang. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Ming Ying Yang
Department of Electrical Engineering and Computer Science
May 18, 2023

Certified by: Marzyeh Ghassemi
Assistant Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by:: Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Evaluating the Impact of Social Determinants of Health on Prediction of Clinical Outcomes in the Intensive Care Unit

by

Ming Ying Yang

Submitted to the Department of Electrical Engineering and Computer Science
on May 18, 2023 in Partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Social determinants of health (SDOH) – the conditions in which people live, grow, and age – play a crucial role in a person’s health and well-being. There is a large, compelling body of evidence in population health studies indicating that a wide range of SDOH is strongly correlated with health outcomes. Yet, a majority of the risk prediction models based on electronic health records (EHR) do not incorporate a comprehensive set of SDOH features as they are often noisy or simply unavailable. Our work links a publicly available EHR database, MIMIC-IV, to well-documented SDOH features. We investigate the impact of such features on common EHR prediction tasks across different patient populations. We find that community-level SDOH features do not enhance the predictive accuracy of a model, but they can improve the model’s calibration and fairness. We further demonstrate that SDOH features are vital for conducting thorough audits of algorithmic biases beyond protective attributes. We hope the new integrated EHR-SDOH database will enable studies on the relationship between community health and individual outcomes and provide new benchmarks to study algorithmic biases beyond race, gender, and age.

Thesis Supervisor: Marzyeh Ghassemi

Title: Assistant Professor of Electrical Engineering and Computer Science

Acknowledgments

First and foremost, I would like to express my deepest gratitude to Professor Marzyeh Ghassemi for her invaluable guidance and support throughout my MEng journey. Her expertise and insightful feedback have been instrumental in shaping this research project.

I would also like to extend my heartfelt appreciation to my collaborators on this project: Dr. Leo Anthony Celi, Dr. Tom Pollard, and Dr. Gloria Hyunjung Kwak. Their expertise, constructive discussions, and willingness to share their knowledge have been invaluable to the success of this project.

Furthermore, I would like to thank the Healthy ML group. The stimulating environment and interdisciplinary discussions within the group have broadened my perspective and made research fun. Special thanks to Hammaad Adams, Aparna Balagopalan, Hyewon Jeong, Qixuan (Alice) Jin, Intae Moon, Vinith Suriyakumar, Yuxin Xiao, and Haoran Zhang for their valuable feedback and constructive review of this work. I am especially appreciative of Hammaad, an incredible research mentor and collaborator.

I have been extremely fortunate to be surrounded by many great friends throughout this journey. They have made my time at MIT truly remarkable, and I am really grateful for their presence in my life.

Lastly, I would like to express my heartfelt appreciation to my parents and my sister who have not only supported me throughout my college career but also my entire life. Their unconditional love, encouragement, and understanding have been a constant source of motivation. I would not be where I am today without them.

Contents

1	Introduction	10
2	Related Work	13
2.1	SDOH in Health Prediction	13
2.2	SDOH in EHR	14
2.3	Fairness and Bias in Healthcare	15
3	Creation of EHR-SDOH Database: MIMIC-IV-SDOH	17
3.1	Public SDOH Databases	18
3.1.1	County Health Rankings (CHR)	18
3.1.2	Social Vulnerability Index (SVI)	18
3.1.3	Social Determinants of Health Database (SDOHD)	18
3.2	EHR-SDOH Integration	19
3.3	Comparison of SDOH Features	20
4	Data and Methods	22
4.1	Data	22
4.1.1	MIMIC-IV-SDOH	23
4.1.2	All of Us Dataset	25
4.2	Models Benchmarked	26
4.3	Evaluation	26
4.3.1	Classification Performance	26
4.3.2	Group Fairness	27

5	Impact of SDOH on Clinical Prediction Tasks	28
5.1	Impact of SDOH on the General Population	28
5.1.1	Limited Improvement in Model Performance at the General Population Level	29
5.1.2	Generalizability of the Finding	30
5.2	Varying Impact of SDOH on Vulnerable Patient Populations	31
5.2.1	Limited Improvement in Model Performance in Vulnerable Pa- tient Populations	31
5.3	SDOH As Fairness Audit Categories	32
5.3.1	SDOH Features Enable More Granular Audits	33
6	Discussion	36
6.1	The Need for Better Data	36
6.2	On More Actionable Audits	37
6.3	Future Work on SDOH and Health Predictions	38
7	Conclusion	39
A	Additional Information on Model Training and Data Pre-processing	40
A.1	Model Training	40
A.2	Binning SDOHD Features	41
A.3	SDOH Features in the All of Us Dataset	41
B	Tables	43

List of Figures

3-1	Comparison of features in the MIMIC-IV-SDOH databases. SDOHD is the most comprehensive and granular SDOH database of the three. Its features are more correlated with race than CHR and SVI indices. Each database emphasizes a set of SDOH domains more than others.	21
4-1	Comparison of selected SDOH features between the MIMIC-IV and the All of Us patient cohorts. Because the All of Us dataset is more geographically diverse, the variation in its SDOH data is much greater than that in MIMIC-IV.	23
5-1	Comparison of the performance of XGBoost classifiers for predicting in-hospital mortality, 30-day readmission, and one-year mortality in the general ICU population. There are no detectable differences between AUROC and AUPRC of models that do not incorporate SDOH features and those that do not. However, when combining SDOH features with both tabular EHR data and discharge notes, we observe significant impacts on ECE, FPR, and recall. We highlight such occurrences with asterisks. The error bars denote the 95% confidence intervals obtained through 1000 bootstrap samples.	29

5-2 Comparison of the FPR of XGBoost classifiers trained on all EHR data (A11) and all EHR data combined with the most helpful SDOH features for that task in all ICU patients in MIMIC-IV. FPR is reported for subgroups defined by race, gender, age, and five SDOH features, which are binned into quartiles. The bin edges are documented in section A.2. The error bars denote the 95% confidence intervals obtained through 1000 bootstrap samples. 34

List of Tables

3.1	Characteristics of the final MIMIC-IV-SDOH tables where d is the number of SDOH features. Note that both the SVI and SDOHD are available at the county level; SDOHD is also available at the zip code level. We only use county-level data to minimize the risk of patients being identified.	19
4.1	Breakdown of feature sets by category	25
5.1	Comparison of model performance for XGBoost classifiers trained with and without SDOH for the task of 30-day readmission in MIMIC-IV and the All of Us dataset. In both datasets, the addition of SDOH has no effect on model performance.	30
5.2	Combinations of EHR and SDOH features that influence performance of the best models for each patient population in terms of AUPRC. We report the number of occurrences in which incorporating SDOH features significantly improves or worsens performance in the form of (# improves/# worsens) across the three prediction tasks (total number of occurrences is 3). Significance is evaluated using a 1000-sample bootstrap hypothesis test at the 5% significance level.	32

B.1	Characteristics of the MIMIC-IV and All of Us cohorts. The MIMIC-IV cohort has six patient populations: all ICU patients, diabetic patients, black diabetic patients, elderly diabetic patients, female diabetic patients, and non-English speaking diabetic patients. The All of Us cohort contains all in-patient hospital stays. N is the number of patients in each group.	43
B.2	Comparison of the performance of models trained with and without SDOH feature to predict in-hospital mortality, 30-day readmission, and one-year mortality for the six patient populations. The evaluation is done only on the best model for each feature set category and task. In general, incorporating SDOH features has a limited impact on model performance. Values in bold indicate significantly better performance compared to the baseline model trained without SDOH, evaluated using a 1000-sample bootstrap hypothesis test at the 5% significance level.	44

Chapter 1

Introduction

The increasing adoption of electronic health records (EHRs) in modern healthcare systems has facilitated the development of machine learning (ML) models to predict the progression of diseases and patient outcomes. Many such models [47, 55, 77] incorporate clinical factors (e.g., labs, vitals, medication, procedures) and basic demographic features (e.g., age, gender, and race) to identify high-risk patients. However, a patient's clinical profile only offers a partial view of all the risk factors that affect their health. Understanding the conditions of their living environment may help to fill in the missing pieces and benefit patients' health and medical care. Human health is affected by many non-clinical factors, commonly known as social determinants of health (SDOH).

The Healthy People 2030 initiative [58], developed by the United States Department of Health and Human Services, describes SDOH as "the conditions in the environments where people are born, live, learn, work, play, worship, and age that affect a wide range of health, functioning, and quality-of-life outcomes and risks." They grouped SDOH into five key domains: (1) economic stability [19, 116], (2) education access and quality [52, 73], (3) health care and quality [63], (4) neighborhood and built environment [74, 90], and (5) social and community context [25, 89].

Across all five domains, SDOH can have either a direct or indirect impact on one's health. At a high level, they can be viewed as individual-level determinants or community-level determinants [21]. The former determinants are specific to a person, and examples include education level, annual income, and family dynam-

ics. Access to individual-level SDOH is limited due to the lack of standardized and validated SDOH screening questions [21] and privacy concerns [87]. In contrast, community-level SDOH measure broader socioeconomic, neighborhood, and environmental characteristics such as unemployment rate, access to public transportation, and air pollution levels. They serve as “community vital sign” [12] that reflect complex societal factors and health disparities that influence one’s health [2, 17].

Population health studies have identified many SDOH to be strongly correlated with acute and chronic conditions [6, 38, 41, 45, 102]. SDOH are also underlying, contributing factors of health disparities (e.g., poverty [32, 46, 128], unequal access to health care [25, 35], low educational attainment [8, 36, 52], and segregation [22, 107]). However, to date, there has been less focus in the ML community to include SDOH in common EHR prediction tasks because many SDOH measures are poorly collected, lack granularity, or are simply unavailable. An American Health Information Management Association (AHIMA) survey [96] finds that most health-care organizations are collecting SDOH data, but they face challenges with a lack of standardization and integration of the SDOH data into EHR and patient distrust in sharing the data. Thus, while SDOH are being increasingly studied in population health [63, 108, 135] and primary care settings [71, 96], data limitations have left the association between SDOH and critical care outcomes largely unexplored.

In this work, we investigate the impact of incorporating SDOH features on common EHR prediction tasks in the intensive care unit (ICU). We first link MIMIC-IV [68], a publicly available EHR database, to external SDOH databases based on patient zip code. We then train models on the common tasks of mortality and readmission risk, evaluating the contribution of SDOH as compared to the EHR data alone. We find that adding SDOH does not improve model accuracy. We do note that, as compared to the EHR data alone, incorporating SDOH can lead to better-calibrated and fairer models in specific patient groups, with varying levels of contribution depending on the population and predictive task. Finally, we illustrate that fairness audits based on both protective attributes and SDOH features help to connect the commonly observed disparities to the underlying mechanisms that drive adverse health outcomes downstream.

Our work makes three main contributions.

- We release a publicly accessible database that combines EHR data with SDOH measures. To the best of our knowledge, this is the first public EHR database that contains structural features that span all five defined SDOH domains. The database will enable new studies on the relationship between community health and individual clinical outcomes.
- We investigate the impact of incorporating SDOH in predictive models across three tasks, three model classes, and six patient populations. We find that the inclusion of SDOH can improve performance for certain vulnerable subgroups.
- We demonstrate that SDOH features enable more fine-grain audits of algorithmic fairness, reporting the FPR parity – the difference in false positive rates (FPR) – across intersectional patient subgroups.

Chapter 2

Related Work

2.1 SDOH in Health Prediction

A number of studies in population health have attempted to assess the impact of social factors on health [2, 18, 84, 106]. There is a large, compelling body of evidence showing that a wide range of SDOH is strongly correlated with health outcomes, such as sepsis [6], heart failure [30], pneumonia [89], cardiovascular disease [45], and diabetes [129, 134]. A particular study found that 40% of deaths in the United States are caused by behavior patterns that could be modified by preventive interventions and suggested that only 10-15% of preventable mortality could be avoided by higher-quality medical care [85]. Other studies have also indicated that the effect of medical care may be more limited than commonly believed [63, 81, 82]. However, there are active controversies regarding the strength of the evidence that suggests a causal relationship between SDOH and well-being. These researchers are increasingly utilizing SDOH to predict individual health outcomes [115, 141].

While several studies have shown that machine learning models can predict individual patient outcomes, such as in-hospital mortality [47, 51, 77, 92, 134] and readmission [49, 77, 110], very few have incorporated SDOH into the models due to the lack of granular and high-quality SDOH data at the individual level.

Due to the limited availability of individual-level SDOH data, many studies are limited to community-level SDOH data [28]. Most found that community-level SDOH do not lead to improvement in model performance [15, 66, 126], partly

due to the low data resolution. In contrast, researchers who are able to access individual-level SDOH generally report improvements in the model’s predictive performance [28, 43, 91]. These studies often focus on a specific outcome for a specific patient group, such as HIV risk assessment [43, 95] and suicide attempts [130, 140]. There have also been studies of model improvements for readmission and mortality prediction in specific subgroups, such as the elderly and obese [138]. One has shown that integrating SDOH into predictive models can improve the fairness of the prediction in underserved heart failure subpopulations [80].

Despite a growing body of SDOH-focused research, the relationship between SDOH and critical care outcomes is unclear. While some have argued that the ICU is not an appropriate setting to collect and identify SDOH, there are several reasons why it could be essential. For example, critical conditions place high demands on the patient and their social network [88, 123]. Social isolation may increase the risk of adverse outcomes, such as mortality [123]. By incorporating SDOH into MIMIC-IV, our work investigates the contribution of community-level SDOH on common EHR prediction tasks in a multi-year cohort in the critical care setting.

2.2 SDOH in EHR

In order for SDOH features to be readily incorporated into risk prediction models, they need to be collected and documented with individual health outcomes. EHRs contain clinical information about patients, such as medical history, vital signs, laboratory data, immunizations, and medications [24]. In the United States, few SDOH features are currently documented in structured EHR data fields due to the lack of adoption of standardized and validated SDOH screening questions [21]. The set of SDOH features available for research use is typically limited to insurance type, preferred language, and smoking and alcohol use, but SDOH information can also be extracted from unstructured EHR data (i.e. clinical notes) [20, 43, 93]. SDOH may also be captured in billing codes [5], but they have not been widely utilized by providers [67].

The integration of SDOH in EHRs is further delayed due to concerns regarding

privacy and misuse [87]. The United States Public Health Services Syphilis Study at Tuskegee (Tuskegee Study) among African Americans [131] and efforts to sterilize American Indians [37, 76] are examples of a dark history of structural inequities in healthcare and unethical medical experimentation against racial and ethnic minorities. As a result, mistrust of the healthcare system and medical research has been well documented among minority groups [16, 33, 42]. The collection and utilization of SDOH require the consent and trust of the patients. Patients who identify with populations that medical establishments and medical researchers historically mistreat might not want to share any personal or sensitive information.

Overall, current EHR-derived SDOH data do not constitute a comprehensive set of SDOH domains. In this study, we link a large, multi-year EHR dataset to public SDOH datasets covering *all five SDOH domains* to comprehensively study the relationship between the community-level SDOH and patient outcomes.

2.3 Fairness and Bias in Healthcare

While much work has been done in algorithmic fairness and bias in health, most of the studies that focused on group-based fairness have only examined bias from the lens of protected attributes, namely age, gender, and race [1, 11, 26, 27, 56, 59, 83, 97, 113, 114, 117, 136, 137]. Recent fairness literature has underscored the importance of measuring biases from multidimensional perspectives, focusing on social processes that produce the biases [53, 62, 112].

There is strong evidence that intersectional social identities are related to a patient’s health outcomes [71, 114]. Capturing social context beyond protected attributes in the form of SDOH is thus vital for this cause. For example, in the primary care setting, researchers have observed a negative correlation between the odds of receiving appropriate prevention and screening and the number of social risk factors experienced by the patient [71]. The more factors a patient was living with, the less likely they were to receive care such as a mammogram screening or vaccinations. This is not something that can be detected through race or gender alone.

Moreover, a recent meta-analysis [127] ranked 47 studies using a self-developed SDOH scoring system based on the type and number of SDOH features used. The researchers found that Black patients had significantly higher prostate cancer-specific mortality (PCSM) than White patients when there was minimal accounting for SDOH. In contrast, studies with greater consideration for SDOH showed the opposite: Black patients had significantly lower PCSM compared to White patients. The findings of this meta-analysis should not be interpreted as suggesting that racial disparities do not exist. Rather, it suggests that there is a significant interaction between race and SDOH, and health equity researchers should incorporate SDOH features into data collection and analyses to better address the long-standing disparities in healthcare [131].

We hope the new integrated MIMIC-IV-SDOH dataset will enable more studies that follow the complex hierarchical system that defined advantaged or disadvantaged subjects in the first place. Our work serves as a first effort, and we demonstrate how SDOH features allow for more granular, actionable algorithmic audits.

Chapter 3

Creation of EHR-SDOH Database: MIMIC-IV-SDOH

The MIT Laboratory for Computational Physiology (LCP) developed and maintains the publicly available Medical Information Mart for Intensive Care (MIMIC), a database on patients admitted to the emergency department (ED) and intensive care units (ICU) at the Beth Israel Deaconess Medical Center (BIDMC) in Boston [69]. The database is used by researchers in over 30 countries for clinical research studies, exploratory analyses, and the development of decision-support tools [103]. The current version, MIMIC-IV, contains detailed, de-identified data associated with over 70,000 ICU stays from 2008 to 2019 and over 400,000 ED stays from 2011 and 2019. Yet, due to the lack of high-quality SDOH data, none of the studies or tools built based on MIMIC account for SDOH measures beyond basic demographics such as insurance, and language. To enable the study of the relationship between community characteristics and individual health outcomes, we create the MIMIC-IV-SDOH database by linking MIMIC-IV to three public SDOH databases (Table 3.1):

1. County Health Rankings (CHR) [34]
2. Social Vulnerability Index (SVI) [23]
3. Social Determinants of Health Database (SDOHD) [4]

This database will be accessible as a Contributor Review database on PhysioNet [48].

3.1 Public SDOH Databases

While there exist other SDOH databases, such as Area Deprivation Index [72] and Atlas of Rural and Small-Town America [125], they are either domain-specific or not frequently updated. Because MIMIC-IV contains ICU stays from 2008 to 2019, we focus on databases with SDOH variables that span multiple years and all five SDOH domains, as defined by Healthy People 2030 [58].

3.1.1 County Health Rankings (CHR)

CHR evaluates counties within each state in the United States based on modifiable health determinants and is updated annually. CHR estimates that clinical care only accounts for 20% of all contributors to long-term health outcomes, specifically the length and quality of life. The remaining 80% stems from health behaviors (30%), physical environment (10%), and social and economic factors (40%) [63].

3.1.2 Social Vulnerability Index (SVI)

Based on data from the American Community Survey (ACS), SVI evaluates social factors across four themes: socioeconomic status, household composition and disability, minority status and language, and housing type and transportation. Although the index was designed to assess community preparedness and resilience in the face of natural hazards, SVI has been used in many population health and health equity studies [10, 70, 78, 121, 135]. For example, communities with higher levels of social vulnerability experienced greater COVID-19 incidence and mortality [64, 70, 121]. Unlike CHR, SVI is available at both the county level and census tract level.

3.1.3 Social Determinants of Health Database (SDOHD)

To incorporate more granular SDOH data into MIMIC-IV, the last database used in the integration is the Social Determinants of Health Database (SDOHD), which is available at the county, census tract, and zip code levels. The database was recently developed to provide a range of well-documented, readily linkable SDOH variables

Table 3.1: Characteristics of the final MIMIC-IV-SDOH tables where d is the number of SDOH features. Note that both the SVI and SDOHD are available at the county level; SDOHD is also available at the zip code level. We only use county-level data to minimize the risk of patients being identified.

SDOH Database	Data Version/Year	Geographic Level	d
CHR	2010-2020	County	106
SVI	2008, 2014, 2016, 2018	Census Tract	162
SDOHD	2009-2020	Census Tract	1329

across domains without having to access multiple source files. SDOHD is curated based on the five key SDOH domains defined by Healthy People 2030: economic stability, education access and quality, health care and quality, neighborhood and built environment, and social and community context. It contains measures from CHR and SVI, making it the most comprehensive database of the three.

3.2 EHR-SDOH Integration

The creation of the integrated MIMIC-IV-SDOH database is carried out in three steps.

Step 1: SDOH Data Acquisition For each SDOH database, we concatenate all datasets released between 2008 and 2020. We map each feature to one of the five SDOH domains and provide detailed documentation.

Step 2: Geographic Crosswalk Although SVI and SDOHD are available at the census tract level, we only use county-level data to minimize the risk of patients being identified. Each patient’s zip code is mapped to a county using the crosswalk files provided by the United States Department of Housing and Urban Development (HUD) [133]. The files contain a residential ratio column, the ratio of residential addresses in the zip-county area to the total number of residential addresses in the entire zip. Because the mapping is many-to-many, the residential ratio is treated as the probability that the patient with zip code z lives in the census tract t or county c , as suggested by the HUD [40]. Note that only patients with Massachusetts zip

codes that are in the HUD crosswalk files are included in the final MIMIC-IV-SDOH dataset.

Step 3: Data Merging MIMIC-IV is merged with each of the three SDOH databases using the geographic location and the SDOH data year closest to the year of admission.

3.3 Comparison of SDOH Features


The demographic features in MIMIC-IV, such as race and gender, are sometimes used as proxies for SDOH features, such as socioeconomic status and health behaviors [27, 111]. We find that many community-level SDOH features are weakly correlated with race in MIMIC-IV. For example, two SDOHD features, the percentage of households that receive food stamps and the percentage of workers taking public transportation, are both weakly and positively correlated with the Black race (Figure 3-1).

Though to a lesser extent, subindices from SVI (e.g., socioeconomic and household composition) and CHR (e.g. health outcomes, quality of life, and social and economic factors) are also weakly associated with race. There are no strong correlations between SDOH features and other tabular features such as labs, risk scores, and Charlson comorbidities.

To better illustrate the type of features in each SDOH database, we manually map each feature to one of the five SDOH domains (Figure 3-1). SDOHD is arguably the most detailed and comprehensive SDOH database out of the three used in the integration and has the largest feature space. As shown in Figure 3-1, the features in SDOHD are predominantly of the economic stability domain. While SVI shares a similar distribution as SDOHD, it emphasizes the neighborhood and the built environment more. Conversely, more than 40% of the features in CHR are related to health outcomes. The types of features in CHR make it well-suited for health predictions, but it has very low resolution, with only 56 distinct sets of values in the integrated MIMIC-IV-CHR table.



(A) Pearson correlation coefficients between basic demographic features and selected features in CHR (top), SVI (middle), and SDOHD (bottom).



(B) Distribution of features in each SDOH database by domain.

Figure 3-1: Comparison of features in the MIMIC-IV-SDOH databases. SDOHD is the most comprehensive and granular SDOH database of the three. Its features are more correlated with race than CHR and SVI indices. Each database emphasizes a set of SDOH domains more than others.

Chapter 4

Data and Methods

Our primary goal is to determine how incorporating SDOH in ML models could impact predictions of acute and longitudinal outcomes. Leveraging the newly created MIMIC-IV-SDOH database, we assess the impact from the perspective of classification performance and group fairness. We also provide a preliminary investigation of the possible mechanisms behind the contributions of SDOH to model performance or the lack thereof.

4.1 Data

In this study, we analyze five patient populations in the MIMIC-IV-SDOH database to assess the impact of SDOH across three tasks. MIMIC-IV data comes from a single EHR system in one geographic location, so the variation in the community-level SDOH features might be too low to be informative. Many past studies that used community-level SDOH features with EHR data from a single hospital or region ended with similar conclusions [28]. To examine the generalizability of our finding, we compare the MIMIC-IV cohort to a patient population in the All of Us Controlled Tier Dataset v6 [120] for the task of 30-day readmission. Unlike MIMIC-IV, which comes from a single hospital in Boston, the All of Us dataset contains patient-level data from more than 35 hospitals across the United States. Because of this difference, the variation in the SDOH data in All of Us is much greater than that in MIMIC-IV (Figure 4-1). In addition, based on the distribution of SDOH features, we note that

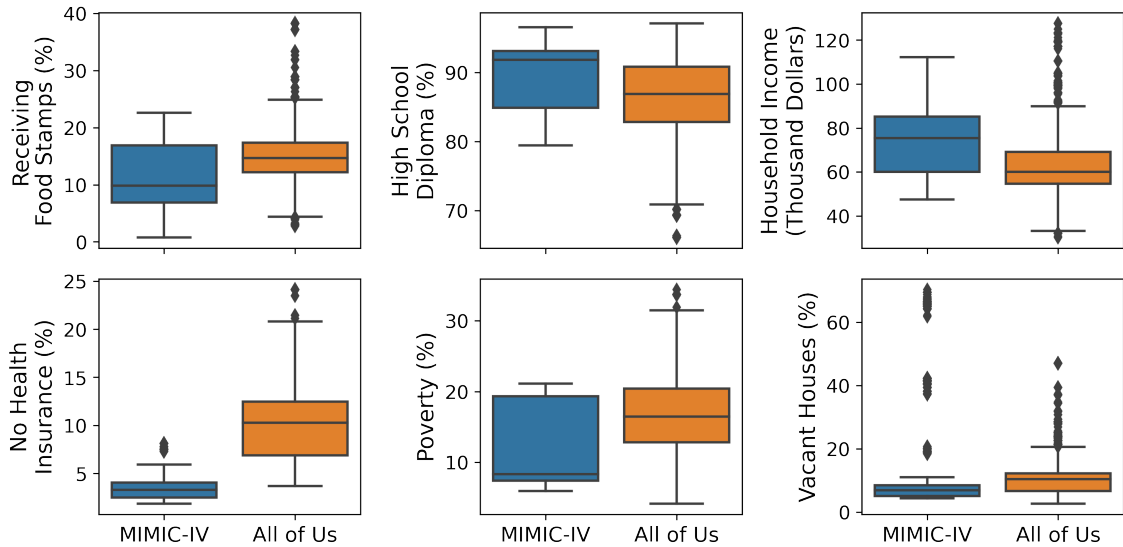


Figure 4-1: Comparison of selected SDOH features between the MIMIC-IV and the All of Us patient cohorts. Because the All of Us dataset is more geographically diverse, the variation in its SDOH data is much greater than that in MIMIC-IV.

the patients in MIMIC-IV are on average more affluent and more educated than the patients in All of Us.

4.1.1 MIMIC-IV-SDOH

Our analysis only includes ICU patients from MIMIC-IV v2.2 who are at least 18 years old and has a hospital length of stay of at least 3 days. The final cohort consists of 42,665 patients and a total of 54,380 admissions.

Task Definition

We focus on three common classification tasks: (1) in-hospital mortality, (2) 30-day readmission, and (3) one-year mortality. Patients who have expired during a stay are excluded from predictions of 30-day readmission and one-year mortality. For 30-day readmission, we only consider non-elective readmissions.

Patient Population Definition

A recent study suggests that the impact of incorporating SDOH in predictive models varies by subpopulation – vulnerable populations like Black patients and the elderly

are likely to benefit more from the inclusion of SDOH [13]. Moreover, several studies have suggested that SDOH are strongly associated with glycemic control [129], as well as diabetic risk, morbidity, and mortality [60]. Diabetic patients also use significantly more healthcare resources compared to patients with other chronic diseases [44]. In fact, they account for 31% of all ICU patients in MIMIC-IV.

Thus, in addition to the cohort of all ICU patients, we evaluate five subgroups:

1. Diabetic patients
2. Black diabetic patients
3. Elderly diabetic patients who are over 75 years old
4. Female diabetic patients
5. Non-English speaking diabetic patients.

On average, all five of these subgroups have more comorbidities compared to the general ICU patients (Table B.1).

Data Pre-processing

To better understand the contribution of different types of features to model performance, we divide the entire dataset into a total of 15 feature sets (Table 4.1) and train separate models on each. These feature sets can be broadly classified into three categories: SDOH features alone, EHR features alone, and SDOH features combined with EHR features.

For the EHR features, we include two different data modalities: tabular data and discharge notes. To enable fair comparison across the three tasks, we use the same tabular features and sections of the discharge notes in all prediction tasks. Tabular features include basic demographics, Charlson comorbidities, labs from the first 24 hours of stay, and risk scores (APSO, SAPS-II, SOFA, and OASIS). The following sections from discharge notes are included: chief complaint, present illness, medical history, medications on admission, allergies, major surgical or invasive procedure, physical exam on admission, pertinent results on admission, and family history.

Table 4.1: Breakdown of feature sets by category

Category	Feature Set
SDOH Only	CHR
	SVI
	SDOHD
EHR Only	Tabular
	Notes
	All EHR (All)
EHR and SDOH	Tabular+CHR
	Tabular+SVI
	Tabular+SDOHD
	Notes+CHR
	Notes+SVI
	Notes+SDOHD
	All+CHR
	All+SVI
	All+SDOHD

Before separating patients into different patient populations, we use median imputation for numerical features before performing standard scaling and constant imputation for categorical features before one-hot encoding. Median imputation is used instead of mean imputation in consideration of skewed data. We also apply principal component analysis (PCA) to reduce the dimensionality of the SDOH data, which is particularly useful as many of the SDOH features are strongly correlated. We retain principal components that explain at least 0.99 of the variance in the data.

Discharge notes are stripped of explicit indicators of in-hospital mortality before being tokenized and lemmatized. Corpus-specific stop words are removed by filtering terms with a document frequency greater than 0.7. Terms with a document frequency less than 0.001 are also removed. Lastly, the notes are converted into a Term Frequency-Inverse Document Frequency (TF-IDF) representation with a vocabulary of size $|V| = 11,751$ words.

4.1.2 All of Us Dataset

Because the All of Us dataset is made up of primarily living participants, we focus on the prediction of 30-day readmission. We only include adult participants who

had an in-patient hospital stay between 2014 and 2021. Unlike MIMIC-IV, the hospital stays in All of Us are not limited to the ICU, so these patients have fewer comorbidities on average (Table B.1).

We exclude patients who stayed less than 3 days in the hospital and didn't have any lab results in the first 24 hours; these are the same labs used in MIMIC-IV. We have 13,324 patients and 21,555 admissions in the final All of Us cohort, representing all 50 US states but Kentucky. More than 50% of the patients come from the Northeast region.

The All of Us dataset only has 7 community-level SDOH features sourced from the 2017 ACS via a three-digit zip code linkage (section A.3). Tabular features in the All of Us dataset are the same as those in MIMIC-IV except for clinical risk scores, which are not available. We apply the same data pre-processing techniques on both datasets. Because the All of Us dataset has no clinical notes, we only train models on three feature sets: (Tabular, SDOH, and Tabular+SDOH).

4.2 Models Benchmarked

We train three types of machine learning models – logistic regression [101], random forest classifier [101], and XGBoost classifier [29] – for each task, patient population, and feature set combination. Each dataset is partitioned into 70:30 train-test splits. To prevent data leakage, no patient appears in both the training set and the test set. Each model is tuned through random hyperparameter search [14] under broad parameter distributions. See section A.1 for additional training details.

4.3 Evaluation

4.3.1 Classification Performance

We evaluate the models in terms of three primary metrics: 1) area under the receiver operating characteristic curve (AUROC), 2) area under the precision-recall curve (AUPRC), and 3) expected calibration error (ECE). While AUROC is a

standard metric to assess accuracy, we include AUPRC to account for class imbalance and ECE to measure the reliability of the prediction. We also use recall as a secondary metric. While threshold selection is complex, cost-dependent, and application-specific, we use a classification threshold of 0.5 for demonstration purposes. 95% confidence intervals are constructed for all metrics by sampling the test set for 1000 bootstrap iterations [39].

4.3.2 Group Fairness

In addition to classification parity, we evaluate the FPR parity – based on the equality of opportunity definition of group fairness [54].

$$\hat{Y} \perp\!\!\!\perp G \mid Y = 0$$

In other words, the probability of the model predicting a negative outcome is independent of group attribute G , conditional on the outcome Y being a true negative.

We examine the differences in TPR across subgroups defined based on the following attributes: (1) race, (2) age, discretized into four bins, (3) gender, (4) median household income, (5) percentage of workers commuting via public transportation, (6) percentage of the population with educational attainment less than high school, (7) percentage of the population receiving food stamps, and (8) the percentage of non-citizens. The SDOH features are discretized into quartiles.

Chapter 5

Impact of SDOH on Clinical Prediction Tasks

The goal of our analysis is to investigate how SDOH impact three common clinical classification tasks. For every task and patient population, we first train a model on each of the 15 feature sets. We then determine the best model for each patient population in terms of AUPRC and report the model performance in Table B.2. Lastly, we provide a preliminary investigation of some of the possible mechanisms behind the impact of SDOH or the lack thereof.

5.1 Impact of SDOH on the General Population

We first examine the impact of SDOH on the general ICU patient population in MIMIC-IV. We find that the inclusion of community-level SDOH in models does not help with predictive performance, measured by AUROC and AUPRC, but can lead to better-calibrated models with a lower FPR. We validate this finding on a more geographically diverse dataset: the All of Us dataset.

5.1.1 Limited Improvement in Model Performance at the General Population Level

First, SDOH features alone, without any EHR data, are not predictive of individual patient outcomes. The mean test AUROC of the XGBoost models, the best models in terms of AUROC, trained on SDOH alone is below 0.60 across all tasks, substantially lower than those trained on tabular EHR features and the TF-IDF representation of discharge notes (Figure 5-1). This is not particularly surprising as most studies that utilize community-level SDOH have arrived at similar conclusions [27]. One possible explanation is that community-level estimates are either imprecise or biased, especially if the within-community variance of a feature is high. Moreover, when a patient is critically ill, information on their upstream risk factors might not be as useful as their current state of health.

Similarly, combining SDOH with tabular EHR data and discharge notes does not improve the AUROC and AUPRC of the model. Again, this trend is observed in all

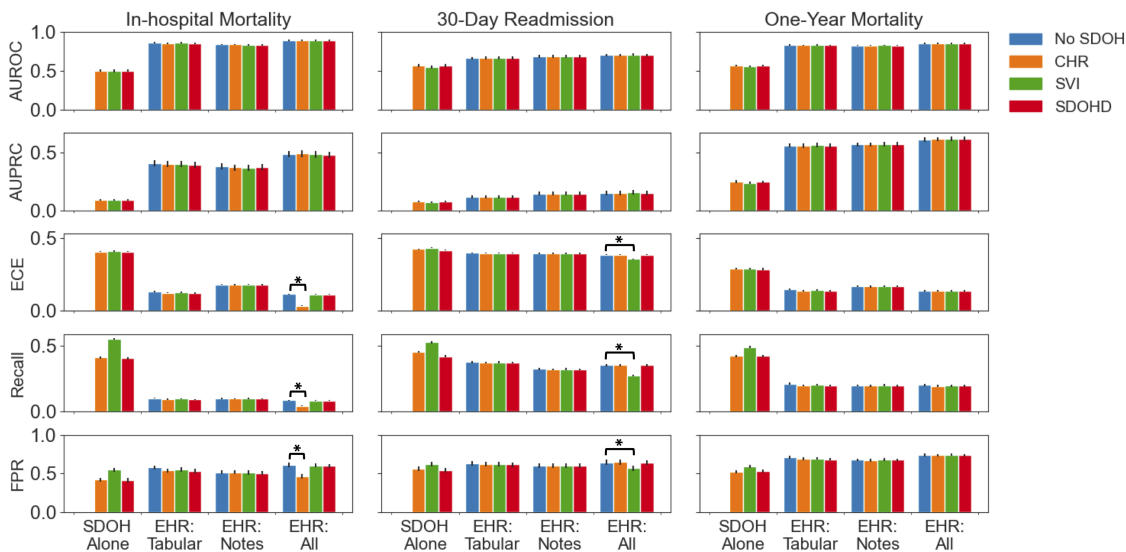


Figure 5-1: Comparison of the performance of XGBoost classifiers for predicting in-hospital mortality, 30-day readmission, and one-year mortality in the general ICU population. There are no detectable differences between AUROC and AUPRC of models that do not incorporate SDOH features and those that do not. However, when combining SDOH features with both tabular EHR data and discharge notes, we observe significant impacts on ECE, FPR, and recall. We highlight such occurrences with asterisks. The error bars denote the 95% confidence intervals obtained through 1000 bootstrap samples.

model classes and SDOH databases. This suggests that the added SDOH features do not provide additional information beyond what is already captured in the EHR. However, SDOH features have some influence on other metrics as such ECE, TPR, and recall. For example, for the task of in-hospital mortality prediction, combining CHR features with all EHR features significantly reduces ECE from 0.11 to 0.03 and FPR from 0.09 to 0.04 (Table B.2). Likewise, for the task of 30-day readmission, combining SVI features with all EHR features significantly reduces ECE from 0.39 to 0.35 and FPR from 0.35 to 0.27. However, these improvements are at the expense of a lower recall.

5.1.2 Generalizability of the Finding

Using the All of Us dataset, we validate that our finding generalizes to a more geographically diverse cohort. Unlike MIMIC-IV, which represents critically ill patient stays at a hospital in Boston, the All of Us cohort includes all in-patient stays across the United States, and many patients do not have any comorbidities (Table B.1). The AUROC of the XGBoost classifiers for All of Us is lower than that in MIMIC due to the lack of detailed, hourly lab and vitals. The AUPRC is higher for All of Us because the 30-day readmission rate is 18% in the All of Us cohort and only 6% in the MIMIC-IV cohort.

Because the two cohorts are drastically different, a direct comparison of model performance is not meaningful, but we can examine the trend in the added value of SDOH. Consistent with our results in MIMIC-IV, we see no significant performance differences in models trained with tabular EHR data and tabular EHR data combined with SDOH data in the All of Us cohort (Table 5.1).

Table 5.1: Comparison of model performance for XGBoost classifiers trained with and without SDOH for the task of 30-day readmission in MIMIC-IV and the All of Us dataset. In both datasets, the addition of SDOH has no effect on model performance.

Feature Set	MIMIC-IV					All of Us				
	AUROC	AUPRC	ECE	FPR	Recall	AUROC	AUPRC	ECE	FPR	Recall
SDOH	0.57	0.08	0.42	0.42	0.54	0.53	0.21	0.30	0.48	0.50
Tabular	0.67	0.11	0.40	0.38	0.63	0.60	0.26	0.27	0.35	0.47
Tabular+SDOH	0.67	0.11	0.40	0.37	0.62	0.60	0.26	0.27	0.34	0.46

5.2 Varying Impact of SDOH on Vulnerable Patient Populations

In this section, we investigate whether including SDOH in predictive models can lead to better performance for specific patient populations. For each task, we compare the best model that incorporates SDOH data to the best baseline model trained on only EHR data.

5.2.1 Limited Improvement in Model Performance in Vulnerable Patient Populations

Similar to the general ICU patients, we find that incorporating SDOH has some impact on model performance in the more vulnerable patient populations. Although predictive performance metrics such as AUROC and AUPRC are largely unaffected, we observe significant improvements in ECE, FPR, or recall in selected models trained on SDOH data. In Table 5.2, we report the aggregated number of occurrences in which incorporating SDOH features significantly improves or worsens model performance across three prediction tasks. See Table B.2 for more granular results.

We find that the added value of SDOH features varies by patient population, prediction task, and the EHR features they are combined with. Even for the same patient population and task, SDOH features are not equally informative or useful. For example, for models trained on all diabetic patients, the only observed performance boost is in recall when CHR features are combined with tabular features. However, the improvement in the recall is at the expense of higher ECE and FPR. In contrast, for models trained on female diabetic patients, incorporating CHR improves model performance in at least one of the three metrics when combined with discharge notes or all EHR data but not tabular EHR data alone.

This varying effect of the inclusion of SDOH by patient population precisely captures why SDOH should be collected and incorporated in analyses. Although individuals in a neighborhood are exposed to the same community-level SDOH,

Table 5.2: Combinations of EHR and SDOH features that influence performance of the best models for each patient population in terms of AUPRC. We report the number of occurrences in which incorporating SDOH features significantly improves or worsens performance in the form of (# improves/# worsens) across the three prediction tasks (total number of occurrences is 3). Significance is evaluated using a 1000-sample bootstrap hypothesis test at the 5% significance level.

Patient Population	Metric	Tabular			Notes			All		
		CHR	SVI	SDOHD	CHR	SVI	SDOHD	CHR	SVI	SDOHD
All Diabetic	ECE (↓)	0/1	0/0	0/0	0/0	0/0	0/1	0/0	0/0	0/0
	FPR (↓)	0/1	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0
	Recall (↑)	1/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0
Black Diabetic	ECE (↓)	1/0	0/0	0/0	0/0	1/0	0/1	0/0	0/1	0/1
	FPR (↓)	1/0	0/1	0/0	0/0	1/0	0/1	0/0	0/1	1/0
	Recall (↑)	0/1	1/0	0/0	0/0	0/0	0/0	0/0	1/0	0/0
Elderly Diabetic	ECE (↓)	0/0	2/0	0/0	0/0	0/0	1/0	0/0	0/0	0/0
	FPR (↓)	0/0	1/0	0/0	0/0	0/0	0/1	0/0	0/0	0/0
	Recall (↑)	0/0	0/0	0/0	0/0	0/0	1/0	0/0	0/0	0/0
Female Diabetic	ECE (↓)	0/0	0/0	0/0	1/0	0/0	0/0	0/0	0/0	0/0
	FPR (↓)	0/0	0/0	0/0	1/0	0/0	0/0	1/0	0/0	0/0
	Recall (↑)	0/0	0/0	0/0	0/1	0/0	0/0	0/0	0/0	0/0
Non-English Speaking Diabetic	ECE (↓)	0/0	1/0	1/0	0/1	0/0	0/0	0/0	0/0	1/0
	FPR (↓)	0/0	0/0	0/0	0/1	0/0	0/0	0/0	0/0	0/0
	Recall (↑)	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0

they have varying social needs [3]. Incorporating SDOH into predictive models may be helpful to identify patients with specific needs and reduce health disparities associated with poor social conditions [2, 7, 27].

5.3 SDOH As Fairness Audit Categories

A 2014 report by the National Academies of Medicine (NAM) argued that the integration of SDOH into the EHR would better enable healthcare providers to address health disparities [31]. Extending on a previous study on the integration of SDOH features and model fairness in patients with heart failure [80], we conduct a thorough audit of FPR parity in all ICU patients using SDOH features in addition to protected attributes such as race, age, and gender. To enable evaluation based on SDOH features, they are binned into quartiles, and the bin edges are documented in section A.2.

In Figure 5-2, we report the FPR of classifiers with the highest AUPRC for each of the three prediction tasks. We focus on the models trained on all EHR data (A11)

and all EHR data combined with the most helpful SDOH features for each task across different subgroups. All+CHR is the best feature set for predicting in-hospital mortality whereas All+SVI is the best feature set for predicting 30-day readmission and one-year mortality.


5.3.1 SDOH Features Enable More Granular Audits

In this setting, a high FPR indicates that the model is overdiagnosing or falsely claiming that the patient is high-risk, which has both medical and economic costs [61]. A high FPR disparity means that members of a protected subgroup would not be given the correct diagnosis or appropriate intervention at the same rate as the other patients.

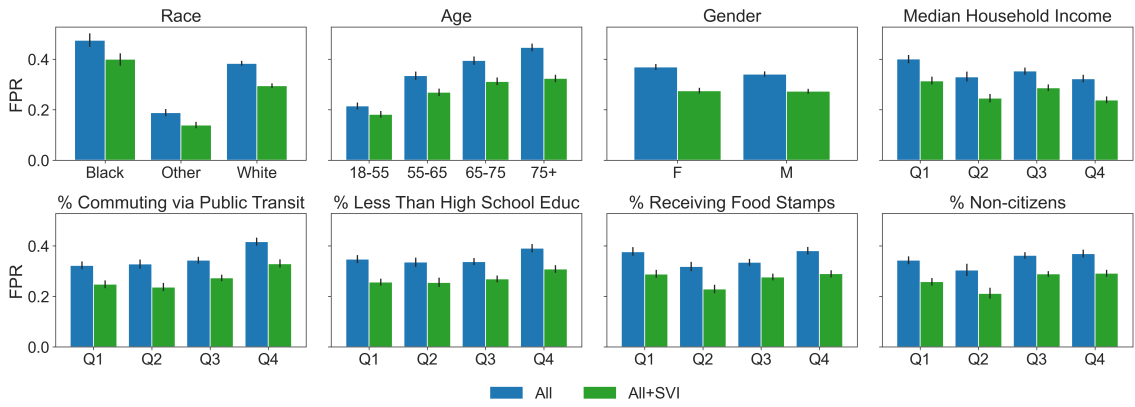
An audit of the FPR based on protective attributes confirms findings from prior work that algorithms exhibit biases against underserved patient populations [114, 136, 137]. We find that patients who are older have higher FPRs across all tasks. For the tasks of predicting in-hospital mortality and 30-day readmission, this FPR disparity becomes smaller when the model accounts for SDOH features.

In addition, we observe significant FPR disparity in the three racial groups and the two gender groups in models trained to predict 30-day readmission (Figure 5-2B). The difference in FPR between female and male patients is also observed in models trained to predict one-year mortality (Figure 5-2C). These differences are among the most commonly reported findings in health disparities research; often, these studies stop there without connecting the observed disparities to mechanisms of systemic biases that drive downstream adverse health outcomes [79]. This is partially due to the lack of additional information on the patients beyond basic demographics.


The fairness audit based on SDOH features provided additional insight and raised more questions. In the tasks of predicting 30-day readmission (Figure 5-2B) and one-year mortality (Figure 5-2C), we observe that the FPR is elevated for patients residing in communities where a larger proportion of individuals commute to work using public transit. We also find FPR disparity in patients belonging to the lower and upper quartiles of median household income (Figure 5-2B) and percentage of non-citizens (Figure 5-2C).



(A) All ICU Patients: In-hospital Mortality



(B) All ICU Patients: 30-Day Readmission



(C) All ICU Patients: One-Year Mortality

Figure 5-2: Comparison of the FPR of XGBoost classifiers trained on all EHR data (All) and all EHR data combined with the most helpful SDOH features for that task in all ICU patients in MIMIC-IV. FPR is reported for subgroups defined by race, gender, age, and five SDOH features, which are binned into quartiles. The bin edges are documented in section A.2. The error bars denote the 95% confidence intervals obtained through 1000 bootstrap samples.

We hypothesize that the FPR disparity is a result of bias propagation, which has been suggested by previous studies [1, 114]. While future work is needed to validate the hypothesis, one interpretation of the FPR disparity between patients in quartiles defined based on the percent of workers commuting via public transit is that patients in the fourth quartile likely do not own a car and hence have higher transportation barriers and limited access to healthcare [32, 35, 46, 94, 118, 122]. Additionally, a lower household income and a higher percentage of non-citizens could represent socioeconomic and linguistic disparities in access to care [25, 35, 50, 122].

Chapter 6

Discussion

6.1 The Need for Better Data

Overall, our analysis validates previous findings that community-level SDOH features do not improve the accuracy of clinical prediction tasks [28] in both a multi-year cohort and a geographically diverse cohort. We expect individual-level SDOH to be better predictors of outcomes, as prior studies that incorporated them all reported significantly improved performance [6, 95, 119]. However, this data are not readily available. For example, although individual-level SDOH can be extracted from participant surveys in the All of Us dataset, less than 15% of the participants have completed the SDOH survey. Moreover, the survey responses were collected only once for each participant, so these survey-based SDOH features may not accurately reflect the lived experience of the respondents beyond the period the survey was conducted. In light of our findings, we call for further efforts to standardize the routine collection of SDOH data and integration into EHR.

The healthcare system plays a vital role in collecting, using, and sharing actionable SDOH data [96]. To facilitate this effort, providers and operations staff across care settings should focus on actions that enhance the standardization and integration of SDOH data. Organizations such as the Office of National Coordinator for Health IT (ONC), the Joint Commission, and Health Level Seven International (HL7) are all leading efforts to further SDOH interoperability and standards [99, 109]. It should also be a focus to provide sufficient training and education

for the staff who are collecting and encoding the data from the patients while adhering to cultural competency, privacy, and confidentiality standards [87].

As the research community awaits access to the more granular EHR-SDOH data, we hope the MIMIC-IV-SDOH database will serve as a starting point for studies on the relationship between community risk factors and patient outcomes and those looking to understand the needs of vulnerable subpopulations.

6.2 On More Actionable Audits

Despite spending a higher percentage of our GDP on medical care expenditures than other developed countries, health outcomes in the United States are among the worst for developed countries [98]. Numerous studies have confirmed the potential of AI in improving health outcomes, but very few tools that were developed have actually helped [9, 132]. A promising direction forward is to look beyond the clinical walls and understand the conditions that affect the health of the people upstream [86].

The growing evidence around the association between SDOH and health outcomes calls for targeted action, but there is a lack of consensus on what interventions would work [86]. Progress in evidence-informed policymaking requires a commitment to enhancing our current understanding of how SDOH affect different populations and ways to measure the effectiveness of interventions targeting specific SDOH domains. Thus, community-level SDOH features are essential for evaluating and monitoring health disparities [108].

By evaluating fairness using intersectional social identities, we could better account for the socially constructed nature of protected attributes such as race and gender. Capturing SDOH provides information on the social processes that created health disparities in the first place [25, 52, 79]. Audits of biases from the lens of SDOH are also more actionable because these features are not social constructs but modifiable factors that can be addressed [108]. Consider transportation, which is one of the SDOH features we used in the fairness audit. Surveys and audits have identified transportation barriers as one of the leading causes of missed or delayed medical appointments, especially in the elderly and those in rural areas [57, 118, 122]. Health

insurance and healthcare delivery organizations are addressing the issue through partnerships with popular ride-share companies to provide non-emergency medical transportation (NEMT) services [104, 124]. These programs have decreased costs [105] and the frequency of urgent care visits [100]. The development of this intervention would not be possible without an understanding of the underlying SDOH and the population affected.

6.3 Future Work on SDOH and Health Predictions

While our work shows that the inclusion of SDOH has minimal impact on three common EHR prediction tasks, they could be more helpful in other tasks and patient groups. Specifically, we did not include any phenotype prediction tasks. Given the associations between SDOH and chronic diseases [6, 30], it is possible that SDOH features are good risk predictors for specific comorbidities. In addition, SDOH could be important to account for in the estimation of treatment effects, which several studies have done using the MIMIC database but without SDOH data [65, 75, 139]. Likewise, although our study utilized three different model classes, they are all relatively simple. Neural networks could potentially uncover more underlying relationships between SDOH and health outcomes [140].

Regardless of the predictive value of SDOH, it is a good idea to account for them in analyses for more granular benchmarking and evaluation of fairness. For example, MIMIC-IV-SDOH can be mapped to MIMIC-CXR, a large database of chest radiographs with radiology reports. There have been many works that focus on group fairness in the field of medical imaging [113, 114, 136], which the inclusion of SDOH could contribute to.

Chapter 7

Conclusion

This work advances our understanding of the impact of SDOH on health prediction. First, we develop a new EHR-SDOH database by linking a popular EHR database, MIMIC-IV, to public community-level SDOH databases. This database can be used to uncover underlying trends between community health and individual health outcomes and provide more benchmarks for evaluating bias and fairness. Second, we demonstrate that incorporating SDOH features in certain vulnerable subgroups can improve model performance. The value of adding SDOH features, however, is dependent on the cohort characteristics and the prediction task. Third, we highlight that algorithmic audits conducted through the lens of SDOH are more comprehensive and actionable. However, the lack of access to high-resolution, individual SDOH data is a limitation of the study. To address this, future work should focus on collecting individual-level SDOH features and accounting for them in analyses to better address patient needs and promote health equity.

Appendix A

Additional Information on Model Training and Data Pre-processing

A.1 Model Training

Due to the class imbalance in all the prediction tasks, we use AUROC for model selection during hyperparameter tuning. The distributions of parameters sampled during the randomized search for logistic regression (`lr`), random forest (`rf`), and XGBoost (`xgb`) classifiers are as followed:

```
lr_param_grid = {
    "C" : [1e-5, 1e-4, 1e-3, 1e-2, 0.1, 1, 5],
    "solver" : ["liblinear"],
}
```

```
rf_param_grid = {
    "n_estimators": [50, 100, 200, 500],
    "max_depth": scipy.stats.randint(2, 10),
    "min_samples_split": scipy.stats.randint(2, 10),
    "min_samples_leaf": scipy.stats.randint(1, 10),
}
```

```
xgb_param_grid = {
```



```
"n_estimators": [50,100,200,500],
"max_depth": scipy.stats.randint(2, 10),
"learning_rate": (0.01,0.05,0.1,0.2,0.3),
"min_child_weight": scipy.stats.randint(2, 10),
"colsample_bytree": [0.5,1],
"subsample" : [0.3,0.6,0.9],
"reg_alpha" : scipy.stats.randint(0, 10),
"reg_lambda": scipy.stats.randint(0, 10),
}
```

A.2 Binning SDOHD Features

The quartile bin edges for SDOH features used in the fairness audit are as followed:

1. Percentage of non-citizens:
0, 0.32, 1.09, 2.54, 30.58
2. Median household income in dollars:
10446, 60698.5, 74902, 92381, 250001
3. Percentage with less than high school education:
0, 4.04, 6.79, 11.64, 67.49
4. Percentage of households receiving food stamps:
0, 4.15, 6.85, 12.1, 78.43
5. Percentage of workers taking public transit:
0, 4.74, 10.5, 20.44, 77.61

A.3 SDOH Features in the All of Us Dataset

The community-SDOH features used include the following:

1. Percentage of households receiving food stamps

2. Percentage of the population with at least a high school diploma
3. Median household income
4. Percentage of the population with no health insurance coverage
5. Percentage of the population with income below the poverty level
6. Percentage of houses that are vacant
7. Deprivation index

Appendix B

Tables

Table B.1: Characteristics of the MIMIC-IV and All of Us cohorts. The MIMIC-IV cohort has six patient populations: all ICU patients, diabetic patients, black diabetic patients, elderly diabetic patients, female diabetic patients, and non-English speaking diabetic patients. The All of Us cohort contains all in-patient hospital stays. N is the number of patients in each group.

Attribute	Subgroup	MIMIC-IV						All of Us
		All ICU (N=42,665)	All Diab. (N=12,651)	Black Diab. (N=1,710)	Elderly Diab. (N=4,520)	Female Diab. (N=5,251)	Non-English Speaking Diab. (N=1,675)	All Inpatient (N=13,324)
Age	17-55	10,136 (23%)	1,806 (14%)	337 (19%)	0 (0%)	681 (12%)	158 (9%)	6,481 (48%)
	55-65	8,773 (20%)	2,759 (21%)	407 (23%)	0 (0%)	1,008 (19%)	332 (19%)	3,134 (23%)
	65-75	10,013 (23%)	3,706 (29%)	473 (27%)	0 (0%)	1,457 (27%)	415 (24%)	2,430 (18%)
	75+	13,742 (32%)	4,380 (34%)	493 (28%)	4,520 (100%)	2,105 (40%)	770 (45%)	1,279 (9%)
Gender	Female	18,677 (43%)	5,251 (41%)	927 (54%)	2,169 (47%)	5,251 (100%)	801 (47%)	8,226 (61%)
	Male	23,988 (56%)	7,400 (58%)	783 (45%)	2,351 (52%)	0 (0%)	874 (52%)	4,800 (36%)
	Other	–	–	–	–	–	–	298 (2%)
Race	White	29,148 (68%)	8,033 (63%)	0 (0%)	3,033 (67%)	3,150 (59%)	455 (27%)	6,168 (46%)
	Black	3,880 (9%)	1,700 (13%)	1,710 (100%)	523 (11%)	920 (17%)	192 (11%)	3,589 (26%)
	Other	9,637 (22%)	2,918 (23%)	0 (0%)	964 (21%)	1,181 (22%)	1,028 (61%)	3,567 (26%)
Insurance Type	Medicaid	2,976 (6%)	763 (6%)	153 (8%)	68 (1%)	370 (7%)	249 (14%)	7,109 (53%)
	Medicare	18,844 (44%)	6,453 (51%)	781 (45%)	3,169 (70%)	2,828 (53%)	721 (43%)	3,270 (24%)
	Other	20,845 (48%)	5,435 (42%)	776 (45%)	1,283 (28%)	2,053 (39%)	705 (42%)	2,945 (22%)
Language	English	38,291 (89%)	11,018 (87%)	1,523 (89%)	3,744 (82%)	4,463 (84%)	0 (0%)	–
	Other	4,374 (10%)	1,633 (12%)	187 (10%)	776 (17%)	788 (15%)	1,675 (100%)	–
Charlson Comorbidity Index	0	2,856 (6%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	5,977 (44%)
	1	3,208 (7%)	301 (2%)	63 (3%)	0 (0%)	110 (2%)	27 (1%)	1,315 (9%)
	2	4,232 (9%)	626 (4%)	106 (6%)	0 (0%)	254 (4%)	56 (3%)	1,240 (9%)
	3	5,151 (12%)	1,039 (8%)	131 (7%)	0 (0%)	397 (7%)	109 (6%)	1,174 (8%)
	4	5,576 (13%)	1,434 (11%)	176 (10%)	173 (3%)	561 (10%)	174 (10%)	855 (6%)
	5	5,329 (12%)	1,793 (14%)	212 (12%)	511 (11%)	777 (14%)	245 (14%)	651 (4%)
	6	4,558 (10%)	1,713 (13%)	215 (12%)	653 (14%)	720 (13%)	230 (13%)	554 (4%)
7+	11,755 (27%)	5,745 (45%)	807 (47%)	3,183 (70%)	2,432 (46%)	834 (49%)	1,558 (11%)	

Table B.2: Comparison of the performance of models trained with and without SDOH feature to predict in-hospital mortality, 30-day readmission, and one-year mortality for the six patient populations. The evaluation is done only on the best model for each feature set category and task. In general, incorporating SDOH features has a limited impact on model performance. Values in bold indicate significantly better performance compared to the baseline model trained without SDOH, evaluated using a 1000-sample bootstrap hypothesis test at the 5% significance level.

Patient Group	Feature Set	In-hospital Mortality					30-Day Readmission					One-Year Mortality				
		AUROC (↑)	AUPRC (↑)	ECE (↓)	FPR (↓)	Recall (↑)	AUROC (↑)	AUPRC (↑)	ECE (↓)	FPR (↓)	Recall (↑)	AUROC (↑)	AUPRC (↑)	ECE (↓)	FPR (↓)	Recall (↑)
All ICU	Tabular	0.86 ³	0.4 ³	0.13 ³	0.1 ³	0.58 ³	0.66 ³	0.12 ³	0.4 ³	0.37 ³	0.58 ³	0.83 ³	0.56 ³	0.15 ³	0.21 ³	0.71 ³
	Tabular+SDOH	0.86 ²	0.4 ²	0.13 ²	0.1 ²	0.56 ²	0.67 ³	0.12 ³	0.39 ³	0.36 ³	0.62 ³	0.83 ²	0.56 ²	0.14 ²	0.2 ²	0.69 ²
	Notes	0.84 ³	0.38 ³	0.18 ³	0.1 ³	0.52 ³	0.69 ³	0.14 ³	0.39 ³	0.32 ³	0.6 ³	0.83 ³	0.57 ³	0.17 ³	0.2 ³	0.68 ³
	Notes+SDOH	0.84 ³	0.37 ³	0.18 ³	0.1 ³	0.51 ³	0.69 ³	0.14 ³	0.39 ³	0.32 ³	0.6 ³	0.83 ²	0.57 ²	0.17 ²	0.2 ²	0.68 ²
	All	0.9 ³	0.49 ³	0.11 ³	0.09 ³	0.61 ³	0.7 ³	0.15 ³	0.39 ³	0.35 ³	0.65 ³	0.85 ³	0.61 ³	0.14 ³	0.2 ³	0.74 ³
	All+SDOH	0.89 ¹	0.49 ¹	0.03¹	0.04¹	0.47 ¹	0.71 ²	0.15 ²	0.35²	0.27²	0.57 ²	0.86 ²	0.62 ²	0.14 ²	0.2 ²	0.74 ²
All Diabetic	Tabular	0.86 ³	0.4 ³	0.09 ³	0.07 ³	0.45 ³	0.65 ³	0.15 ³	0.38 ³	0.37 ³	0.55 ³	0.78 ³	0.52 ³	0.13 ³	0.23 ³	0.62 ³
	Tabular+SDOH	0.85 ²	0.39 ²	0.08 ²	0.06 ²	0.42 ²	0.64 ²	0.14 ²	0.38 ²	0.36 ²	0.58 ²	0.78 ¹	0.52 ¹	0.19 ¹	0.32 ¹	0.74¹
	Notes	0.82 ³	0.33 ³	0.21 ³	0.18 ³	0.65 ³	0.67 ³	0.17 ³	0.37 ³	0.3 ³	0.55 ³	0.78 ³	0.51 ³	0.14 ³	0.19 ³	0.58 ³
	Notes+SDOH	0.81 ¹	0.32 ¹	0.21 ¹	0.18 ¹	0.63 ¹	0.66 ³	0.17 ³	0.39 ³	0.3 ³	0.55 ³	0.78 ¹	0.52 ¹	0.13 ¹	0.18 ¹	0.56 ¹
	All	0.89 ³	0.47 ³	0.07 ³	0.05 ³	0.48 ³	0.68 ³	0.18 ³	0.37 ³	0.3 ³	0.54 ³	0.81 ³	0.57 ³	0.11 ³	0.2 ³	0.65 ³
	All+SDOH	0.89 ²	0.47 ²	0.07 ²	0.05 ²	0.47 ²	0.67 ²	0.18 ²	0.37 ²	0.31 ²	0.52 ²	0.81 ²	0.57 ²	0.11 ²	0.19 ²	0.64 ²
Black Diabetic	Tabular	0.83 ³	0.36 ³	0.27 ³	0.22 ³	0.68 ³	0.59 ³	0.14 ³	0.37 ³	0.19 ³	0.31 ³	0.74 ³	0.49 ³	0.18 ³	0.27 ³	0.59 ³
	Tabular+SDOH	0.83 ²	0.35 ²	0.26 ²	0.2 ²	0.67 ²	0.59 ²	0.16 ²	0.35 ²	0.38 ²	0.56²	0.72 ¹	0.46 ¹	0.08¹	0.17¹	0.41 ¹
	Notes	0.76 ³	0.15 ³	0.16 ³	0.06 ³	0.19 ³	0.6 ³	0.17 ³	0.3 ³	0.01 ³	0.02 ³	0.75 ³	0.53 ³	0.21 ³	0.09 ³	0.35 ³
	Notes+SDOH	0.76 ²	0.19 ²	0.05²	0.01²	0.07 ²	0.57 ³	0.15 ³	0.26 ³	0.17 ³	0.2³	0.74 ¹	0.52 ¹	0.21 ¹	0.08 ¹	0.32 ¹
	All	0.83 ³	0.34 ³	0.27 ³	0.21 ³	0.68 ³	0.58 ³	0.15 ³	0.26 ³	0.17 ³	0.24 ³	0.79 ³	0.59 ³	0.21 ³	0.13 ³	0.52 ³
	All+SDOH	0.82 ²	0.33 ²	0.26 ²	0.21 ²	0.66 ²	0.6 ³	0.18 ³	0.37 ³	0.06³	0.14 ³	0.78 ²	0.58 ²	0.27 ²	0.26 ²	0.66²
Elderly Diabetic	Tabular	0.79 ³	0.38 ³	0.16 ³	0.06 ³	0.35 ³	0.63 ³	0.14 ³	0.37 ³	0.36 ³	0.58 ³	0.73 ³	0.55 ³	0.14 ³	0.31 ³	0.64 ³
	Tabular+SDOH	0.8 ²	0.37 ²	0.08²	0.06 ²	0.36 ²	0.63 ¹	0.16 ¹	0.37 ¹	0.35 ¹	0.53 ¹	0.73 ²	0.54 ²	0.08²	0.25²	0.55 ²
	Notes	0.76 ³	0.29 ³	0.09 ³	0.03 ³	0.17 ³	0.57 ³	0.1 ³	0.37 ³	0.03 ³	0.05 ³	0.72 ³	0.54 ³	0.09 ³	0.23 ³	0.51 ³
	Notes+SDOH	0.76 ³	0.29 ³	0.09 ³	0.03 ³	0.18 ³	0.55 ³	0.1 ³	0.33³	0.2 ³	0.24³	0.72 ²	0.55 ²	0.09 ²	0.23 ²	0.53 ²
	All	0.84 ³	0.43 ³	0.05 ³	0.04 ³	0.31 ³	0.61 ³	0.12 ³	0.25 ³	0.18 ³	0.33 ³	0.76 ³	0.59 ³	0.07 ³	0.22 ³	0.56 ³
	All+SDOH	0.84 ¹	0.43 ¹	0.05 ¹	0.04 ¹	0.31 ¹	0.6 ¹	0.12 ¹	0.25 ¹	0.2 ¹	0.36 ¹	0.76 ²	0.59 ²	0.07 ²	0.21 ²	0.58 ²
Female Diabetic	Tabular	0.84 ³	0.4 ³	0.26 ³	0.22 ³	0.73 ³	0.61 ³	0.13 ³	0.37 ³	0.36 ³	0.55 ³	0.75 ³	0.44 ³	0.2 ³	0.32 ³	0.7 ³
	Tabular+SDOH	0.83 ²	0.38 ²	0.26 ²	0.22 ²	0.71 ²	0.57 ¹	0.12 ¹	0.35 ¹	0.34 ¹	0.47 ¹	0.75 ³	0.45 ³	0.22 ³	0.35 ³	0.71 ³
	Notes	0.76 ³	0.28 ³	0.07 ³	0.02 ³	0.12 ³	0.62 ³	0.15 ³	0.4 ³	0.15 ³	0.27 ³	0.75 ³	0.44 ³	0.21 ³	0.3 ³	0.68 ³
	Notes+SDOH	0.76 ²	0.27 ²	0.07 ²	0.01 ²	0.11 ²	0.61 ¹	0.15 ¹	0.39 ¹	0.13 ¹	0.23 ¹	0.74 ¹	0.44 ¹	0.12¹	0.15¹	0.41 ¹
	All	0.86 ³	0.44 ³	0.04 ³	0.03 ³	0.38 ³	0.62 ³	0.16 ³	0.39 ³	0.16 ³	0.26 ³	0.77 ³	0.49 ³	0.11 ³	0.18 ³	0.53 ³
	All+SDOH	0.86 ²	0.44 ²	0.04 ²	0.03 ²	0.37 ²	0.62 ¹	0.14 ¹	0.38 ¹	0.08¹	0.17 ¹	0.77 ³	0.49 ³	0.1 ³	0.17 ³	0.52 ³
Non-English Speaking Diabetic	Tabular	0.8 ³	0.34 ³	0.27 ³	0.24 ³	0.68 ³	0.61 ³	0.12 ³	0.36 ³	0.27 ³	0.46 ³	0.8 ³	0.55 ³	0.13 ³	0.13 ³	0.51 ³
	Tabular+SDOH	0.79 ²	0.35 ²	0.26 ²	0.23 ²	0.71 ²	0.58 ³	0.13 ³	0.3³	0.31 ³	0.48 ³	0.79 ²	0.54 ²	0.07²	0.14 ²	0.51 ²
	Notes	0.74 ³	0.27 ³	0.05 ³	0.02 ³	0.08 ³	0.63 ³	0.11 ³	0.29 ³	0.18 ³	0.27 ³	0.71 ³	0.48 ³	0.13 ³	0.18 ³	0.46 ³
	Notes+SDOH	0.75 ¹	0.26 ¹	0.05 ¹	0.01 ¹	0.06 ¹	0.56 ¹	0.09 ¹	0.37 ¹	0.32 ¹	0.38 ¹	0.72 ²	0.51 ²	0.13 ²	0.18 ²	0.49 ²
	All	0.82 ³	0.38 ³	0.09 ³	0.11 ³	0.44 ³	0.61 ³	0.12 ³	0.36 ³	0.27 ³	0.46 ³	0.77 ³	0.57 ³	0.12 ³	0.17 ³	0.55 ³
	All+SDOH	0.81 ²	0.39 ²	0.1 ²	0.11 ²	0.41 ²	0.58 ³	0.13 ³	0.3³	0.31 ³	0.48 ³	0.77 ²	0.56 ²	0.12 ²	0.18 ²	0.56 ²

¹ Trained on CHR, ² Trained on SVI, ³ Trained on SDOHD

Bibliography

- [1] Hammaad Adam, Ming Ying Yang, Kenrick Cato, Ioana Baldini, Charles Senteio, Leo Anthony Celi, Jiaming Zeng, Moninder Singh, and Marzyeh Ghassemi. Write It Like You See It: Detectable Differences in Clinical Notes by Race Lead to Differential Model Recommendations. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 7–21, Oxford United Kingdom, July 2022. ACM.
- [2] Nancy E. Adler, M. Maria Glymour, and Jonathan Fielding. Addressing Social Determinants of Health and Health Inequalities. *JAMA*, 316(16):1641–1642, October 2016.
- [3] Agency for Healthcare Research and Quality. About SDOH in Healthcare, February 2020.
- [4] Agency for Healthcare Research and Quality. Social Determinants of Health Database, August 2022.
- [5] American Hospital Association. ICD-10-CM Coding for Social Determinants of Health. Technical report, American Hospital Association, January 2022.
- [6] Fatemeh Amrollahi, Supreeth P Shashikumar, Angela Meier, Lucila Ohno-Machado, Shamim Nemati, and Gabriel Wardi. Inclusion of social determinants of health improves sepsis readmission prediction models. *Journal of the American Medical Informatics Association*, 29(7):1263–1270, July 2022.
- [7] Anne Andermann. Taking action on the social determinants of health in clinical practice: a framework for health professionals. *CMAJ : Canadian Medical Association Journal*, 188(17-18):E474–E483, December 2016.
- [8] Shervin Assari, Brianna Preiser, and Marisa Kelly. Education and Income Predict Future Emotional Well-Being of Whites but Not Blacks: A Ten-Year Cohort. *Brain Sciences*, 8(7):122, June 2018.
- [9] Avi Goldfarb and Florenta Teodoridis. Why is AI adoption in health care lagging?, March 2022.
- [10] Rosevine A. Azap, Anghela Z. Paredes, Adrian Diaz, J. Madison Hyer, and Timothy M. Pawlik. The association of neighborhood social vulnerability with surgical textbook outcomes among patients undergoing hepatopancreatic surgery. *Surgery*, 168(5):868–875, November 2020.

- [11] Aparna Balagopalan, Haoran Zhang, Kimia Hamidieh, Thomas Hartvigsen, Frank Rudzicz, and Marzyeh Ghassemi. The Road to Explainability is Paved with Bias: Measuring the Fairness of Explanations. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1194–1206, June 2022. arXiv:2205.03295 [cs].
- [12] Andrew W Bazemore, Erika K Cottrell, Rachel Gold, Lauren S Hughes, Robert L Phillips, Heather Angier, Timothy E Burdick, Mark A Carrozza, and Jennifer E DeVoe. “Community vital signs” : incorporating geocoded social determinants into electronic records to promote patient and population health. *Journal of the American Medical Informatics Association*, 23(2):407–412, March 2016.
- [13] Anas Belouali, Haibin Bai, Kanimozhi Raja, Star Liu, Xiyu Ding, and Hadi Kharrazi. Impact of social determinants of health on improving the LACE index for 30-day unplanned readmission prediction. *JAMIA open*, 5(2):oac046, July 2022.
- [14] James Bergstra and Yoshua Bengio. Random Search for Hyper-Parameter Optimization. *Journal of Machine Learning Research*, 13(10):281–305, 2012.
- [15] Nrupen A. Bhavsar, Aijing Gao, Matthew Phelan, Neha J. Pagidipati, and Benjamin A. Goldstein. Value of Neighborhood Socioeconomic Status in Predicting Risk of Outcomes in Studies That Use Electronic Health Record Data. *JAMA Network Open*, 1(5):e182716, September 2018.
- [16] Willie Boag, Harini Suresh, L. Celi, Peter Szolovits, and M. Ghassemi. Racial Disparities and Mistrust in End-of-Life Care. August 2018.
- [17] Luisa N. Borrell, Jennifer R. Elhawary, Elena Fuentes-Afflick, Jonathan Witonsky, Nirav Bhakta, Alan H.B. Wu, Kirsten Bibbins-Domingo, José R. Rodríguez-Santana, Michael A. Lenoir, James R. Gavin, Rick A. Kittles, Noah A. Zaitlen, David S. Wilkes, Neil R. Powe, Elad Ziv, and Esteban G. Burchard. Race and Genetic Ancestry in Medicine — A Time for Reckoning with Racism. *New England Journal of Medicine*, 384(5):474–480, February 2021.
- [18] Paula Braveman and Laura Gottlieb. The Social Determinants of Health: It’s Time to Consider the Causes of the Causes. *Public Health Reports*, 129(1_suppl2):19–31, January 2014. Publisher: SAGE Publications Inc.
- [19] Ann-Sylvia Brooker and Joan M. Eakin. Gender, class, work-related stress and health: toward a power-centred approach. *Journal of Community & Applied Social Psychology*, 11(2):97–109, 2001. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/casp.620>.
- [20] Jeremiah R. Brown, Iben M. Ricket, Ruth M. Reeves, Rashmee U. Shah, Christine A. Goodrich, Glen Gobbel, Meagan E. Stabler, Amy M. Perkins, Freneka

- Minter, Kevin C. Cox, Chad Dorn, Jason Denton, Bruce E. Bray, Ramkiran Gouripeddi, John Higgins, Wendy W. Chapman, Todd MacKenzie, and Michael E. Matheny. Information Extraction From Electronic Health Records to Predict Readmission Following Acute Myocardial Infarction: Does Natural Language Processing Using Clinical Notes Improve Prediction of Readmission? *Journal of the American Heart Association*, 11(7):e024198, April 2022.
- [21] Michael N. Cantor and Lorna Thorpe. Integrating Data On Social Determinants Of Health Into Electronic Health Records. *Health Affairs*, 37(4):585–590, April 2018. Publisher: Health Affairs.
- [22] Patrick M. Carter and Marc A. Zimmerman. Addressing Residential Segregation as a Social Determinant of Health. *JAMA Network Open*, 5(6):e2222619, June 2022.
- [23] CDC/ATSDR. CDC/ATSDR Social Vulnerability Index (SVI), October 2022.
- [24] Centers for Medicare & Medicaid Services. Electronic Health Records, December 2021.
- [25] Cindy D. Chang. Social Determinants of Health and Health Disparities Among Immigrants and their Children. *Current Problems in Pediatric and Adolescent Health Care*, 49(1):23–30, January 2019.
- [26] Irene Chen, Fredrik D. Johansson, and David Sontag. Why Is My Classifier Discriminatory?, December 2018. arXiv:1805.12002 [cs, stat].
- [27] Irene Y. Chen, Peter Szolovits, and Marzyeh Ghassemi. Can AI Help Reduce Disparities in General Medical and Mental Health Care? *AMA Journal of Ethics*, 21(2):167–179, February 2019. Publisher: American Medical Association.
- [28] Min Chen, Xuan Tan, and Rema Padman. Social determinants of health in electronic health records and their impact on analysis and risk prediction: A systematic review. *Journal of the American Medical Informatics Association: JAMIA*, 27(11):1764–1773, November 2020.
- [29] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA, August 2016. Association for Computing Machinery.
- [30] M. H. Chin and L. Goldman. Correlates of early hospital readmission or death in patients with congestive heart failure. *The American Journal of Cardiology*, 79(12):1640–1644, June 1997.
- [31] Committee on the Recommended Social and Behavioral Domains and Measures for Electronic Health Records, Board on Population Health and Public Health Practice, and Institute of Medicine. *Capturing Social and Behavioral Domains and Measures in Electronic Health Records: Phase 2*. National Academies Press (US), Washington (DC), 2015.

- [32] Richard A. Cooper, Matthew A. Cooper, Emily L. McGinley, Xiaolin Fan, and J. Thomas Rosenthal. Poverty, Wealth, and Health Care Utilization: A Geographic Assessment. *Journal of Urban Health*, 89(5):828–847, October 2012.
- [33] Giselle Corbie-Smith, Stephen B. Thomas, and Diane Marie M. St. George. Distrust, Race, and Research. *Archives of Internal Medicine*, 162(21):2458–2463, November 2002.
- [34] County Health Rankings & Roadmaps. Massachusetts Data and Resources, 2022.
- [35] Sergio Cristancho, D. Marcela Garces, Karen E. Peters, and Benjamin C. Mueller. Listening to rural Hispanic immigrants in the Midwest: a community-based participatory assessment of major barriers to health care access and use. *Qualitative Health Research*, 18(5):633–646, May 2008.
- [36] David Cutler and Adriana Lleras-Muney. Education and Health: Evaluating Theories and Evidence. Technical Report w12352, National Bureau of Economic Research, Cambridge, MA, July 2006.
- [37] Sally M Davis and Raymond Reid. Practicing participatory research in American Indian communities. *The American Journal of Clinical Nutrition*, 69(4):755S–759S, April 1999.
- [38] Angela G. E. M de Boer, Wouter Wijker, and Hanneke C. J. M de Haes. Predictors of health care utilization in the chronically ill: a review of the literature. *Health Policy*, 42(2):101–115, November 1997.
- [39] Thomas J. DiCiccio and Bradley Efron. Bootstrap confidence intervals. *Statistical Science*, 11(3):189–228, September 1996. Publisher: Institute of Mathematical Statistics.
- [40] Alexander Din and Ron Wilson. Crosswalking ZIP Codes to Census Geographies: Geoprocessing the U.S. Department of Housing & Urban Development’s ZIP Code Crosswalk Files. *Cityscape*, 22(1):293–314, 2020. Publisher: US Department of Housing and Urban Development.
- [41] Aline Dugravot, Aurore Fayosse, Julien Dumurgier, Kim Bouillon, Tesnim Ben Rayana, Alexis Schnitzler, Mika Kivimaki, Séverine Sabia, and Archana Singh-Manoux. Social inequalities in multimorbidity, frailty, disability, and transitions to mortality: a 24-year follow-up of the Whitehall II cohort study. *The Lancet Public Health*, 5(1):e42–e50, January 2020. Publisher: Elsevier.
- [42] Annette Dula. African American Suspicion of the Healthcare System Is Justified: What Do We Do about It? *Cambridge Quarterly of Healthcare Ethics*, 3(3):347–357, 1994. Publisher: Cambridge University Press.

- [43] Daniel J. Feller, Jason Zucker, Michael T. Yin, Peter Gordon, and Noémie Elhadad. Using Clinical Notes and Natural Language Processing for Automated HIV Risk Assessment. *Journal of Acquired Immune Deficiency Syndromes (1999)*, 77(2):160–166, February 2018.
- [44] Lior Fuchs, Catherine E. Chronaki, Shinhyuk Park, Victor Novack, Yael Baumfeld, Daniel Scott, Stuart McLennan, Daniel Talmor, and Leo Celi. ICU admission characteristics and mortality rates among elderly and very elderly patients. *Intensive Care Medicine*, 38(10):1654–1661, October 2012.
- [45] Bruna Galobardes, George Davey Smith, and John W. Lynch. Systematic Review of the Influence of Childhood Socioeconomic Circumstances on Risk for Cardiovascular Disease in Adulthood. *Annals of Epidemiology*, 16(2):91–104, February 2006.
- [46] Maria G. Garcia Popa-Lisseanu, Anthony Greisinger, Marsha Richardson, Kimberly J. O’Malley, Namieta M. Janssen, Donald M. Marcus, Jasmeet Tagore, and Maria E. Suarez-Almazor. Determinants of treatment adherence in ethnically diverse, economically disadvantaged patients with rheumatic disease. *The Journal of Rheumatology*, 32(5):913–919, May 2005.
- [47] Marzyeh Ghassemi, Tristan Naumann, Finale Doshi-Velez, Nicole Brimmer, Rohit Joshi, Anna Rumshisky, and Peter Szolovits. Unfolding Physiological State: Mortality Modelling in Intensive Care Units. *KDD : proceedings / International Conference on Knowledge Discovery & Data Mining. International Conference on Knowledge Discovery & Data Mining*, 2014:75–84, August 2014.
- [48] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C. K. Peng, and H. E. Stanley. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation*, 101(23):E215–220, June 2000.
- [49] Sara Nouri Golmaei and Xiao Luo. DeepNote-GNN: predicting hospital readmission using clinical notes and patient network. In *Proceedings of the 12th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics*, BCB ’21, pages 1–9, New York, NY, USA, August 2021. Association for Computing Machinery.
- [50] Daniel J. Gottlieb, Alexa S. Beiser, and George T. O’Connor. Poverty, Race, and Medication Use Are Correlates of Asthma Hospitalization Rates: A Small Area Analysis in Boston. *Chest*, 108(1):28–35, July 1995.
- [51] Paulina Grnarova, Florian Schmidt, Stephanie L. Hyland, and Carsten Eickhoff. Neural Document Embeddings for Intensive Care Patient Mortality Prediction, December 2016. arXiv:1612.00467 [cs].
- [52] Robert A Hahn and Benedict I Truman. Education Improves Public Health and Promotes Health Equity. *International journal of health services*, 45(4):657–678, January 2015.

- [53] Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. Towards a critical race methodology in algorithmic fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, pages 501–512, New York, NY, USA, January 2020. Association for Computing Machinery.
- [54] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [55] Hrayr Harutyunyan, Hrant Khachatrian, David C. Kale, Greg Ver Steeg, and Aram Galstyan. Multitask learning and benchmarking with clinical time series data. *Scientific Data*, 6(1):96, June 2019. Number: 1 Publisher: Nature Publishing Group.
- [56] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness Without Demographics in Repeated Loss Minimization. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1929–1938. PMLR, July 2018. ISSN: 2640-3498.
- [57] Health Research & Educational Trust. Social Determinants of Health Series: Transportation and the Role of Hospitals, November 2017.
- [58] Healthy People 2030. Social Determinants of Health - Healthy People 2030 | health.gov, August 2020.
- [59] Paul L. Hebert, Elizabeth A. Howell, Edwin S. Wong, Susan E. Hernandez, Seppo T. Rinne, Christine A. Sulc, Emily L. Neely, and Chuan-Fen Liu. Methods for Measuring Racial Differences in Hospitals Outcomes Attributable to Disparities in Use of High-Quality Hospital Care. *Health Services Research*, 52(2):826–848, April 2017.
- [60] Felicia Hill-Briggs, Patti L. Ephraim, Elizabeth A. Vraney, Karina W. Davidson, Renee Pekmezaris, Debbie Salas-Lopez, Catherine M. Alfano, and Tiffany L. Gary-Webb. Social Determinants of Health, Race, and Diabetes Population Health Improvement: Black/African Americans as a Population Exemplar. *Current Diabetes Reports*, 22(3):117–128, 2022.
- [61] Jerome R. Hoffman and Richelle J. Cooper. Overdiagnosis of Disease: A Modern Epidemic. *Archives of Internal Medicine*, 172(15):1123–1124, August 2012.
- [62] Anna Lauren Hoffmann. Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society*, 22(7):900–915, June 2019. Publisher: Routledge _eprint: <https://doi.org/10.1080/1369118X.2019.1573912>.
- [63] Carlyn M. Hood, Keith P. Gennuso, Geoffrey R. Swain, and Bridget B. Catlin. County Health Rankings: Relationships Between Determinant Factors and

- Health Outcomes. *American Journal of Preventive Medicine*, 50(2):129–135, February 2016.
- [64] Carrie R. Howell, Li Zhang, Nengjun Yi, Tapan Mehta, Andrea L. Cherrington, and W. Timothy Garvey. Associations between cardiometabolic disease severity, social determinants of health (SDoH), and poor COVID-19 outcomes. *Obesity (Silver Spring, Md.)*, 30(7):1483–1494, July 2022.
- [65] Douglas J. Hsu, Mengling Feng, Rishi Kothari, Hufeng Zhou, Kenneth P. Chen, and Leo A. Celi. The Association Between Indwelling Arterial Catheters and Mortality in Hemodynamically Stable Patients With Respiratory Failure: A Propensity Score Analysis. *Chest*, 148(6):1470–1476, December 2015.
- [66] Mehdi Jamei, Aleksandr Nisnevich, Everett Wetchler, Sylvia Sudat, and Eric Liu. Predicting all-cause risk of 30-day hospital readmission using artificial neural networks. *PloS One*, 12(7):e0181173, 2017.
- [67] Jessica L Maksut, Carla J Hodge, Charlayne D Van, Amaya Razmi, and Meagan T Khau. Utilization of Z Codes for Social Determinants of Health among Medicare Fee-for-Service Beneficiaries, 2019. Technical Report 24, Centers for Medicare and Medicaid Services, September 2021.
- [68] Alistair Johnson, Lucas Bulgarelli, Tom Pollard, Steven Horng, Leo Anthony Celi, and Roger Mark. MIMIC-IV, June 2022. Version Number: 2.0 Type: dataset.
- [69] Alistair E W Johnson, David J Stone, Leo A Celi, and Tom J Pollard. The MIMIC Code Repository: enabling reproducibility in critical care research. *Journal of the American Medical Informatics Association*, 25(1):32–39, January 2018.
- [70] Monita Karmakar, Paula M. Lantz, and Renuka Tipirneni. Association of Social and Demographic Factors With COVID-19 Incidence and Death Rates in the US. *JAMA Network Open*, 4(1):e2036462, January 2021.
- [71] Alan Katz, Dan Chateau, Jennifer E. Enns, Jeff Valdivia, Carole Taylor, Randy Walld, and Scott McCulloch. Association of the Social Determinants of Health With Quality of Primary Care. *The Annals of Family Medicine*, 16(3):217–224, May 2018. Publisher: The Annals of Family Medicine Section: Original Research.
- [72] Amy J.H. Kind and William R. Buckingham. Making Neighborhood-Disadvantage Metrics Accessible — The Neighborhood Atlas. *New England Journal of Medicine*, 378(26):2456–2458, June 2018. Publisher: Massachusetts Medical Society _eprint: <https://doi.org/10.1056/NEJMp1802313>.
- [73] Patrick M. Krueger, Melanie K. Tran, Robert A. Hummer, and Virginia W. Chang. Mortality Attributable to Low Levels of Education in the United States. *PloS One*, 10(7):e0131809, 2015.

- [74] Daniel J. Kruger, Thomas M. Reischl, and Gilbert C. Gee. Neighborhood Social Conditions Mediate the Association Between Physical Deterioration and Mental Health. *American Journal of Community Psychology*, 40(3-4):261–271, December 2007.
- [75] Peng Lan, Ting-Ting Wang, Hang-Yang Li, Ru-Shuang Yan, Wei-Chao Liao, Bu-Wen Yu, Qian-Qian Wang, Ling Lin, Kong-Han Pan, Yun-Song Yu, and Jian-Cang Zhou. Utilization of echocardiography during septic shock was associated with a decreased 28-day mortality: a propensity score-matched analysis of the MIMIC-III database. *Annals of Translational Medicine*, 7(22):662, November 2019.
- [76] Jane Lawrence. The Indian Health Service and the Sterilization of Native American Women. *American Indian Quarterly*, 24(3):400–419, 2000. Publisher: University of Nebraska Press.
- [77] Li-wei Lehman, Mohammed Saeed, William Long, Joon Lee, and Roger Mark. Risk Stratification of ICU Patients Using Topic Models Inferred from Unstructured Progress Notes. *AMIA Annual Symposium Proceedings*, 2012:505–511, November 2012.
- [78] Erica Adams Lehnert, Grete Wilt, Barry Flanagan, and Elaine Hallisey. Spatial exploration of the CDC’s Social Vulnerability Index and heat-related health outcomes in Georgia. *International Journal of Disaster Risk Reduction*, 46:101517, June 2020.
- [79] Elle Lett, Emmanuella Asabor, Sourik Beltrán, Ashley Michelle Cannon, and Onyebuchi A. Arah. Conceptualizing, Contextualizing, and Operationalizing Race in Quantitative Health Sciences Research. *Annals of Family Medicine*, 20(2):157–163, 2022.
- [80] Yikuan Li, Hanyin Wang, and Yuan Luo. Improving Fairness in the Prediction of Heart Failure Length of Stay and Mortality by Integrating Social Determinants of Health. *Circulation: Heart Failure*, 15(11):e009473, November 2022. Publisher: American Heart Association.
- [81] J. P. Mackenbach. The contribution of medical care to mortality decline: McKeeown revisited. *Journal of Clinical Epidemiology*, 49(11):1207–1213, November 1996.
- [82] J. P. Mackenbach, K. Stronks, and A. E. Kunst. The contribution of medical care to inequalities in health: differences between socio-economic groups in decline of mortality from conditions amenable to medical intervention. *Social Science & Medicine (1982)*, 29(3):369–376, 1989.
- [83] Sarah B. Maness, Laura Merrell, Erika L. Thompson, Stacey B. Griner, Nolan Kline, and Christopher Wheldon. Social Determinants of Health and Health Disparities: COVID-19 Exposures and Mortality Among African American People in the United States. *Public Health Reports*, 136(1):18–22, January 2021. Publisher: SAGE Publications Inc.

- [84] Michael Marmot and Richard Wilkinson. *Social Determinants of Health*. OUP Oxford, October 2005. Google-Books-ID: AmwiS8HZeRIC.
- [85] J. Michael McGinnis. Actual Causes of Death in the United States. *JAMA: The Journal of the American Medical Association*, 270(18):2207, November 1993.
- [86] J. Michael McGinnis, Pamela Williams-Russo, and James R. Knickman. The Case For More Active Policy Attention To Health Promotion. *Health Affairs*, 21(2):78–93, March 2002. Publisher: Health Affairs.
- [87] Deven McGraw. *Privacy Concerns Related to Inclusion of Social and Behavioral Determinants of Health in Electronic Health Records*. National Academies Press (US), January 2015. Publication Title: Capturing Social and Behavioral Domains and Measures in Electronic Health Records: Phase 2.
- [88] Joanne McPeake, Leanne Boehm, Elizabeth Hibbert, Katrina Hauschildt, Rita Bakhru, Anthony Bastin, Brad Butcher, Tammy Eaton, Wendy Harris, Aluko Hope, James Jackson, Annie Johnson, Janet Kloos, Karen Korzick, Judith McCartney, Joel Meyer, Ashley Montgomery-Yates, Tara Quasim, Andrew Slack, Dorothy Wade, Mary Still, Giora Netzer, Ramona O. Hopkins, Mark E. Mikkelsen, Theodore Iwashyna, Kimberley Haines, and Carla Sevin. Modification of social determinants of health by critical illness and consequences of that modification for recovery: an international qualitative study. *BMJ Open*, 12(9):e060454, September 2022. Publisher: British Medical Journal Publishing Group Section: Intensive care.
- [89] Jennifer Meddings, Heidi Reichert, Shawna N. Smith, Theodore J. Iwashyna, Kenneth M. Langa, Timothy P. Hofer, and Laurence F. McMahon. The Impact of Disability and Social Determinants of Health on Condition-Specific Readmissions beyond Medicare Risk Adjustments: A Cohort Study. *Journal of General Internal Medicine*, 32(1):71–80, January 2017.
- [90] Clifford S. Mitchell, Junfeng (Jim) Zhang, Torben Sigsgaard, Matti Jantunen, Paul J. Liroy, Robert Samson, and Meryl H. Karol. Current State of the Science: Health Effects and Indoor Environmental Quality. *Environmental Health Perspectives*, 115(6):958–964, June 2007.
- [91] Todd D. Molfenter, Abhik Bhattacharya, and David H. Gustafson. The roles of past behavior and health beliefs in predicting medication adherence to a statin regimen. *Patient Preference and Adherence*, 6:643–651, 2012.
- [92] Intae Moon, Stefan Groha, and Alexander Gusev. SurvLatent ODE : A Neural ODE based time-to-event model with competing risks for longitudinal data improves cancer-associated Venous Thromboembolism (VTE) prediction, August 2022. arXiv:2204.09633 [cs, stat].
- [93] Amol S. Navathe, Feiran Zhong, Victor J. Lei, Frank Y. Chang, Margarita Sordo, Maxim Topaz, Shamkant B. Navathe, Roberto A. Rocha, and Li Zhou.

- Hospital Readmission and Social Risk Factors Identified from Physician Notes. *Health Services Research*, 53(2):1110–1136, April 2018.
- [94] John P. Ney, David N. van der Goes, Marc R. Nuwer, Lonnie Nelson, and Matthew A. Eccher. Continuous and routine EEG in intensive care. *Neurology*, 81(23):2002–2008, December 2013.
- [95] Ank E. Nijhawan, Lisa R. Metsch, Song Zhang, Daniel J. Feaster, Lauren Gooden, Mamta K. Jain, Robrina Walker, Shannon Huffaker, Michael J. Muga-vero, Petra Jacobs, Wendy S. Armstrong, Eric S. Daar, Meg Sullivan, Carlos del Rio, and Ethan A. Halm. Clinical and Sociobehavioral Prediction Model of 30-Day Hospital Readmissions Among People with HIV and Substance Use Disorder: Beyond Electronic Health Record Data. *Journal of acquired immune deficiency syndromes (1999)*, 80(3):330–341, March 2019.
- [96] NORC at the University of Chicago and American Health Information Management Association. Social Determinants of Health Data: Survey Results on the Collection, Integration, and Use. Technical report, February 2023.
- [97] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, October 2019. Publisher: American Association for the Advancement of Science.
- [98] OECD. *Health at a Glance 2021: OECD Indicators*. Health at a Glance. OECD, November 2021.
- [99] Office of the National Coordinator for Health Information Technology. Social Determinants of Health (SDOH) Clinical Decision Support (CDS) Feasibility Brief. Technical report, February 2023.
- [100] PatientEngagementHIT. Lyft Fills Medical Transportation Gaps for One-Third of Riders, January 2019.
- [101] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Pas-sos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [102] Stephen Pfohl, Ben Marafino, Adrien Coulet, Fatima Rodriguez, Latha Palaniappan, and Nigam H. Shah. Creating Fair Models of Atherosclerotic Cardiovascular Disease Risk. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '19, pages 271–278, New York, NY, USA, January 2019. Association for Computing Machinery.
- [103] Tom J. Pollard and Leo Anthony Celi. Enabling Machine Learning in Critical Care. *ICU management & practice*, 17(3):198–199, 2017.

- [104] Brian Powers, Scott Rinefort, and Sachin H. Jain. Shifting Non-Emergency Medical Transportation To Lyft Improves Patient Experience And Lowers Costs. *Health Affairs Forefront*, September 2018.
- [105] Brian W. Powers, Scott Rinefort, and Sachin H. Jain. Nonemergency Medical Transportation: Delivering Care in the Era of Lyft and Uber. *JAMA*, 316(9):921–922, September 2016.
- [106] WHO Healthy Cities Project. *Social Determinants of Health: The Solid Facts*. World Health Organization, 2003. Google-Books-ID: QDFzqNZZHLMC.
- [107] PRRAC. "Racial and Ethnic Residential Segregation as a Root Social Determinant of Public Health and Health Inequity: A Persistent Public Health Challenge in the United States" by Robert A. Hahn (April-June 2017 P&R Issue), June 2017.
- [108] Jonathan Purtle, Rachel Peters, Jennifer Kolker, and Ana V. Diez Roux. Uses of Population Health Rankings in Local Policy Contexts: A Multisite Case Study. *Medical Care Research and Review*, 76(4):478–496, August 2019.
- [109] Andrea Ribick and Megan Moriarty. New HL7® FHIR® Accelerator Project Aims to Improve Interoperability of Social Determinants of Health Data. August 2019.
- [110] A. Rumshisky, M. Ghassemi, T. Naumann, P. Szolovits, V. M. Castro, T. H. McCoy, and R. H. Perlis. Predicting early psychiatric readmission with natural language processing of narrative discharge summaries. *Translational Psychiatry*, 6(10):e921, October 2016.
- [111] Eliane Rössli, Selen Bozkurt, and Tina Hernandez-Boussard. Peeking into a black box, the fairness and generalizability of a MIMIC-III benchmarking model. *Scientific Data*, 9(1):24, January 2022. Number: 1 Publisher: Nature Publishing Group.
- [112] Andrew D. Selbst, Danah Boyd, Sorelle Friedler, Suresh Venkatasubramanian, and Janet Vertesi. Fairness and Abstraction in Sociotechnical Systems, August 2018.
- [113] Laleh Seyyed-Kalantari, Guanxiong Liu, Matthew McDermott, Irene Y. Chen, and Marzyeh Ghassemi. CheXclusion: Fairness gaps in deep chest X-ray classifiers, October 2020. arXiv:2003.00827 [cs, eess, stat].
- [114] Laleh Seyyed-Kalantari, Haoran Zhang, Matthew B. A. McDermott, Irene Y. Chen, and Marzyeh Ghassemi. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature Medicine*, 27(12):2176–2182, December 2021. Number: 12 Publisher: Nature Publishing Group.
- [115] Efrat Shadmi, Natalie Flaks-Manov, Moshe Hoshen, Orit Goldman, Haim Bitterman, and Ran D. Balicer. Predicting 30-day readmissions with preadmission electronic health record data. *Medical Care*, 53(3):283–289, March 2015.

- [116] M. Shain and D. M. Kramer. Health promotion in the workplace: framing the concept; reviewing the evidence. *Occupational and Environmental Medicine*, 61(7):643–648, 585, July 2004.
- [117] Vinith M. Suriyakumar, Marzyeh Ghassemi, and Berk Ustun. When Personalization Harms: Reconsidering the Use of Group Attributes in Prediction, February 2023. arXiv:2206.02058 [cs, stat].
- [118] Samina T. Syed, Ben S. Gerber, and Lisa K. Sharp. Traveling Towards Disease: Transportation Barriers to Health Care Access. *Journal of community health*, 38(5):976–993, October 2013.
- [119] Paul Y Takahashi, Euijung Ryu, Janet E Olson, Erin M Winkler, Matthew A Hathcock, Ruchi Gupta, Jeff A Sloan, Jyotishman Pathak, Suzette J Bielski, and James R Cerhan. Health behaviors and quality of life predictors for risk of hospitalization in an electronic health record-linked biobank. *International Journal of General Medicine*, 8:247–254, August 2015.
- [120] The All of Us Research Program Investigators. The “All of Us” Research Program. *New England Journal of Medicine*, 381(7):668–676, August 2019.
- [121] Renuka Tipirneni, Monita Karmakar, Megan O’Malley, Hallie C. Prescott, and Vineet Chopra. Contribution of Individual- and Neighborhood-Level Social, Demographic, and Health Factors to COVID-19 Hospitalization Outcomes. *Annals of Internal Medicine*, 175(4):505–512, April 2022.
- [122] Samuel D. Towne, Xiaojun Liu, Rui Li, Matthew Lee Smith, Jay E. Maddock, Anran Tan, Samah Hayek, Shira Zelber-Sagi, Xiaoqing Jiang, Haotian Ruan, and Zhaokang Yuan. Social and Structural Determinants of Health Inequities: Socioeconomic, Transportation-Related, and Provincial-Level Indicators of Cost-Related Forgone Hospital Care in China. *International Journal of Environmental Research and Public Health*, 18(11):6113, June 2021.
- [123] Alison E. Turnbull, Danielle Groat, Victor D. Dinglas, Narjes Akhlaghi, Somnath Bose, Valerie Banner-Goodspeed, Mustafa Mir-Kasimov, Carla M. Sevin, James C. Jackson, Sarah Beesley, Ramona O. Hopkins, Dale M. Needham, Samuel M. Brown, Elise Caraker, Sai Phani Sree Cherukuri, Naga Preethi Kadiri, Tejaswi Kalva, Mounica Koneru, Pooja Kota, Emma Maelian Lee, Mazin Ali Mahmoud, Albahi Malik, Roozbeh Nikooie, Darin Roberts, Sriharsha Singu, Parvaneh Vaziri, Katie Brown, Austin Daw, Mardee Merrill, Rilee Smith, Ellie Hirshberg, Jorie Butler, Benjamin Hoenig, Maria Karamourtopoulos, Margaret Hays, Rebecca Abel, Craig High, Emily Beck, Brent Armbruster, Darrin Applegate, Melissa Fergus, Naresh Kumar, Megan Roth, Susan Mogan, Rebecca Abel, Andrea De Souza Licht, Isabel Londono, Julia Larson, Krystal Capers, Andrew Toksoz-Exley, and Julia Crane. Perceived Social Support among Acute Respiratory Failure Survivors in a Multicenter Prospective Cohort Study. *Annals of the American Thoracic Society*, 19(11):1930–1933, November 2022. Publisher: American Thoracic Society - AJRCCM.

- [124] Unite Us. Lyft and Unite Us Team Up to Reduce Transportation Barriers and Improve Access to Health and Social Care, March 2020.
- [125] U.S. Department of Agriculture Economic Research Service. Atlas of Rural and Small-Town America, December 2022.
- [126] Joshua R. Vest and Ofir Ben-Assuli. Prediction of emergency department revisits using area-level social determinants of health measures and health information exchange information. *International Journal of Medical Informatics*, 129:205–210, September 2019.
- [127] Randy A. Vince, Jr, Ralph Jiang, Merrick Bank, Jake Quarles, Milan Patel, Yilun Sun, Holly Hartman, Nicholas G. Zaorsky, Angela Jia, Jonathan Shoag, Robert T. Dess, Brandon A. Mahal, Kristian Stensland, Nicholas W. Eyrich, Mariana Seymore, Rebecca Takele, Todd M. Morgan, Matthew Schipper, and Daniel E. Spratt. Evaluation of Social Determinants of Health and Prostate Cancer Outcomes Among Black and White Patients: A Systematic Review and Meta-analysis. *JAMA Network Open*, 6(1):e2250416, January 2023.
- [128] Philip Vutien, Rucha Shah, Karen Ma, Nasir Saleem, and Joshua Melson. Utilization of Census Tract-Based Neighborhood Poverty Rates to Predict Non-adherence to Screening Colonoscopy. *Digestive Diseases and Sciences*, 64(9):2505–2513, September 2019.
- [129] Rebekah J. Walker, Mulugeta Gebregziabher, Bonnie Martin-Harris, and Leonard E. Egede. Relationship between social determinants of health and processes and outcomes in adults with type 2 diabetes: validation of a conceptual framework. *BMC Endocrine Disorders*, 14(1):82, October 2014.
- [130] Colin G. Walsh, Jessica D. Ribeiro, and Joseph C. Franklin. Predicting suicide attempts in adolescents with longitudinal clinical data and machine learning. *Journal of Child Psychology and Psychiatry, and Allied Disciplines*, 59(12):1261–1270, December 2018.
- [131] Harriet A. Washington. *Medical Apartheid: The Dark History of Medical Experimentation on Black Americans from Colonial Times to the Present*. Doubleday, 2006. Google-Books-ID: q0bMejttqgEC.
- [132] Will Douglas Heaven. Hundreds of AI tools have been built to catch covid. None of them helped., July 2021.
- [133] Ron Wilson. Understanding and Enhancing the U.S. Department of Housing and Urban Development’s ZIP Code Crosswalk Files. *Cityscape: A Journal of Policy Development and Research*, 20:277–294, 2018.
- [134] Jiancheng Ye, Liang Yao, Jiahong Shen, Rethavathi Janarthanam, and Yuan Luo. Predicting mortality in critically ill patients with diabetes using machine learning and clinical notes. *BMC Medical Informatics and Decision Making*, 20(11):295, December 2020.

- [135] Chia-Yuan Yu, Ayoung Woo, Christopher T. Emrich, and Biyuan Wang. Social Vulnerability Index and obesity: An empirical study in the US. *Cities*, 97:102531, February 2020.
- [136] Haoran Zhang, Natalie Dullerud, Karsten Roth, Lauren Oakden-Rayner, Stephen Robert Pfohl, and Marzyeh Ghassemi. Improving the Fairness of Chest X-ray Classifiers, March 2022. arXiv:2203.12609 [cs, eess].
- [137] Haoran Zhang, Amy X. Lu, Mohamed Abdalla, Matthew McDermott, and Marzyeh Ghassemi. Hurtful Words: Quantifying Biases in Clinical Contextual Word Embeddings, March 2020. arXiv:2003.11515 [cs, stat].
- [138] Yongkang Zhang, Yiye Zhang, Evan Sholle, Sajjad Abedian, Marianne Sharko, Meghan Reading Turchioe, Yiyuan Wu, and Jessica S Ancker. Assessing the impact of social determinants of health on predictive models for potentially avoidable 30-day readmission or death. *PLoS One*, 15(6):e0235064, 2020.
- [139] Zhongheng Zhang, Kun Chen, and Hongying Ni. Calcium supplementation improves clinical outcome in intensive care unit patients: a propensity score matched analysis of a large clinical database MIMIC-II. *SpringerPlus*, 4:594, 2015.
- [140] Le Zheng, Oliver Wang, Shiyang Hao, Chengyin Ye, Modi Liu, Minjie Xia, Alex N. Sabo, Liliana Markovic, Frank Stearns, Laura Kanov, Karl G. Sylvester, Eric Widen, Doff B. McElhinney, Wei Zhang, Jiayu Liao, and Xuefeng B. Ling. Development of an early-warning system for high-risk patients for suicide attempt using deep learning and electronic health records. *Translational Psychiatry*, 10(1):1–10, February 2020. Number: 1 Publisher: Nature Publishing Group.
- [141] Ezgi Özyılmaz, Özlem Özkan Kuşçu, Emre Karakoç, Aslı Boz, Gülşah Orhan Traşçı, Rengin Güzel, and Gülşah Seydaoğlu. Worse pre-admission quality of life is a strong predictor of mortality in critically ill patients. *Turkish Journal of Physical Medicine and Rehabilitation*, 68(1):19–29, March 2022.