Chasing Your Long Tails: Differentially Private Prediction in Health Care Settings

Vinith M. Suriyakumar, Nicolas Papernot, Anna Goldenberg, Marzyeh Ghassemi

vinith@cs.toronto.edu

University of Toronto, Vector Institute

ABSTRACT

Machine learning models in health care are often deployed in settings where it is important to protect patient privacy. In such settings, methods for *differentially private* (DP) learning provide a general-purpose approach to learn models with privacy guarantees. Modern methods for DP learning ensure privacy through the addition of calibrated noise. The resulting privacy-preserving models are unable to learn too much information about the *tails* of a data distribution, resulting in a loss of accuracy that can disproportionately affect small groups. In this paper, we study the effects of DP learning in health care. We use state-of-the-art methods for DP learning to train privacy-preserving models in clinical prediction tasks, including x-ray classification of images and mortality prediction in time series data. We use these models to perform a comprehensive empirical investigation of the tradeoffs between privacy, utility, robustness to dataset shift and fairness. Our results highlight lesser-known limitations of methods for DP learning in health care, models that exhibit steep tradeoffs between privacy and utility, and models whose predictions are disproportionately influenced by large demographic groups in the training data. We discuss the costs and benefits of differentially private learning in health care with open directions for differential privacy, machine learning and health care.

CCS CONCEPTS

• **Applied computing** → **Health informatics**; • **Security and privacy** → **Usability in security and privacy**; • **Computing methodologies** → **Machine learning**.

KEYWORDS

machine learning, health care, privacy, fairness, robustness

ACM Reference Format:

Vinith M. Suriyakumar, Nicolas Papernot, Anna Goldenberg, Marzyeh Ghassemi. 2021. Chasing Your Long Tails: Differentially Private Prediction in Health Care Settings. In *Conference on Fairness, Accountability, and Transparency (FAccT '21)*, March 3–10, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 33 pages. <https://doi.org/10.1145/3442188.3445934>

This work is licensed under a Creative Commons Attribution International 4.0 License.

FAccT '21, March 3–10, 2021, Virtual Event, Canada

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8309-7/21/03.

<https://doi.org/10.1145/3442188.3445934>

1 INTRODUCTION

The potential for machine learning to learn clinically relevant patterns in health care has been demonstrated across a wide variety of tasks [36, 80, 92, 100]. However, machine learning models are susceptible to privacy attacks [31, 84] that allow malicious entities with access to these models to recover sensitive information, e.g., HIV status or zip code, of patients who were included in the training data. Others have shown that anonymized electronic health records (EHR) can be re-identified using simple “linkages” with public data [90], and that neural models trained on EHR are susceptible to membership inference attacks [49, 84].

Differential privacy (DP) has been proposed as a leading technique to minimize re-identification risk through linkage attacks [24, 65], and is being used to collect personal data by the 2020 US Census [42], user statistics in iOS and MacOS by Apple [91], and Chrome Browser data by Google [68]. DP is an algorithm-level guarantee used in machine learning [22], where an algorithm is said to be differentially private if its output is statistically indistinguishable when applied to two input datasets that differ by only one record in the dataset. DP learning focuses with increasing intensity on learning the “body” of a targeted distribution as the desired level of privacy increases. Techniques such as differentially private stochastic gradient descent (DP-SGD) [2] and objective perturbation [8, 66] have been developed to efficiently train models with DP guarantees, but introduce a *privacy-utility tradeoff* [32]. This tradeoff has been well-characterized in computer vision [74], and tabular data [47, 84] but have not yet been characterized in health care datasets. Further, DP learning has asymptotic theoretical guarantees about generalization that have been established [50, 69]. There have not been any theoretical or empirical investigations on connections between differential privacy and out-of-distribution generalization which we term *privacy-robustness tradeoffs*. Finally, more “unique” minority data may not be well-characterized by DP, leading to a noted *privacy-fairness tradeoff* in vision [5, 27] and natural language settings [5].

To date there has not been a robust characterization of utility, privacy, robustness, and fairness tradeoffs for DP models in health care settings. Patient health and care are often highly individualized with a heavy “tail” due to the complexity of illness and treatment variation [45], and any loss of model utility in a deployed model is likely to hinder delivered care [93]. Privacy-robustness tradeoffs may also be high cost in health care, as the data distribution is constantly changing in response to new conditions [15], clinical practice shifts [44], and underlying EHR systems changing [67]. Privacy-fairness tradeoffs are perhaps the most pernicious concern in health care as there are well-documented prejudices in health care [12]. Importantly, the data of patients from minority groups also often lie even further in the data distribution tails because lack

of access to care can impact patients' EHR presence [30], and leads to small sample sizes of non-white patients [11].

In this work, we investigate the feasibility of using DP methods to train models for health care tasks. We characterize the impact of DP learning in both linear and neural models on 1) accuracy, 2) robustness, and 3) fairness. First, we establish the privacy-utility tradeoffs within two health care datasets (NIH Chest X-Ray data [99], and MIMIC-III EHR data [48]) as compared to two vision datasets (MNIST [59] and Fashion-MNIST [102]). We find that DP models have severe privacy-utility tradeoffs in the MIMIC-III EHR setting, using three common tasks — mortality, long-length of stay (LOS), and an intervention (vasopressor) onset [41, 98]. Second, we investigate the impact of DP on robustness to dataset shifts in EHR data. Because medical data often contains dataset shifts over time [33], we create a realistic yearly model training scenario and evaluate the robustness of DP models under these shifts. Finally, we investigate the impact of DP on fairness in two ways: loss of performance and loss of influence. We define loss of performance through the standard, and often competing, group fairness metrics [40, 52, 53] of performance gap, parity gap, recall gap, and specificity gap. We examine fairness further by looking at loss of minority data importance with influence functions [57]. In our audits, we focus on loss of population minority influence, e.g., importance of Black patient data, and label minority influence, e.g., importance of positive class patient data, across levels of privacy (low to high) and levels of dataset shift (least to most malignant).

Across our experiments we find that DP learning algorithms are *not* well-suited for off-the-shelf use in health care, outlining a series of technical challenges for the DP community. First, DP models have significantly more severe privacy-utility tradeoffs in the MIMIC-III EHR setting, and this tradeoff is proportional to the size of the tails in the data. This tradeoff holds even in larger datasets such as NIH Chest X-Ray. We further find that DP learning does not increase model robustness in the presence of small or large dataset shifts except for one instance. This observation warrants further empirical and theoretical investigations on connections between differential privacy and out of distribution generalization. Finally, we do not find a significant drop in standard group fairness definitions, unlike other domains [5], likely due to the dominating effect of utility loss. We do, however, find a large drop in minority class influence. Specifically, we show that Black training patients lose “helpful” influence on Black test patients. Finally, based on our empirical observations, we outline a series of open problems that future work should address to make DP learning feasible in health care settings.

1.1 Contributions

In this work, we evaluate the impact of DP learning on linear and neural networks across three tradeoffs: privacy-utility, privacy-robustness and privacy-fairness. Our analysis contributes to a call for ensuring that privacy mechanisms equally protect all individuals [25]. We present the following contributions:

- **Privacy-utility tradeoffs scale sharply with tail length.** We find that DP has particularly strong tradeoffs as tasks have fewer positive examples, resulting in unusable classifier performance.

Further, increasing the dataset size does not improve utility tradeoffs in our health care tasks.

- **There is no correlation between privacy and robustness in EHR shifts.** We show that DP generally does not improve shift robustness, with the mortality task as one exception. Despite this, we find no correlation between increasing privacy and improved shift robustness in our tasks, most likely due to the poor utility tradeoffs. Thus, given the one instance we believe further theoretical and empirical research should explore connections between differential privacy and out-of-distribution generalization.
- **DP gives unfair influence to majority groups that is hard to detect with standard measures of group fairness.** We show that increasing privacy does not result in disparate impact for minority groups across multiple protected attributes and standard group fairness definitions because the privacy-utility tradeoff is so extreme. We use influence functions to demonstrate that the inherent group privacy property of DP results in large losses of influence for minority groups across patient class label, and patient ethnicity labels.

2 RELATED WORK

2.1 Differential Privacy

DP provides much stronger privacy guarantees over methods such as k-anonymity [89] and t-closeness [60], to a number of privacy attacks such as reconstruction, tracing, linkage, and differential attacks [24]. The outputs of DP analyses are resistant to attacks based on auxiliary information, meaning they cannot be made less private [23]. Such benefits have made DP a leading method for ensuring privacy in consumer data settings [42, 68, 91]. Further, theoretical analyses have demonstrated improved generalization guarantees for out of sample data [50]. Inspired by these theoretical guarantees, we are interested in investigating empirically if differential privacy improves out of distribution generalization. Our goal is to motivate future theoretical analyses connecting differential privacy and domain generalization. Other theoretical analyses demonstrate that a model that is both private and approximately fair can exist in finite sample access settings. However, they show that it is impossible to achieve DP and exact fairness with non-trivial accuracy [19]. Further theoretical analyses show the unavoidable tradeoff between privacy and utility on long-tailed datasets [28]. This is empirically shown in DP-SGD which has disparate impact on complex minority groups in vision and NLP [5, 27].

Differential Privacy in Health Care. Prior work on DP in machine learning for health care has focused on the distributed setting, where multiple hospitals collaborate to learn a model [6, 76]. This work has shown that DP learning leads to a loss in model performance defined by area under the receiver operator characteristic (AUROC). We instead focus on analyzing the tradeoffs between privacy, robustness, and fairness, with an emphasis on the impact that DP has on subgroups.

2.2 Utility, Robustness, and Fairness in Health Care

Utility Needs in Health Care Tasks. Machine learning in health care is intended to support clinicians in their decision making,

DATASET	DATA TYPE	OUTCOME VARIABLE	n	d	CLASSIFICATION TASK	TAIL SIZE	PROTECTED ATTRIBUTES	EVALUATION
HEALTH CARE								
mimic_mortality	TIME SERIES	IN-ICU MORTALITY	21,877	(24,69)	BINARY	LARGE	ETHNICITY	U,R, F
mimic_los_3	TIME SERIES	LENGTH OF STAY > 3 DAYS	21,877	(24,69)	BINARY	SMALL	ETHNICITY	U,R, F
mimic_intervention	TIME SERIES	VASOPRESSOR ADMINISTRATION	21,877	(24,69)	MULTICLASS (4)	SMALL	ETHNICITY	U,R, F
NIH_chest_x_ray	IMAGING	MULTILABEL DISEASE PREDICTION	112,120	(256,256)	MULTICLASS MULTILABEL (14)	LARGEST	SEX	U,F
VISION BASELINES								
mnist	IMAGING	NUMBER CLASSIFICATION	60,000	(28,28)	MULTICLASS (10)	NONE	N/A	U
fashion_mnist	IMAGING	CLOTHING CLASSIFICATION	60,000	(28,28)	MULTICLASS (10)	NONE	N/A	U

Table 1: We analyze tradeoffs in two vision baseline datasets and two health care datasets. We use three prediction tasks in MIMIC-III with different tail sizes and focus our utility (U), robustness (R), and fairness (F) analyses on these tasks. Finally, we choose NIH Chest X-Ray which is a larger dataset with the largest tail to examine whether increasing the dataset size has an impact on utility and fairness tradeoffs.

which suggests that models need to perform similarly to physicians [20]. The specific metric is dependent on the task as high positive predictive value may be preferred over high negative predictive value [55]. In this work, we focus on predictive accuracy as AUROC and AUPRC, characterizing this loss as privacy levels increase.

Robustness to Dataset Shift. The effect of dataset shift has been studied in non-DP health care settings, demonstrating that model performance often deteriorates when the data distribution is non-stationary [21, 51, 88]. Recent work has demonstrated that performance deteriorates rapidly on patient LOS and mortality prediction tasks in the MIMIC-III EHR dataset, when trained on past years, and applied to a future year [67]. We focus on this setting for a majority of our experiments, leveraging year-to-year changes in population as small dataset shifts, and a change in EHR software between 2008 and 2009 as a large dataset shift.

Group Fairness. Disparities exist between white and Black patients, resulting in health inequity in the U.S.A [70, 72]. Further, even the use of some sensitive data like ethnicity in medical practice is contentious [97], and has been called into question in risk scores, for instance in estimating kidney function [26, 62].

Much work has described the ability of machine learning models to exacerbate disparities between protected groups [11]; even state-of-the-art chest X-Ray classifiers demonstrate diagnostic disparities between sex, ethnicity, and insurance type [83]. We leverage recent work in measuring the group fairness of machine learning models for different statistical definitions [40] in supervised learning.

We complement these standard metrics by also examining loss of data importance through influence functions [57]; influence functions have also been extended to approximate the effects of subgroups on a model’s prediction [58]. They demonstrate that memorization is required for small generalization error on long tailed distributions [28].

3 DATA

Details of each data source and prediction task are shown in Table 1. The four datasets are intentionally of different sizes, with respective tasks that represent distributions with and without long tails.

3.1 Vision Baselines

We use MNIST [59] and FashionMNIST [102] to demonstrate the benchmark privacy-utility tradeoffs in non-health settings with no tails. We use the NIH Chest X-Ray dataset [99] (112,120 images, details in Appendix B.2) to benchmark privacy-utility tradeoffs in a medically based, but still vision-focused, setting with the longest tails of all of our tasks.

3.2 MIMIC-III Time Series EHR Data

For the remainder of our analyses on privacy-robustness and privacy-fairness, we use the MIMIC-III database [48]—a publicly available anonymized EHR dataset of intensive care unit (ICU) patients (21,877 unique patient stays, details in Appendix B.1). We focus on two binary prediction tasks of predicting (1) ICU mortality (class imbalanced), (2) LOS greater than 3 days (class balanced) and choose one multiclass prediction tasks of predicting intervention onset for (3) vasopressor administration (class balanced) [41, 98].

Source of Distribution Shift. In MIMIC-III, there is a known source of dataset shift after 2008 due to a transition in the EHR used [1]. There are also smaller shifts in non-transition years as the patient distribution is non-stationary [67].

4 METHODOLOGY

We use both DP-SGD and objective perturbation across three different privacy levels to evaluate the impact that DP learning has on utility and robustness to dataset shift. Given the worse utility and robustness tradeoffs using objective perturbation, we focus our subsequent fairness analyses on DP-SGD in health care settings.

4.1 Model Classes

Vision Baselines. We use different convolutional neural network architectures for the MNIST and FashionMNIST prediction tasks based on prior work [74]. We use DenseNet-121 pretrained on ImageNet for the NIH Chest X-Ray experiments [83].

MIMIC EHR Tasks. For the MIMIC-III health care tasks analyses, we choose one linear model and one neural network per task, based on the best baselines, trained without privacy, outlined in prior work creating benchmarks for the MIMIC-III dataset [98]. For

binary prediction tasks we use logistic regression (LR) [17] and gated recurrent unit with decay (GRU-D) [10]. For our multiclass prediction task, we use LR and 1D convolutional neural networks.

4.2 Differentially Private Training

We train models without privacy guarantees using stochastic gradient descent (SGD). DP models are trained with DP-SGD [2], which is the de-facto approach for both linear models and neural networks. We choose not to train models using PATE [73], because it requires access to public data for semi-supervised learning and this is unrealistic in health care settings. In the Appendix, we provide results for models trained using objective perturbation [8, 66] which provides $(\epsilon, 0)$ -DP. It is only applicable to our linear models. We focus on DP-SGD due to its more optimal theoretical guarantees [46] regarding privacy-utility tradeoffs, and objective perturbation's limited applicability to linear models. The modifications made to SGD involve clipping gradients computed on a per-example basis to have a maximum ℓ_2 norm, and then adding Gaussian noise to these gradients before applying parameter updates [2] (Appendix E.1).

We choose three different levels of privacy to measure the effect of increasing levels of privacy by varying levels of epsilon. We selected these levels based on combinations of the noise level, clipping norm, number of samples, and number of epochs. Our three privacy levels are: None, Low (Clip Norm = 5.0, Noise Multiplier = 0.1), and High (Clip Norm = 1.0, Noise Multiplier = 1.0). We provide a detailed description of training setup in terms of hyperparameters and infrastructure in Appendix F.

4.3 Privacy Metrics

We measure DP using the ϵ bound derived analytically using the Renyi DP accountant for DP-SGD. Larger values of ϵ reflect lower privacy. Note that the privacy guarantees reported for each model are underestimates because they do not include the privacy loss due to hyperparameter searches [9, 61].

5 PRIVACY-UTILITY TRADEOFFS

We analyze the privacy-utility tradeoff by training linear and neural models with DP learning. We analyze performance across three privacy levels for the vision, MIMIC-III and NIH Chest X-Ray datasets. The privacy-utility tradeoffs for these datasets and tasks have *not* been characterized yet. Our work provides a benchmark for future work on evaluating DP learning.

Experimental Setup. We train both linear and neural models on the tabular MIMIC-III tasks. We train deep neural networks on NIH Chest X-Ray image tasks and the vision baseline tasks. We first analyze the effect that increased tail length in MIMIC-III has on the privacy-utility tradeoff. Next, we compare whether linear or neural models have better privacy-utility tradeoffs. Finally, we use the NIH Chest X-Ray dataset to evaluate if increasing dataset size, while keeping similar tail sizes, results in better tradeoffs.

Time Series Utility Metrics. For MIMIC-III, we average the model AUROC across all shifted test sets to quantitatively measure the utility tradeoff. We measure the privacy-utility tradeoff based on the difference in performance metrics as the level of privacy increases. The average performance across years is used because

it incorporates the performance variability between each of the years due to dataset shift. Results for AUPRC for MIMIC-III can be found in Appendix H.1. Both our AUROC and AUPRC results show extreme utility tradeoffs in health care tasks. Both metrics are commonly used to evaluate clinical performance of diagnostic tests [37].

Imaging Utility Metrics. For the NIH Chest X-Ray experiments, the task we experiment on is multiclass multilabel disease prediction. We average the AUROC across all 14 disease labels. For the MNIST and FashionMNIST vision baselines, the task we experiment on is multiclass prediction (10 labels for both) where we evaluate using accuracy.

5.1 Health Care Tasks Have Steep Utility Tradeoffs

We compare the privacy-utility tradeoffs in Table 2. DP-SGD generally has a negative impact on model utility. The extreme tradeoffs in MIMIC-III mortality prediction, and NIH Chest X-Ray diagnosis exemplify the information DP-SGD loses from the tails, because the positive cases are in the long tails of the distribution. There is a 22% and 26% drop in the AUROC between no privacy and high privacy settings for mortality prediction for LR and GRUD respectively. There is a 35% drop in AUROC between the no privacy and high privacy settings for the NIH Chest X-Ray prediction task which has a much longer tail than mortality prediction. Our results for objective perturbation show worse utility tradeoffs than those presented by DP-SGD (Appendix G.1).

5.2 Linear Models Have Better Privacy-Utility Tradeoffs

Across all three prediction tasks in the MIMIC-III dataset we find that linear models have better tradeoffs in the presence of long tails. This is likely due to two issues: small generalization error in neural networks often requires memorization in long tail settings [28, 104] and gradient clipping introduces more bias as the number of model parameters increases [14, 85].

5.3 Long Tails Are Still Difficult to Overcome With Larger Datasets

By definition of differential privacy, privacy-utility tradeoffs can be improved with larger datasets [95] and less noise is required to achieve the same privacy guarantee with more data. We find that the NIH Chest X-Ray dataset also has extreme tradeoffs. Despite its larger size, the dataset's positive labels in long tails are similarly lost. Additionally, we downsample the MIMIC-III Mortality task according to different percentages and find it still has similar privacy-utility tradeoffs to when the model is trained on the entire training set (Table 19).

6 PRIVACY-ROBUSTNESS TRADEOFFS

A potential motivation for using DP despite extreme utility tradeoffs are the recent theoretical generalization guarantees [50]. These guarantees hold for out of sample data that is still from the same distribution as the training sample which differs from our setting. In our setting, we empirically analyze whether there is motivation

VISION BASELINES				
DATASET	MODEL	NONE (ϵ, δ)	Low (ϵ, δ)	HIGH (ϵ, δ)
MNIST	CNN	98.83 ± 0.06 ($\infty, 0$)	98.58 ± 0.06 ($2.6 \cdot 10^5$)	93.78 ± 0.25 (2.01)
FASHIONMNIST	CNN	87.92 ± 0.19 ($\infty, 0$)	87.90 ± 0.16 ($2.6 \cdot 10^5$)	79.53 ± 0.10 (2.01)
MIMIC-III				
TASK	MODEL	NONE (ϵ, δ)	Low (ϵ, δ)	HIGH (ϵ, δ)
MORTALITY	LR	0.82 ± 0.03 ($\infty, 0$)	0.76 ± 0.05 ($3.50 \cdot 10^5, 10^{-5}$)	0.60 ± 0.04 ($3.54, 10^{-5}$)
	GRUD	0.79 ± 0.03 ($\infty, 0$)	0.59 ± 0.09 ($1.59 \cdot 10^5, 10^{-5}$)	0.53 ± 0.03 ($2.65, 10^{-5}$)
LENGTH OF STAY > 3	LR	0.69 ± 0.02 ($\infty, 0$)	0.66 ± 0.03 ($3.50 \cdot 10^5, 10^{-5}$)	0.60 ± 0.04 ($3.54, 10^{-5}$)
	GRUD	0.67 ± 0.03 ($\infty, 0$)	0.63 ± 0.02 ($1.59 \cdot 10^5, 10^{-5}$)	0.61 ± 0.03 ($2.65, 10^{-5}$)
INTERVENTION ONSET (VASO)	LR	0.90 ± 0.03 ($\infty, 0$)	0.87 ± 0.03 ($1.63 \cdot 10^7, 10^{-5}$)	0.77 ± 0.05 ($0.94, 10^{-5}$)
	CNN	0.88 ± 0.04 ($\infty, 0$)	0.86 ± 0.02 ($5.95 \cdot 10^7, 10^{-5}$)	0.68 ± 0.04 ($0.66, 10^{-5}$)
NIH CHEST X-RAY				
METRIC	MODEL	NONE (ϵ, δ)	Low (ϵ, δ)	HIGH (ϵ, δ)
AVERAGE AUC	DENSENET-121	0.84 ± 0.00 ($\infty, 0$)	0.51 ± 0.01 ($1.74 \cdot 10^5, 10^{-6}$)	0.49 ± 0.00 ($0.84, 10^{-6}$)
BEST AUC	DENSENET-121	0.98 ± 0.00 (HERNIA)	0.54 ± 0.04 (EDEMA)	0.54 ± 0.05 (PLEURAL THICKENING)
WORST AUC	DENSENET-121	0.72 ± 0.00 (INFILTRATION)	0.48 ± 0.02 (FIBROSIS)	0.47 ± 0.02 (PLEURAL THICKENING)

Table 2: Health care tasks have a significant tradeoff between the High and Low or None setting. The tradeoff is better in tasks with small tails (length of stay and intervention onset), and worst in tasks such as mortality and NIH Chest X-Ray with long tails. We provide the ϵ, δ guarantees in parentheses, where ϵ represents the privacy loss (lower is better) and δ represents the probability that the guarantee does not hold (lower is better).

Figure 1: The effect of DP learning on robustness to non-stationarity and dataset shift. One instance of increased robustness in the 2009 column for mortality prediction in the high privacy setting (A), but this does not hold across all tasks and models. Performance drops in the 2009 column for LOS in both LR and GRU-D (B), and a much worse drop in the high privacy CNN for intervention prediction (C).

for theoretical analyses about the effect of DP on out-of-distribution generalization guarantees. We investigate the impact of DP to mitigating dataset shift for time series MIMIC-III tasks by analyzing model performance across years of care. We first record generalization as the difference in performance when a model is trained and tested on data drawn from p , versus performance on a shifted test

set drawn from q) and the malignancy of the shift. We then measure the malignancy of the yearly shifts using a domain classifier. Finally we perform a Pearson's correlation test [86] between the model's generalization capacity and the shift malignancy.

Experimental Setup. We analyze the robustness of DP models to dataset shift in the MIMIC-III health care tasks. We use year-to-year variation in hospital practices as a small shifts, and a change in EHR software between 2008-2009 as a source of major dataset shift. We define robustness as the difference in test accuracy between in-distribution and out-distribution data. For instance, to measure model robustness from the 2006 to 2007, we would 1) *train* a model on data from 2006, 2) *test* the model on data from 2006, and 3) test the same model on data from 2007. The difference in these two test accuracies is the 2006-2007 model robustness.

Robustness Metrics. To measure the impact of DP-SGD on robustness to dataset shift, we measure the malignancy of yearly shifts from 2002 to 2012 for the MIMIC-III dataset. We then measure the correlation between malignancy of yearly shift and model performance. As done by others we use We use a binary domain classifier (model class is chosen best on data type) trained to discriminate between in-domain p and out-domain q . The malignancy of the dataset shift is proportional to how difficult it is to train on p and perform well in q [78]. Other methods such as multiple univariate hypothesis testing or multivariate hypothesis testing assume that the data is i.i.d [78]. A full procedure is given in Appendix A, with complete significance and malignancies for each year in Appendix D.

6.1 DP Does Not Impart Robustness to Yearly EHR Data Shift

While we expect that DP will be more robust to dataset shift across all tasks and models, we find that model performance drops when the EHR shift occurs (2008-2009) across all privacy levels and tasks (Fig. 1). We note one exception: high privacy models are more stable in the mortality task during more malignant shifts (2007-2009) (Fig. 1).¹ Despite this, we find that there are no significant correlations between model robustness and privacy level (Table 13).

Our analyses find that the robustness guarantees that DP provides do not hold in a large, tabular EHR setting. We note that the privacy-utility tradeoff from Section 5 is too extreme in health care to conclusively understand the effect on model robustness.

7 PRIVACY-FAIRNESS TRADEOFFS

Prior work has demonstrated that DP learning has disparate impact on complex minority groups in vision and NLP [5]. We expect similar disparate impacts on patient minority groups in the MIMIC-III and NIH Chest X-Ray datasets, based on known disparities in treatment and health care delivery [70, 72]. We evaluate disparities based on four standard group fairness definitions, and on loss of minority patient influence.

We focus on the disparities between white and Black patients in MIMIC-III, based on prior work showing classifier variation in the low number of Black patients [11]. We focus on male and female patients in NIH Chest X-Ray based on prior work exposing disparities in chest x-ray classifier performance between these two groups [83].

Group Fairness Experimental Setup and Metrics. We measure fairness according to four standard group fairness definitions:

¹We did not observe this improvement when training with objective perturbation (Appendix G.2).

performance gap, parity gap, recall gap, and specificity gap [40]. The performance gap for our health care tasks is the difference in AUROC between the selected subgroups. The remaining three definitions of fairness for binary prediction tasks are presented in Appendix A.3.

Influence Experimental Setup and Metrics. We use influence functions to measure the relative influence of training points on test set performance (equations in Appendix A.4). Influences above 0 are *helpful* in minimizing the test loss for the test patient in that column, and influences below 0 are *harmful* in minimizing the test loss for that patient. Our influence function method [57] assumes a smooth, convex loss function with respect to model parameters, and is therefore only valid for LR. We focus group privacy analyses on the LR model in no and high privacy settings for mortality prediction.

First, we aim to confirm that the gradient clipping in DP-SGD bounds the influence of all training patients on the loss of all test patients. For the utility tradeoff, we measure the group influence that the training patients of each label group has on the loss of each test patient. For the robustness tradeoff, we measure the individual influence of all training patients on the loss of test patients between the least malignant and most malignant dataset shifts. Finally, we measure the group influence of training patients in each ethnicity on the white test patients and Black test patients separately.

7.1 DP Has No Impact on Group Fairness on Average, But Reduces Variance Over Time

To measure the average fairness gap in MIMIC-III, we average group fairness measures across all years of care. In the NIH Chest X-ray data we average across all disease labels.

We find that DP-SGD has little impact on any tested fairness definitions in both MIMIC-III (Table 14) and NIH Chest X-Ray, likely due to the high privacy-utility tradeoff. DP-SGD does confer lower variance in fairness measures on MIMIC-III tasks over time (Appendix H.3.1).

7.2 DP Learning Gives Unfair Influence to Majority Groups

We find that DP-SGD reduces the influence of all training points on individual test points (Fig 2) because gradient clipping tightly bounds influence of all training points across test points.

Influence-Utility Tradeoff. We find the worst privacy-utility tradeoff in the mortality task. Non-DP models find the patients who died to be the most helpful in predictions of mortality (Fig. 3 and Table 3). However, because positive labels, i.e., death, are rare, DP models focus influence on patients who survived, resulting in unfair over-influence.

Influence-Robustness Tradeoff. We see improved robustness in LR for the mortality task which has the most malignant dataset shift (2008-2009) (Figure 1).

We find that the variance of the influence is fairly low for non-DP models during lower malignancy shifts. In more malignant shifts, the variance of the influence is high with many training points being harmful (Fig. 2). This is likely due to gradient clipping reducing

Figure 2: DP bounds the individual influence of training patients on the loss of test patients (A) which improved robustness for mortality prediction between the least malignant shift in 2007 (B) and the most malignant in 2009 (C). Individual influence of training data in the no privacy setting on 100 test patients with highest influence variance. Each column on the x-axis is an individual test patient. A unique colour is plotted per column/test patient for ease of assessment. The influence value of each patient in the training set on a specific test point is plotted as a point in that patient's column. Influence of training points is bounded in the high privacy setting (red dotted line).

PRIVACY LEVEL	AVERAGE SURVIVED INFLUENCE	AVERAGE DIED INFLUENCE	MOST HELPFUL GROUP	MOST HARMFUL GROUP INFLUENCE
NONE	-1.07 ± 7.25	2.28 ± 6.91	DIED	SURVIVED
Low	-0.34 ± 0.95	0.03 ± 0.18	SURVIVED	SURVIVED
HIGH	-0.14 ± 4.69	0.04 ± 1.34	SURVIVED	SURVIVED

Table 3: Group influence summary statistics of training data by class label in all privacy levels for all test patients. Privacy changes the most helpful group the patients who died (minority) to the patients who survived (majority). DP learning minimizes the helpful influence of minority groups resulting in worse utility.

influence variance and is entangled with the poor privacy-utility tradeoff H.2.

Influence-Fairness Tradeoff. We approximate the collective group influence of different ethnicities in the training set on the test loss in Fig. 10 and Table 4. We show that group privacy results in white patients having a more significant influence, both helpful and harmful, on test patients in the high privacy setting.

8 DISCUSSION

8.1 On Utility, Robustness and Trust in Clinical Prediction Tasks

Poor Utility Impacts Trust. While some reduced utility in long tail tasks are known [28], the extreme tradeoffs that we observe in Table 2 are much worse than expected. Machine learning can only support the decision making processes of clinicians if there is clinical trust. If models do not perform well as, or better than,

Figure 3: Group influence of training data by class label in no privacy (A) and high privacy (B) settings on 100 test patients with highest influence variance. In the no privacy setting, patients who died have a helpful influence despite being a minority class. High privacy gives the majority group the most influence due to the group privacy guarantee.

WHITE TEST PATIENTS				
PRIVACY LEVEL	AVERAGE WHITE INFLUENCE	AVERAGE BLACK INFLUENCE	MOST HELPFUL ETHNICITY	MOST HARMFUL ETHNICITY
NONE	0.29 ± 2.40	0.71 ± 1.40	WHITE	WHITE
Low	-0.22 ± 0.70	-0.03 ± 0.17	WHITE	WHITE
HIGH	-0.11 ± 3.94	0.03 ± 1.35	WHITE	WHITE
BLACK TEST PATIENTS				
PRIVACY LEVEL	AVERAGE WHITE INFLUENCE	AVERAGE BLACK INFLUENCE	MOST HELPFUL ETHNICITY	MOST HARMFUL ETHNICITY
NONE	0.48 ± 1.39	0.44 ± 2.19	BLACK	WHITE
Low	-0.23 ± 0.75	-0.03 ± 0.18	WHITE	WHITE
HIGH	-0.40 ± 4.10	0.12 ± 1.45	WHITE	WHITE

Table 4: Group influence summary statistics across all privacy levels for white (majority) and Black (minority) training patients on both white and Black test patients in MIMIC-III. Privacy changes the most helpful group from Black patients to the majority white patients and minimizes their helpful influence. This needs careful consideration as the use of ethnicity is still being investigated in medical practice.

clinicians once we include privacy, there is little reason to trust them [93].

Importance of Model Robustness. Despite the promising theoretical transferability guarantees of DP, the results in Fig 1 and Table H.2 demonstrate these do not transfer in our health care setting. While we explored changes in EHR software as dataset shift, there are many other known shifts in healthcare data, e.g., practice modifications due to reimbursement policy changes [56], or changing clinical needs in public health emergencies such as COVID-19 [15]. If models do not maintain their utility after dataset shifts, catastrophic silent failures could occur in deployed models [67].

8.2 Majority Group Influence Is Harmful in Health Care

We show in Figure 3 that the tails of the label distribution are minority-rich, results in poor mortality prediction performance under DP. Prior work in evaluating model fairness in health care has focused on standard group fairness definitions [77]. However,

these definitions do not provide a detailed understanding of model fairness under reduced utility. Other work has shown that large utility loss can “wash out” fairness impacts [27]. Our work demonstrates that DP learning does harm group fairness in such “washed out” poor utility settings by giving majority groups (e.g., those that survived, and white patients) the most influence on predictions across all subgroups.

Why Influence Matters. Disproportionate assignment of influence is an important problem. Differences in access, practice, or recording reflect societal biases [79, 81], and models trained on biased data may exhibit unfair performance in populations due to this underlying variation [13]. Further, while patients with the same diagnosis are usually more helpful for estimating prognosis in practice [18], labels in health care often lack precision or, in some cases, may be unreliable [71]. In this setting, understanding what factors are consistent high-influence in patient phenotypes is an important task [38, 103].

Loss Of Black Influence. Ethnicity is currently used in medical practice as a factor in many risk scores, where different risk

profiles are assumed for patients of different races [62]. However, the validity of this stratification has recently been called into question by the medical community [26]. Prior work has established the complexity of treatment variation in practice, as patient care plans are highly individualized, e.g., in a cohort of 250 million patients, 10% of diabetes and depression patients and almost 25% of hypertension patients had a unique treatment pathway [45]. Thus having the white patients become the most influential in Black patients predictions may not be desirable.

Anchoring Influence Loss in Systemic Injustice. Majority over-influence is prevalent in medical settings, and has direct impact on the survival of patients. Many female and minority patients receive worse care and have worse outcomes because clinicians base their symptomatic evaluations on white and/or male patients [34, 35]. Further, randomized control trials (RCTs) are an important tool that 10-20% of treatments are based on [63]. However, prior work has shown that RCTs have notorious exclusive criteria for inclusion; in one salient example, only 6% of asthmatic patients would have been eligible to enroll in the RCT that resulted in their treatments [94]. RCTs tend to be comprised of white, male patients, resulting in their data determining what is an effective treatment [43]. By removing influence from women, Hispanics, and Blacks, naive machine learning practices can exacerbate systemic injustices [12].

There are ongoing efforts to improve representation of the population in RCTs, shifting away from the majority having majority influence on treatments [87]. Researchers using DP should follow suit, and work to reduce the disparate impact on influence to ensure that it does not perpetuate this existing bias in health care. One solution is to start measuring individual example privacy loss [29] instead of a conservative worst bound across all patients. Currently, DP-SGD uses constant gradient clipping for all examples to ensure this constant worst bound. Instead, individual privacy accounting can help support adaptive gradient clipping for each example which may help to reduce the disparate impact DP-SGD has on influence. We also encourage future privacy-fairness tradeoff analyses to include loss of influence as a standard metric, especially where the utility tradeoff is extreme.

8.3 Less Privacy In Tails Is Not an Option

The straightforward solution to the long tail issue is to “provide less or no privacy for the tails” [54]. In some settings this solution may make sense. In health care, this solution could amplify existing systemic biases against minority subgroups, and minority mistrust of medical institutions. For example, Black mothers in the US are most likely to be mistreated, dying in childbirth at a rate three times higher than white women [7]. In this setting, it is not ethical choose between a “non-private” prediction that will potentially leak unwanted information, e.g., prior history of abortion, and a “private” prediction that will deliver lower quality care.

8.4 On the Costs and Benefits of Privacy in Health Care

Privacy Issues With Health Care Data. Most countries have regulations that define the protective measures to maintain patient data privacy. In North America, these laws are defined by the

Health Insurance Portability and Accountability Act (HIPAA) [3] in the US and Personal Information Protection and Electronic Documents Act (PIPEDA) [4] in Canada. These laws are governed by the General Data Protection Regulation (GDPR) in the EU. Recent work has shown that HIPAA’s privacy regulations and standards such as anonymizing data are not sufficient to prevent advanced re-identification of data [64]. In one instance, researchers were able to re-identify individuals’ faces from MRIs using facial recognition software [82]. Privacy attacks such as these demonstrate the fear of health care data loss.

8.5 Open Problems for DP in Health Care

While health care has been cited as an important motivation for the development of DP [8, 22, 23, 74, 96, 101], our work demonstrates that there are a number of open areas of research to make it well suited for machine learning for healthcare. The theoretical benefits of DP apply in extremely large data collection settings, such as the successful deployment of DP US Census data storage. We highlight potential areas of research that both the DP and machine learning communities should focus on to make DP usable in health care data:

- (1) **Connections Between Differential Privacy and Domain Generalization** In our robustness to dataset shift experiments, we observe one instance where differential privacy improves model robustness. Dataset shift can be viewed as a type of domain generalization. Thus, an interesting direction would be to see how the generalization guarantees can be extended to domain generalization.
- (2) **Boosting Utility and Fairness of Differential Privacy Using Post-Processing** A number of recent post-processing techniques in algorithmic fairness have been used to convert classifiers into those that satisfy equalized odds and multicalibration. A differentially private implementation of the equalized odds post-processing was used to improve the fairness of DP learning. Further post-processing techniques should be examined as a way to boost the utility of differentially private learning.
- (3) **Private Collaborative Learning** In this work we investigate the single hospital setting which often means that there is limited data. An important direction in the health care setting is to develop private and confidential learning algorithms which allow hospitals to learning a single model while ensuring the privacy of their own data. This can be especially helpful in settings of class imbalance if one hospital has more data about a subpopulation that is infrequent at another hospital.
- (4) **Practical Deployments of Differential Privacy in Health Care** More interdisciplinary work between the differential privacy community, biostatisticians, and epidemiologists are needed to build differentially private algorithms which provide good utility based on the inherent challenges of health care data.

9 CONCLUSION

In this work, we investigate the feasibility of using DP-SGD to train models for health care prediction tasks. We find that DP-SGD is not well-suited to health care prediction tasks in its current formulation. First, we demonstrate that DP-SGD loses important information about minority classes (e.g., dying patients, minority

ethnicities) that lie in the tails of the dat distribution. The theoretical robustness guarantees of DP-SGD do not apply to the dataset shifts we evaluated. We show that DP learning disparately impacts group fairness when looking at loss of influence for majority groups. We show this disparate impact occurs even when standard measures of group fairness show no disparate impact due to poor utility. This imposed asymmetric valuation of data by the model requires careful thought, because the appropriate use of class membership labels in medical settings in an active topic of discussion and debate. Finally, we propose open areas of research to improve the usability of DP in health care settings. Future work should target modifying DP-SGD, or creating novel DP learning algorithms, that can learn from data distribution tails effectively, without compromising privacy.

ACKNOWLEDGMENTS

We would like to acknowledge the following funding sources: New Frontiers in Research Fund - NFRFE-2019-00844. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute www.vectorinstitute.ai/partners. Thank you to the MIT Laboratory of Computational Physiology for facilitating year of care access to the MIMIC-III database. Finally, we would like to thank Nathan Ng, Taylor Killian, Victoria Cheng, Varun Chandrasekaran, Sindhu Gowda, Laleh Seyyed-Kalantari, Berk Ustun, Shalmali Joshi, Natalie Dullerud, Shrey Jain, and Sicong (Sheldon) Huang for their helpful feedback.

REFERENCES

- [1] [n.d.]. MIMIC. <https://mimic.physionet.org/>, note = Accessed: 2020-09-30.
- [2] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (Vienna, Austria) (CCS '16)*. ACM, New York, NY, USA, 308–318.
- [3] Accountability Act. 1996. Health insurance portability and accountability act of 1996. *Public law* 104 (1996), 191.
- [4] Privacy Act. 2000. Personal Information Protection and Electronic Documents Act. *Department of Justice, Canada. Full text available at <http://laws.justice.gc.ca/en/P-8.6/text.html>* (2000).
- [5] Eugene Bagdasaryan, Omid Poursaeed, and Vitaly Shmatikov. 2019. Differential privacy has disparate impact on model accuracy. In *Advances in Neural Information Processing Systems*. 15453–15462.
- [6] Brett K Beaulieu-Jones, William Yuan, Samuel G Finlayson, and Zhiwei Steven Wu. 2018. Privacy-preserving distributed deep learning for clinical data. *arXiv preprint arXiv:1812.01484* (2018).
- [7] Cynthia J Berg, Hani K Atrash, Lisa M Koonin, and Myra Tucker. 1996. Pregnancy-related mortality in the United States, 1987–1990. *Obstetrics & Gynecology* 88, 2 (1996), 161–167.
- [8] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. 2011. Differentially private empirical risk minimization. *Journal of Machine Learning Research* 12, Mar (2011), 1069–1109.
- [9] Kamalika Chaudhuri and Staal A Vinterbo. 2013. A stability-based validation procedure for differentially private machine learning. In *Advances in Neural Information Processing Systems*. 2652–2660.
- [10] Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. 2018. Recurrent Neural Networks for Multivariate Time Series with Missing Values. *Scientific Reports* 8, 1 (April 2018), 6085. <https://doi.org/10.1038/s41598-018-24271-9>
- [11] Irene Chen, Fredrik D Johansson, and David Sontag. 2018. Why is my classifier discriminatory?. In *Advances in Neural Information Processing Systems*. 3539–3550.
- [12] Irene Y. Chen, Emma Pierson, Sherri Rose, Shalmali Joshi, Kadija Ferryman, and Marzyeh Ghassemi. 2020. Ethical Machine Learning in Health Care. *arXiv:2009.10576 [cs.CV]*
- [13] Irene Y Chen, Peter Szolovits, and Marzyeh Ghassemi. 2019. Can AI Help Reduce Disparities in General Medical and Mental Health Care? *AMA Journal of Ethics* 21, 2 (2019), 167–179.
- [14] Xiangyi Chen, Zhiwei Steven Wu, and Mingyi Hong. 2020. Understanding Gradient Clipping in Private SGD: A Geometric Perspective. *arXiv preprint arXiv:2006.15429* (2020).
- [15] Joseph Paul Cohen, Paul Morrison, Lan Dao, Karsten Roth, Tim Q Duong, and Marzyeh Ghassemi. 2020. Covid-19 image data collection: Prospective predictions are the future. *arXiv preprint arXiv:2006.11988* (2020).
- [16] R Dennis Cook and Sanford Weisberg. 1980. Characterizations of an empirical influence function for detecting influential cases in regression. *Technometrics* 22, 4 (1980), 495–508.
- [17] David R Cox. 1972. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)* 34, 2 (1972), 187–202.
- [18] Peter Croft, Douglas G Altman, Jonathan J Deeks, Kate M Dunn, Alastair D Hay, Harry Hemingway, Linda LeResche, George Peat, Pablo Perel, Steffen E Petersen, et al. 2015. The science of clinical practice: disease diagnosis or patient prognosis? Evidence about “what is likely to happen” should shape clinical practice. *BMC medicine* 13, 1 (2015), 20.
- [19] Rachel Cummings, Varun Gupta, Dhamma Kimpara, and Jamie Morgenstern. 2019. On the compatibility of privacy and fairness. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*. 309–315.
- [20] Thomas Davenport and Ravi Kalakota. 2019. The potential for artificial intelligence in healthcare. *Future healthcare journal* 6, 2 (2019), 94.
- [21] Sharon E Davis, Thomas A Lasko, Guanhua Chen, Edward D Siew, and Michael E Matheny. 2017. Calibration drift in regression and machine learning models for acute kidney injury. *Journal of the American Medical Informatics Association* 24, 6 (2017), 1052–1061.
- [22] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.
- [23] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3–4 (2014), 211–407.
- [24] Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. 2017. Exposed! a survey of attacks on private data. (2017).
- [25] Michael D Ekstrand, Rezvan Joshaghani, and Hoda Mehrpouyan. 2018. Privacy for all: Ensuring fair and equitable privacy protections. In *Conference on Fairness, Accountability and Transparency*. 35–47.
- [26] Nwamaka Denise Eneanya, Wei Yang, and Peter Philip Reese. 2019. Reconsidering the consequences of using race to estimate kidney function. *Jama* 322, 2 (2019), 113–114.
- [27] Tom Farrand, Fatemehsadat Miresheghallah, Sahib Singh, and Andrew Trask. 2020. Neither Private Nor Fair: Impact of Data Imbalance on Utility and Fairness in Differential Privacy. *arXiv preprint arXiv:2009.06389* (2020).
- [28] Vitaly Feldman. 2020. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 954–959.
- [29] Vitaly Feldman and Tijana Zrnic. 2020. Individual Privacy Accounting via a Renyi Filter. *arXiv preprint arXiv:2008.11193* (2020).
- [30] Kadija Ferryman and Mikaela Pitcan. 2018. Fairness in precision medicine. *Data & Society* (2018).
- [31] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. 1322–1333.
- [32] Quan Geng, Wei Ding, Ruiqi Guo, and Sanjiv Kumar. 2020. Tight Analysis of Privacy and Utility Tradeoff in Approximate Differential Privacy. In *International Conference on Artificial Intelligence and Statistics*. 89–99.
- [33] Marzyeh Ghassemi, Tristan Naumann, Peter Schulam, Andrew L Beam, and Rajesh Ranganath. 2018. Opportunities in machine learning for healthcare. *arXiv preprint arXiv:1806.00388* (2018).
- [34] Brad N Greenwood, Seth Carnahan, and Laura Huang. 2018. Patient–physician gender concordance and increased mortality among female heart attack patients. *Proceedings of the National Academy of Sciences* 115, 34 (2018), 8569–8574.
- [35] Brad N Greenwood, Rachel R Hardeman, Laura Huang, and Aaron Sojourner. 2020. Physician–patient racial concordance and disparities in birthing mortality for newborns. *Proceedings of the National Academy of Sciences* 117, 35 (2020), 21194–21200.
- [36] Varun Gulshan, Lily Peng, Marc Coram, Martin C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, Ramasamy Kim, Rajiv Raman, Philip C Nelson, Jessica L Mega, and Dale R Webster. 2016. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. *JAMA* 316, 22 (Dec. 2016), 2402–2410.
- [37] Karimollah Hajian-Tilaki. 2013. Receiver operating characteristic (ROC) curve analysis for medical diagnostic test evaluation. *Caspian journal of internal medicine* 4, 2 (2013), 627.
- [38] Yoni Halpern, Steven Horng, Youngduck Choi, and David Sontag. 2016. Electronic medical record phenotyping using the anchor and learn framework. *Journal of the American Medical Informatics Association* 23, 4 (2016), 731–740.

- [39] Han-JD. 2019. Gated Recurrent Unit with a Decay mechanism for Multivariate Time Series with Missing Values.. <https://github.com/Han-JD/GRU-D>
- [40] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*. 3315–3323.
- [41] Hrayr Harutyunyan, Hrant Khachatryan, David C. Kale, Greg Ver Steeg, and Aram Galstyan. 2017. Multitask Learning and Benchmarking with Clinical Time Series Data. *arXiv:1703.07771 [cs, stat]* (March 2017). <http://arxiv.org/abs/1703.07771> arXiv: 1703.07771.
- [42] Michael B. Hawes. 2020. Implementing Differential Privacy: Seven Lessons From the 2020 United States Census. *Harvard Data Science Review* (30 4 2020). <https://doi.org/10.1162/99608f92.353c6f99> <https://hdsr.mitpress.mit.edu/pub/dgg03v06>.
- [43] Asefeh Heiat, Cary P Gross, and Harlan M Krumholz. 2002. Representation of the elderly, women, and minorities in heart failure clinical trials. *Archives of internal medicine* 162, 15 (2002).
- [44] Diana Herrera-Perez, Alyson Haslam, Tyler Crain, Jennifer Gill, Catherine Livingston, Victoria Kaestner, Michael Hayes, Dan Morgan, Adam S Cifu, and Vinay Prasad. 2019. Meta-Research: A comprehensive review of randomized clinical trials in three medical journals reveals 396 medical reversals. *Elife* 8 (2019), e45183.
- [45] G. Hripsak, P.B. Ryan, J.D. Duke, N.H. Shah, R.W. Park, V. Huser, M.A. Suchard, M.J. Schuemie, F.J. DeFalco, A. Perotte, et al. 2016. Characterizing treatment pathways at scale using the OHDSI network. *Proceedings of the National Academy of Sciences* 113, 27 (2016), 7329–7336.
- [46] Matthew Jagielski, Jonathan Ullman, and Alina Oprea. 2020. Auditing Differentially Private Machine Learning: How Private is Private SGD? *arXiv preprint arXiv:2006.07709* (2020).
- [47] Bargav Jayaraman and David Evans. 2019. Evaluating differentially private machine learning in practice. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 1895–1912.
- [48] Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific Data* 3, 1 (May 2016), 1–9. <https://doi.org/10.1038/sdata.2016.35>
- [49] James Jordon, Daniel Jarrett, Jinsung Yoon, Tavian Barnes, Paul Elbers, Patrick Thoral, Ari Ercole, Cheng Zhang, Danielle Belgrave, and Mihaela van der Schaar. 2020. Hide-and-Seek Privacy Challenge. *arXiv preprint arXiv:2007.12087* (2020).
- [50] Christopher Jung, Katrina Ligett, Seth Neel, Aaron Roth, Saeed Sharif-Malvajerdi, and Moshe Shenfeld. 2019. A New Analysis of Differential Privacy’s Generalization Guarantees. *arXiv preprint arXiv:1909.03577* (2019).
- [51] Kenneth Jung and Nigam H Shah. 2015. Implications of non-stationarity on predictive modeling using EHRs. *Journal of biomedical informatics* 58 (2015), 168–174.
- [52] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*. 2564–2572.
- [53] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2019. An empirical study of rich subgroup fairness for machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 100–109.
- [54] Michael Kearns, Aaron Roth, Zhiwei Steven Wu, and Grigory Yaroslavtsev. 2015. Privacy for the protected (only). *arXiv preprint arXiv:1506.00242* (2015).
- [55] Christopher J Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. 2019. Key challenges for delivering clinical impact with artificial intelligence. *BMC medicine* 17, 1 (2019), 195.
- [56] Robert Kocher, Ezekiel J Emanuel, and Nancy-Ann M DeParle. 2010. The Affordable Care Act and the future of clinical medicine: the opportunities and challenges. *Annals of internal medicine* 153, 8 (2010), 536–539.
- [57] Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 1885–1894.
- [58] Pang Wei W Koh, Kai-Siang Ang, Hubert Teo, and Percy S Liang. 2019. On the accuracy of influence functions for measuring group effects. In *Advances in Neural Information Processing Systems*. 5255–5265.
- [59] Yann LeCun, Corinna Cortes, and CJ Burges. 2010. MNIST handwritten digit database. (2010).
- [60] Ninghui Li, Tiancheng Li, and Suresh Venkatasubramanian. 2007. t-closeness: Privacy beyond k-anonymity and l-diversity. In *2007 IEEE 23rd International Conference on Data Engineering*. IEEE, 106–115.
- [61] Jingcheng Liu and Kunal Talwar. 2019. Private selection from private candidates. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 298–309.
- [62] Toni Martin. 2011. The color of kidneys. *American Journal of Kidney Diseases* 58, 5 (2011), A27–A28.
- [63] J Michael McGinnis, Leigh Stuckhardt, Robert Saunders, Mark Smith, et al. 2013. *Best care at lower cost: the path to continuously learning health care in America*. National Academies Press.
- [64] Liangyuan Na, Cong Yang, Chi-Cheng Lo, Fangyuan Zhao, Yoshimi Fukuoka, and Anil Aswani. 2018. Feasibility of reidentifying individuals in large national physical activity data sets from which protected health information has been removed with use of machine learning. *JAMA network open* 1, 8 (2018), e186040–e186040.
- [65] A Narayanan and V Shmatikov. 2008. Robust de-anonymization of large sparse datasets [Netflix]. In *IEEE Symposium on Research in Security and Privacy*, Oakland, CA.
- [66] Seth Neel, Aaron Roth, Giuseppe Vietri, and Zhiwei Steven Wu. 2019. Differentially private objective perturbation: Beyond smoothness and convexity. *arXiv preprint arXiv:1909.01783* (2019).
- [67] Bret Nestor, Matthew B A McDermott, Willie Boag, Gabriela Berner, Tristan Naumann, Michael C Hughes, Anna Goldenberg, and Marzyeh Ghassemi. 2019. Feature Robustness in Non-stationary Health Records: Caveats to Deployable Model Performance in Common Clinical Machine Learning Tasks.
- [68] Thông T Nguyễn, Xiaokui Xiao, Yin Yang, Siu Cheung Hui, Hyejin Shin, and Junbum Shin. 2016. Collecting and analyzing data from smart device users with local differential privacy. *arXiv preprint arXiv:1606.05053* (2016).
- [69] Kobbi Nissim and Uri Stemmer. 2015. On the generalization properties of differential privacy. *CoRR, abs/1504.05800* (2015).
- [70] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [71] K.J. O’malley, K.F. Cook, M.D. Price, K.R. Wildes, J.F. Hurdle, and C.M. Ashton. 2005. Measuring diagnoses: ICD code accuracy. *Health Services Research* 40, 5p2 (2005), 1620–1639.
- [72] Jennifer M Orsi, Helen Margellos-Anast, and Steven Whitman. 2010. Black-white health disparities in the United States and Chicago: a 15-year progress analysis. *American journal of public health* 100, 2 (2010), 349–356.
- [73] Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian Goodfellow, and Kunal Talwar. 2016. Semi-supervised Knowledge Transfer for Deep Learning from Private Training Data. (Oct. 2016). [arXiv:1610.05755 \[stat.ML\]](https://arxiv.org/abs/1610.05755)
- [74] Nicolas Papernot, Steve Chien, Shuang Song, Abhradeep Thakurta, and Úlfar Erlingsson. 2020. Making the Shoe Fit: Architectures, Initializations, and Tuning for Learning with Privacy. <https://openreview.net/forum?id=rJg851rYwH>
- [75] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (Oct. 2011), 2825–2830. <http://jmlr.csail.mit.edu/papers/v12/pedregosa11a.html>
- [76] Stephen R Pfohl, Andrew M Dai, and Katherine Heller. 2019. Federated and Differentially Private Learning for Electronic Health Records. *arXiv preprint arXiv:1911.05861* (2019).
- [77] Stephen R Pfohl, Agata Foryciarz, and Nigam H Shah. 2020. An Empirical Characterization of Fair Machine Learning For Clinical Risk Prediction. *arXiv preprint arXiv:2007.10306* (2020).
- [78] Stephan Rabanser, Stephan Günnemann, and Zachary Lipton. 2019. Failing loudly: an empirical study of methods for detecting dataset shift. In *Advances in Neural Information Processing Systems*. 1394–1406.
- [79] Alvin Rajkomar, Michaela Hardt, Michael D Howell, Greg Corrado, and Marshall H Chin. 2018. Ensuring fairness in machine learning to advance health equity. *Annals of internal medicine* 169, 12 (2018), 866–872.
- [80] Alvin Rajkomar, Eyal Oren, Kai Chen, Andrew M Dai, Nissan Hajaj, Michaela Hardt, Peter J Liu, Xiaobing Liu, Jake Marcus, Mimi Sun, et al. 2018. Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine* 1, 1 (2018), 18.
- [81] Sherri Rose. 2018. Machine learning for prediction in electronic health data. *JAMA network open* 1, 4 (2018), e181404–e181404.
- [82] Christopher G Schwarz, Walter K Kremers, Terry M Therneau, Richard R Sharp, Jeffrey L Gunter, Prashanthi Vemuri, Arvin Arani, Anthony J Spychalla, Kejal Kantarci, David S Knopman, et al. 2019. Identification of anonymous MRI research participants with face-recognition software. *New England Journal of Medicine* 381, 17 (2019), 1684–1686.
- [83] Laleh Seyyed-Kalantari, Guanxiong Liu, Matthew McDermott, and Marzyeh Ghassemi. 2020. CheXclusion: Fairness gaps in deep chest X-ray classifiers. *arXiv preprint arXiv:2003.00827* (2020).
- [84] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3–18.
- [85] Shuang Song, Om Thakkar, and Abhradeep Thakurta. 2020. Characterizing private clipped gradient descent on convex generalized linear problems. *arXiv preprint arXiv:2006.06783* (2020).
- [86] Stephen M Stigler. 1989. Francis Galton’s account of the invention of correlation. *Statist. Sci.* (1989), 73–79.
- [87] Karien Stronks, Nicolien F Wieringa, and Anita Hardon. 2013. Confronting diversity in the production of clinical evidence goes beyond merely including

- under-represented groups in clinical trials. *Trials* 14, 1 (2013), 1–6.
- [88] Adarsh Subbaswamy, Peter Schulam, and Suchi Saria. 2018. Preventing failures due to dataset shift: Learning predictive models that transport. *arXiv preprint arXiv:1812.04597* (2018).
 - [89] Latanya Sweeney. 2002. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, 05 (2002), 557–570.
 - [90] Latanya Sweeney. 2015. Only you, your doctor, and many others may know. *Technology Science* 2015092903, 9 (2015), 29.
 - [91] Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang. 2017. Privacy loss in apple's implementation of differential privacy on macos 10.12. *arXiv preprint arXiv:1709.02753* (2017).
 - [92] Nenad Tomašev, Xavier Glorot, Jack W Rae, Michal Zielinski, Harry Askham, Andre Saraiva, Anne Mottram, Clemens Meyer, Suman Ravuri, Ivan Protsyuk, Alistair Connell, Cian O Hughes, Alan Karthikesalingam, Julien Cornebise, Hugh Montgomery, Geraint Rees, Chris Laing, Clifton R Baker, Kelly Peterson, Ruth Reeves, Demis Hassabis, Dominic King, Mustafa Suleyman, Trevor Back, Christopher Nielson, Joseph R Ledsam, and Shakir Mohamed. 2019. A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature* 572, 7767 (Aug. 2019), 116–119.
 - [93] Eric J Topol. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nature medicine* 25, 1 (2019), 44–56.
 - [94] Justin Travers, Suzanne Marsh, Mathew Williams, Mark Weatherall, Brent Caldwell, Philippa Shirlcliffe, Sarah Aldington, and Richard Beasley. 2007. External validity of randomised controlled trials in asthma: to whom do the results of the trials apply? *Thorax* 62, 3 (2007), 219–223.
 - [95] Salil Vadhan. 2017. The complexity of differential privacy. In *Tutorials on the Foundations of Cryptography*. Springer, 347–450.
 - [96] Giuseppe Vietri, Borja Balle, Akshay Krishnamurthy, and Zhiwei Steven Wu. 2020. Private Reinforcement Learning with PAC and Regret Guarantees. *arXiv preprint arXiv:2009.09052* (2020).
 - [97] Darshali A Vyas, Leo G Eisenstein, and David S Jones. 2020. Hidden in Plain Sight—Reconsidering the Use of Race Correction in Clinical Algorithms.
 - [98] Shirley Wang, Matthew B. A. McDermott, Geeticka Chauhan, Michael C. Hughes, Tristan Naumann, and Marzyeh Ghassemi. 2019. MIMIC-Extract: A Data Extraction, Preprocessing, and Representation Pipeline for MIMIC-III. *arXiv:1907.08322 [cs, stat]* (July 2019). <http://arxiv.org/abs/1907.08322> arXiv: 1907.08322.
 - [99] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers. 2017. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2097–2106.
 - [100] Denny Wu, Hirofumi Kobayashi, Charles Ding, Lei Cheng, and Keisuke Goda Marzyeh Ghassemi. 2019. Modeling the Biological Pathology Continuum with HSIC-regularized Wasserstein Auto-encoders. *arXiv:1901.06618 [cs, stat]* (Jan. 2019). <http://arxiv.org/abs/1901.06618> arXiv: 1901.06618.
 - [101] Xi Wu, Fengang Li, Arun Kumar, Kamalika Chaudhuri, Somesh Jha, and Jeffrey Naughton. 2017. Bolt-on differential privacy for scalable stochastic gradient descent-based analytics. In *Proceedings of the 2017 ACM International Conference on Management of Data*. 1307–1322.
 - [102] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017).
 - [103] S. Yu, Y. Ma, J. Grönsbell, T. Cai, A.N. Ananthakrishnan, V.S. Gainer, S.E. Churchill, P. Szolovits, S.N. Murphy, I.S. Kohane, et al. 2017. Enabling phenotypic big data with PheNorm. *Journal of the American Medical Informatics Association* 25, 1 (2017), 54–60.
 - [104] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2016. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530* (2016).