

Fair Classification via Transformer Neural Networks: Case Study of an Educational Domain

Modar Sulaiman
University of Tartu
Tartu, Estonia
modar.sulaiman@ut.ee

Kallol Roy
University of Tartu
Tartu, Estonia
kallol.roy@ut.ee

ABSTRACT

Educational technologies nowadays increasingly use data and Machine Learning (ML) models. This gives the students, instructors, and administrators support and insights for the optimum policy. However, it is well acknowledged that ML models are subject to bias, which raises concerns about the fairness, bias, and discrimination of using these automated ML algorithms in education and its unintended and unforeseen negative consequences. The contribution of bias during the decision-making comes from datasets used for training ML models and the model architecture. This paper presents a preliminary investigation of the fairness of transformer neural networks on the two tabular datasets: Law School and Student-Mathematics. In contrast to classical ML models, the transformer-based models transform these tabular datasets into a richer representation while solving the classification task. We use different fairness metrics for evaluation and check the trade-off between fairness and accuracy of the transformer-based models over the tabular datasets. Empirically, our approach shows impressive results regarding the trade-off between fairness and performance on the Law School dataset.

Keywords

Fairness; Classification; Educational Data; Bias; Representation Learning; Machine Learning; Transformer

1. INTRODUCTION

Automated decision-making with ML models in education is increasingly used to aid and support teachers, educators, and other stakeholders for optimal policy formulation. Though this method holds immense potential to improve prediction accuracy, the outcomes of the ML models show unfair results to some groups (e.g., an underrepresented minority group) of society. For example, the ML models show unfair approval for student loans for African students or predict lower bar exam success of students of low socioeconomic group [5].

In general, the notion of unfairness in ML is broadly categorized as follows [20], (i) Disparate treatment: where the ML model classifies differently (unfairly) people with the same values of non-sensitive features but different values of sensitive features, (ii) Disparate impact: where the ML model classifies that benefits (or hurts) people who are sharing the value of a sensitive feature vector more frequently than the other group, (iii) Disparate mistreatment: where the ML model achieves different classification accuracy for groups of people sharing different values of a sensitive feature. Our paper mainly investigated the fairness of classification using transformer models on tabular data in the educational domain. We used tabular non-textual data to train the transformer networks and thus verify their ability to make fair predictions in the classification tasks in two scenarios: first, training transformer networks without any bias mitigation method, and second, considering a bias mitigation method (Section 3) while training transformer networks. The representation of non-textual data in a rich representation via the transformer model is one of the main strengths of the transformer models. It follows the distributional hypothesis: a word is characterized by the company it keeps.

In our work, we showed that there exists a transformer model, namely the SAINT model, which achieves perfect group fairness without requiring any explicit debiasing method. In addition, one of the transformer models, the Tab model (with fairness constraint), improves fairness for the protected group at a negligible cost in terms of accuracy compared to the other models on the Law School dataset. Furthermore, we indicated a case using the Student-Mathematics dataset, where we do not recommend using the transformer models to mitigate bias in the final predictions. Finally, we demonstrated the possibility of empirically achieving a slight trade-off between performance and fairness using transformer models.

To the best of our knowledge, our work is the first study that explores the fairness ability of transformer-based models on tabular data in educational and other domains—in addition to considering the use of fairness constraint to train transformer-based models on tabular data in particular.

2. RELATED WORK

A comprehensive studies of algorithmic fairness in education is reported in [13]. They have investigated how discrimination emerges in automated systems and how it can be mitigated through studying and calibrating: the measurement

of input data, model learning, and output presentation. A comparative study of different fairness algorithms on multiple datasets are reported [6]).

The Law School Admission Council (LSAC) National Longitudinal Study reported discrimination against Africans (and other minorities) examinees during the bar passage examinations [19]. Recent approaches of using transformer-based models are showcased in [8, 11, 17] and in Section 4 we provide details of those transformer-based models. However, in contrast to the previous works [8, 11, 17] which only considered the performance of transformer-based models, we validate the fairness capability and performance of transformer-based models simultaneously. A flurry of interesting and insightful research on fairness constraints are investigated recently by various research groups [20, 15, 1]. Several authors addressed the problem of trade-off between fairness and accuracy [16, 3] by explaining why a trade-off exists on a given biased dataset and demonstrating the possibility of empirically improving fairness without sacrificing accuracy.

3. FAIR CLASSIFICATION THEORY

Fairness of machine learning examines the disparate treatment, disparate impact, or disparate mistreatment of different groups in a population, which are devised based on sensitive attributes, denoted by z , such as race, etc. Groups are divided into privileged groups (favorable treatment) and unprivileged groups (unfavorable treatment). Without loss of generality, let $z = 0$ be the unprivileged group and $z = 1$ be the privileged group. Privileged population groups would be more likely to receive favorable treatment than unprivileged population groups by a biased model.

In general, the binary classification task can be formulated as finding the optimum mapping function \mathcal{T} using input features $x \in \mathcal{R}^d$ and class labels $y \in (0, 1)$. For example, the label $y = 1$ could represent the application loan is getting accepted and $y = 0$ otherwise. This decision boundary-based classifier \mathcal{T} predicts a label y , by minimizing the cross-entropy loss function over a training set:

$$\arg \min_w \mathcal{L}_w W \quad (1)$$

$$\mathcal{L}(\mathcal{T}(x) = \hat{y}, y) = -(y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})) \quad (2)$$

where \hat{y} is the predicted label and w is the set of model weights. However, the classifier can be trained fairly using additional constraint for fairness guarantees of unprivileged group [1, 20]. The fairness constrained optimization problem can be expressed as the following:

$$\begin{aligned} & \arg \min_w \mathcal{L}_w \\ & \text{subject to } P(\cdot | z = 0) = P(\cdot | z = 1) \end{aligned} \quad (3)$$

where the constraint in (3) can be formulated as a condition of *no disparate treatment* (4), *no disparate impact* (5) or *no disparate mistreatment* (6, 7, 8) as shown below.

$$P(\hat{y}|x, z) = P(\hat{y}|x) \quad (4)$$

$$P(\hat{y} = 1 | z = 0) = P(\hat{y} = 1 | z = 1) \quad (5)$$

$$P(\hat{y} \neq y | z = 0) = P(\hat{y} \neq y | z = 1) \quad (6)$$

$$P(\hat{y} \neq y | y = -1, z = 0) = P(\hat{y} \neq y | y = -1, z = 1) \quad (7)$$

$$P(\hat{y} \neq y | y = 1, z = 0) = P(\hat{y} \neq y | y = 1, z = 1) \quad (8)$$

We consider only the condition of *no disparate treatment* (4) as the fairness constraint in our experiments. However, finding an optimum fair classifier defined by the constrained optimization problem (3) is non-trivial. The desired fair classifier (satisfying the constraints) may be of a non-convex boundary-based type; thus, finding the optimum weights w in some cases is challenging.

4. PROPOSED MODELS

This section highlights the models that we used in our work. Most of the transformer-based models are adaptations of the transformer architecture [18], and we examine them for fair classification in the tabular datasets in the educational domain.

The first transformer-based model that is explicitly designed for tabular data is Tab-Transformer [11]. For the sake of brevity, we refer to the Tab-Transformer model as *Tab* model in our paper. The Tab model uses attention mechanism to embed only categorical features in the tabular data. A very similar model to the Tab model is FT-Transformer (*Feature Tokenizer + Transformer*) model [8], or *FT* model, for short. The FT model is also designed specifically for tabular data. However, the FT model transforms all features (categorical and numerical) to embeddings where the continuous features are projected into a d -dimensional space before passing them through the transformer encoder.

Our paper shows that the most critical transformer-based model regarding fairness in the tabular dataset is the *SAINT* model, *Self-Attention and Intersample Attention Transformer* [17]. The architecture of the SAINT model itself plays the role of a regularizer for achieving fairness. Each block of SAINT model consists of two attention blocks: self-attention transformer block [18] and intersample attention (a type of row attention) transformer block. The intersample attention block of SAINT model computes attention across samples. In this approach, the features of a given data point interact with each other, then data points interact with each other using entire rows. Consequently, the SAINT model implicitly satisfies the condition (4) since it provides a better contextual representation than the other models for the data points and thus, does not need to add a fairness constraint to the objective function in some cases.

The last transformer-based model we used is Perceiver [12]. Perceiver model is designed to be architecture agnostic of the nature of the input data. It handles arbitrary configurations of different modalities (audio, images, and text). In our experiment we check the performance and fairness of Perceiver model in tabular data modality. In addition to the previous transformer-based models, we used one of the classical ML models, Logistic Regression (LR) model. LR is one of the common statistical analysis methods to predict a binary outcome and used in different works when studying fairness of ML.

5. DATASETS AND FAIRNESS METRICS

In this section, we describe the datasets used in our experiments and define the fairness metrics for assessing the models' fairness.

5.1 Datasets

We used two tabular datasets in our experiments: Law School (LSAC) and Student-Mathematics datasets. For both datasets,

we consider using the same features and processing¹ that are used to conduct the experiments in the survey [14].

The Law School dataset was developed by a Law School Admission Council (LSAC) survey across 163 law schools in the United States in 1991 [19]. The dataset is investigated in a variety of studies and is currently hosted in the database of Project SEAPHE². After cleaning and processing the law school data, we got 20,798 samples of students. We used 12 attributes (3 categorical, 3 binary and 6 numerical attributes) for Law School dataset in experiments. The model is used to predict whether or not an examiner will pass the bar exam. We label 'Non-white' to be the minority and unprivileged group in our experiments, 16% of the Law School dataset.

In the Student-Mathematics dataset³, the task is to predict the performance of secondary school students in mathematics subject [2]. The dataset contains 395 samples of students with 33 features (4 categorical, 13 binary and 16 numerical attributes). The female students in the Student-Mathematics dataset are the majority group, with about 52.7% of the dataset.

5.2 Fairness Metrics

Our paper focused on the disparate treatment of population groups (privileged and unprivileged), so we used different group fairness metrics to determine the bias in the model's final predictions.

We investigated three group fairness metrics in total: Absolute Between-ROC Area (*ABROC*) [7], Equal Opportunity Difference (*EOD*) [9] and Statistical Parity Difference (*SPD*) [4]. We interpret that the more unfair a model, the greater the difference in fairness metric values of each subgroup.

The *ABROCA* statistic is based on the Receiver Operating Characteristics (ROC) curve, and measures the absolute value of the area between the baseline group ROC curve ROC_b and those of one or more comparison groups ROC_c . Whereas *EOD* measures the difference between True Positive Rates (TPR) for unprivileged and privileged groups. Finally, *SPD* measures the difference between the probability of unprivileged group gets favorable prediction and the probability of privileged group gets favorable prediction.

6. EXPERIMENTS AND RESULTS

This section presents the results of our experiments and illustrates the trade-offs between fairness and performance. We randomly split the Law school and Student-Mathematics datasets into training and test sets for our experiments. Each test set was chosen to have 30% of the samples and the training sets to contain the remaining samples from each dataset.

Table 1 shows that all the transformer-based models (without fairness constraint) have comparable performance on the Law School dataset. Nonetheless, the Tab model (without fairness constraint) has shown marginally better performance scores on the Law School dataset than the other transformer models, with accuracy of 0.90016 and F1 score

¹<https://github.com/tailequy/fairness-dataset/tree/main/experiments>

²<http://www.seaphe.org/databases.php>

³<https://archive.ics.uci.edu/ml/datasets/student+performance>

Dataset: Law School. Protected attribute: race.				
Model	F1	Accuracy	SPD	EOD
LR	0.94984	0.90721	0.189538	0.082670
FT	0.94504	0.89839	-0.215906	-0.124452
Tab	0.94664	0.90016	-0.112048	-0.049809
Perceiver	0.94590	0.89919	-0.151387	-0.081128
SAINT	0.94299	0.89214	0	0

Dataset: Student-Mathematics. Protected attribute: sex.				
Model	F1	Accuracy	SPD	EOD
LR	0.91111	0.93277	0.153193	-0.005847
SAINT	0.76041	0.61344	0	0

Table 1: Results of each model without applying any fairness constraint.

Model	F1	Accuracy	SPD	EOD
LR	0.94643	0.89983	-0.100764	-0.043942
FT	0.94305	0.89230	0.001328	0.000617
Tab	0.94371	0.89342	0	0
Perceiver	0.94237	0.89102	0	0
SAINT	0.94012	0.88701	0	0

Table 2: Results of each model with fairness constraints on Law School. Protected attribute: race.

of 0.94664. Moreover, Table 1 indicates that the LR model has better performance than the Tab model by a very small margin, with an accuracy improvement by +0.00705 and F1-score improvement by +0.0032. Furthermore, Table 1 and Figure 1 show that the SAINT model outperforms the other models in terms of fairness without requiring any explicit debiasing method. This ensures that the candidates are treated equally based on different races (protected attribute in the Law School dataset) when predicting whether or not the candidate will pass the bar exam. In short, the SAINT model emerges as an ideal candidate for correctly identifying successful students at equal rates for different race subgroups.

However, transformer-based models are nonlinear with a large number of model parameters. This requires a large enough dataset to train [10] for achieving a good performance. Thus, the number of samples of Student-Mathematics dataset (only 395 samples) is not enough for training these transformer neural networks. As a result, the LR model outperforms the SAINT model on the Student-Mathematics dataset, with an accuracy of 0.93277 and F1-score of 0.91111. Nevertheless, same as in the Law School dataset, the SAINT model outperforms the LR model on the Student-Mathematics dataset based on the fairness metrics (SPD, EOD) without requiring any explicit debiasing technique.

Using the mentioned bias mitigation method in Section 3, Table 2 and Figure 1 show that the transformer-based models (with fairness constraint) successfully limit the bias in the final predictions on the Law School dataset at a minimal cost in terms of performance. Intriguingly, the Tab model (with fairness constraint) shows better performance and fairness than the other transformer models, with an accuracy of 0.89342, F1-score of 0.94371, and perfect group fairness. Furthermore, although the SAINT model showed its usefulness in achieving perfect group fairness without re-

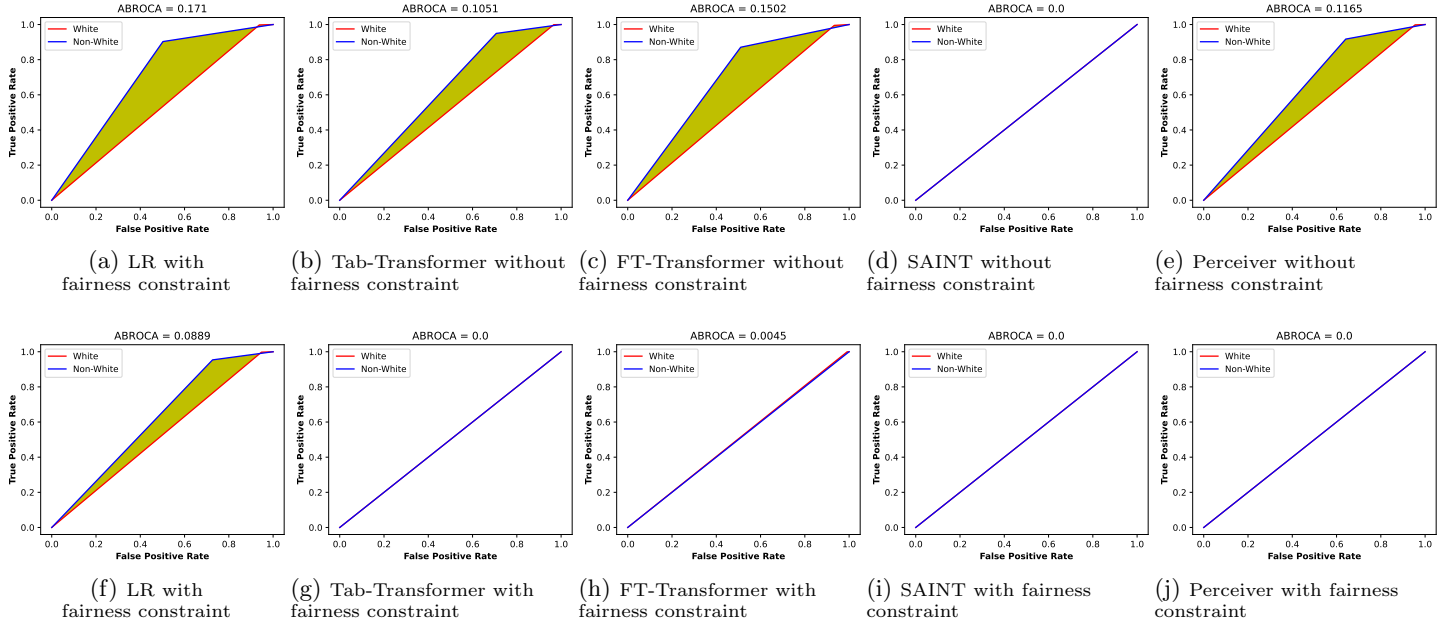


Figure 1: ABROCA of each model on Law School Data. Figures (a, b, c, d, e) show ABROCA of each model without applying fairness constraints. Figures (f, g, h, i, j) show ABROCA of each model when applying fairness constraints.

quiring any explicit debiasing method, the Tab model (with fairness constraint) outperformed the SAINT model (without fairness constraint) in terms of fairness and performance simultaneously, with an accuracy improvement by $+0.00128$ and F1-score improvement by $+0.00072$.

Additionally, the LR model showed two drawbacks: (i) it hardly shows a noticeable better performance than the other transformer models on the Law School dataset, and (ii) Figure 1 indicates that the LR model tends to fail to converge to a complete fair solution, with ABROCA of 0.171 for the LR model (without fairness constraint) and ABROCA of 0.0889 for the LR model (with fairness constraint).

In general, previous results in Table 1 and Table 2 show that when we use a large enough dataset, such as the Law School dataset, there is a slight trade-off between the performance and fairness using the transformer models (with fairness constraint). Additionally, Table 1 indicates that the SAINT model approximates the ideal distribution of the given dataset, which has a negligible accuracy-fairness trade-off.

7. CONCLUSION

Improving fairness while training the model and preserving its sound performance is challenging in ML. The prior works have focused on enhancing fairness in the final predictions using different bias mitigation techniques in classical ML models. Still, as we show, transformer models have some advantages when we use them to study fairness over the tabular dataset. In particular, we showed that the SAINT model achieves perfect group fairness without requiring any explicit debiasing method. Additionally, we showed that the Tab model (with fairness constraint) improves fairness for the protected group in the Law School dataset at a negligible cost in terms of accuracy compared to the other models. Our critical insight is that when trying to improve the fair-

ness in final results using transformer models, it is valuable to check if the tabular dataset is large enough and compatible with the number of the parameters in the transformer model.

Consequently, we believe that this insight and use of transformer models for fairness over tabular datasets provide a foundation for pursuing the fairness of artificial intelligence in educational and other domains.

For future work, we will comprehensively study the empirical performance of the transformer models for fair classification in different tabular datasets. In addition, we will employ the ideas of transfer learning and self-supervised learning to solve the problem of training the transformer-based models on a small dataset.

ACKNOWLEDGMENTS

This work has received funding by the European Social Fund via IT Academy programme.

8. REFERENCES

- [1] A. Agarwal, A. Beygelzimer, M. Dudík, J. Langford, and H. Wallach. A reductions approach to fair classification. In *International Conference on Machine Learning*, pages 60–69. PMLR, 2018.
- [2] P. Cortez and A. M. G. Silva. Using data mining to predict secondary school student performance. 2008.
- [3] S. Dutta, D. Wei, H. Yueksel, P.-Y. Chen, S. Liu, and K. Varshney. Is there a trade-off between fairness and accuracy? a perspective using mismatched hypothesis testing. In *International Conference on Machine Learning*, pages 2803–2813. PMLR, 2020.
- [4] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science*

- conference, pages 214–226, 2012.
- [5] M. Fei and D.-Y. Yeung. Temporal models for predicting student dropout in massive open online courses. In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, pages 256–263. IEEE, 2015.
 - [6] S. A. Friedler, C. Scheidegger, S. Venkatasubramanian, S. Choudhary, E. P. Hamilton, and D. Roth. A comparative study of fairness-enhancing interventions in machine learning, 2018.
 - [7] J. Gardner, C. Brooks, and R. Baker. Evaluating the fairness of predictive student models through slicing analysis. In *Proceedings of the 9th international conference on learning analytics & knowledge*, pages 225–234, 2019.
 - [8] Y. Gorishniy, I. Rubachev, V. Khrulkov, and A. Babenko. Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems*, 34, 2021.
 - [9] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.
 - [10] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. d. L. Casas, L. A. Hendricks, J. Welbl, A. Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
 - [11] X. Huang, A. Khetan, M. Cvitkovic, and Z. Karnin. Tabtransformer: Tabular data modeling using contextual embeddings. *arXiv preprint arXiv:2012.06678*, 2020.
 - [12] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira. Perceiver: General perception with iterative attention. In *International Conference on Machine Learning*, pages 4651–4664. PMLR, 2021.
 - [13] R. F. Kizilcec and H. Lee. Algorithmic fairness in education. *CoRR*, abs/2007.05443, 2020.
 - [14] T. Le Quy, A. Roy, V. Iosifidis, W. Zhang, and E. Ntoutsi. A survey on datasets for fairness-aware machine learning. *WIREs Data Mining and Knowledge Discovery*, Mar. 2022.
 - [15] S. Liu and L. N. Vicente. Accuracy and fairness trade-offs in machine learning: A stochastic multi-objective approach. *arXiv preprint arXiv:2008.01132*, 2020.
 - [16] K. T. Rodolfa, H. Lamba, and R. Ghani. Empirical observation of negligible fairness–accuracy trade-offs in machine learning for public policy. *Nature Machine Intelligence*, 3(10):896–904, 2021.
 - [17] G. Somepalli, M. Goldblum, A. Schwarzschild, C. B. Bruss, and T. Goldstein. Saint: Improved neural networks for tabular data via row attention and contrastive pre-training. *arXiv preprint arXiv:2106.01342*, 2021.
 - [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
 - [19] L. F. Wightman. Lsac national longitudinal bar passage study. lsac research report series. 1998.
 - [20] M. B. Zafar, I. Valera, M. Gomez-Rodriguez, and K. P. Gummadi. Fairness constraints: A flexible approach for fair classification. *The Journal of Machine Learning Research*, 20(1):2737–2778, 2019.