

# TRANSITIONING FROM REAL TO SYNTHETIC DATA: QUANTIFYING THE BIAS IN MODEL

Aman Gupta, Deepak Bhatt & Anubha Pandey

Mastercard, India

{aman.gupta, deepak.bhatt, anubha.pandey}@mastercard.com

## ABSTRACT

With advent of generative modelling techniques, synthetic data and its use has penetrated across various domains from unstructured data such as image, text to structured dataset modelling healthcare outcome, risk decisioning in financial domain and many more. It overcomes various challenges such as limited training data, class imbalance, restricted access to dataset owing to privacy issues. To ensure trained model used for automated decisioning purposes makes fair decision there exist prior work to quantify and mitigate those issues. This study aims to establish trade-off between bias and fairness in the models trained using synthetic data. Variants of synthetic data generation techniques were studied to understand bias amplification including differentially private generation schemes. Through experiments on a tabular dataset, we demonstrate there exist varying level of bias impact on models trained using synthetic data. Techniques generating less correlated feature performs well as evident through fairness metrics with 94%, 82%, and 88% relative drop in DPD (demographic parity difference), EoD (equality of odds) and EoP (equality of opportunity) respectively, and 24% relative improvement in DRP (demographic parity ratio) with respect to the real dataset. We believe the outcome of our research study will help data science practitioners understand the bias in use of synthetic data.

**Keywords:** Fairness, Bias, Differential Privacy, Synthetic Dataset, Class Imbalance

## 1 INTRODUCTION

Machine learning systems are being adopted in every aspect of our daily life. Amazon and Netflix uses it for movie recommendations, recommending products customized based on user preference. Among many, recruiting industry is one such example leveraging AI to find potential candidates. Owing to its capability to reduce manual hours and assist in automated decision-making, they are being widely deployed and integrated across various use-cases such as public use, policy-making, predictive policing, etc. On the other hand, though machine learning systems are quite successful yet it found that “Robots are racist and sexist, just like the people who created it<sup>1</sup>”, claimed that “crime rate in color people are higher<sup>2</sup>” and “dark skins are unattractive<sup>3</sup>”, “recommending less qualified male candidates higher than more qualified female candidates on job portal” (Lahoti et al., 2019). As machine learning models learn what humans teach them, therefore, like people, machine learning systems are vulnerable to amplify existing human biases and societal stereotypes when they make decisions (Wang et al., 2019). One of the factors that may influence the fairness of machine learning models is change in model outcome with protected attributes like gender, race, religion, sexual orientation, economic condition, etc (Mehrabi et al., 2019). Recent research found that even if we

---

<sup>1</sup><https://www.theguardian.com/commentisfree/2017/apr/20/robots-racist-sexist-people-machines-ai-language>

<sup>2</sup><https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

<sup>3</sup><https://www.theguardian.com/technology/2016/sep/08/beauty-contest-doesnt-like-black-people>

don't use the protected attributes for the model training, information about the race, gender, and age are implicitly encoded into intermediate representations of the model (Elazar & Goldberg, 2018).

Various bias mitigation techniques exist such as: (1) Pre-processing: transforming data representation before training (d'Alessandro et al., 2017) (Bellamy et al., 2018), (2) In-processing: modify the existing algorithms during the training process to remove discrimination, and (3) Post-processing: transform the model output to improve prediction fairness (Berk et al., 2017). The effectiveness of these methods have been tested on a real dataset; however, synthetic data is gaining traction in light of limited data access arising due to privacy and compliance. Recent research aims to study how biases are amplified, for example, in compressed deep neural networks (Hooker et al., 2020). Another study estimates the impact of bias in adopting differentially private mechanisms to protect the model from learning sensitive attributes (Bagdasaryan & Shmatikov, 2019).

In this study, we are particularly interested in bias amplification when models are trained using synthetic data. Generative adversarial networks (GANs) have become a popular choice for synthetic data generation (Goodfellow et al., 2014). Several variants of GANs are available today that are more stable and capable of generating realistic samples. In this paper, we have used state-of-the-art GAN architecture for tabular data generation: CTGAN (Xu et al., 2019), Gaussian Copula GAN (Masarotto et al., 2012), and Copula GAN<sup>4</sup>. Further to generate differentially private synthetic data, PATE-GAN (Jordon et al., 2018) is used.

The major contributions are: (1) demonstrating the impact of bias on models trained using synthetic data (2) Studied how differentially private synthetic data affect the model accuracy and the fairness (3) compared and contrast the key differences in model trained using synthetic data and its differentially private version.

## 2 EXPERIMENTAL RESULTS

### 2.1 FAIRNESS METRICS

There exists a myriad of notions in the literature to quantify fairness. Each measure emphasizes different aspects of fairness. Fairness in machine learning measures the degree of disparate treatment for different groups (e.g., female vs. male), or individual fairness, emphasizing similar individuals should be treated similarly. Evaluation metrics to measure group fairness are (i) Demographic Parity Difference (DPD) (Verma & Rubin, 2018). (ii) Demographic Parity Ratio (DPR) (iii) Equality of Odds (EoD) (iv) Equality of Opportunity (EoP) (Hardt et al., 2016). Please find the details of the metrics in the supplementary. We have used FairLearn library<sup>5</sup> to access fairness metrics. Some cut-off values have been defined in the literature for these metrics, suggesting whether models are fair or not. For DPD, EoD, and EoP, if the absolute value is smaller than 0.1, the model can be considered fair, whereas for DPR fairness range for this metric is between 0.8 and 1.25.

### 2.2 DATASET DESCRIPTION

We have used Adult Dataset<sup>6</sup> for our analysis, a well-known example, which is widely used in the fairness literature to predict whether the income of an individual exceeds  $\$50K/yr$  based on census data. The Adult dataset has about 48K records with a binary label indicating a salary  $\leq \$50K$  or  $\geq \$50K$ , with a class imbalance of [76%, 24%], where 76% of the observations are earning  $\leq \$50K$ . There are 14 attributes: 8 categorical and 6 continuous attributes. There are two attributes corresponding to the race and sex of an individual that can be used to define the protected groups.

<sup>4</sup>[https://en.wikipedia.org/wiki/Copula%28probability\\_theory%29](https://en.wikipedia.org/wiki/Copula%28probability_theory%29)

<sup>5</sup><https://fairlearn.github.io/>

<sup>6</sup><http://archive.ics.uci.edu/ml>

Table 1: Fairness Report

	DPD	DPR	EoD	EoP
Real Data	0.18	0.80	0.11	0.09
Balanced Data	0.29	0.64	0.20	0.20
CTGAN	0.05	0.95	0.04	0.042
GaussianCopula	0.13	0.86	0.16	0.074
CopulaGAN	0.01	0.99	0.02	0.01
PATE-GAN	0.45	0.46	0.39	0.39

Table 2: How frequently model is predicting an individual earning  $\geq \$50K$  for the groups (male and female) in test set.

	Male (7677)	Female (3639)
Ground Truth	2346 (30%)	412 (11%)
Real Data	2125 (27%)	347 (9%)
Balanced Data	3702 (48%)	699 (19%)
CTGAN	736 (9.5%)	170 (4%)
GaussianCopula	1325 (17%)	151 (4%)
CopulaGAN	305 (3.9%)	121 (3.3%)
PATE-GAN	4790 (62%)	646 (17.7%)

### 2.3 BIAS INDUCED FROM THE DATASET

The model’s performance in terms of group fairness metrics discussed above is presented in Table 1. We can see that the values in Table 1 indicate that the model is unfair as DPD and EoD are greater than 0.1. To further investigate the cause of the bias, we visualize the correlation matrix of the dataset’s attributes and found a high correlation between some of the allowed variables (non-protected) and gender (protected variables) as illustrated in Figure 1b. Next, we trained the Random Forest classifier on the balanced dataset by maintaining the same class ratio. However, from Table 1 we can observe that the balanced dataset has no impact as the model remains unfair. From Figure 1b we can observe non-protected attributes still have a high correlation with protected attributes.

Further, we analyze individual fairness in the model outcome. Here, we evaluate how frequently the model predicts an individual as ”earning more than \$50k per annum” irrespective of gender. Table 2 shows the performance in terms of individual fairness. We can observe that 30% of Male and 11% Female individuals earn more than \$50k per annum. Hence the dataset has more instances of male earning more than \$50k per annum than females. When we train the model on the Adult dataset after removing the gender (protected attribute) information, the proportion of predicting high-earning individuals for males is higher than for females. The model trained on real dataset predicts 27% male and 9% female individuals as earning more than \$50k per annum. However, the model trained on balanced dataset predicts 48% male and 19% female individuals as earning more than \$50k per annum. This difference in predicting high-earning individuals between both the groups increases even after training the model on the class balanced dataset.

### 2.4 IMPACT OF SYNTHETIC DATA ON BIAS AND FAIRNESS

We further study the impact of synthetic data on the model’s fairness. To generate the synthetic data for our experiments, we have used Synthetic Data Vault (SDV) library<sup>7</sup>. SDV is an open-source library to help end-users easily generate synthetic Data for different data modalities, including single table, multi-table, and time-series data. The SDV Python package includes a comprehensive set of generative models and evaluation framework that facilitates the task of generating evaluating the quality of Synthetic Data. From the SDV library, we have used CTGAN, GaussianCopula (Masarotto et al., 2012), and Copula model for data generation. We also explored the effect of differentially private algorithms for data generation. We implemented the PATE-GAN framework from the paper (Jordon et al., 2018) with  $\epsilon = 2$  privacy budget, where smaller  $\epsilon$  implies greater privacy.

<sup>7</sup><https://sdv.dev/>

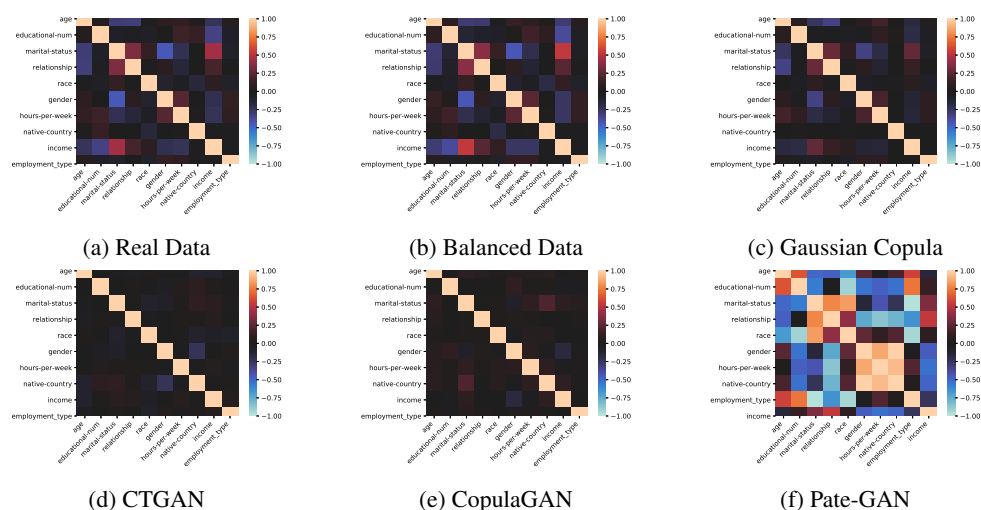


Figure 1: Represents the heatmap of the correlation matrix in synthetic data obtained from different generative models. Each heatmap of the correlation matrix shows a correlation coefficient between the variables. The higher the color variation in a cell, the higher is the relationship between the variables.

Table 1 provides details of performance on synthetic data in terms of fairness metrics. The synthetic data generated by methods like GaussianCopula, CTGAN, and CopulaGAN, significantly reduces the model’s disparity. The improvement in DPD and EoP indicates the model correctly classifies the same proportion of positive outcomes in both groups. However, a differentially private version of synthetic data (PATE-GAN) amplifies bias and unfairness in the model. This behavior can be explained by the correlation of the attributes in the synthetic data shown in Figure 1. In the synthetic data generated from CTGAN and CopulaGAN, all the attributes are weakly correlated and loosely dependent upon protected attributes (gender). In PATE-GAN, the attributes are highly correlated.

Next, we evaluate the model performance in terms of individual fairness as shown in Table 2. We find that training the model on synthetic data from CTGAN and CopulaGAN, the proportion of predicting high-earning individuals is almost the same for both groups. In contrast, PATE-GAN model predicts the higher-earning individuals more frequently for the male group. The disparities in outcomes for PATE-GAN is due to random noise added in their differentially private mechanism. Based on all the observations it is evident that the model trained on weakly correlated data is fairer and less biased. Hence an effective way to improve the fairness of the machine learning systems by introducing the use of less correlated synthetic data.

### 3 CONCLUSION

With the widespread adoption of synthetic data for being compliant and preserving privacy, the study evaluates its impact on bias amplification. Specifically, we studied disparity in machine learning model output when trained using synthetic data. Further, a differentially private version of GANs known as PATE-GAN is also studied in quantifying the impact of bias amplification. Experimental evaluation reveals that differentially private versions of the synthetic dataset possess more disparity as evident from fairness metrics. This effect is attributed to high cross-correlation in generated features. The outcome of the study will ensure that synthetic data is carried in responsible ways while accounting for the model disparity. We encourage future research directions to evaluate the impact of the privacy budget on model disparity. This further opens up the avenue to explore the trisection of utility, privacy, and fairness in synthetic datasets.

## REFERENCES

- Eugene Bagdasaryan and Vitaly Shmatikov. Differential privacy has disparate impact on model accuracy. *arXiv preprint arXiv:1905.12101*, 2019.
- Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, et al. Ai fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias. *arXiv preprint arXiv:1810.01943*, 2018.
- Richard Berk, Hoda Heidari, Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409*, 2017.
- Brian d’Alessandro, Cathy O’Neil, and Tom LaGatta. Conscientious classification: A data scientist’s guide to discrimination-aware classification. *Big data*, 5(2):120–134, 2017.
- Yanai Elazar and Yoav Goldberg. Adversarial removal of demographic attributes from text data. *arXiv preprint arXiv:1808.06640*, 2018.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.
- Moritz Hardt, Eric Price, and Nathan Srebro. Equality of opportunity in supervised learning. *arXiv preprint arXiv:1610.02413*, 2016.
- Sara Hooker, Nyalleng Moorosi, Gregory Clark, Samy Bengio, and Emily Denton. Characterising bias in compressed models. *arXiv preprint arXiv:2010.03058*, 2020.
- James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. Pate-gan: Generating synthetic data with differential privacy guarantees. In *International Conference on Learning Representations*, 2018.
- Preethi Lahoti, Krishna P Gummadi, and Gerhard Weikum. ifair: Learning individually fair data representations for algorithmic decision making. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pp. 1334–1345. IEEE, 2019.
- Guido Masarotto, Cristiano Varin, et al. Gaussian copula marginal regression. *Electronic Journal of Statistics*, 6:1517–1549, 2012.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*, 2019.
- Sahil Verma and Julia Rubin. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (Fairware)*, pp. 1–7. IEEE, 2018.
- Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5310–5319, 2019.
- Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. Modeling tabular data using conditional gan. *arXiv preprint arXiv:1907.00503*, 2019.