



HAL
open science

An Empirical Analysis of Fairness Notions under Differential Privacy *

Anderson Santana de Oliveira, Caelin Kaplan, Khawla Mallat, Tanmay
Chakraborty

► **To cite this version:**

Anderson Santana de Oliveira, Caelin Kaplan, Khawla Mallat, Tanmay Chakraborty. An Empirical Analysis of Fairness Notions under Differential Privacy *. PPAI 2023 - Fourth AAAI Workshop on Privacy-Preserving Artificial Intelligence, Feb 2023, Washington D.C., United States. hal-04387685

HAL Id: hal-04387685

<https://inria.hal.science/hal-04387685v1>

Submitted on 11 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

An Empirical Analysis of Fairness Notions under Differential Privacy*

Anderson Santana de Oliveira,¹ Caelin Kaplan,² Khawla Mallat¹ Tanmay Chakraborty³

¹ SAP

² SAP and INRIA

³ SAP and Eurecom

firstname.lastname@sap.com

Abstract

Recent works have shown that selecting an optimal model architecture suited to the differential privacy setting is necessary to achieve the best possible utility for a given privacy budget using differentially private stochastic gradient descent (DP-SGD) (Tramèr and Boneh 2020; Cheng et al. 2022). In light of these findings, we empirically analyse how different fairness notions, belonging to distinct classes of statistical fairness criteria (independence, separation and sufficiency), are impacted when one selects a model architecture suitable for DP-SGD, optimized for utility. Using standard datasets from ML fairness literature, we show using a rigorous experimental protocol, that by selecting the optimal model architecture for DP-SGD, the differences across groups concerning the relevant fairness metrics (demographic parity, equalized odds and predictive parity) more often decrease or are negligibly impacted, compared to the non-private baseline, for which optimal model architecture has also been selected to maximize utility. These findings challenge the understanding that differential privacy will necessarily exacerbate unfairness in deep learning models trained on biased datasets.

Introduction

Bias is frequently present in data used for machine learning. Today, systems responsible for controlling high stakes aspects of our societies, such as hiring, healthcare, social benefit allocations, or yet admissions to higher education, may produce unfair outcomes for underprivileged groups. This concern has motivated a considerable amount of research effort (see Barocas, Hardt, and Narayanan 2019; Chouldechova and Roth 2018).

Research on fairness has mainly been concentrated on defining fairness notions, creating interventions on ML training to achieve a given notion of fairness, and applying the approach to a variety of datasets in a supervised learning setting (see Feldman et al. 2015; Hardt, Price, and Srebro 2016; Zafar et al. 2017a,b). However, fairness is not the only concern that ML models face. They are also subject to membership inference attacks (MIAs), where an adversary can infer if some dataset he crafted or collected has been used

during a given ML model’s training (Hilprecht, Härterich, and Bernau 2019; Shokri et al. 2017).

Researchers have considered how defending against these attacks, most notably using differential privacy (Bagdasaryan, Poursaeed, and Shmatikov 2019; Xu, Du, and Wu 2020; Noe, Herskind, and Søgaard 2022), impacts fairness. These studies focus on the fact that the accuracy reduction introduced by differential privacy is more accentuated for underprivileged groups. Alternatively, we show that privacy protection does not always have a negative effect when it comes to disparities across groups.

Recent work has demonstrated that it is necessary to select an optimal model architecture when training a machine learning model with differentially private stochastic gradient descent (DP-SGD) (Abadi et al. 2016), rather than simply adding DP-SGD to the best identified non-private model (Tramèr and Boneh 2020; Cheng et al. 2022; De et al. 2022).

Building on these results, we show empirically that differential privacy does not systematically amplify disparities across groups, when we examine certain fairness notions. We measured selected indicators supporting fairness notions that belong to opposing statistical non-discrimination criteria: independence, separation and sufficiency, using a rigorous experimental protocol.

By using datasets typically examined in ML fairness research, we observed that by selecting the optimal model architecture with DP-SGD, either the fairness metrics are negligibly impacted or (more often) the differences across groups decrease, when compared to the non-private baseline, for which optimal model architecture and hyperparameter search has also been conducted for reaching maximum utility (represented by the ROC AUC score).

We selected ROC AUC score because it is robust to class imbalance and also independent of a decision threshold. We observed that the ROC AUC score on the test set remains close for the best private and non-private models. The fact that we evaluate utility only, allows us to compare to previous research, which simply incorporate a DP-SGD optimizer to the best non-private baseline architecture that have found that DP-SGD disparately impacts model’s accuracy.

Another important differentiation in this work is that we did not exclude minorities from the analysis, nor outliers, while it may be a widespread practice in data science and

*Accepted for oral presentation at the The Fourth AAAI Workshop on Privacy-Preserving Artificial Intelligence (PPAI-23) Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

in (machine learning) research. We measure the fairness disparities in realistic situation where the fairness indicators on all minority groups and their intersections, as we consider two protected attributes for each dataset (race and sex).

We draw our conclusions after performing exhaustive grid-search, for the private and non-private settings (including activation function, optimizer, number of hidden layers, dropout probability, learning rate, ...). The private grid search procedure uses randomization, consuming privacy budget. We measure key group fairness metrics using k-fold cross validation for each dataset to measure equalized odds, predictive parity and demographic parity differences.

The remainder of this paper is organized as follows: The Section Background brings the definitions on fairness notions and differential privacy required for understanding the contributions here. In Sec. Related Work, we discuss several papers about the interplay between fairness and differential privacy. Sec. Methodology describes the configurable model architecture we used and the privacy-preserving model selection process. Next, the Sec. Experimental Evaluation presents the main findings. Finally Sec. Conclusion closes the paper.

Background

This section provides an overview of the main concepts necessary to understand the present paper.

Differential Privacy

Differential Privacy is the flagship definition of data privacy. Differential Privacy ensures that approximately nothing can be learned about an individual whether she/he participates or not in a database. Differential Privacy defines a constraint on the processing of data such that the output of two adjacent databases is approximately the same. More formally:

Definition 1 *A randomized algorithm M gives (ϵ, δ) -differential privacy, if for all databases X_1 and X_2 differing on at most one element, and all $S \subseteq \text{Range}(M)$, $\Pr[M(X_1) \in S] \leq \exp(\epsilon) \times \Pr[M(X_2) \in S] + \delta$ (Dwork and Roth 2014).*

This condition encapsulates the crucial notion of indistinguishability of the results of a database manipulation whether someone participates or not in that database. The confidence parameter, ϵ , is referred to as the privacy budget. Note that as ϵ decreases, the output of the mechanism becomes more private. The values of interest for δ are less than the inverse of any polynomial in the size of the database. Any function, M , that satisfies the Differential Privacy condition can be used to generate data that guarantees the privacy of the individuals in the database. One interesting property of Differential Privacy is composability. If M and M' are ϵ and ϵ' differentially private, respectively, then their composition, $M \circ M'$, is $(\epsilon + \epsilon')$ -differentially private (Dwork 2008). This property allows a user to create and combine many different mechanisms to achieve differential privacy.

Rényi Differential Privacy Rényi Differential Privacy (RDP) (Mironov 2017) is a generalization of pure

differential privacy that uses the Rényi divergence of order λ between two distributions ($D(M(X_1))$ and $D(M(X_2))$), such that ϵ becomes a function $\epsilon(\lambda)$, rather than simply a scalar value. It offers simple and precise composition property, allowing to better control how much privacy budget is spent.

DP-SGD Abadi et al. (2016) developed an algorithm to insert differential privacy into the training procedure of a deep learning model. This method requires truncating the computed gradients and adding random noise to obfuscate the influence of a single input data on the final model. The training process can be defined as follows: first, a set of inputs, called lots, is processed by the model and the loss is computed; next, an array, composed of the gradients for each weight, is calculated with respect to the loss; then, the gradients are clipped, returning the minimum value between the norm of the gradients and an upper bound defined as a hyperparameter of the differential privacy mechanism; finally, Noise, coming from a normal distribution, is added to the sum of clipped gradients with a variance proportional to the aforementioned upper bound and a noise multiplier parameter. The result is a sanitized version of the gradients, in which the influence of the input data is bounded, guaranteeing privacy. Consequently, the resulting model will have less accuracy than a comparable baseline model, trained without noise addition.

Fairness Notions

In this section, we briefly define the group fairness notions that we analyze in this work. We selected notions from opposing classes of statistical non-discrimination criteria: independence, separation, sufficiency (Barocas, Hardt, and Narayanan 2019), which are properties about the joint distributions of sensitive attributes, the prediction target variable, and the outcome of a classifier, formally taking the form of conditional independence statements among the aforementioned variables (Castelnuovo et al. 2022). Independence is represented in our work by demographic parity, separation by equalized odds, while sufficiency is interpreted as predictive parity. We do not consider whether these notions are adequate for the use cases discussed. Instead, we focus on the effects of differential privacy on fairness indicators related to these notions.

Demographic Parity A predictor that satisfies the demographic parity constraint should yield that each subgroup receives a positive outcome at equal rates. Demographic parity can be expressed as follows:

$$P(\hat{Y} = 1 | A = a) = P(\hat{Y} = 1 | A = b), \quad \forall a, b \in A$$

where \hat{Y} is the predicted label and A is the protected attribute indicating the subgroup. The notion of demographic parity helps prevent the reinforcement of historical biases and supports the underprivileged groups in the short term due to the progressive enforcement of a positive feedback loop. However, ensuring demographic parity focuses only on the final outcome and not on the equality of treatment,

which can lead to a problem called laziness in (Makhlouf, Zhioua, and Palamidessi 2021; Carey and Wu 2022), or self-fulfilling prophecy in (Dwork et al. 2012), where a trained model selects the true positives of the privileged groups while selecting randomly (i.e., with a coin toss) subjects from the underprivileged groups as long as the number of subjects selected from each subgroup is valid. In addition, demographic parity could be used in an inappropriate context, where the disparity is truly present but is not related to a protected attribute (Dwork et al. 2012).

In this paper, we use demographic parity difference, which is defined as the difference between the highest and the lowest group-level selection rate. A demographic parity difference of 0 means that all groups have the same selection rate.

Equalized Odds A predictor that satisfies the equalized odds constraint should predict each of the labels equally well for all the subgroups (Hardt, Price, and Srebro 2016). Equalized odds can be expressed as follows:

$$P(\hat{Y} = 1|Y = y, A = a) = P(\hat{Y} = 1|Y = y, A = b) \\ \forall y \in 0, 1 \wedge \forall a, b \in A$$

where A is the protected attribute, \hat{Y} is the predicted label, and Y is the ground truth label. Unlike demographic parity, equalized odds can ensure the equality of treatment, which eliminates the laziness issue above-mentioned. However, the notion of equalized odds would not help dealing with bias in the long term. It does not take into consideration potential bias outside of the model, which can lead to enforcing the bias, particularly in cases where an extreme disproportionality is present between different subgroups.

In this paper, we are interested in evaluating the maximum equalized odds difference across groups. This is given by whichever is greater between the true positive rate difference and false positive rate difference between two groups. An equalized odds difference of 0 means that all groups have the same true positive, true negative, false positive, and false negative rates.

Predictive Parity Predictive Parity regards the error parity among groups who are assigned the same outcome:

$$P(Y = 1|A = a, \hat{Y} = 1) = P(Y = 1|A = b, \hat{Y} = 1) \\ \forall a, b \in A$$

That is, the model should have the same precision across sensitive groups. Predictive rate parity is achieved if the positive predictive values in the subgroups are close to each other. Precision stands for the number of the true positives divided by the total number of examples predicted positive within a group. As for the previous notions, we will assess maximum predictive parity difference across groups, meaning that a difference of 0 mean that the classifier achieves the same precision score for all groups.

It is important to note that independence, separation and sufficiency cannot hold at the same time in conjunction with one another (see Barocas, Hardt, and Narayanan 2019,

Chapter 3): it means that when one assess datasets like the ones we use in the present work, where the sensitive attribute is not independent from the target label, imposing more than one fairness criteria at the same time will only leave degenerate solutions.

Related Work

In this section, we discuss some papers approaching fairness concerns under differential privacy. In (Bagdasaryan, Poursaeed, and Shmatikov 2019), several datasets were used to study the impact of DP-SGD on underrepresented groups. Training on the data of the underrepresented subgroups produces larger gradients. Thus, clipping more substantially biases these gradient values and limits the influence of their data on the model. They also show that random noise addition has the biggest impact on the underrepresented inputs. However, in some of the experiments, the group membership is confused with the class labels, and minorities are introduced artificially via under-sampling. The impact of differential privacy on a given fairness notion is not considered.

Pujol et al. (2020) studies fair allotment and decision rule problems using differential privacy: it shows that the noise brought by strong privacy guarantees, unsurprisingly, induces disparate decision errors concerning minorities. The empirical analysis uses extremely small privacy budgets (in the order of 10^{-5}).

The paper (Xu, Du, and Wu 2021) introduces a DP-SGD approach that assigns an adaptive gradient clipping norm per group in order to obtain a uniform accuracy loss per group (or per class). The experiments do not include how other fairness notions (demographic parity, equalized odds, etc.) are impacted by the approach.

In (Tran et al. 2021) evaluates excessive risk as the fairness notion when protecting privacy using the ensemble method PATE (Papernot et al. 2018). They show how the different model parameters affect the excessive risk difference across groups. They also show possible mitigation to the disparate effects using soft labels.

In (Jagielski et al. 2019), optimization procedures for achieving equalized odds fairness under differential privacy are proposed. The work poses the theoretical bounds between accuracy, privacy, and equalized odds. The experimental results show a clear negative correlation between privacy budget and fairness violation. Ganey, Oprisanu, and De Cristofaro (2022) demonstrate differential privacy’s disparate impact on generative models regarding minority/majority groups either by reducing or increasing the gap between classes or demographic groups in the generated synthetic data, with respect to the real data for both their proportions in the data, and accuracy of downstream classification tasks.

Cheng et al. (2022) showed that model utility (measured using test accuracy) is significantly improved when a specific search for the best network architecture is performed during model training with DP-SGD. The reader needs to be aware that the privacy budget values, reported in this work (Cheng et al. 2022), do not seem to account for the potential privacy leakage that occurs during the network architecture search itself in (Papernot and Steinke 2022).

In the present paper, we do not undersample any of the classes or demographic groups. We utilize the group intersections between gender and race, generating all subgroups in the existing datasets. Like most of the works mentioned in this section, we optimize the model training for utility only, not seeking to minimize the disparities in the baseline non-private nor for the privacy preserving models. This allows us to compare to these related works. Furthermore, we consider three different fairness notions, each corresponding to one of the statistical fairness criteria (demographic parity for independence, equalized odds for separation and predictive parity for sufficiency): we discovered that the best performing models, resulting from a hyperparameter search under differential privacy constraints, often reduces the disparities across groups as compared to the non-private baseline, and if not, it maintains similar levels of disparities.

Methodology

Our motivation is to understand how privacy protection impacts fairness in real-world settings. Therefore, we selected on datasets analyzed in the ML fairness literature that contain data gathered from use cases where fairness and privacy are genuine issues. Unlike previous work, we do not limit our analysis to a set of artificially constructed groups within the dataset. Instead, we leave the datasets intact and analyze the impact on fairness to all existing groups.

We defined a template for a feed-forward network, composed of an input block, a variable number of hidden blocks, and a final softmax layer. The input and hidden blocks contain each a linear, a normalization and a dropout layer (see Fig. 1 in the Appendix). Using this architecture, we explored multiple configurations ranges:

- Learning rate: 1^{-4} , 1^{-3} , 1^{-2} , 1^{-1}
- Dropout probability: 0.0, 0.1, 0.2
- Number of hidden blocks: 1 to 3
- Batch size: 256, 512
- Activation functions: Relu, Tanh
- Optimizers: Adam, SGD
- Maximum gradient norm (only relevant for differential): 0.01, 0.1, 1

We performed a 5-fold cross validation hyperparameter search on each dataset to find the best hyperparameter configuration with the goal of maximizing the overall model AUC score (area under the ROC curve). This metric is appropriate because it is insensitive to class imbalance and the decision threshold, as well as comprises False Negative and False Positive rates. Remark that We we optimise the utility without considering fairness interventions - so our usage of the word “best” stands uniquely for robust AUC score.

We performed the search for each of the datasets under four different setups: (i) With and without the protected attributes (gender and race); (ii) With and without differential privacy, where we applied DP-SGD using *Opacus* (Yousefpour et al. 2021). We performed hyperparameter search with the help of the *Ray tune* library¹. We have run all experiments on 2 GPUs on a NVIDIA DGX server.

¹<https://docs.ray.io/en/latest/tune/index.html>

For the search of the best model under differential privacy, we set the same privacy budget for all datasets ($\epsilon = 5.0$) using DP-SGD, for a single training run. This means that the total privacy spent can be roughly estimated to be $\epsilon = 20$ using the basic composition theorem of differential privacy, because we perform 5-fold cross-validation (each training data split will be used four times for training, once for validation). Furthermore, the randomized hyperparameter search also implies that additional privacy budget is used. Based on (Papernot and Steinke 2022, Theorem 6), which amounts to additional $\epsilon = 2$. Finally, we train the model for computing the final average metrics, adding up $\epsilon = 5$. Thus, we can make a conservative estimation of the total privacy budget of our procedure to be $\epsilon = 27$ for the results we report in this paper.

Next, we trained models for each of the datasets using the best configuration for three settings: *Baseline*, without privacy protection; *Baseline+DP* which is the baseline model configuration with the addition of differential privacy; and *Best DP Model*). It is important to remark that the *Baseline+DP* case, we did not adjust any other hyperparameters, and used the privacy budget $\epsilon = 27$. Using the default model with differential privacy approximates what has occurred in most previous research, where differential privacy was added to the baseline models with minor adjustments, such as batch size or learning rate.

We used a hold-out test set (i.e., not used during the 5-fold cross validation phase) to collect all the average metric values on 10 model training runs, using the *Fairlearn* library (Kulshrestha and Safro 2021).

Experimental Evaluation

Datasets

We used real-world datasets containing demographic information, widely used in machine learning fairness research, all of them present biases of different kinds in the input data. Unlike other works, which analysed fairness aspects using only two demographic groups based on sex, we kept all the original race categories present in all datasets, because we want to understand how all identified minorities are actually affected by these classifiers.

ACS Employment The ACS employment dataset (Ding et al. 2021) is derived from the American Community Survey released by the US Census bureau, where the feature informing if a given individual was employed in the year of the data collection is transformed into the prediction target. Here, we use 10% of the survey results for the state of California in 2018. It includes features such as educational attainment, marital status, citizenship, employment status of parents, among others.

ACS Income The ACS income dataset (Ding et al. 2021) was created as an alternative to the UCI Adult dataset. The task is to predict if a person’s income is above \$50k/yr. As for the ACS Employment dataset, we use a 10% sample of data from the ACS survey results for the state of California in 2018.

Dataset	Metric	Baseline	Baseline + DP	Best DP Model
ACS Emp.	Overall AUC	0.8837 \pm 0.0011	0.8110 \pm 0.0062	0.8702 \pm 0.0013
	AUC difference	0.3401 \pm 0.0875	0.3134 \pm 0.0653	0.2073 \pm 0.0480
	Demographic parity difference	0.4383 \pm 0.0805	0.5317 \pm 0.1561	0.3154 \pm 0.0359
	Equalized odds difference	0.5884 \pm 0.1360	0.6623 \pm 0.1701	0.2874 \pm 0.0534
	Precision difference	0.5534 \pm 0.0998	0.4833 \pm 0.1296	0.2968 \pm 0.0674
ACS Inc.	Overall AUC	0.8878 \pm 0.0011	0.8155 \pm 0.0045	0.8820 \pm 0.0008
	AUC difference	0.2546 \pm 0.0569	0.3498 \pm 0.0867	0.2225 \pm 0.0219
	Demographic parity difference	0.4223 \pm 0.0613	0.6105 \pm 0.1234	0.2556 \pm 0.0490
	Equalized odds difference	0.4360 \pm 0.0780	0.8499 \pm 0.1656	0.3756 \pm 0.0019
	Precision difference	0.3550 \pm 0.0518	0.7882 \pm 0.2009	0.4032 \pm 0.0271
LSAC	Overall AUC	0.8343 \pm 0.0029	0.7755 \pm 0.0125	0.7962 \pm 0.0077
	AUC difference	0.0435 \pm 0.0056	0.0422 \pm 0.0142	0.0575 \pm 0.0128
	Demographic parity difference	0.3064 \pm 0.0653	0.2007 \pm 0.0437	0.1687 \pm 0.0151
	Equalized odds difference	0.2548 \pm 0.0862	0.2853 \pm 0.0467	0.1975 \pm 0.0722
	Precision difference	0.1688 \pm 0.0485	0.2202 \pm 0.0061	0.2197 \pm 0.0082
Adult	Overall AUC	0.9056 \pm 0.0011	0.8476 \pm 0.0073	0.9005 \pm 0.0009
	AUC difference	0.1264 \pm 0.0249	0.2226 \pm 0.0556	0.1841 \pm 0.0599
	Demographic parity difference	0.2750 \pm 0.0155	0.4683 \pm 0.0612	0.2375 \pm 0.0207
	Equalized odds difference	0.7845 \pm 0.0492	0.8005 \pm 0.0472	0.8000 \pm 0.0000
	Precision difference	0.9400 \pm 0.0966	0.7567 \pm 0.1933	0.8000 \pm 0.0000
Compas	Overall AUC	0.6895 \pm 0.0041	0.5349 \pm 0.0359	0.6863 \pm 0.0030
	AUC difference	0.1162 \pm 0.0273	0.1295 \pm 0.0390	0.0824 \pm 0.0264
	Demographic parity difference	0.5101 \pm 0.0209	0.3322 \pm 0.1331	0.3694 \pm 0.0230
	Equalized odds difference	0.5592 \pm 0.0476	0.3905 \pm 0.1356	0.3726 \pm 0.0375
	Precision difference	0.3347 \pm 0.0749	0.3009 \pm 0.1052	0.3168 \pm 0.0467

Table 1: Mean and standard deviation values for 10 training runs for the best model configurations with and without privacy. The difference metrics report only the highest different for any two groups for each dataset. The total privacy budget is $\epsilon = 27$ for the *Baseline + DP* and *Best DP* models.

LSAC The Law School dataset (Wightman 1998) from the law school admissions council’s national longitudinal bar passage study to predict whether a candidate would pass the bar exam. It consists of law school admission records, including gender, race, features like family income, and others derived from the GPA. The classification target is the binary feature “isPassBar”.

Adult The UCI Adult dataset (Dua and Graff 2017) contains US census income survey records. We use the binarized “income” feature as the target variable for our classification task to predict if an individual’s income is above \$50k/yr.

COMPAS The COMPAS dataset (Angwin et al. 2016) for recidivism risk prediction consists of records comprising convicted persons criminal history, including gender and race as demographic features. We use the ground truth of whether the offender was re-arrested (binary) as the target variable for classification. Notice that the positive outcome (True) is the undesirable result. Moreover, the majority group (Black males) obtains higher positive rates.

We transformed all categorical attributes using one-hot encoding and standardized all features vectors to have zero mean and unit variance.

Main Results

Table 1 shows our results on holdout test data for five distinct datasets with the following metrics: overall AUC (for all groups of individuals, the higher the better), maximum AUC difference across groups, maximum demographic parity difference across groups, maximum equalized odds difference across groups, and maximum predictive parity difference (for these last four metrics, the lower the better). Here, we use a differential privacy budget (epsilon) equivalent to 27 for each privacy-preserving version. That is, we explicitly set $\epsilon = 27$ for the *Baseline + DP* models, whereas we set it to $\epsilon = 5.0$ for the *Best DP model* settings, and account for the privacy budget spent during the hyperparameter search phase. We set the decision threshold to 0.5, which is not relevant for the AUC scores, but which influences the other metrics.

In the reported results, we decided to omit the case where the protected attributes were not considered during training. Dwork et al. (Dwork et al. 2012) show that removing the sensitive attributes does not help mitigate bias because these features are correlated with other proxy features. Additionally, by omitting the sensitive attributes, we observed in our experiments that the AUC score declined significantly for all datasets and in all settings (i.e., with or without privacy).

Due to the decline in AUC score, assessing the impact on fairness notions is not useful because the model’s utility is not satisfactory.

In Table 1, the *Baseline* column reports the metric values for the best model configuration found during the grid search for the setting without privacy protection optimized for maximum overall AUC. The *Baseline + DP* corresponds to the best baseline model configuration, now trained with DP-SGD, and finally, the *Best DP Model* column lists the metric values for the model search was conducted while training the model with DP-SGD. Consequently, the latter has different architecture and hyperparameters than the previous ones, but they have been trained on the same data, with the same objective function.

The baseline column confirms the existence of biases in the datasets. There is no single scale for measuring sufficiency, separation and independence, but the measured differences across groups are sufficiently large to imagine that in most use cases these level of disparity would not be acceptable, and that fairness interventions would be required to reduce disparities according to the fairness notion of choice for a given dataset. Here we want to observe how these fairness metrics are affected by the use of DP-SGD. Group specific metrics are found in the appendix for each of the datasets.

We can observe that the *Best DP Model* outperforms the *Baseline + DP* for basically all metrics on all datasets. This is not so surprising, in particular, because we do not perform any additional fine-tuning (e.g. adjusting the learning rate) for the *Baseline + DP* models. On the other hand, what is more remarkable is the fact that in many cases, differential privacy maintained a very good utility when compared to the baseline models - with slight decreases in overall AUC scores - while decreasing the disparities with respect to two or more of the fairness metrics. The best examples can be observed for the ACS Employment and ACS Income datasets, where the Overall AUC is basically equal to their respective non-private baselines, and all differences across groups have drastically decreased. On the other side of the spectrum, we do not observe substantial decreases in disparities for the Adult dataset, but differential privacy does not seem to exacerbate biases. The regularization effect of noise addition and gradient clipping brought by DP-SGD is likely decreasing the tendency of the neural network to overfit the majority groups in most of the datasets. However, a specific analysis is needed for each dataset because the group distributions are too diverse to make general statements.

These results show that differential privacy will not systematically magnify disparities when compared to the non-private baseline. One needs to reconsider the findings in previous research that have exposed significantly worse utility and increased unfairness, which basically correspond to the *Baseline + DP* setting, where all indicators are consistently degraded.

Therefore, our experiments indicate that in addition to protecting against privacy attacks, differential privacy does not exacerbate unequal access to opportunities, even when opposing fairness notions are under consideration. Since the influence of each training input data is bounded, the decision

frontiers shift, resulting in positive outcomes assigned more often to the underprivileged (or minority) groups. These conclusions are supported by further fairness indicators for each subgroup in the Cartesian product between the sensitive attributes for sex and gender on each dataset (see the Appendix).

Conclusions

Previous research has insisted on the fact that differential privacy adversely impacts accuracy for underprivileged groups. While this observation may be the case when adding differential privacy to an existing non-private model configuration, by using a model resulting from a hyperparameter search under differential privacy constraints, one can observe that the resulting model often displays less disparities for fairness notions belonging to opposing statistical non-discrimination criteria, such as demographic parity, equalized odds and predictive parity, when compared to its non-privacy preserving counterpart found following the same model architecture and hyperparameter search procedures.

References

- Abadi, M.; Chu, A.; Goodfellow, I. J.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep Learning with Differential Privacy. In Weippl, E. R.; Katzenbeisser, S.; Kruegel, C.; Myers, A. C.; and Halevi, S., eds., *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, 308–318. ACM.
- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2016. Machine bias. In *Ethics of Data and Analytics*, 254–264. Propublica.
- Bagdasaryan, E.; Poursaeed, O.; and Shmatikov, V. 2019. Differential Privacy Has Disparate Impact on Model Accuracy. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 15453–15462.
- Barocas, S.; Hardt, M.; and Narayanan, A. 2019. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- Carey, A. N.; and Wu, X. 2022. The Causal Fairness Field Guide: Perspectives From Social and Formal Sciences. *Frontiers Big Data*, 5: 892837.
- Castelnovo, A.; Crupi, R.; Greco, G.; Regoli, D.; Penco, I. G.; and Cosentini, A. C. 2022. A clarification of the nuances in the fairness metrics landscape. *Scientific Reports*, 12(1): 1–21.
- Cheng, A.; Wang, J.; Zhang, X. S.; Chen, Q.; Wang, P.; and Cheng, J. 2022. DPNAS: Neural Architecture Search for Deep Learning with Differential Privacy. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(6): 6358–6366.
- Chouldechova, A.; and Roth, A. 2018. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810*.

- De, S.; Berrada, L.; Hayes, J.; Smith, S. L.; and Balle, B. 2022. Unlocking High-Accuracy Differentially Private Image Classification through Scale. *CoRR*, abs/2204.13650.
- Ding, F.; Hardt, M.; Miller, J.; and Schmidt, L. 2021. Retiring Adult: New Datasets for Fair Machine Learning. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y. N.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 6478–6490.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository.
- Dwork, C. 2008. Differential Privacy: A Survey of Results. In Agrawal, M.; Du, D.; Duan, Z.; and Li, A., eds., *Theory and Applications of Models of Computation, 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings*, volume 4978 of *Lecture Notes in Computer Science*, 1–19. Springer.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. S. 2012. Fairness through awareness. In Goldwasser, S., ed., *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, 214–226. ACM.
- Dwork, C.; and Roth, A. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4): 211–407.
- Feldman, M.; Friedler, S. A.; Moeller, J.; Scheidegger, C.; and Venkatasubramanian, S. 2015. Certifying and Removing Disparate Impact. In Cao, L.; Zhang, C.; Joachims, T.; Webb, G. I.; Margineantu, D. D.; and Williams, G., eds., *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10-13, 2015*, 259–268. ACM.
- Ganev, G.; Oprisanu, B.; and De Cristofaro, E. 2022. Robin Hood and Matthew Effects: Differential Privacy Has Disparate Impact on Synthetic Data. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 6944–6959. PMLR.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of Opportunity in Supervised Learning. In Lee, D. D.; Sugiyama, M.; von Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, 3315–3323.
- Hilprecht, B.; Härterich, M.; and Bernau, D. 2019. Monte Carlo and Reconstruction Membership Inference Attacks against Generative Models. *Proc. Priv. Enhancing Technol.*, 2019(4): 232–249.
- Jagielski, M.; Kearns, M. J.; Mao, J.; Oprea, A.; Roth, A.; Sharifi-Malvajerdi, S.; and Ullman, J. R. 2019. Differentially Private Fair Learning. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 3000–3008. PMLR.
- Kulshrestha, A.; and Safro, I. 2021. FAIRLEARN: Configurable and Interpretable Algorithmic Fairness. *arXiv preprint arXiv:2111.08878*.
- Makhlouf, K.; Zhioua, S.; and Palamidessi, C. 2021. Machine learning fairness notions: Bridging the gap with real-world applications. *Inf. Process. Manag.*, 58(5): 102642.
- Mironov, I. 2017. Rényi Differential Privacy. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, 263–275. IEEE Computer Society.
- Noe, F.; Herskind, R.; and Søgaard, A. 2022. Exploring the Unfairness of DP-SGD Across Settings. *CoRR*, abs/2202.12058.
- Papernot, N.; Song, S.; Mironov, I.; Raghunathan, A.; Talwar, K.; and Erlingsson, Ú. 2018. Scalable Private Learning with PATE. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Papernot, N.; and Steinke, T. 2022. Hyperparameter Tuning with Rényi Differential Privacy. In *International Conference on Learning Representations*.
- Pujol, D.; McKenna, R.; Kuppam, S.; Hay, M.; Machanavajjhala, A.; and Miklau, G. 2020. Fair decision making using privacy-protected data. In Hildebrandt, M.; Castillo, C.; Celis, L. E.; Ruggieri, S.; Taylor, L.; and Zanfir-Fortuna, G., eds., *FAT* '20: Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27-30, 2020*, 189–199. ACM.
- Shokri, R.; Stronati, M.; Song, C.; and Shmatikov, V. 2017. Membership Inference Attacks Against Machine Learning Models. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, 3–18. IEEE Computer Society.
- Tramèr, F.; and Boneh, D. 2020. Differentially Private Learning Needs Better Features (or Much More Data). *CoRR*, abs/2011.11660.
- Tran, C.; Dinh, M. H.; Beiter, K.; and Fioretto, F. 2021. A Fairness Analysis on Private Aggregation of Teacher Ensembles. *CoRR*, abs/2109.08630.
- Wightman, L. F. 1998. LSAC National Longitudinal Bar Passage Study. LSAC Research Report Series. *LSAC Research Report Series*.
- Xu, D.; Du, W.; and Wu, X. 2020. Removing Disparate Impact of Differentially Private Stochastic Gradient Descent on Model Accuracy. *CoRR*, abs/2003.03699.
- Xu, D.; Du, W.; and Wu, X. 2021. Removing Disparate Impact on Model Accuracy in Differentially Private Stochastic Gradient Descent. In Zhu, F.; Ooi, B. C.; and Miao, C., eds., *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14-18, 2021*, 1924–1932. ACM.
- Yousefpour, A.; Shilov, I.; Sablayrolles, A.; Testuggine, D.; Prasad, K.; Malek, M.; Nguyen, J.; Ghosh, S.; Bharadwaj, A.; Zhao, J.; Cormode, G.; and Mironov, I. 2021. Opaucus: User-Friendly Differential Privacy Library in PyTorch. *arXiv preprint arXiv:2109.12298*.

Zafar, M. B.; Valera, I.; Gomez-Rodriguez, M.; and Gummadi, K. P. 2017a. Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment. In Barrett, R.; Cummings, R.; Agichtein, E.; and Gabrilovich, E., eds., *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017*, 1171–1180. ACM.

Zafar, M. B.; Valera, I.; Gomez-Rodriguez, M.; and Gummadi, K. P. 2017b. Fairness Constraints: Mechanisms for Fair Classification. In Singh, A.; and Zhu, X. J., eds., *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, volume 54 of *Proceedings of Machine Learning Research*, 962–970. PMLR.

Appendix




Figure 1: Representation of the predefined model architecture: the number of hidden blocks is part of the grid search.

Group distributions in the datasets

The tables here display the number of instances in the training data for each group, for all datasets among training and hold out test data.

Intersection Metrics

In this section we present several group metrics for all datasets. The decision threshold has been 0.5 for all threshold dependent metrics.

Race	Sex	Train	Test
Alaska Native	Female	3	1
American Indian	Female	109	53
	Male	82	35
Amerindian Tribes	Female	37	15
	Male	33	16
Asian	Female	2199	909
	Male	1956	844
Black	Female	644	312
	Male	625	283
Native Hawaiian	Female	33	16
	Male	39	20
Some Other Race	Female	1520	683
	Male	1538	672
Two or More Races	Female	688	286
	Male	700	267
White	Female	8345	3520
	Male	7966	3433

Table 2: ACS Employment group sizes

Race	Sex	Train	Test
Alaska Native	Female	2	0
American Indian	Female	97	47
	Male	75	56
Amerindian Tribes	Female	25	12
	Male	31	10
Asian	Female	2178	965
	Male	1887	847
Black	Female	618	283
	Male	645	290
Native Hawaiian	Female	45	14
	Male	52	21
Some Other Race	Female	1493	675
	Male	1629	697
Two or More Races	Female	692	259
	Male	725	309
White	Female	8277	3499
	Male	8046	3381

Table 3: ACS Income group sizes

Race	Sex	Train	Test
Black	Female	753	314
	Male	518	205
Other	Female	887	415
	Male	1072	451
White	Female	6493	2816
	Male	8862	3765

Table 4: LSAC group sizes

Race	Sex	Train	Test
Amer-Indian-Eskimo	Female	125	60
	Male	214	71
Asian-Pac-Islander	Female	379	138
	Male	703	299
Black	Female	1637	671
	Male	1655	722
Other	Female	113	42
	Male	181	70
White	Female	9065	3962
	Male	20117	8618

Table 5: Adult group sizes

Race	Sex	Train	Test
Black	Female	453	199
	Male	2113	931
Other	Female	131	45
	Male	639	249
White	Female	391	176
	Male	1322	564

Table 6: Compas group sizes

Metric	Group	Baseline	Baseline + DP	Best DP Model
Accuracy	American Indian-Female	0.7615 ± 0.0698	0.6146 ± 0.0688	0.6830 ± 0.0149
Accuracy	American Indian-Male	0.8154 ± 0.0671	0.7686 ± 0.0755	0.7571 ± 0.0179
Accuracy	Amerindian Tribes-Female	0.7077 ± 0.0487	0.7143 ± 0.1650	0.7636 ± 0.0469
Accuracy	Amerindian Tribes-Male	0.7643 ± 0.0894	0.7000 ± 0.0704	0.7692 ± 0.0000
Accuracy	Asian-Female	0.7615 ± 0.0090	0.7063 ± 0.0091	0.7473 ± 0.0052
Accuracy	Asian-Male	0.8251 ± 0.0082	0.7704 ± 0.0080	0.8355 ± 0.0047
Accuracy	Black-Female	0.7502 ± 0.0140	0.7287 ± 0.0230	0.8171 ± 0.0055
Accuracy	Black-Male	0.8365 ± 0.0138	0.7349 ± 0.0250	0.8354 ± 0.0077
Accuracy	Native Hawaiian-Female	0.6462 ± 0.1214	0.6933 ± 0.1142	0.8636 ± 0.0000
Accuracy	Native Hawaiian-Male	0.8375 ± 0.0437	0.8300 ± 0.0888	0.8000 ± 0.0471
Accuracy	Some Other Race-Female	0.7521 ± 0.0088	0.7285 ± 0.0050	0.7725 ± 0.0058
Accuracy	Some Other Race-Male	0.8477 ± 0.0079	0.8084 ± 0.0074	0.8612 ± 0.0081
Accuracy	Two or More Races-Female	0.8026 ± 0.0181	0.7907 ± 0.0100	0.8139 ± 0.0057
Accuracy	Two or More Races-Male	0.8911 ± 0.0124	0.8132 ± 0.0106	0.8920 ± 0.0037
Accuracy	White-Female	0.7743 ± 0.0026	0.7025 ± 0.0064	0.7624 ± 0.0018
Accuracy	White-Male	0.8283 ± 0.0052	0.7724 ± 0.0066	0.8283 ± 0.0018
Precision	American Indian-Female	0.6486 ± 0.0889	0.5160 ± 0.0839	0.5924 ± 0.0158
Precision	American Indian-Male	0.8531 ± 0.1051	0.8285 ± 0.1193	0.6816 ± 0.0213
Precision	Amerindian Tribes-Female	0.4150 ± 0.0747	0.6933 ± 0.2779	0.8893 ± 0.0588
Precision	Amerindian Tribes-Male	0.9650 ± 0.0568	0.4886 ± 0.1907	0.7143 ± 0.0000
Precision	Asian-Female	0.7185 ± 0.0136	0.6305 ± 0.0084	0.6626 ± 0.0055
Precision	Asian-Male	0.7978 ± 0.0087	0.7103 ± 0.0093	0.8048 ± 0.0047
Precision	Black-Female	0.6805 ± 0.0208	0.5546 ± 0.0256	0.7496 ± 0.0090
Precision	Black-Male	0.7437 ± 0.0232	0.5884 ± 0.0247	0.7313 ± 0.0072
Precision	Native Hawaiian-Female	0.6855 ± 0.0944	0.6640 ± 0.1147	0.8333 ± 0.0000
Precision	Native Hawaiian-Male	0.8086 ± 0.0479	0.8271 ± 0.0969	0.7848 ± 0.0155
Precision	Some Other Race-Female	0.6797 ± 0.0123	0.6241 ± 0.0046	0.6649 ± 0.0057
Precision	Some Other Race-Male	0.8146 ± 0.0073	0.7457 ± 0.0103	0.8217 ± 0.0057
Precision	Two or More Races-Female	0.7170 ± 0.0233	0.6398 ± 0.0152	0.6713 ± 0.0092
Precision	Two or More Races-Male	0.8303 ± 0.0166	0.6430 ± 0.0194	0.7942 ± 0.0065
Precision	White-Female	0.7132 ± 0.0044	0.6032 ± 0.0069	0.6755 ± 0.0026
Precision	White-Male	0.8045 ± 0.0065	0.7100 ± 0.0072	0.8087 ± 0.0062
TNR	American Indian-Female	0.8059 ± 0.0623	0.5800 ± 0.1860	0.6290 ± 0.0228
TNR	American Indian-Male	0.9292 ± 0.0623	0.8222 ± 0.1589	0.6704 ± 0.0324
TNR	Amerindian Tribes-Female	0.7100 ± 0.0316	0.6333 ± 0.2919	0.7333 ± 0.1405
TNR	Amerindian Tribes-Male	0.9000 ± 0.1610	0.6692 ± 0.2085	0.7143 ± 0.0000
TNR	Asian-Female	0.7206 ± 0.0230	0.5112 ± 0.0224	0.6251 ± 0.0126
TNR	Asian-Male	0.7654 ± 0.0139	0.5591 ± 0.0228	0.7357 ± 0.0076
TNR	Black-Female	0.7538 ± 0.0265	0.6707 ± 0.0460	0.8053 ± 0.0106
TNR	Black-Male	0.8435 ± 0.0180	0.6506 ± 0.0439	0.8093 ± 0.0060
TNR	Native Hawaiian-Female	0.4200 ± 0.1751	0.6625 ± 0.2503	0.8182 ± 0.0000
TNR	Native Hawaiian-Male	0.7875 ± 0.0604	0.8364 ± 0.1271	0.6250 ± 0.0000
TNR	Some Other Race-Female	0.7355 ± 0.0180	0.6016 ± 0.0194	0.6648 ± 0.0085
TNR	Some Other Race-Male	0.7864 ± 0.0096	0.6851 ± 0.0213	0.8069 ± 0.0079
TNR	Two or More Races-Female	0.8150 ± 0.0192	0.7764 ± 0.0302	0.7790 ± 0.0102
TNR	Two or More Races-Male	0.8620 ± 0.0162	0.7832 ± 0.0256	0.8784 ± 0.0058
TNR	White-Female	0.7551 ± 0.0096	0.5605 ± 0.0182	0.6946 ± 0.0070
TNR	White-Male	0.8011 ± 0.0090	0.6374 ± 0.0141	0.7930 ± 0.0107

Table 7: ACS Employment: Mean and standard deviation values for 10 training runs for the best model configurations with and without differential privacy for several metrics regarding all subgroups between race and gender. More metrics for this dataset in Table8.

Metric	Group	Baseline	Baseline + DP	Best DP Model
TPR	American Indian-Female	0.6778 ± 0.1610	0.6687 ± 0.2166	0.7591 ± 0.0220
TPR	American Indian-Male	0.6333 ± 0.1305	0.7118 ± 0.2191	0.8636 ± 0.0000
TPR	Amerindian Tribes-Female	0.7000 ± 0.1892	0.7750 ± 0.3810	0.7750 ± 0.0527
TPR	Amerindian Tribes-Male	0.7273 ± 0.0958	0.7667 ± 0.3258	0.8333 ± 0.0000
TPR	Asian-Female	0.8080 ± 0.0227	0.9221 ± 0.0189	0.8961 ± 0.0150
TPR	Asian-Male	0.8821 ± 0.0165	0.9578 ± 0.0107	0.9198 ± 0.0050
TPR	Black-Female	0.7450 ± 0.0322	0.8500 ± 0.0400	0.8339 ± 0.0134
TPR	Black-Male	0.8236 ± 0.0150	0.8850 ± 0.0467	0.8796 ± 0.0190
TPR	Native Hawaiian-Female	0.7875 ± 0.1324	0.7286 ± 0.2970	0.9091 ± 0.0000
TPR	Native Hawaiian-Male	0.8875 ± 0.0710	0.8222 ± 0.2468	0.9167 ± 0.0786
TPR	Some Other Race-Female	0.7751 ± 0.0216	0.9015 ± 0.0318	0.9218 ± 0.0102
TPR	Some Other Race-Male	0.9070 ± 0.0114	0.9331 ± 0.0192	0.9171 ± 0.0172
TPR	Two or More Races-Female	0.7820 ± 0.0343	0.8202 ± 0.0736	0.8822 ± 0.0078
TPR	Two or More Races-Male	0.9314 ± 0.0209	0.8810 ± 0.0478	0.9186 ± 0.0133
TPR	White-Female	0.7994 ± 0.0157	0.8923 ± 0.0175	0.8535 ± 0.0114
TPR	White-Male	0.8569 ± 0.0104	0.9110 ± 0.0116	0.8632 ± 0.0112
Selection Rate	American Indian-Female	0.3615 ± 0.0683	0.5171 ± 0.1878	0.5321 ± 0.0173
Selection Rate	American Indian-Male	0.2872 ± 0.0590	0.4371 ± 0.1741	0.5694 ± 0.0179
Selection Rate	Amerindian Tribes-Female	0.3846 ± 0.0513	0.6000 ± 0.3144	0.6364 ± 0.0606
Selection Rate	Amerindian Tribes-Male	0.5929 ± 0.0757	0.4684 ± 0.2386	0.5385 ± 0.0000
Selection Rate	Asian-Female	0.5270 ± 0.0211	0.6945 ± 0.0188	0.6099 ± 0.0126
Selection Rate	Asian-Male	0.5661 ± 0.0129	0.7149 ± 0.0151	0.6198 ± 0.0041
Selection Rate	Black-Female	0.4519 ± 0.0253	0.4980 ± 0.0417	0.4575 ± 0.0104
Selection Rate	Black-Male	0.3932 ± 0.0116	0.5421 ± 0.0390	0.4464 ± 0.0083
Selection Rate	Native Hawaiian-Female	0.7077 ± 0.0873	0.5200 ± 0.2471	0.5455 ± 0.0000
Selection Rate	Native Hawaiian-Male	0.5500 ± 0.0493	0.4600 ± 0.1630	0.7000 ± 0.0471
Selection Rate	Some Other Race-Female	0.4788 ± 0.0175	0.6112 ± 0.0243	0.5810 ± 0.0072
Selection Rate	Some Other Race-Male	0.5662 ± 0.0070	0.6222 ± 0.0189	0.5495 ± 0.0105
Selection Rate	Two or More Races-Female	0.4086 ± 0.0171	0.4176 ± 0.0433	0.4447 ± 0.0085
Selection Rate	Two or More Races-Male	0.4712 ± 0.0133	0.4202 ± 0.0308	0.3909 ± 0.0075
Selection Rate	White-Female	0.4846 ± 0.0120	0.6334 ± 0.0169	0.5394 ± 0.0087
Selection Rate	White-Male	0.5203 ± 0.0082	0.6332 ± 0.0112	0.5372 ± 0.0108
ROC Auc	American Indian-Female	0.7418 ± 0.0870	0.6244 ± 0.0658	0.6941 ± 0.0145
ROC Auc	American Indian-Male	0.7812 ± 0.0760	0.7670 ± 0.0775	0.7670 ± 0.0162
ROC Auc	Amerindian Tribes-Female	0.7050 ± 0.0949	0.7042 ± 0.1449	0.7542 ± 0.0686
ROC Auc	Amerindian Tribes-Male	0.8136 ± 0.1023	0.7179 ± 0.0857	0.7738 ± 0.0000
ROC Auc	Asian-Female	0.7643 ± 0.0088	0.7167 ± 0.0088	0.7606 ± 0.0054
ROC Auc	Asian-Male	0.8237 ± 0.0082	0.7585 ± 0.0087	0.8277 ± 0.0048
ROC Auc	Black-Female	0.7494 ± 0.0143	0.7603 ± 0.0154	0.8196 ± 0.0057
ROC Auc	Black-Male	0.8336 ± 0.0129	0.7678 ± 0.0233	0.8445 ± 0.0097
ROC Auc	Native Hawaiian-Female	0.6038 ± 0.1257	0.6955 ± 0.1167	0.8636 ± 0.0000
ROC Auc	Native Hawaiian-Male	0.8375 ± 0.0437	0.8293 ± 0.0992	0.7708 ± 0.0393
ROC Auc	Some Other Race-Female	0.7553 ± 0.0090	0.7515 ± 0.0077	0.7933 ± 0.0059
ROC Auc	Some Other Race-Male	0.8467 ± 0.0079	0.8091 ± 0.0073	0.8620 ± 0.0082
ROC Auc	Two or More Races-Female	0.7985 ± 0.0201	0.7983 ± 0.0239	0.8306 ± 0.0043
ROC Auc	Two or More Races-Male	0.8967 ± 0.0127	0.8321 ± 0.0156	0.8985 ± 0.0053
ROC Auc	White-Female	0.7773 ± 0.0038	0.7264 ± 0.0058	0.7740 ± 0.0027
ROC Auc	White-Male	0.8290 ± 0.0052	0.7742 ± 0.0065	0.8281 ± 0.0018

Table 8: ACS Employment: Mean and standard deviation values for 10 training runs for the best model configurations with and without differential privacy for several metrics regarding all subgroups between race and gender.

Metric	Group	Baseline	Baseline + DP	Best DP Model
Accuracy	American Indian-Female	0.7121 ± 0.0434	0.7420 ± 0.0394	0.7814 ± 0.0196
Accuracy	American Indian-Male	0.8636 ± 0.0262	0.7323 ± 0.0505	0.7605 ± 0.0083
Accuracy	Amerindian Tribes-Female	0.8000 ± 0.0493	0.6000 ± 0.1042	0.7308 ± 0.0405
Accuracy	Amerindian Tribes-Male	0.7167 ± 0.0315	0.8250 ± 0.1344	0.8700 ± 0.0483
Accuracy	Asian-Female	0.7709 ± 0.0056	0.6838 ± 0.0063	0.7775 ± 0.0054
Accuracy	Asian-Male	0.8307 ± 0.0053	0.7741 ± 0.0051	0.8580 ± 0.0039
Accuracy	Black-Female	0.8067 ± 0.0133	0.7209 ± 0.0056	0.7925 ± 0.0059
Accuracy	Black-Male	0.8221 ± 0.0128	0.7701 ± 0.0282	0.8183 ± 0.0056
Accuracy	Native Hawaiian-Female	0.7381 ± 0.0645	0.7733 ± 0.1891	0.8588 ± 0.0304
Accuracy	Native Hawaiian-Male	0.8765 ± 0.0585	0.7652 ± 0.0622	0.9000 ± 0.0323
Accuracy	Some Other Race-Female	0.7401 ± 0.0094	0.7137 ± 0.0056	0.7658 ± 0.0023
Accuracy	Some Other Race-Male	0.8592 ± 0.0054	0.8111 ± 0.0080	0.8554 ± 0.0046
Accuracy	Two or More Races-Female	0.8204 ± 0.0153	0.7599 ± 0.0158	0.8063 ± 0.0044
Accuracy	Two or More Races-Male	0.8765 ± 0.0082	0.8537 ± 0.0135	0.8663 ± 0.0026
Accuracy	White-Female	0.7704 ± 0.0025	0.7038 ± 0.0098	0.7759 ± 0.0023
Accuracy	White-Male	0.8337 ± 0.0038	0.7783 ± 0.0046	0.8447 ± 0.0014
Precision	American Indian-Female	0.6203 ± 0.0361	0.7015 ± 0.0713	0.6241 ± 0.0248
Precision	American Indian-Male	0.8072 ± 0.0595	0.6598 ± 0.0831	0.6789 ± 0.0094
Precision	Amerindian Tribes-Female	0.7561 ± 0.0372	0.3751 ± 0.2904	0.5982 ± 0.0282
Precision	Amerindian Tribes-Male	0.6183 ± 0.0247	0.5067 ± 0.4133	0.7250 ± 0.0403
Precision	Asian-Female	0.7212 ± 0.0090	0.6068 ± 0.0061	0.6896 ± 0.0054
Precision	Asian-Male	0.8005 ± 0.0055	0.7051 ± 0.0073	0.8287 ± 0.0044
Precision	Black-Female	0.7261 ± 0.0157	0.6286 ± 0.0115	0.7264 ± 0.0094
Precision	Black-Male	0.6692 ± 0.0143	0.6266 ± 0.0374	0.7198 ± 0.0064
Precision	Native Hawaiian-Female	0.7276 ± 0.0460	0.8000 ± 0.4216	0.6827 ± 0.0425
Precision	Native Hawaiian-Male	0.9549 ± 0.0390	0.5415 ± 0.1072	0.7929 ± 0.0553
Precision	Some Other Race-Female	0.6577 ± 0.0143	0.6113 ± 0.0087	0.6403 ± 0.0016
Precision	Some Other Race-Male	0.8289 ± 0.0078	0.7432 ± 0.0097	0.8093 ± 0.0025
Precision	Two or More Races-Female	0.7178 ± 0.0200	0.6210 ± 0.0173	0.6932 ± 0.0053
Precision	Two or More Races-Male	0.7755 ± 0.0138	0.7547 ± 0.0223	0.7725 ± 0.0032
Precision	White-Female	0.7000 ± 0.0062	0.6037 ± 0.0099	0.6901 ± 0.0024
Precision	White-Male	0.8075 ± 0.0049	0.7165 ± 0.0054	0.8268 ± 0.0028
TNR	American Indian-Female	0.7300 ± 0.0422	0.8033 ± 0.0838	0.7586 ± 0.0230
TNR	American Indian-Male	0.8821 ± 0.0478	0.6778 ± 0.1304	0.6619 ± 0.0151
TNR	Amerindian Tribes-Female	0.7125 ± 0.0604	0.6600 ± 0.2633	0.6250 ± 0.0000
TNR	Amerindian Tribes-Male	0.5300 ± 0.0483	0.8833 ± 0.1766	0.8571 ± 0.0000
TNR	Asian-Female	0.7400 ± 0.0167	0.4738 ± 0.0195	0.6454 ± 0.0086
TNR	Asian-Male	0.7431 ± 0.0101	0.5744 ± 0.0184	0.7615 ± 0.0085
TNR	Black-Female	0.8139 ± 0.0128	0.6442 ± 0.0322	0.7580 ± 0.0130
TNR	Black-Male	0.8114 ± 0.0112	0.7195 ± 0.0460	0.7913 ± 0.0064
TNR	Native Hawaiian-Female	0.5667 ± 0.0820	1.0000 ± 0.0000	0.8083 ± 0.0403
TNR	Native Hawaiian-Male	0.8000 ± 0.1721	0.7471 ± 0.1002	0.8400 ± 0.0516
TNR	Some Other Race-Female	0.7508 ± 0.0200	0.5974 ± 0.0291	0.6631 ± 0.0027
TNR	Some Other Race-Male	0.8040 ± 0.0116	0.7020 ± 0.0181	0.7704 ± 0.0033
TNR	Two or More Races-Female	0.8206 ± 0.0168	0.7529 ± 0.0230	0.7677 ± 0.0053
TNR	Two or More Races-Male	0.8627 ± 0.0103	0.8574 ± 0.0200	0.8481 ± 0.0029
TNR	White-Female	0.7547 ± 0.0104	0.5747 ± 0.0202	0.7074 ± 0.0041
TNR	White-Male	0.7902 ± 0.0075	0.6462 ± 0.0125	0.8172 ± 0.0042
TPR	American Indian-Female	0.6846 ± 0.1330	0.6500 ± 0.1810	0.8286 ± 0.0369
TPR	American Indian-Male	0.8313 ± 0.0836	0.8077 ± 0.1462	0.8824 ± 0.0000
TPR	Amerindian Tribes-Female	0.8875 ± 0.0922	0.4800 ± 0.4022	0.9000 ± 0.1054
TPR	Amerindian Tribes-Male	0.9500 ± 0.0645	0.6500 ± 0.4743	0.9000 ± 0.1610
TPR	Asian-Female	0.8080 ± 0.0218	0.9224 ± 0.0155	0.9341 ± 0.0055
TPR	Asian-Male	0.9078 ± 0.0116	0.9632 ± 0.0087	0.9364 ± 0.0089
TPR	Black-Female	0.7951 ± 0.0261	0.8264 ± 0.0350	0.8374 ± 0.0101
TPR	Black-Male	0.8458 ± 0.0389	0.8639 ± 0.0393	0.8617 ± 0.0123
TPR	Native Hawaiian-Female	0.8667 ± 0.0805	0.5750 ± 0.3545	0.9800 ± 0.0632
TPR	Native Hawaiian-Male	0.8929 ± 0.0505	0.8167 ± 0.2659	1.0000 ± 0.0000
TPR	Some Other Race-Female	0.7237 ± 0.0283	0.8746 ± 0.0364	0.9241 ± 0.0076
TPR	Some Other Race-Male	0.9122 ± 0.0099	0.9287 ± 0.0168	0.9371 ± 0.0089
TPR	Two or More Races-Female	0.8200 ± 0.0346	0.7733 ± 0.0581	0.8702 ± 0.0069
TPR	Two or More Races-Male	0.9028 ± 0.0090	0.8464 ± 0.0417	0.8981 ± 0.0067
TPR	White-Female	0.7921 ± 0.0116	0.8791 ± 0.0127	0.8672 ± 0.0066
TPR	White-Male	0.8771 ± 0.0084	0.9132 ± 0.0127	0.8722 ± 0.0041

Table 9: ACS Income: Mean and standard deviation values for 10 training runs for the best model configurations with and without differential privacy for several metrics regarding all subgroups between race and gender. More metrics in Table 10

Metric	Group	Baseline	Baseline + DP	Best DP Model
Selection Rate	American Indian-Female	0.4333 ± 0.0701	0.3780 ± 0.1183	0.4326 ± 0.0196
Selection Rate	American Indian-Male	0.3773 ± 0.0548	0.5258 ± 0.1282	0.5816 ± 0.0083
Selection Rate	Amerindian Tribes-Female	0.5875 ± 0.0604	0.3867 ± 0.2945	0.5769 ± 0.0405
Selection Rate	Amerindian Tribes-Male	0.6833 ± 0.0457	0.2500 ± 0.2125	0.3700 ± 0.0483
Selection Rate	Asian-Female	0.5087 ± 0.0182	0.7116 ± 0.0167	0.6197 ± 0.0053
Selection Rate	Asian-Male	0.6029 ± 0.0096	0.7017 ± 0.0132	0.6237 ± 0.0079
Selection Rate	Black-Female	0.4193 ± 0.0121	0.5539 ± 0.0331	0.5004 ± 0.0106
Selection Rate	Black-Male	0.3929 ± 0.0157	0.4851 ± 0.0370	0.4591 ± 0.0066
Selection Rate	Native Hawaiian-Female	0.6810 ± 0.0504	0.3067 ± 0.1891	0.4235 ± 0.0372
Selection Rate	Native Hawaiian-Male	0.7706 ± 0.0434	0.4000 ± 0.1293	0.4750 ± 0.0323
Selection Rate	Some Other Race-Female	0.4379 ± 0.0213	0.6006 ± 0.0317	0.5680 ± 0.0042
Selection Rate	Some Other Race-Male	0.5612 ± 0.0093	0.6015 ± 0.0156	0.5903 ± 0.0050
Selection Rate	Two or More Races-Female	0.4082 ± 0.0175	0.4279 ± 0.0317	0.4723 ± 0.0040
Selection Rate	Two or More Races-Male	0.4006 ± 0.0066	0.3821 ± 0.0238	0.4242 ± 0.0035
Selection Rate	White-Female	0.4746 ± 0.0107	0.6178 ± 0.0152	0.5391 ± 0.0046
Selection Rate	White-Male	0.5441 ± 0.0070	0.6305 ± 0.0117	0.5276 ± 0.0039
ROC Auc	American Indian-Female	0.7073 ± 0.0573	0.7267 ± 0.0589	0.7936 ± 0.0217
ROC Auc	American Indian-Male	0.8567 ± 0.0326	0.7427 ± 0.0497	0.7721 ± 0.0075
ROC Auc	Amerindian Tribes-Female	0.8000 ± 0.0493	0.5700 ± 0.1229	0.7625 ± 0.0527
ROC Auc	Amerindian Tribes-Male	0.7400 ± 0.0327	0.7667 ± 0.2144	0.8786 ± 0.0805
ROC Auc	Asian-Female	0.7740 ± 0.0061	0.6981 ± 0.0058	0.7898 ± 0.0052
ROC Auc	Asian-Male	0.8254 ± 0.0052	0.7688 ± 0.0055	0.8489 ± 0.0038
ROC Auc	Black-Female	0.8045 ± 0.0151	0.7353 ± 0.0043	0.7977 ± 0.0053
ROC Auc	Black-Male	0.8286 ± 0.0190	0.7917 ± 0.0245	0.8265 ± 0.0064
ROC Auc	Native Hawaiian-Female	0.7167 ± 0.0642	0.7875 ± 0.1773	0.8942 ± 0.0336
ROC Auc	Native Hawaiian-Male	0.8464 ± 0.0968	0.7819 ± 0.1088	0.9200 ± 0.0258
ROC Auc	Some Other Race-Female	0.7373 ± 0.0104	0.7360 ± 0.0065	0.7936 ± 0.0031
ROC Auc	Some Other Race-Male	0.8581 ± 0.0054	0.8153 ± 0.0079	0.8537 ± 0.0046
ROC Auc	Two or More Races-Female	0.8203 ± 0.0182	0.7631 ± 0.0237	0.8190 ± 0.0046
ROC Auc	Two or More Races-Male	0.8828 ± 0.0077	0.8519 ± 0.0179	0.8731 ± 0.0032
ROC Auc	White-Female	0.7734 ± 0.0023	0.7269 ± 0.0086	0.7873 ± 0.0026
ROC Auc	White-Male	0.8336 ± 0.0038	0.7797 ± 0.0046	0.8447 ± 0.0014

Table 10: ACS Income: Mean and standard deviation values for 10 training runs for the best model configurations with and without differential privacy for several metrics regarding all subgroups between race and gender.

Metric	Group	Baseline	Baseline + DP	Best DP Model
Accuracy	Black-Female	0.7659 ± 0.0155	0.7331 ± 0.0087	0.7379 ± 0.0078
Accuracy	Black-Male	0.7420 ± 0.0310	0.6722 ± 0.0100	0.7005 ± 0.0181
Accuracy	Other-Female	0.8272 ± 0.0098	0.7937 ± 0.0058	0.8099 ± 0.0110
Accuracy	Other-Male	0.8344 ± 0.0052	0.8144 ± 0.0033	0.8233 ± 0.0072
Accuracy	White-Female	0.8656 ± 0.0030	0.8332 ± 0.0009	0.8426 ± 0.0085
Accuracy	White-Male	0.8826 ± 0.0033	0.8521 ± 0.0026	0.8645 ± 0.0070
Precision	Black-Female	0.7706 ± 0.0234	0.7039 ± 0.0108	0.7108 ± 0.0067
Precision	Black-Male	0.7354 ± 0.0513	0.6311 ± 0.0063	0.6488 ± 0.0142
Precision	Other-Female	0.8241 ± 0.0119	0.7861 ± 0.0059	0.8006 ± 0.0109
Precision	Other-Male	0.8530 ± 0.0102	0.8117 ± 0.0032	0.8216 ± 0.0087
Precision	White-Female	0.8803 ± 0.0047	0.8316 ± 0.0010	0.8486 ± 0.0142
Precision	White-Male	0.9011 ± 0.0043	0.8513 ± 0.0025	0.8685 ± 0.0109
TNR	Black-Female	0.5960 ± 0.0753	0.3766 ± 0.0399	0.4040 ± 0.0239
TNR	Black-Male	0.6594 ± 0.0936	0.3865 ± 0.0409	0.4135 ± 0.0358
TNR	Other-Female	0.4384 ± 0.0516	0.2741 ± 0.0276	0.3357 ± 0.0471
TNR	Other-Male	0.5071 ± 0.0490	0.3116 ± 0.0176	0.3571 ± 0.0417
TNR	White-Female	0.4164 ± 0.0294	0.1039 ± 0.0067	0.2197 ± 0.0933
TNR	White-Male	0.4694 ± 0.0281	0.1290 ± 0.0182	0.2509 ± 0.0761
TPR	Black-Female	0.8768 ± 0.0625	0.9658 ± 0.0145	0.9558 ± 0.0141
TPR	Black-Male	0.8147 ± 0.0493	0.9239 ± 0.0466	0.9532 ± 0.0101
TPR	Other-Female	0.9710 ± 0.0140	0.9858 ± 0.0038	0.9851 ± 0.0052
TPR	Other-Male	0.9425 ± 0.0207	0.9805 ± 0.0062	0.9773 ± 0.0074
TPR	White-Female	0.9668 ± 0.0063	0.9976 ± 0.0008	0.9830 ± 0.0112
TPR	White-Male	0.9652 ± 0.0061	0.9965 ± 0.0013	0.9871 ± 0.0073
Selection Rate	Black-Female	0.6901 ± 0.0663	0.8306 ± 0.0240	0.8137 ± 0.0162
Selection Rate	Black-Male	0.5927 ± 0.0652	0.7785 ± 0.0432	0.7815 ± 0.0170
Selection Rate	Other-Female	0.8605 ± 0.0224	0.9157 ± 0.0096	0.8986 ± 0.0152
Selection Rate	Other-Male	0.8308 ± 0.0274	0.9080 ± 0.0084	0.8942 ± 0.0150
Selection Rate	White-Female	0.8963 ± 0.0101	0.9789 ± 0.0017	0.9457 ± 0.0261
Selection Rate	White-Male	0.8928 ± 0.0092	0.9756 ± 0.0038	0.9475 ± 0.0186
ROC Auc	Black-Female	0.7364 ± 0.0150	0.6712 ± 0.0138	0.6799 ± 0.0093
ROC Auc	Black-Male	0.7370 ± 0.0338	0.6552 ± 0.0087	0.6834 ± 0.0191
ROC Auc	Other-Female	0.7047 ± 0.0214	0.6300 ± 0.0125	0.6604 ± 0.0222
ROC Auc	Other-Male	0.7248 ± 0.0149	0.6461 ± 0.0068	0.6672 ± 0.0182
ROC Auc	White-Female	0.6916 ± 0.0122	0.5507 ± 0.0031	0.6014 ± 0.0412
ROC Auc	White-Male	0.7173 ± 0.0119	0.5628 ± 0.0087	0.6190 ± 0.0346

Table 11: LSAC: Mean and standard deviation values for 10 training runs for the best model configurations with and without differential privacy for several metrics regarding all subgroups between race and gender.

Metric	Group	Baseline	Baseline + DP	Best DP Model
Accuracy	Amer-Indian-Eskimo- Female	0.9767 ± 0.0086	0.9283 ± 0.0158	0.9667 ± 0.0000
Accuracy	Amer-Indian-Eskimo- Male	0.8718 ± 0.0300	0.8085 ± 0.0426	0.8803 ± 0.0120
Accuracy	Asian-Pac-Islander- Female	0.8935 ± 0.0060	0.8290 ± 0.0591	0.8906 ± 0.0053
Accuracy	Asian-Pac-Islander- Male	0.7849 ± 0.0100	0.7074 ± 0.0381	0.7759 ± 0.0071
Accuracy	Black- Female	0.9651 ± 0.0021	0.9544 ± 0.0126	0.9620 ± 0.0019
Accuracy	Black- Male	0.8824 ± 0.0050	0.8115 ± 0.0286	0.8747 ± 0.0033
Accuracy	Other- Female	0.9738 ± 0.0075	0.9357 ± 0.0490	0.9548 ± 0.0176
Accuracy	Other- Male	0.8357 ± 0.0154	0.7643 ± 0.0772	0.8229 ± 0.0100
Accuracy	White- Female	0.9256 ± 0.0013	0.9031 ± 0.0049	0.9245 ± 0.0008
Accuracy	White- Male	0.8076 ± 0.0021	0.7239 ± 0.0189	0.8058 ± 0.0012
Precision	Amer-Indian-Eskimo- Female	0.9400 ± 0.0966	0.5100 ± 0.3950	0.8000 ± 0.0000
Precision	Amer-Indian-Eskimo- Male	0.6079 ± 0.2475	0.3556 ± 0.2669	0.5952 ± 0.0471
Precision	Asian-Pac-Islander- Female	0.7273 ± 0.0817	0.3902 ± 0.1816	0.6736 ± 0.0457
Precision	Asian-Pac-Islander- Male	0.7240 ± 0.0228	0.5604 ± 0.0477	0.7594 ± 0.0242
Precision	Black- Female	0.6979 ± 0.0407	0.6245 ± 0.1845	0.6120 ± 0.0234
Precision	Black- Male	0.7373 ± 0.0217	0.5165 ± 0.0610	0.7188 ± 0.0135
Precision	Other- Female	0.0000 ± 0.0000	0.0250 ± 0.0791	0.0000 ± 0.0000
Precision	Other- Male	0.6801 ± 0.0573	0.4691 ± 0.1071	0.6381 ± 0.0426
Precision	White- Female	0.7211 ± 0.0138	0.6234 ± 0.0504	0.7225 ± 0.0067
Precision	White- Male	0.7183 ± 0.0132	0.5429 ± 0.0245	0.7264 ± 0.0064
TNR	Amer-Indian-Eskimo- Female	0.9945 ± 0.0088	0.9836 ± 0.0200	0.9818 ± 0.0000
TNR	Amer-Indian-Eskimo- Male	0.9738 ± 0.0207	0.8967 ± 0.0713	0.9508 ± 0.0000
TNR	Asian-Pac-Islander- Female	0.9817 ± 0.0086	0.8933 ± 0.0884	0.9767 ± 0.0066
TNR	Asian-Pac-Islander- Male	0.8854 ± 0.0128	0.7116 ± 0.0936	0.9197 ± 0.0122
TNR	Black- Female	0.9895 ± 0.0023	0.9812 ± 0.0183	0.9823 ± 0.0015
TNR	Black- Male	0.9515 ± 0.0054	0.8635 ± 0.0601	0.9500 ± 0.0037
TNR	Other- Female	0.9976 ± 0.0077	0.9561 ± 0.0512	0.9780 ± 0.0180
TNR	Other- Male	0.9625 ± 0.0056	0.8143 ± 0.1345	0.9643 ± 0.0000
TNR	White- Female	0.9711 ± 0.0027	0.9655 ± 0.0128	0.9723 ± 0.0011
TNR	White- Male	0.8883 ± 0.0112	0.7139 ± 0.0450	0.8973 ± 0.0044
TPR	Amer-Indian-Eskimo- Female	0.7800 ± 0.0632	0.3200 ± 0.2700	0.8000 ± 0.0000
TPR	Amer-Indian-Eskimo- Male	0.2500 ± 0.1354	0.2700 ± 0.2058	0.4500 ± 0.0850
TPR	Asian-Pac-Islander- Female	0.3056 ± 0.0472	0.4000 ± 0.2219	0.3167 ± 0.0527
TPR	Asian-Pac-Islander- Male	0.5881 ± 0.0156	0.6990 ± 0.1016	0.4941 ± 0.0211
TPR	Black- Female	0.4781 ± 0.0331	0.4188 ± 0.1744	0.5563 ± 0.0247
TPR	Black- Male	0.5846 ± 0.0171	0.5875 ± 0.1136	0.5500 ± 0.0166
TPR	Other- Female	0.0000 ± 0.0000	0.1000 ± 0.3162	0.0000 ± 0.0000
TPR	Other- Male	0.3286 ± 0.0768	0.5643 ± 0.2368	0.2571 ± 0.0499
TPR	White- Female	0.5754 ± 0.0175	0.4224 ± 0.0823	0.5565 ± 0.0069
TPR	White- Male	0.6289 ± 0.0228	0.7461 ± 0.0459	0.6032 ± 0.0071
Selection Rate	Amer-Indian-Eskimo- Female	0.0700 ± 0.0105	0.0417 ± 0.0379	0.0833 ± 0.0000
Selection Rate	Amer-Indian-Eskimo- Male	0.0577 ± 0.0215	0.1268 ± 0.0858	0.1056 ± 0.0120
Selection Rate	Asian-Pac-Islander- Female	0.0558 ± 0.0123	0.1449 ± 0.1001	0.0616 ± 0.0115
Selection Rate	Asian-Pac-Islander- Male	0.2746 ± 0.0100	0.4271 ± 0.0927	0.2201 ± 0.0135
Selection Rate	Black- Female	0.0328 ± 0.0032	0.0379 ± 0.0243	0.0434 ± 0.0018
Selection Rate	Black- Male	0.1494 ± 0.0059	0.2215 ± 0.0697	0.1442 ± 0.0051
Selection Rate	Other- Female	0.0024 ± 0.0075	0.0452 ± 0.0520	0.0214 ± 0.0176
Selection Rate	Other- Male	0.0957 ± 0.0166	0.2614 ± 0.1472	0.0800 ± 0.0100
Selection Rate	White- Female	0.0917 ± 0.0042	0.0791 ± 0.0203	0.0885 ± 0.0015
Selection Rate	White- Male	0.2726 ± 0.0146	0.4292 ± 0.0444	0.2584 ± 0.0052
ROC Auc	Amer-Indian-Eskimo- Female	0.8873 ± 0.0310	0.6518 ± 0.1280	0.8909 ± 0.0000
ROC Auc	Amer-Indian-Eskimo- Male	0.6119 ± 0.0717	0.5834 ± 0.0783	0.7004 ± 0.0425
ROC Auc	Asian-Pac-Islander- Female	0.6436 ± 0.0211	0.6467 ± 0.0841	0.6467 ± 0.0243
ROC Auc	Asian-Pac-Islander- Male	0.7367 ± 0.0101	0.7053 ± 0.0280	0.7069 ± 0.0086
ROC Auc	Black- Female	0.7338 ± 0.0161	0.7000 ± 0.0807	0.7693 ± 0.0124
ROC Auc	Black- Male	0.7680 ± 0.0085	0.7255 ± 0.0287	0.7500 ± 0.0077
ROC Auc	Other- Female	0.4988 ± 0.0039	0.5280 ± 0.1550	0.4890 ± 0.0090
ROC Auc	Other- Male	0.6455 ± 0.0381	0.6893 ± 0.0791	0.6107 ± 0.0250
ROC Auc	White- Female	0.7732 ± 0.0077	0.6939 ± 0.0355	0.7644 ± 0.0032
ROC Auc	White- Male	0.7586 ± 0.0062	0.7300 ± 0.0094	0.7503 ± 0.0016

Table 12: Adult: Mean and standard deviation values for 10 training runs for the best model configurations with and without differential privacy for several metrics regarding all subgroups between race and gender.

Metric	Group	Baseline	Baseline + DP	Best DP Model
Accuracy	Black-Female	0.6628 ± 0.0122	0.5387 ± 0.0533	0.6558 ± 0.0216
Accuracy	Black-Male	0.6245 ± 0.0092	0.5172 ± 0.0387	0.6103 ± 0.0071
Accuracy	Other-Female	0.7178 ± 0.0150	0.6289 ± 0.0777	0.7000 ± 0.0189
Accuracy	Other-Male	0.6265 ± 0.0157	0.5253 ± 0.0422	0.6378 ± 0.0105
Accuracy	White-Female	0.6847 ± 0.0203	0.5682 ± 0.0444	0.6506 ± 0.0180
Accuracy	White-Male	0.6289 ± 0.0072	0.5051 ± 0.0404	0.6266 ± 0.0151
Precision	Black-Female	0.5803 ± 0.0216	0.4145 ± 0.0444	0.5578 ± 0.0305
Precision	Black-Male	0.6640 ± 0.0065	0.5862 ± 0.0369	0.6705 ± 0.0060
Precision	Other-Female	0.3300 ± 0.0740	0.2969 ± 0.1149	0.3537 ± 0.0489
Precision	Other-Male	0.5953 ± 0.0289	0.4523 ± 0.0452	0.6005 ± 0.0216
Precision	White-Female	0.6178 ± 0.0422	0.4290 ± 0.0534	0.5488 ± 0.0353
Precision	White-Male	0.5893 ± 0.0094	0.4513 ± 0.0365	0.5870 ± 0.0195
TNR	Black-Female	0.7861 ± 0.0202	0.6041 ± 0.1355	0.7303 ± 0.0277
TNR	Black-Male	0.5013 ± 0.0111	0.4361 ± 0.1050	0.5619 ± 0.0113
TNR	Other-Female	0.9000 ± 0.0152	0.7147 ± 0.1127	0.8382 ± 0.0208
TNR	Other-Male	0.7887 ± 0.0223	0.5739 ± 0.1063	0.7613 ± 0.0292
TNR	White-Female	0.8436 ± 0.0226	0.6536 ± 0.0766	0.8000 ± 0.0329
TNR	White-Male	0.7184 ± 0.0100	0.5028 ± 0.1645	0.7193 ± 0.0160
TPR	Black-Female	0.4675 ± 0.0237	0.4351 ± 0.1092	0.5377 ± 0.0336
TPR	Black-Male	0.7137 ± 0.0139	0.5759 ± 0.0883	0.6454 ± 0.0114
TPR	Other-Female	0.1545 ± 0.0439	0.3636 ± 0.1660	0.2727 ± 0.0429
TPR	Other-Male	0.4112 ± 0.0159	0.4607 ± 0.0999	0.4738 ± 0.0202
TPR	White-Female	0.4197 ± 0.0328	0.4258 ± 0.0595	0.4015 ± 0.0229
TPR	White-Male	0.5149 ± 0.0145	0.5081 ± 0.1352	0.5085 ± 0.0247
Selection Rate	Black-Female	0.3121 ± 0.0181	0.4111 ± 0.1206	0.3734 ± 0.0212
Selection Rate	Black-Male	0.6234 ± 0.0095	0.5709 ± 0.0874	0.5583 ± 0.0090
Selection Rate	Other-Female	0.1133 ± 0.0164	0.3044 ± 0.1084	0.1889 ± 0.0189
Selection Rate	Other-Male	0.2972 ± 0.0130	0.4410 ± 0.0962	0.3398 ± 0.0244
Selection Rate	White-Female	0.2551 ± 0.0170	0.3761 ± 0.0601	0.2756 ± 0.0259
Selection Rate	White-Male	0.3842 ± 0.0096	0.5020 ± 0.1497	0.3809 ± 0.0130
ROC Auc	Black-Female	0.6268 ± 0.0122	0.5196 ± 0.0376	0.6340 ± 0.0220
ROC Auc	Black-Male	0.6075 ± 0.0087	0.5060 ± 0.0395	0.6036 ± 0.0070
ROC Auc	Other-Female	0.5273 ± 0.0226	0.5392 ± 0.0787	0.5555 ± 0.0238
ROC Auc	Other-Male	0.6000 ± 0.0151	0.5173 ± 0.0388	0.6175 ± 0.0082
ROC Auc	White-Female	0.6317 ± 0.0215	0.5397 ± 0.0385	0.6008 ± 0.0148
ROC Auc	White-Male	0.6166 ± 0.0075	0.5055 ± 0.0280	0.6139 ± 0.0157

Table 13: Compas: Mean and standard deviation values for 10 training runs for the best model configurations with and without differential privacy for several metrics regarding all subgroups between race and gender.