

Exploring the Unfairness of DP-SGD Across Settings

Frederik Noe, Rasmus Herskind and Anders Søgaard
University of Copenhagen

Abstract

End users and regulators require private and fair artificial intelligence models, but previous work suggests these objectives may be at odds. We use the CivilComments to evaluate the impact of applying the *de facto* standard approach to privacy, DP-SGD, across several fairness metrics. We evaluate three implementations of DP-SGD: for dimensionality reduction (PCA), linear classification (logistic regression), and robust deep learning (Group-DRO). We establish a negative, logarithmic correlation between privacy and fairness in the case of linear classification and robust deep learning. DP-SGD had no significant impact on fairness for PCA, but upon inspection, also did not seem to lead to private representations.

Introduction

The vast majority of artificial intelligence research projects focus on learning accurate models, but in practice, real-life application of machine learning models require balancing performance with the need to prevent discrimination against protected demographic subgroups and satisfying privacy principles. Both needs are prerequisites for employing machine learning models at scale in many domains, and the lack of fairness and privacy in many widely used models threatens the public perception of artificial intelligence. Unfortunately, the needs are seemingly at odds (Ekstrand, Joshaghani, and Mehrpouyan 2018; Cummings et al. 2019; Bagdasaryan, Poursaeed, and Shmatikov 2019; Farrand et al. 2020; Chang and Shokri 2021; Agarwal 2021). Common fairness and privacy objectives are in conflict, because privacy-preserving algorithms often disproportionately affect members of minority classes (Farrand et al. 2020).¹ The exact nature of this trade-off has not been analyzed, however, and perhaps more surprisingly, there has been no evaluations of whether learning algorithms that were developed with group fairness in mind, such as Sagawa et al. (2020), are less sensitive to optimization with DP-SGD.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Note this is a different trade-off than the fairness-privacy trade-off which results from the need for collecting sensitive data to learn fair models (Veale and Binns 2017).

Toxic	Comment Text	Male	Female	LGBTQ	White	Black	...	Christian
0	I applaud your father. He was a good man! We need more like him.	1	0	0	0	0	...	0
0	As a Christian, I will not be patronizing any of those businesses.	0	0	0	0	0	...	1
0	What do Black and LGBT people have to do with bicycle licensing?	0	0	1	0	1	...	0
0	Government agencies track down foreign baddies and protect law-abiding white citizens. How many shows does that describe?	0	0	0	1	0	...	0
1	Maybe you should learn to write a coherent sentence so we can understand WTF your point is.	0	0	0	0	0	...	0

Figure 1: A few datapoints from CivilComments including demographic identities (Koh et al. 2021)

Contributions We evaluate the impact on fairness (Unequal Risk) of using DP-SGD optimization, across three settings: PCA, logistic regression and Group-DRO (Sagawa et al. 2020). For the baseline condition, we confirm past results and establish a negative correlation between fairness and privacy. We further show that this negative correlation is *logarithmic*. Moreover, we show that DP-SGD has little impact on fairness in the PCA setting. This, however, seems to be because differentially private dimensionality reduction is insufficient to guarantee the privacy of downstream classifiers.

Fairness and Privacy

We represent each data point as $z = (x, g, y) \in \mathcal{X} \times \mathcal{G} \times \mathcal{Y}$, with $g \in \mathcal{G}$ encoding its protected attributes (e.g., race, gender). Let \mathcal{D}_y^g denote the distribution of data with protected attribute g and label y . Several definitions of fairness exist (Williamson and Menon 2019), but we use (generalized) approximately constant conditional (equal) risk (Donini et al. 2018):

Definition 0.1 (δ -Unequal Risk). Let $\ell^{g_i}(\theta) = \mathbb{E}[\ell(\theta(x), y) | g = g_i]$ be the risk of the samples in the group defined by g_i , and $\delta \in [0, 1]$. We say that a model θ exhibits δ -Unequal Risk if for any two values of g , say g_i and g_j , $|\ell^{g_i}(\theta) - \ell^{g_j}(\theta)| < \delta$.

Note that if ℓ coincides with the performance metric of a task, and $\delta = 0$, this is identical to performance or classification parity (Yuan et al. 2021).² δ directly measures what

²Such a notion of fairness can be derived from John Rawls' theory on distributive justice and stability, treating model performance as a resource to be allocated. Rawls' *difference principle*, maximizing the welfare of the worst-off group, is argued to lead to stability and mobility in society at large (Rawls 1971). Performance or clas-

is sometimes called Rawlsian *min-max fairness* (Bertsimas, Farias, and Trichakis 2011). Recall the standard definition of (ϵ, ϕ) -privacy³ is as follows:

Definition 0.2 ((ϵ, ϕ) -privacy). θ is (ϵ, ϕ) -private iff $\Pr[\theta(\mathcal{X})] \leq \exp(\epsilon) \times \Pr[\theta(\mathcal{X}')] + \phi$ for any two distributions, \mathcal{X} and \mathcal{X}' , different at most in one row.

Differential privacy ensures that an algorithm will generate similar outputs for similar training data sets. Note the multiplicative bound $\exp(\epsilon)$ and the additive bound ϕ serve different roles: The ϕ term represents the possibility that a few data points are not governed by the multiplicative bound, which controls the level of privacy (rather than its scope). Note that it also follows directly that if $\epsilon = 0$ and $\phi = 0$, absolute privacy is required, leading θ to be independent of the data. In our experiments, we assume $\phi = 0$ and look at the impact of DP-SGD training across different values of ϵ .

Agarwal (2021) shows that a $(\epsilon, 0)$ -private and fully fair model – using equalized odds as the definition of fairness – will be unable to learn anything. To see this, remember that a fully private model is independent of the data and unable to learn from correlations between input and output. If θ is, in addition, required to be fair, it is thereby required to be fair for all distributions, which prevents θ from encoding any prior beliefs about the output distribution. Equal odds is often less interesting for NLP applications than equal risk, but note this finding generalizes straight-forwardly to equalized risk, and to approximate fairness (since even for finite distributions, we can define a $\delta > 0$, such that preserving absolute privacy leads to a constant θ).

Theorem 1. For sufficiently small values of ϵ , a fully $(\epsilon, 0)$ -private model θ that is also δ -fair, will have trivial performance.

Proof. This follows directly from the above. \square

We focus on δ -Unequal Risk below, but for completeness include alternative definitions. While δ -Unequal Risk only considers the difference between extremes, δ -Variance considers the overall F_1 -score variance across the groups, with $\delta = \sigma^2$ defined as the mean of the squared distances of ℓ^{g_i} from their mean $\bar{\ell}^g$, i.e., $\frac{\sum_{i=1}^n (\ell^{g_i} - \bar{\ell}^g)^2}{n}$. Often Unequal Risk is contrasted with Equal Odds, i.e., the idea that fairness requires similar positive rates across groups. Zafar et al. (2017) introduce the so-called *p%-rule*, a variant of Equal Odds, which we will also adopt below. In the binary case, we say a model θ exhibits *p*-Equal Odds if

$$p = \min_{i,j} \min \left(\frac{\mathbb{E}[\theta(\mathcal{X}^{g_i})]}{\mathbb{E}[\theta(\mathcal{X}^{g_j})]}, \frac{\mathbb{E}[\theta(\mathcal{X}^{g_j})]}{\mathbb{E}[\theta(\mathcal{X}^{g_i})]} \right)$$

Note that if the *p*%-rule returns 1.0, θ exhibits 1.0-Equal Odds, which amounts to identical positive rates across

sification parity has, however, been argued to suffer from statistical limitations in (Corbett-Davies and Goel 2018), which remind us that when risk distributions differ, standard error metrics are poor proxies of individual equity. This is known as the problem of infra-marginality. Note, however, that this argument does not apply to binary classification problems.

³We talk about (ϵ, ϕ) -privacy rather than (ϵ, δ) -privacy, to avoid confusion with δ -fairness.




Figure 2: Baseline performance (F_1) across groups

groups. It is easy to show that $p \in [0, 1]$, with higher values suggesting fairness under Equal Odds. We modify the *p*%-rule, which serves as a compromise between Unequal Risk and Equal Odds.

Definition 0.3 (Modified *p*%-rule). We say a model θ exhibits *p*-Equal Correct Priors if

$$p = \min_{i,j} \min \left(\frac{\mathbb{E}[\theta(\mathcal{X}^{g_i})]}{\mathbb{E}[\mathcal{Y}^{g_i}]}, \frac{\mathbb{E}[\theta(\mathcal{X}^{g_j})]}{\mathbb{E}[\mathcal{Y}^{g_j}]}, \frac{\mathbb{E}[\theta(\mathcal{X}^{g_j})]}{\mathbb{E}[\theta(\mathcal{X}^{g_i})]}, \frac{\mathbb{E}[\theta(\mathcal{X}^{g_i})]}{\mathbb{E}[\theta(\mathcal{X}^{g_j})]} \right)$$

This definition says the divergence from the positive class prior should be no different across any two groups. Note how this, unlike δ -Unequal Risk, says *nothing* about the actual risk of misclassification. It is merely a weighted version of *p*-Equal Odds, which takes the true priors into account. We report results for δ -Variance and the *p*%-rule in the Appendix, and focus on δ -Unequal Risk and the Modified *p*%-rule here.

Experiments

Our experiments evaluate the impact of differential privacy (Dwork et al. 2006) on the fairness of different machine learning architectures. Differentially private stochastic gradient descent (DP-SGD) is the *de facto* baseline approach to learning private models, and we limit ourselves to variants of DP-SGD. Vanilla DP-SGD limits the influence of training samples by (i) clipping the per-batch gradient where its norm exceeds a pre-determined clipping bound C , and by (ii) adding Gaussian noise \mathcal{N} characterized by a noise scale σ to the aggregated per-sample gradients. We control this influence with a privacy budget ϵ , where lower values for ϵ indicates a more strict level of privacy. DP-SGD has remained popular, among other things because it generalizes to iterative training procedures (McMahan et al. 2018), and supports tighter bounds using the Rényi method (Mironov 2017a). We evaluate the impact of differentially private training on fairness across dimensionality reduction, linear classification, and deep learning.

Bertweet and DistilBert Transformer architectures are often trained with masked language modeling objectives to learn representations that are useful for downstream NLP

tasks (Devlin et al. 2019a; Liu et al. 2019). We use two pre-trained, Transformer-based language models: (a) Bertweet (Nguyen, Vu, and Nguyen 2020) is trained on English tweets with the same hyper-parameters as the original BERT-base model (Devlin et al. 2019b). DistilBert (Sanh et al. 2020) is induced by distillation of BERT (Devlin et al. 2019a) to obtain a more efficient and smaller version of it.

Principal Component Analysis Principal components are linear combinations or mixtures of the initial variables and internally uncorrelated to minimize the information loss. We rely on the IBM Differential Privacy Library (Diffprivlib)⁴ (Holohan et al. 2019) for DP-PCA, which relies on Imtiaz and Sarwate (2016), who add Wishart noise to this reconstruction to obtain ϵ -privacy. Their algorithm is shown to approximate the underlying data well, especially for $\epsilon \geq 1$. Alternatives are discussed in ?. We apply DP-PCA to Bertweet representations of CivilComments and pass these to a three-layer feed-forward network. We varied ϵ -values in $[0.1, 10]$ with a step-size of 0.5, and performed grid search for the hyper-parameters of the feed-forward network. Its dimensionality was given (768, 64, and 32 in the outermost layer), but all other parameters were optimized for validation performance: We trained for 2979 epochs, with learning rate $1e^{-06}$, MSE loss, AdamW optimization, and a batch size of 32.

Logistic Regression We also evaluate the impact of DP-SGD on linear classification. In logistic regression, the weighted sum of model parameters and input is transformed to probabilities by the logistic function, with parameters fit to minimize logistic loss with DP-SGD. We train 400 models with $\epsilon \in \{0.1, 0.2, 0.3, 39.9, 40\}$. Again, we rely on Bertweet and Diffprivlib.

Group Distributionally Robust Optimization We also consider a deep architecture optimized for group fairness, namely GroupDRO (Sagawa et al. 2020), which relies on distributionally robust optimization (Hashimoto et al. 2018) to minimize the worst-case loss over the groups in the training data. We use DP-SGD in Opacus⁵ (AI 2020) to implement differentially private GroupDRO. DP-SGD was generalized to deep learning in Abadi et al. (2016); Opacus uses Renyi Differential Privacy (RDP) (Mironov 2017b).

Data We use Civil Comments (Koh et al. 2021), a toxicity detection dataset in which comments are annotated for toxicity and *the demographics of the target of toxicity* (see Fig. 1). The task is a binary classification task. There are eight target groups:⁶ LGBTQ, male, female, Christian, Muslim, other religions, Black, and White. For simplicity, we

⁴<https://github.com/IBM/differential-privacy-library>

⁵<https://github.com/pytorch/opacus>

⁶We cite the names used in the publicly available dataset. It is unclear to what extent the names impacted annotations. Some seem to represent communities that can in part be referred to by other, more general or more specific, terms, e.g., LGBT, LGBTQ+, the Rainbow Community. By using this data, we do not mean to imply that these terms are useful for referring to groups of citizens. We use CivilComments simply because it is the most widely distributed NLP dataset for fairness studies.

only use data from these groups: Prior work has shown that toxicity classifiers pick up on biases in the training data and spuriously associate toxicity with the mention of certain demographics (Borkan et al. 2019). Thus, for classifying the toxicity of online comments, the setup shown in Figure 1 ensures that we are able to keep track of fairness by analysing scores across different groups.

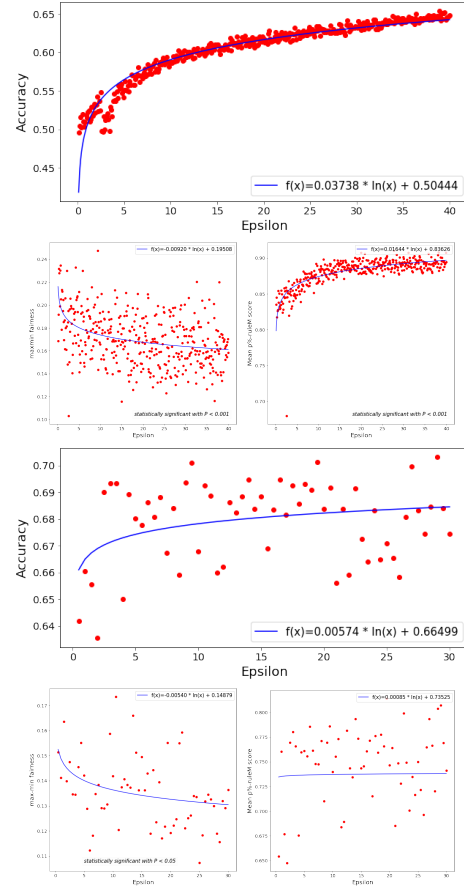


Figure 3: **Top (Linear Classification):** Logarithmic correlations between ϵ (x) and accuracy (y) (**first row**), ϵ and Unequal Risk (**second, left**), and ϵ and the modified p%-rule (**second, right**). **Bottom (Deep GroupDRO):** Logarithmic correlation between ϵ (x) and accuracy (y) (**third**), ϵ and Unequal Risk (**fourth, left**), and the modified p%-rule (**fourth, right**).

Results

The baseline F1 scores across the eight groups are plotted in Figure 2. Scores significantly across the eight groups, with lowest toxicity detection scores for minority religious groups and Christians, and highest toxicity detection scores for blacks and whites, reportedly being the prototypical targets for American and European racist slur (Croom 2015). Generally, PCA reduces performance slightly, while our GroupDRO architecture performs slightly better than our baseline with fixed Bertweet embeddings. We are mainly in-

terested in how overall performance and fairness is affected by DP-SGD.

Figure 3 (**second row**) shows the correlation between the degree of approximate privacy and Unequal Risk for linear classification: The more private a model, the less fair. This correlation contrasts with the positive correlation between ϵ and accuracy (**first**): The more private a model, the worse it performs. Note that both of these correlations are *significant* ($p < 0.001$) and *logarithmic* and have similar coefficients.

GroupDRO with DistilBert exhibits similar correlations; see Figure 3. While odds become logarithmically more equal with higher ϵ values for linear models, odds exhibit high variance for GroupDRO and are not significantly correlated with ϵ . The fact that GroupDRO exhibit a similar, negative, logarithmic correlation between ϵ and Unequal Risk, is particularly interesting, since GroupDRO assumes access to group information during training and uses this information to learn a model with Equal Risk. Under sufficient privacy guarantees, the approach is ineffective, however.

In contrast, with PCA, we see a weak, linear correlation between ϵ and performance: The less privacy, the better performance, again, but differences are minimal (absolute drops in accuracy within a 0.5 margin). We also see a weak, negative, linear correlation between ϵ and Unequal Risk: The less privacy, the more fairness, but this correlation is *not* significant. At first sight, it may seem promising that differentially private dimensionality reduction does not lead to Unequal Risk. The same pattern is seen across all the fairness measures. However, this is unfortunately *not* a case of getting the best of both worlds: Our feed-forward network seems to be able to recover private information from our dimensionality-reduced representations. To see why this may happen, consider that the algorithm presented in Imtiaz and Sarwate (2016) is only approximate and not providing strict guarantees. It adds stochastic noise to each dimension independently, and only provides expected guarantees. The more dimensions that carry signal, the higher the chance that private information is recoverable. The sensitivity of differentially private dimensionality reduction has been observed before (Aggarwal 2005). We hypothesize that our feed-forward network in a way similar to spectral filtering (Kargupta et al. 2003) can filter off the random noise by implicitly analyzing eigenstates.

Discussion and Conclusion

The trade-off between fairness and privacy is of great concern to the machine learning community, but previous work has not suggested promising research directions for how to best solve this dilemma. Our results go beyond previous work on the trade-offs between fairness and privacy (Bagdasaryan, Poursaeed, and Shmatikov 2019; Agarwal 2021) in evaluating this trade-off across multiple machine learning settings and in establishing *logarithmic* correlations between fairness and privacy across two different settings: linear classification and robust deep learning. Finally, we also provide an example of a successful *attack* against differentially private dimensionality reduction by training a three-layer feed-forward network on the learned representations of DP-PCA.

Uniyal et al. (2021) presented a proposal, which exhibited a better trade-off, empirically. Liu et al. (2021) presented an approach to fairness and privacy that relies on a self-adaptive mechanism and dynamically adjusts instance influence in each class depending on the theoretical bias-variance bounds.⁷ Fioretto, Tran, and Hentenryck (2021) proposes various solutions, including output perturbations, linearizations, learning piece-wise linear proxy-functions, or fairness payment. These methods have several drawbacks, and most are not generally applicable. Piece-wise linear proxy-functions and GroupDRO assume that demographics are available during training, which limits their application severely. In summary, finding robustly applicable approaches to simultaneously optimizing for fairness and privacy remains an open problem. Our contributions are (a) showing the *logarithmic* nature of this correlation, suggesting the usefulness of maximizing the second order gradient of this trade-off; (b) showing this holds even for group-robust deep learning; and (c) showing, by example, that seemingly good results for *both* fairness and privacy may be the result of implicit privacy attacks, when deep classifiers exploit privacy leaks.

Based on a sequence of experiments with toxicity detection, relying on large-scale pretrained language models, this paper established the logarithmic nature of the relationship between fairness and privacy, in the form of DP-SGD, and showed that this logarithmic nature generalizes to models trained to optimize for fairness (with GroupDRO). Moreover, we saw that when applied to dimensionality reduction, DP-SGD did not seem to violate fairness. This anomaly was explained by leakage of private information, however. We discussed possible research directions for jointly learning private and fair machine learning models, a topic we believe to be of great importance to the future of artificial intelligence.

References

- Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
- Agarwal, S. 2021. Trade-Offs between Fairness and Privacy in Machine Learning. In *IJCAI 2021 Workshop on AI for Social Good*.
- Aggarwal, C. C. 2005. On k-Anonymity and the Curse of Dimensionality. In *Proceedings of the 31st International Conference on Very Large Data Bases, VLDB '05*, 901–909. VLDB Endowment. ISBN 1595931546.
- AI, F. 2020. Opacus.
- Bagdasaryan, E.; Poursaeed, O.; and Shmatikov, V. 2019. Differential Privacy Has Disparate Impact on Model Accuracy. In Wallach, H.; Larochelle, H.; Beygelzimer, A.;

⁷The paper was rejected, though, because it did not prove reasonable privacy guarantees, and because it relied on non-standard definitions of fairness. See <https://openreview.net/forum?id=IqVB8e0DIUd>

- d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Bertsimas, D.; Farias, V. F.; and Trichakis, N. 2011. The Price of Fairness. *Oper. Res.*, 59(1): 17–31.
- Borkan, D.; Dixon, L.; Sorensen, J.; Thain, N.; and Vasserman, L. 2019. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion Proceedings of The 2019 World Wide Web Conference*.
- Chang, H.; and Shokri, R. 2021. On the Privacy Risks of Algorithmic Fairness. arXiv:2011.03731.
- Corbett-Davies, S.; and Goel, S. 2018. The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning. arXiv:1808.00023.
- Croom, A. M. 2015. Slurs, Stereotypes, and in-Equality: A Critical Review of “How Epithets and Stereotypes Are Racially Unequal”. *Language Sciences*, 52: 139–154.
- Cummings, R.; Gupta, V.; Kimpura, D.; and Morgenstern, J. 2019. On the compatibility of privacy and fairness. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*, 309–315.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019a. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019b. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805.
- Donini, M.; Oneto, L.; Ben-David, S.; Shawe-Taylor, J. S.; and Pontil, M. 2018. Empirical Risk Minimization Under Fairness Constraints. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the Third Conference on Theory of Cryptography, TCC'06*, 265–284. Berlin, Heidelberg: Springer-Verlag. ISBN 3540327312.
- Ekstrand, M. D.; Joshaghani, R.; and Mehrpouyan, H. 2018. Privacy for All: Ensuring Fair and Equitable Privacy Protections. In Friedler, S. A.; and Wilson, C., eds., *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, 35–47. New York, NY, USA: PMLR.
- Farrand, T.; Mireshghallah, F.; Singh, S.; and Trask, A. 2020. Neither Private Nor Fair: Impact of Data Imbalance on Utility and Fairness in Differential Privacy. In *Proceedings of the 2020 Workshop on Privacy-Preserving Machine Learning in Practice*, 15–19.
- Fioretto, F.; Tran, C.; and Hentenryck, P. V. 2021. Decision Making with Differential Privacy under a Fairness Lens. arXiv:2105.07513.
- Hashimoto, T. B.; Srivastava, M.; Namkoong, H.; and Liang, P. 2018. Fairness Without Demographics in Repeated Loss Minimization. arXiv:1806.08010.
- Holohan, N.; Braghin, S.; Aonghusa, P. M.; and Levacher, K. 2019. Diffprivlib: The IBM Differential Privacy Library. arXiv:1907.02444.
- Imtiaz, H.; and Sarwate, A. D. 2016. Symmetric matrix perturbation for differentially-private principal component analysis. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2339–2343.
- Kargupta, H.; Datta, S.; Wang, Q.; and Sivakumar, K. 2003. On the Privacy Preserving Properties of Random Data Perturbation Techniques. In *Proceedings of the Third IEEE International Conference on Data Mining, ICDM '03*, 99. USA: IEEE Computer Society. ISBN 0769519784.
- Koh, P. W.; Sagawa, S.; Marklund, H.; Xie, S. M.; Zhang, M.; Balsubramani, A.; Hu, W.; Yasunaga, M.; Phillips, R. L.; Gao, I.; Lee, T.; David, E.; Stavness, I.; Guo, W.; Earnshaw, B. A.; Haque, I. S.; Beery, S.; Leskovec, J.; Kundaje, A.; Pierson, E.; Levine, S.; Finn, C.; and Liang, P. 2021. WILDS: A Benchmark of in-the-Wild Distribution Shifts. arXiv:2012.07421.
- Liu, W.; Wang, X.; Lu, X.; Cheng, J.; Jin, B.; Wang, X.; and Zha, H. 2021. Fair Differential Privacy Can Mitigate the Disparate Impact on Model Accuracy.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- McMahan, B.; Andrew, G.; Mironov, I.; Papernot, N.; Kairouz, P.; Chien, S.; and Úlfar Erlingsson. 2018. A General Approach to Adding Differential Privacy to Iterative Training Procedures. Workshop on Privacy Preserving Machine Learning (NeurIPS 2018).
- Mironov, I. 2017a. Rényi Differential Privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, 263–275.
- Mironov, I. 2017b. Rényi Differential Privacy. *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*.
- Nguyen, D. Q.; Vu, T.; and Nguyen, A. T. 2020. BERTweet: A pre-trained language model for English Tweets. arXiv:2005.10200.
- Rawls, J. 1971. *A Theory of Justice*. Cambridge, Massachusetts: Belknap Press of Harvard University Press, 1 edition. ISBN 0-674-88014-5.
- Sagawa, S.; Koh, P. W.; Hashimoto, T. B.; and Liang, P. 2020. Distributionally Robust Neural Networks for Group Shifts: On the Importance of Regularization for Worst-Case Generalization. arXiv:1911.08731.
- Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2020. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv:1910.01108.
- Uniyal, A.; Naidu, R.; Kotti, S.; Singh, S.; Kenfack, P. J.; Mireshghallah, F.; and Trask, A. 2021. DP-SGD vs PATE: Which Has Less Disparate Impact on Model Accuracy? arXiv:2106.12576.

Veale, M.; and Binns, R. 2017. Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data & Society*, 4(2): 2053951717743530.

Williamson, R.; and Menon, A. 2019. Fairness risk measures. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 6786–6797. PMLR.

Yuan, M.; Kumar, V.; Ahmad, M. A.; and Teredesai, A. 2021. Assessing Fairness in Classification Parity of Machine Learning Models in Healthcare. arXiv:2102.03717.

Zafar, M. B.; Valera, I.; Rodriguez, M. G.; and Gummadi, K. P. 2017. Fairness Constraints: Mechanisms for Fair Classification.