ON THE IMPACT OF MACHINE LEARNING RANDOMNESS ON GROUP FAIRNESS

Anonymous authors
Paper under double-blind review

ABSTRACT

Statistical measures for group fairness in machine learning reflect the gap in performance of algorithms across different groups. These measures, however, exhibit a high variance, between different training instances, that makes them unreliable for empirical evaluation of fairness. What is the cause of this variance, and how can we reduce it? We investigate the impact of different sources of randomness in machine learning on group fairness. We show that the variance in group fairness measures is mainly due to the high volatility of the learning process on under-represented groups, which itself is largely caused by the stochasticity of data order during training. Based on these findings, we show how to manipulate group level accuracy (i.e. model fairness), with high efficiency and negligible impact on the overall predictive power of the model, by changing the data order.

1 INTRODUCTION

Machine learning models are shown to manifest and escalate historical prejudices and biases present in their training data (Crawford, 2013; Barocas & Selbst, 2016; Zhao et al., 2017; Abbasi et al., 2019). Understanding these biases and the following ethical considerations has led to the rise of fair machine learning research (Chouldechova & Roth, 2018; Caton & Haas, 2020; Mehrabi et al., 2021). Recent work in fair deep learning have observed a trend of high variance in fairness measures across multiple training runs (Qian et al., 2021; Amir et al., 2021; Sellam et al., 2021; Soares et al., 2022), usually attributed to non-determinism in training. These results have challenged the legitimacy of existing claims in the literature (Soares et al., 2022), and have even disputed the effectiveness of various bias mitigation techniques (Amir et al., 2021; Sellam et al., 2021). Thus, a reliable extraction of fairness trends in a model requires accounting for the high variance to avoid lottery winners (see Figure 1).

The naive solution of executing a large number of identical training runs to capture the overall variance creates a huge computational demand, and discourages the examination of biases in several rapidly growing forefronts of machine learning research by increasing the resource requirements. Therefore, understanding the actual cause of the high variance in the fairness measures is critical. To the best of our knowledge, we are the first to study fairness variance beyond trivially executing multiple identical training runs. More specifically, we show the following:

- We show that the trends of fairness variance observed in literature are dominated by random data reshuffling during training, which causes high fairness variance between epochs even within a single training run, while the non-determinism in weight initialization has minimal influence.
- We extract an empirical relationship between group representation and instability in group level performance, highlighting higher vulnerability of minority to changing model behavior. Our results attribute the high fairness variance to lower prediction stability for under-represented subgroups.
- We demonstrate an immediate dominance of the data order on model fairness. A model's fairness is predictable, based on only the most recent training points, irrespective of preceding model behavior.
- Based on this information, we propose to use the fairness variance across epochs as a proxy to study the changing model fairness across multiple training runs, thus reducing the computational requirements by a significant margin.
- Finally, we manipulate group level performances (i.e., model fairness) by changing the data order, with a relatively minor impact on the overall accuracy. This manipulation can improve fairness as well as reverse the effects of several bias mitigation algorithms within a single training epoch.

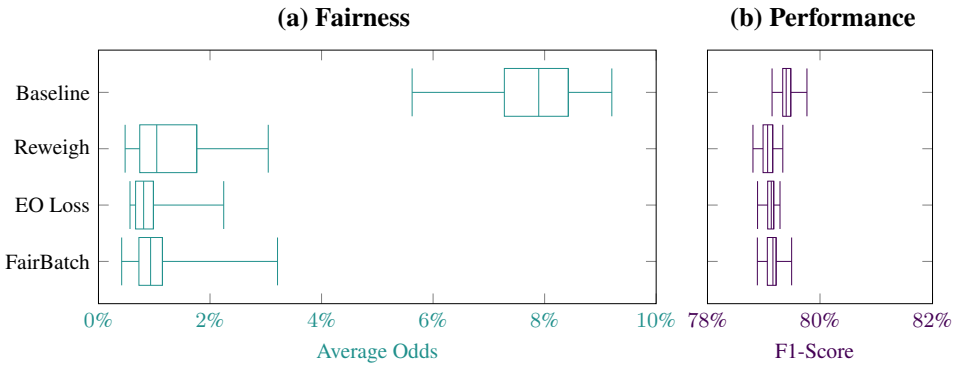

Figure 1: **(a)** Fairness (average odds) has a high variance across identical runs due to non-determinism in training. This variance persists even with several state-of-the-art bias mitigation algorithms (*Reweighing (Kamiran & Calders, 2012); Equalized Odds Loss (Fukuchi et al., 2020); FairBatch (Roh et al., 2020)*), and reliable comparison between these methods can only be made after capturing the overall model behavior, due to intersecting ranges. **(b)** The performance score (F1 score), however, has significantly smaller range of variance and doesn't face similar issues as fairness.

2 Experimental Setup

Datasets and Models We will conduct our investigation on ACSIncome and ACSEmployment tasks of the Folktables dataset (Ding et al., 2021), and binary classification of the 'smiling' label in CelebA dataset (Liu et al., 2015), with gender (Female vs. Male) as the sensitive attribute for all datasets. For CelebA, input features are obtained by passing the input image through a pre-trained and frozen ResNet-50 backbone and extracting the output feature vector after the final average pooling layer. We train a feed forward network on input features with a single hidden layer consisting of 64 neurons and *ReLU* activation, and train the model with *cross-entropy (CE)* loss objective for a total of 300 epochs, in all our experiments. We use train-val-test split of $0.7 : 0.1 : 0.2$, and maintain the same split throughout all our experiments, i.e. we do not consider potential non-determinism introduced due to changing train-val-test splits. All our evaluations are performed on the test split. We will focus primarily on the ACSIncome dataset in the main text, while additional experiments on CelebA and ACSEmployment are included in the appendix.

Fairness Metrics and Variance Fairness in machine learning has been interpreted widely into a multitude of definitions. In this paper, we will focus on average odds (AO), empirically interpreted as the average disparity between *true positive rates* and *false positive rates* of various groups. For model predictions R, true labels Y, and sensitive attributes A, average odds can be defined as,

$$AverageOdds := \frac{1}{2} \sum_{y=\{0,1\}} |\mathbb{P}(R = 1|Y = y, A = 0) - \mathbb{P}(R = 1|Y = y, A = 1)| \tag{1}$$

At the heart of our work is the study of fairness variance across model checkpoints. Unless otherwise specified, variance across multiple training runs refers to the variance across final checkpoints at epoch 300 of each training run. Similarly, variance across epochs refers to variance in a single training run across checkpoints at the end of every epoch for the last 200 epochs of training.

Non-Determinism in Model Training We set manual seeds at various intermediate locations in our code to control non-determinism in model training. We refer to the seed set right before building the neural model as the weight initialization seed, which influences the randomness in sampling the weight values. Similarly, we refer to the seed set right before the first training data shuffling as random reshuffling seed, which influences the data order that will be used as reference for the rest of the training. During training, we simply set the epoch number as seed right before reshuffling the reference data order at every epoch. As we can change the reference data order by changing the random reshuffling seed, this setup allows us to control non-determinism in data order throughout training with a single random seed. Fixing both weight initialization and random reshuffling seed

allows us to deterministically replicate model training, while changing both seeds simultaneously is similar to the discussion of non-determinism present in literature. Pseudo code detailing the definition of various random seeds is provided in Appendix A.

3 THE SOURCES OF VARIANCE

We start by moving past the trivial observation that different training runs lead to different outcomes, and investigate the underlying cause of high fairness variance by studying the intricacies of the training dynamics.

3.1 WEIGHT INITIALIZATION AND RANDOM RESHUFFLING

Neural model training under widely adapted mini-batch stochastic gradient descent (SGD) has two main sources of non-determinism, i.e. weight initialization and random reshuffling. To decouple their impact on fairness variance, we perform the following experiment. We execute three sets of 50 unique training runs each while , *(i)* allowing for both sources of randomness, *(ii)* changing only the weight initialization while keeping the random reshuffling fixed, and *(iii)* changing only the random reshuffling while keeping the weight initialization fixed. For ACSIncome, we collect average odds (AO) and F1 score of all training runs in Figure 2 (see Figure 10 for the other datasets).

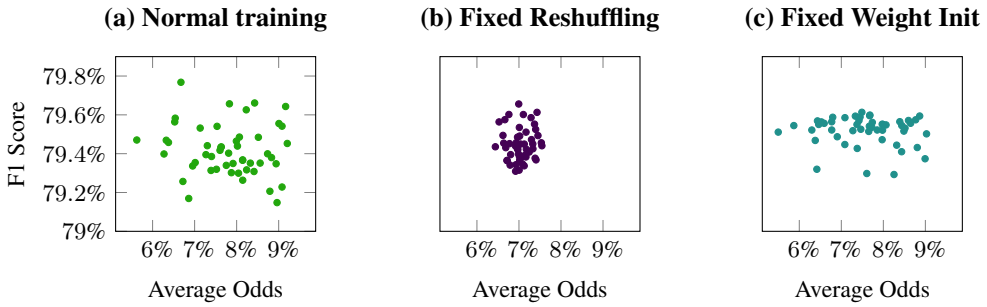

Figure 2: Average Odds (AO) variance by **(a)** allowing for both sources of randomness simultaneously, represents the instability in fairness scores as noted by existing literature. We see a significant drop in variance if we change only **(b)** the weight initialization while keeping the random reshuffling fixed. However, we observe high range of fairness variance by changing only **(c)** the random reshuffling, even for a fixed weight initialization. These results suggest random reshuffling of the data order as the dominant source of variance in fairness, with little influence from weight initialization.

The large range of fairness scores reported by allowing for both sources of randomness represents the variance observed in existing literature. Interestingly, when these sources are examined separately, the variance due to random reshuffling is equivalently large but the variance due to weight initialization drops significantly. It is clear that fairness variance originates from data order shuffling, while randomness in weight initialization has minimal impact. To further probe the difference between these two sources of non-determinism, we plot the median, inter-quartile range and overall range of average odds (AO) across the complete set of 50 training runs for the last 200 epochs in the two isolated settings from the previous experiment, i.e., changing only the weight initialization and only the random reshuffling, respectively. The resulting plots are collected in Figure 3.

We find high correlation in fairness scores across training runs with changing weight initialization (average pairwise pearson coefficient ≈ 0.91), despite high variance between epochs as observed in Figure 3(a). This behavior supports our observations of low variance in fairness scores at the final epoch in Figure 2. Furthermore, there is a lack of any reasonable correlation between fairness scores of training runs with changing random reshuffling (average pairwise pearson coefficient ≈ 0.13), with noticeably high variance across both changing epochs and training runs in Figure 3(b). Interestingly, the high fairness variance across epochs inside a single training run in Figure 3(a) closely matches the variance that we observe across multiple training runs in Figure 3(b).

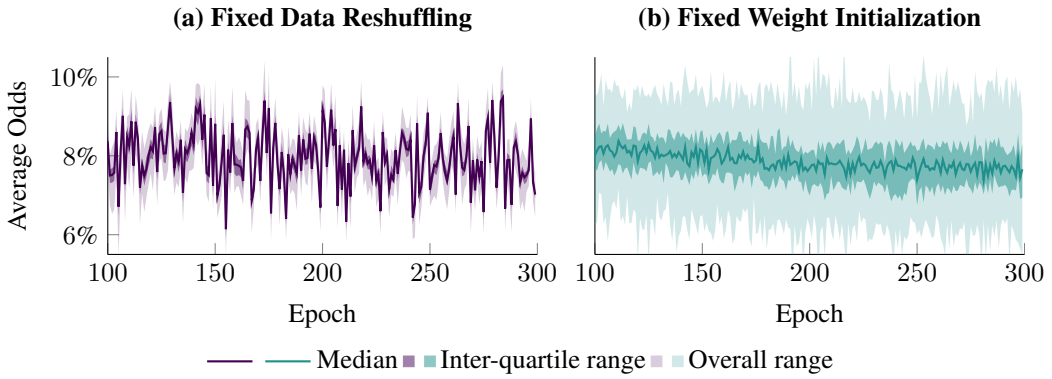

Figure 3: Median, inter-quartile range, and overall range of average odds across 50 different training runs while changing only the weight initialization or the random reshuffling respectively. **(a)** Changing only the weight initialization while keeping random reshuffling fixed implies that for any particular epoch, all training runs are fed the same data order. Despite different initializations, these models have very little variance across training runs, but high variance across epochs. This highlights the dominant impact of random reshuffling on model fairness. **(b)** Changing only the random reshuffling further supports our claim, as we see high variance even across training runs.

3.2 Changing Predictions and Data Distribution

In previous experiments, we observed high variance in model's predictive behavior even between consecutive epochs. The changing predictive behavior of neural models is not surprising, and has been studied extensively in literature (Toneva et al., 2018; Kirkpatrick et al., 2017; Jagielski et al., 2022; Tirumala et al., 2022). While these changing predictions maintain overall balance and keep a stable average performance, they can still have disparate impact on individual groups, the exact characteristics of which are less known. In this section, we study this instability by investigating individual data points which change their predictions. More specifically, we plot the percentage of unique data points from each group that changed their prediction at least once between epochs 100 and k, where we gradually increase the value of k. The results are collected in Figure 4.

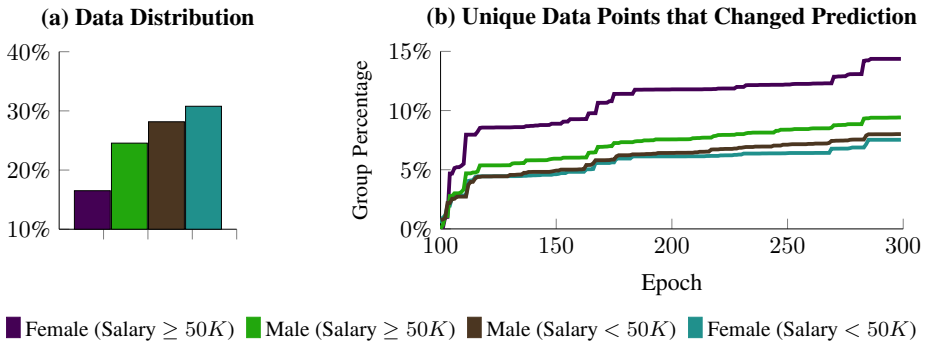

Figure 4: **(a)** The underlying data distribution of ACSIncome shows positive labels from group Female as an under-represented minority. **(b)** Total percentage of unique data points from each subgroup that change prediction across epochs follow the opposite order of their representation in the training data. These results highlight subgroups with least representation, i.e. positive labels from group Female in this setting, being the most vulnerable to changing predictions.

Clearly, the trends in percentage of unique data points with changing predictions mimics the distribution of each group in the original dataset, i.e. the groups which are represented the least are the most vulnerable to changing model behavior. In ACSIncome dataset, positive labels from group Female is severely under represented, and consequently can be seen to have almost twice the percentage of unique examples with changing predictions than any other group in the dataset. We provide additional

results with similar trends for different under-represented groups for CelebA and ACSEmployment in Appendix B.2. These results highlight higher vulnerability to changing model predictions of groups with smaller representation in the training dataset, which in turn results in high fairness variance.

4 Data Order and Fairness

The previous set of experiments suggest the existence of an intricate relationship between the data order and resulting model fairness. In this section, we dig deeper into this relationship and exploit it to approximate fairness variance without separately training multiple models.

4.1 Immediate Impact of Data Order

A single epoch of SGD comprises of learning on a sequence of mini-batches from the training data, changing the model parameters one gradient update at a time. The data order during training governs the order of gradient updates and thus its impact on resulting model behavior is quite obvious (Bengio et al., 2009). However, our investigation into fairness variance suggests that this impact might be more immediate than expected, capable of stably predicting fairness scores irrespective of the model's preceding behavior. To test this, we design an experiment as follows. We sample 1000 unique checkpoints randomly from last 200 epochs of 50 different training runs while allowing both forms of training non-determinism simultaneously. These checkpoints represent a variety of models with different weight initialization, training data order and even number of training epochs. We then train each of these checkpoints for exactly one epoch on a common randomly chosen data order. We collect fairness variance across checkpoints before and after this single epoch of training in Figure 5.

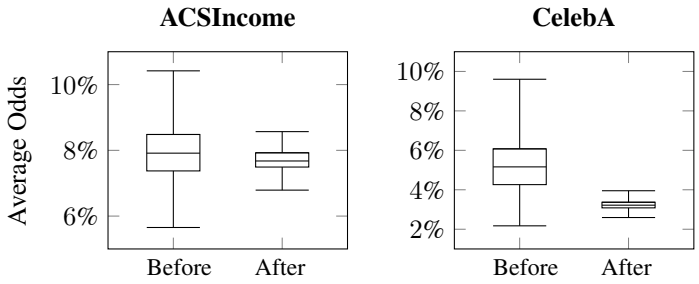

Figure 5: (*Before*) Multiple checkpoints taken from training runs with changing weight initialization, training data order, and even number of training epochs, show a high range of fairness variance (i.e., average odds variance), as expected. *(After)* However, these fairness scores are stabilized significantly by training these checkpoints for only a single epoch on a common randomly chosen data order. This shows the immediate impact of data order on model fairness.

Before training on the same data order for a single epoch, these checkpoints represent the complete range of fairness variance previously noted. However, after only a single epoch of training on a common data order, these models have all moved towards the same fairness score with significantly lower variance. This highlights the immediate impact of data order of model fairness, which is stable based on only the data order of the most recent epoch. An existing line of research suggests the presence of no energy valleys in deep learning loss landscape between minimas of separately trained models (Draxler et al., 2018; Garipov et al., 2018), which could explain the immediate impact of data order, as various checkpoints might be connected in the loss landscape.

Epochs are pseudo constructs, used to record multiple passes through the data, and thus we expect the influence of data order on model fairness to have a granular manifestation as the few most recent gradient updates of the model. Furthermore, while we have shown that the fairness variance is indeed reduced, we are yet to show that it is also predictable. To achieve both simultaneously, we extend the experiment above as follows. We start by extracting the data order corresponding to the epochs achieving best and worst fairness scores on the validation set in a standard training run. Next, we train these checkpoints on exactly b batches, taken from the suffix of data order with best and worst fairness scores as described above. This allows us to obtain models which have experienced the same most recent b gradient updates. For example, at $b = 0$, this setup involves no training and we simply

report the variance across the original set of checkpoints, while at $b = 1070$ (i.e., total number of batches in an epoch for ACSIncome), this setup will repeat the previous experiment. We test for $b = \{0, 20, 50, 100, 200\}$ and the results are collected in Figure 6.

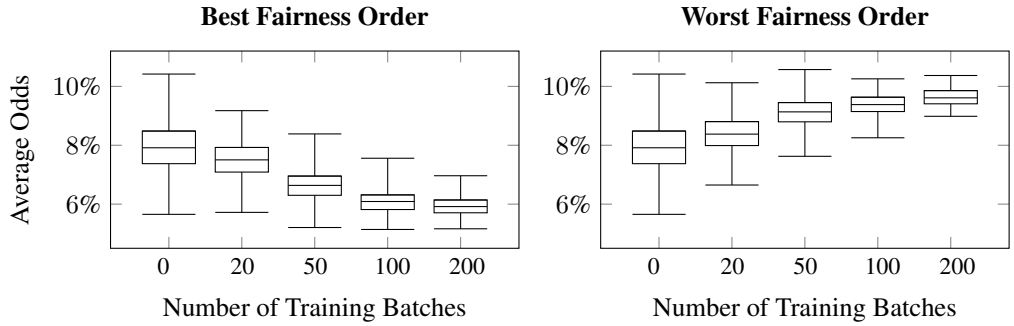

Figure 6: Fairness scores (average odds) stabilize as the number of most recent common training batches increase, highlighting the impact of these gradient updates on model fairness. Moreover, it stabilizes to low or high fairness scores based on the corresponding data order from the epoch with **best** or **worst** fairness score, showing the predictability of model fairness for any given data order.

As the number of fixed batches b increases, the model fairness gets more stable, and already at $b = 200$, the model fairness is as stable as seen in previous set of experiments. This shows that it is indeed the most recent gradient updates seen by the model which dominate its fairness scores, irrespective of model behavior prior to those updates. Moreover, the trends of stability move towards expected fairness based on the choice of original data order taken from best or worst fairness epoch as described earlier, which helps us show the predictability of the impact of data order.

4.2 Capturing Fairness Variance in a Single Training Run

We now return to our original problem of capturing fairness variance without wasting computing resources on a large number of training runs. We saw stability and predictability in fairness scores based on the data order, and thus fairness variance across multiple training runs is simply the randomness in data order across their last epochs. As these data orders are obtained by random shuffling, their distribution across multiple training runs should result in the same empirical range of fairness as their distribution across epochs. Thus, we propose evaluating intermediate checkpoints in a single training run as a proxy for reporting fairness variance across multiple training runs.

We test this by first plotting the distribution of fairness scores for sampling across multiple training runs (three different stopping epochs) and sampling across epochs (three different training runs) in Figure 7(a). Furthermore, we also test the quality of black swans, i.e. the best models captured, as a function of number of unique training runs and number of epochs evaluated per training run. For a data point $epochs = e; seeds = s$, we perform s unique training runs, and evaluate the model for last e epochs in each training run, thus accumulating a total of $e * s$ checkpoints. These checkpoints are then used to calculate some quality measure, and averaged over 50 repeats to compensate for randomness in choosing s training runs. The experiment is performed for all combinations of $e \in [1, 50]$ and $s \in [1, 50]$. We use two different quality measures for a set of checkpoints as described above, *(i)* the best fairness achieved across all checkpoints, and *(ii)* the Hausdorff distance (Birsan & Tiba, 2005) of the pareto-front (including both fairness and F1 scores) from the best achievable pareto-front, i.e., for $e = 50; s = 50$. The results are collected in Figure 7(b).

As we can notice from the plot in Figure 7(a), the empirical distribution of fairness across multiple training runs closely matches the distribution across a single training run. We even performed the Kolmogorov–Smirnov (KS) test (Massey Jr, 1951) to match sampling across multiple runs (at epoch 300) and sampling across epochs (for a single training run with seed 0) and found the maximum difference in empirical CDF to be only 0.07, with a p-value of 0.712, i.e. the probability that both set of checkpoints were sampled from the same underlying distribution. Moreover, the quality of black swans in Figure 7(b) for both quality measures suggests similar trends. For example, when comparing the two extremes, i.e. $e = 1; s = 50$ and $e = 50; s = 1$, we see similar results, with latter

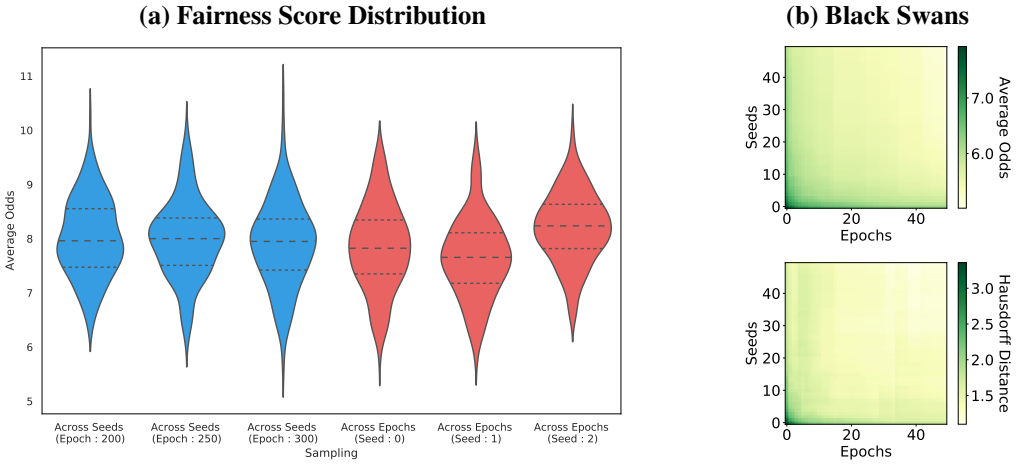

Figure 7: **(a)** Fairness scores (average odds) across multiple training runs and across epochs in a single training run have similar empirical distributions. Thus, studying this distribution across epochs provides a highly efficient alternative to executing multiple training runs. **(b)** Quality of black swans (i.e. extremely rare checkpoints) improves with more checkpoints collected, either in terms of fairness (the lowest achievable average odds score) or the trade-off between overall performance and fairness (the hausdorff distance to the best pareto front). This improvement occurs at the same rate, irrespective of sampling checkpoints across multiple random seeds (x-axis) or multiple epochs in a single seed (y-axis). But, sampling across epochs is significantly cheaper.

requiring 50 times less computation. This shows that the commonly used method to capture fairness variance in literature (i.e., $e = 1; s = 50$) is highly inefficient use of computing resources, and one can achieve the same by simply observing model behavior across a single training run. With these experiments, we show direct benefits of evaluating multiple epochs in a single training run, saving huge amounts of resources and time in capturing the overall model variance.

5 Manipulating Group Level Accuracy with Data Order

In previous sections, we explored the instability in model fairness and took advantage of our insights to reduce computational requirements for studying this variance. However, the predictability of fairness scores based on the most recent batches implies the possibility that we should be able to finely manipulate the fairness scores, or more precisely the group level accuracy values, by simply controlling the data order for a single epoch. We propose that the data distribution seen by the model in the most recent batches can help us nudge the group level accuracy behavior as desired. To test this hypothesis, we perform the following experiment. We start by defining the ratio between different groups, to indicate the distribution that will dominate the model behavior. For example, a ratio of $1 : 10$ between two groups will distinctly favor the latter group, while a ratio of $10 : 1$ will favor the first group. We then form batches in the exact ratio as provided until we run out of data points (which will happen for any ratio that is not the exact distribution of the original dataset). The excess data points are then shuffled randomly and place at the prefix of data order.

Similar to previous experiment settings, we pick the last epoch checkpoint from 50 unique training runs while allowing for both forms of non-determinism (i.e. weight initialization and random reshuffling), and then train them all for a single epoch on our custom data order. For our first experiment, we fix all other ratios and only change the ratio between positive and negative labels in group Female. We record the group level accuracy for all subgroups separately, along with the overall model accuracy. We repeat the experiment by now changing only the ratio between positive and negative labels in group Male. The results are collected in Figure 8.

It is clear that by manipulating the data distribution seen by the model in the most recent gradient updates, we can control the group level accuracy of the model. While the overall accuracy drops noticeably for extreme ratios, it does not change much in the middle despite significant variance in group level performances. This further strengthens our claim on the dominance of data order

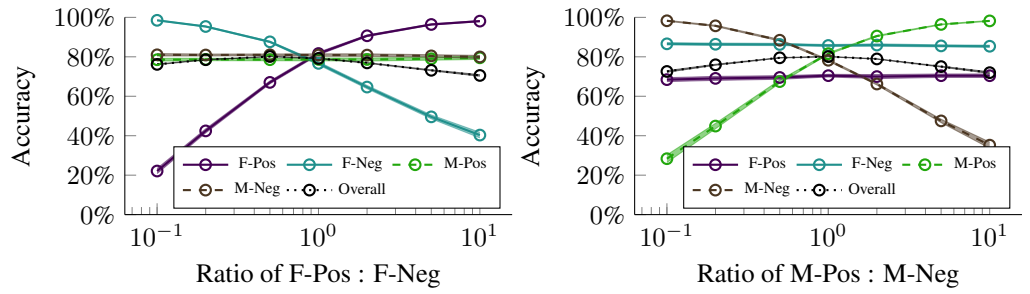

Figure 8: We show how to control group level accuracy by changing data distribution of the most recent gradient updates, tested separately for (a) ratio between positive and negative labels for group Female, and (b) ratio between positive and negative labels for group Male, while keeping other ratios fixed. In only a single epoch of training, we are able to manipulate the group level accuracy trade-off, with relatively small impact on overall accuracy.

on model fairness and group level accuracy. *Fairbatch* (Roh et al., 2020), a state-of-the-art bias mitigation algorithm, follows a similar formulation as our custom data order. They create batches with fixed ratio between groups, and the ratio is optimized at every epoch to counter the existing bias in the model. Our results not only explain the success of their algorithm, but also state that instead of regularly adapting the distribution to compensate for the model bias, one can directly use the final desired distribution and the model will adapt to it immediately.

Finally, we extract two special cases from the above experiments, i.e. ratios of $1:1$ and $1:3$ between positive and negative labels of group Female (for ACSIncome dataset), and call them *EqualOrder* and *AdvOrder* respectively. While *EqualOrder* achieves competitive performances with state-of-the-art bias mitigation methods, *AdvOrder* is capable of derailing the fairness scores of various mitigation algorithms. Note that we can force even worse fairness gaps than achieved by *AdvOrder* by pushing the ratio to its extreme, however that will also noticeably impact the model's overall accuracy. We perform experiments with three unique setups, standard training and two commonly used bias mitigation algorithms. This includes, (i) *Reweighing* (Kamiran & Calders, 2012), a data pre-processing algorithm which weighs every label-group pair based on its representation in the overall dataset, and (ii) *Equalized Odds Loss* (Fukuchi et al., 2020), an in-processing loss function to nudge the model towards fair predictions. We incorporate both *EqualOrder* and *AdvOrder* in these setups to see the influence of data order. The results are collected in Figure 9.

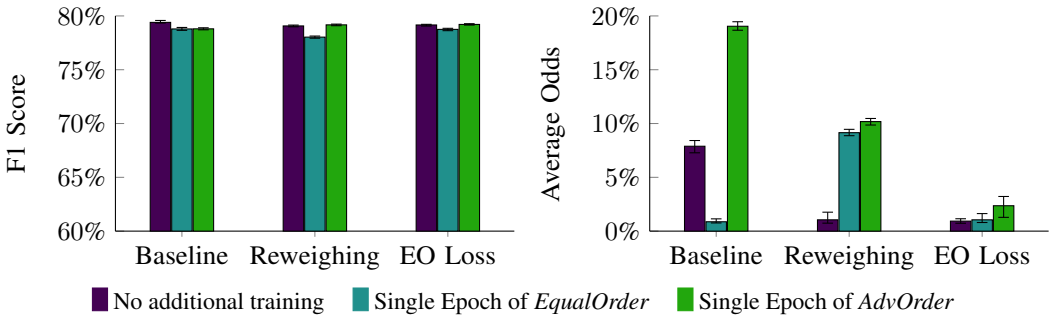

Figure 9: Under baseline training setup, by changing only the data order for just a single epoch of training, *EqualOrder* gets competitive performance to commonly used bias mitigation methods. Similarly, *AdvOrder* gets significantly worse fairness than even the baseline by simply creating an adversarial data order. For Reweighing, we see an increase in bias even with *EqualOrder*, as the beneficial weighing of minority creates bias towards the majority. On the other hand, even reweighing is not enough to counter the effects of *AdvOrder*, which achieves fairness scores close to baseline. Finally, equalized odds loss is capable of dynamically adapting to the model's changing predictive behavior, yet we still observe noticeable increase in bias under *AdvOrder*.

By simply training with *EqualOrder* for a single epoch, the baseline models achieve competitive fairness scores to state-of-the-art bias mitigation algorithms. *Reweighing* suffers from an unexpected high bias under *EqualOrder*, as the combination of ideally distributed data order along with increase in the weight of minority data points pushes the model towards significant unfair behavior against the majority. On the other hand, using *AdvOrder* in the baseline setting can further push the model bias, emphasizing the adversarial power of data ordering. Both *Reweighing* and *Equalized Odds Loss* also suffers similar fates, but with controlled damage in latter due to the loss function regularly adapting to the degrading behavior. These results together cement the effectiveness of manipulating group level accuracy by controlling the data order for just a single epoch of training.

6 Related Work

Variance in Fair Machine Learning Recently, there has been a growing awareness of high variance and instability in fair deep learning, associated with non-determinism in model training or underlying implementations (Amir et al., 2021; Sellam et al., 2021; Soares et al., 2022; Qian et al., 2021), and the challenges it creates in trusting existing results. Soares et al. (2022) refute various trends in literature associated with fairness of large language models by simply executing multiple identical training runs, while Sellam et al. (2021) focus on the impact of randomness in pre-training paradigm and find further increase in instability. Amir et al. (2021) even revisit state-of-the-art bias mitigation in clinical texts and show a lack of statistically significant improvement after accounting for non-determinism in training. While existing literature focuses on exploring the impact of fairness variance, we instead focus on investigating the source of this non-determinism. Furthermore, we propose to move away from the practise of executing multiple training runs to capture this variance and instead provide a computationally efficient proxy.

Data Order and Example Forgetting Controlling the data order during training has gained noticeable success in improving convergence speeds (Bengio et al., 2009; Shah et al., 2020; Soviany et al., 2022; Mohtashami et al., 2022), while its adversarial capabilities have also been explored (Shumailov et al., 2021). Even under randomly shuffled data order, neural models are known to undergo changes in predictive behavior during training (Toneva et al., 2018; Kirkpatrick et al., 2017; Jagielski et al., 2022; Tirumala et al., 2022). We extend the discussion on non-determinism in model behavior to model fairness, and investigate its connections with data order. We focus specifically on the instability in predictions made by the learning model and not its internal state, as we are concerned about the disparity in group level decisions which highlights the encoded bias.

7 Conclusion

Fairness variance in deep learning has raised concerns regarding the reliability of existing results in literature. In this paper, we took a closer look at this variance and proposed a proxy to training multiple models by studying fairness scores inside a single training run, which exploits a surprisingly dominant impact of data order on model behavior. We even manipulated the data order to minutely control group level accuracy without losing overall accuracy in only a single epoch on training. In our work, we focused only on the discrete decisions made by the model, as we were investigating the impact of non-determinism in model training on its fairness. However, further extensions of this discussion to trends in the internal state of the learned model can reveal even granular characteristics, and has potential application in similar fields of research suffering from high variance.

8 Reproducibility Statement

Our work studies various sources of non-determinism in neural model training and its impact on model fairness. Thus, it is important to address the implementation choices we have made to allow easy reproducibility. We provide a minimal psuedo code in Appendix A to clearly define various random seeds used in our setup to control training non-determinism. The complete training and evaluation pipeline is deterministic outside the choice of these random seeds, i.e., one can perfectly replicate the training by simply fixing these seeds. We also provide our code to train models and perform several evaluations done in the paper as supplementary material to allow easy adaptation for future research.

REFERENCES

Mohsen Abbasi, Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. Fairness in representation: quantifying stereotyping as a representational harm. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, pp. 801–809. SIAM, 2019.

Silvio Amir, Jan-Willem van de Meent, and Byron C Wallace. On the impact of random seeds on the fairness of clinical classifiers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 3808–3823, 2021.

Solon Barocas and Andrew D Selbst. Big data's disparate impact. *Calif. L. Rev.*, 104:671, 2016.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.

T Birsan and Dan Tiba. One hundred years since the introduction of the set distance by dimitrie pompeiu. In *IFIP Conference on System Modeling and Optimization*, pp. 35–39. Springer, 2005.

Simon Caton and Christian Haas. Fairness in machine learning: A survey. *arXiv preprint arXiv:2010.04053*, 2020.

Alexandra Chouldechova and Aaron Roth. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810*, 2018.

Kate Crawford. The hidden biases in big data. *Harvard business review*, 1(4), 2013.

Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair machine learning. *Advances in Neural Information Processing Systems*, 34, 2021.

Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. Essentially no barriers in neural network energy landscape. In *International conference on machine learning*, pp. 1309–1318. PMLR, 2018.

Akihiko Fukuchi, Yoko Yabe, and Masashi Sode. Fairtorch. `https://github.com/wbawakate/fairtorch`, 2020.

Timur Garipov, Pavel Izmailov, Dmitrii Podoprikhin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018.

Matthew Jagielski, Om Thakkar, Florian Tramèr, Daphne Ippolito, Katherine Lee, Nicholas Carlini, Eric Wallace, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, et al. Measuring forgetting of memorized training examples. *arXiv preprint arXiv:2207.00099*, 2022.

Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems*, 33(1):1–33, 2012.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114 (13):3521–3526, 2017.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

Frank J Massey Jr. The kolmogorov-smirnov test for goodness of fit. *Journal of the American statistical Association*, 46(253):68–78, 1951.

Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.

Amirkeivan Mohtashami, Sebastian Stich, and Martin Jaggi. Characterizing & finding good data orderings for fast convergence of sequential gradient methods. *arXiv preprint arXiv:2202.01838*, 2022.

Shangshu Qian, Viet Hung Pham, Thibaud Lutellier, Zeou Hu, Jungwon Kim, Lin Tan, Yaoliang Yu, Jiahao Chen, and Sameena Shah. Are my deep learning systems fair? an empirical study of fixed-seed training. *Advances in Neural Information Processing Systems*, 34, 2021.

Yuji Roh, Kangwook Lee, Steven Euijong Whang, and Changho Suh. Fairbatch: Batch selection for model fairness. In *International Conference on Learning Representations*, 2020.

Thibault Sellam, Steve Yadlowsky, Ian Tenney, Jason Wei, Naomi Saphra, Alexander D'Amour, Tal Linzen, Jasmijn Bastings, Iulia Raluca Turc, Jacob Eisenstein, et al. The multiberts: Bert reproductions for robustness analysis. In *International Conference on Learning Representations*, 2021.

Vatsal Shah, Xiaoxia Wu, and Sujay Sanghavi. Choosing the sample with lowest loss makes sgd robust. In *International Conference on Artificial Intelligence and Statistics*, pp. 2120–2130. PMLR, 2020.

Ilia Shumailov, Zakhar Shumaylov, Dmitry Kazhdan, Yiren Zhao, Nicolas Papernot, Murat A Erdogdu, and Ross J Anderson. Manipulating sgd with data ordering attacks. *Advances in Neural Information Processing Systems*, 34:18021–18032, 2021.

Ioana Baldini Soares, Dennis Wei, Karthikeyan Natesan Ramamurthy, Moninder Singh, and Mikhail Yurochkin. Your fairness may vary: Pretrained language model fairness in toxic text classification. In *Annual Meeting of the Association for Computational Linguistics*, 2022.

Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. Curriculum learning: A survey. *International Journal of Computer Vision*, pp. 1–40, 2022.

Kushal Tirumala, Aram H Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. Memorization without overfitting: Analyzing the training dynamics of large language models. *arXiv preprint arXiv:2205.10770*, 2022.

Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. An empirical study of example forgetting during deep neural network learning. In *International Conference on Learning Representations*, 2018.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 2979–2989, 2017.

A NON-DETERMINISM IN MODEL TRAINING

We provide a minimal pseudo-code to explain our control over training non-determinism and define weight initialization and random reshuffling seeds.

```
....
torch.manual_seed(weight_initialization_seed)
model = MLPModel() # Initialize Neural Model
...
train_data, .. = load_dataset() # load data with deterministic order
train_data = shuffle(train_data, seed=random_reshuffling_seed)
...
trainloader = torch.utils.data.DataLoader(train_data, ..., shuffle=True)
...
for epoch in range(total_epochs):
    ...
    torch.manual_seed(epoch)
    for batch in trainloader:
        ...
```

B ADDITIONAL EXPERIMENTS FOR VARIANCE BEYOND RANDOM SEEDS

B.1 WEIGHT INITIALIZATION AND RANDOM RESHUFFLING

We provide additional results in Figure 10, 11 on CelebA and ACSEmployment dataset to show the dominance of random reshuffling on model fairness.

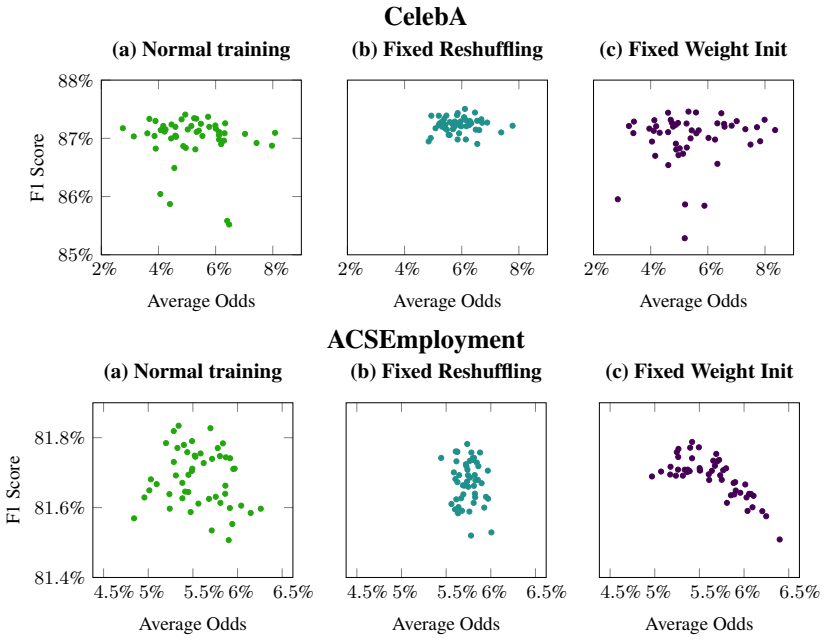

Figure 10: Additional experiments on CelebA and ACSEmployment datasets reveal similar trends as seen in Figure 2. Random reshuffling of data order is the dominant source of variance in fairness scores, with very little influence from the weight initialization.

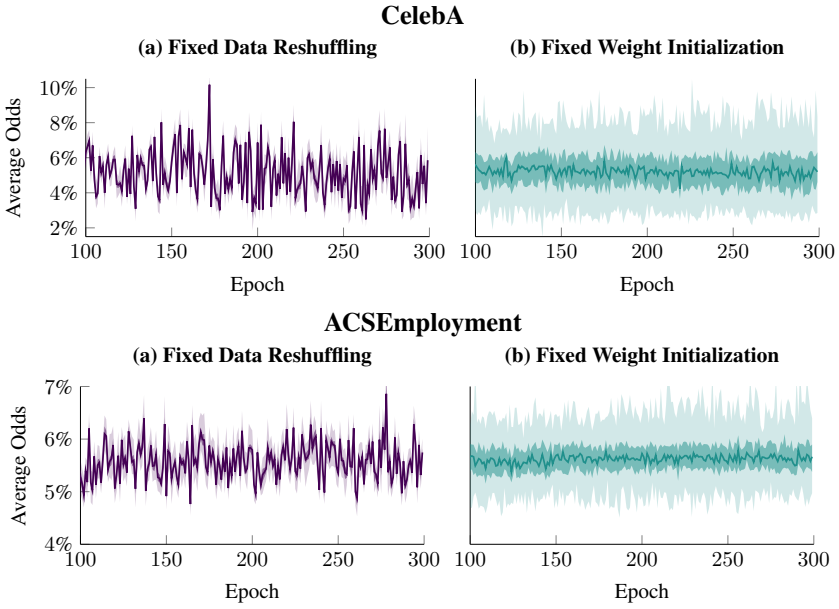

Figure 11: Additional experiments on CelebA and ACSEmployment datasets reveal similar trends as seen in Figure 3. These results further highlight the dominant impact of random reshuffling on model fairness.

B.2 Changing Predictions and Data Distribution

We show the relationship between data distribution and changing predictions for CelebA and ACSEmployment in Figure 12.

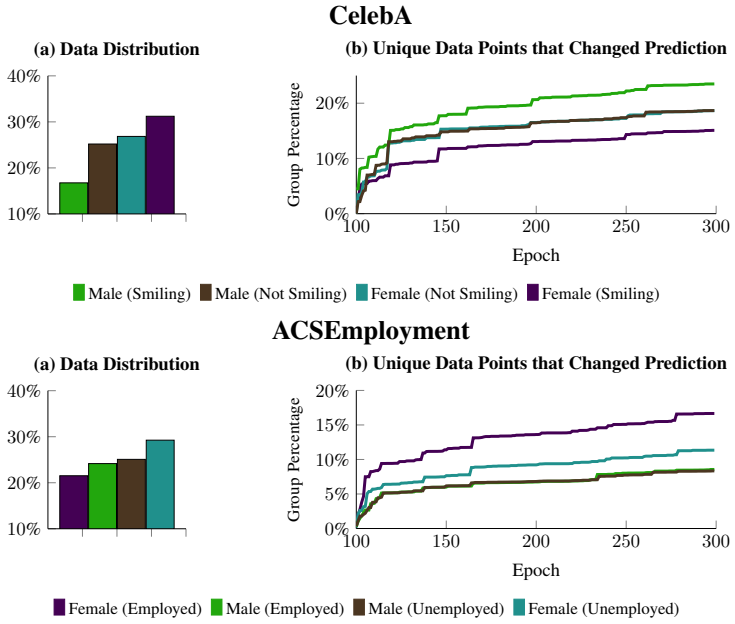

Figure 12: Additional experiments on CelebA and ACSEmployment datasets reveal similar trends as seen in Figure 4. These results highlight subgroups with least representation being the most vulnerable to changing predictions.

We provide ablation study for Figure 6 and show that it is indeed the batches from data order prefix that govern the predictability of model fairness.

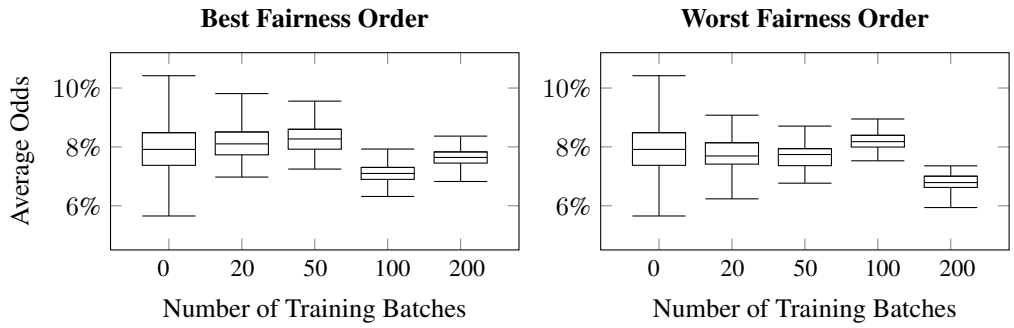

Figure 13: Additional experiments by choosing b batches from the prefix for Figure 6. The results still maintain stability as b increases but now they stabilize to random fairness scores instead of best and worst fairness scores as seen in Figure 6. Clearly, its the most recent batches in that order which truly governs the model fairness.

C Additional Experiments for Data Order and Fairness

C.1 Capturing Variance in a Single Training Run

We show the empirical similarity of distribution across multiple training runs and multiple epochs in a single training run for additional datasets CelebA and ACSIncome in Figure 14.

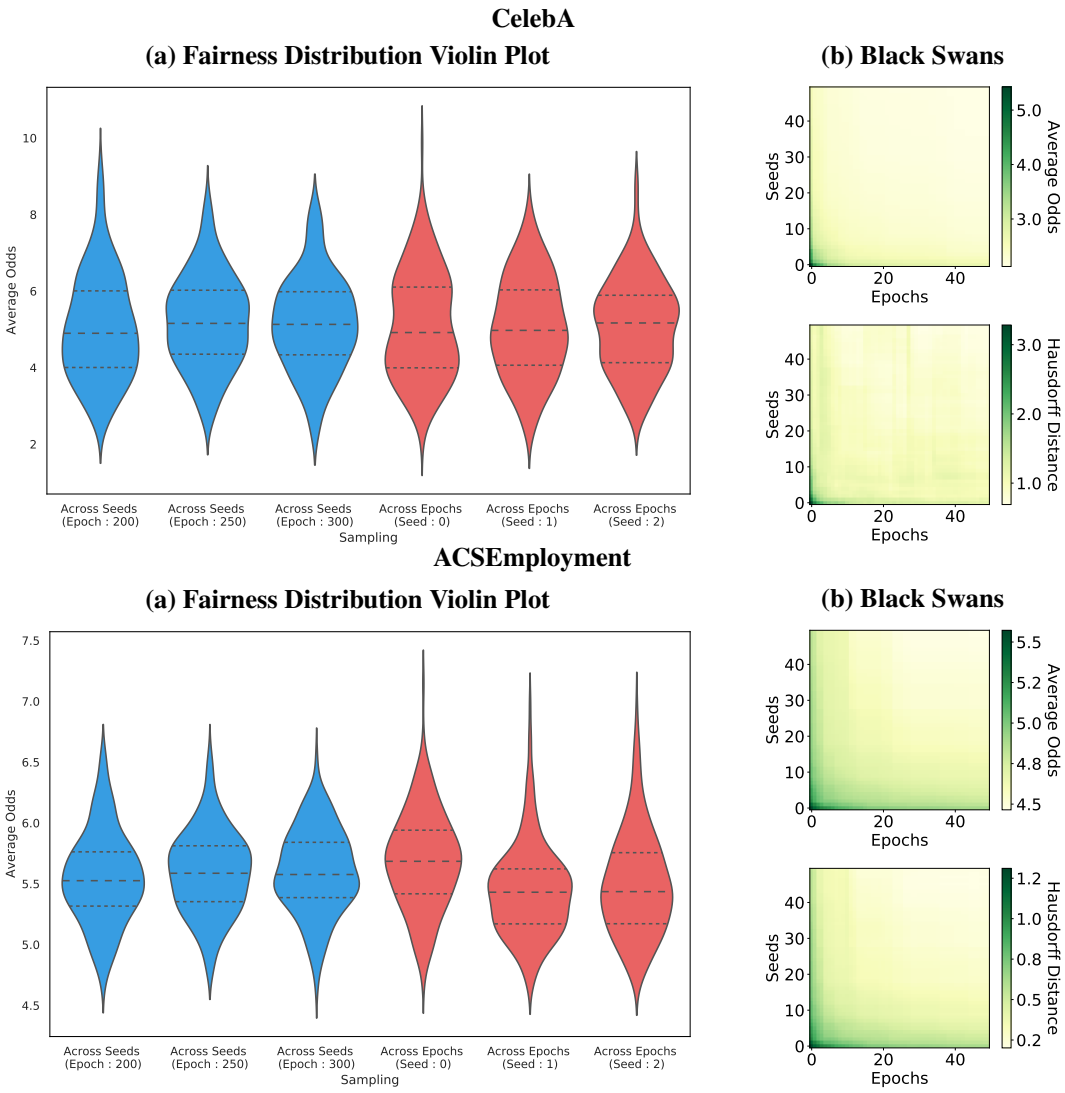

Figure 14: Additional experiments on CelebA and ACSEmployment datasets reveal similar trends as seen in Figure 7. Fairness scores (average odds) across multiple training runs and across epochs in a single training run have similar empirical distributions. Thus, studying this distribution across epochs provides a highly efficient alternative to executing multiple training runs.

D Additional Experiments for Manipulating Group Level Accuracy with Data Order

We provide additional results in Figure 15, 16, highlighting the predictability of model fairness based on the data order, and its impact compared against various state-of-the-art baselines.

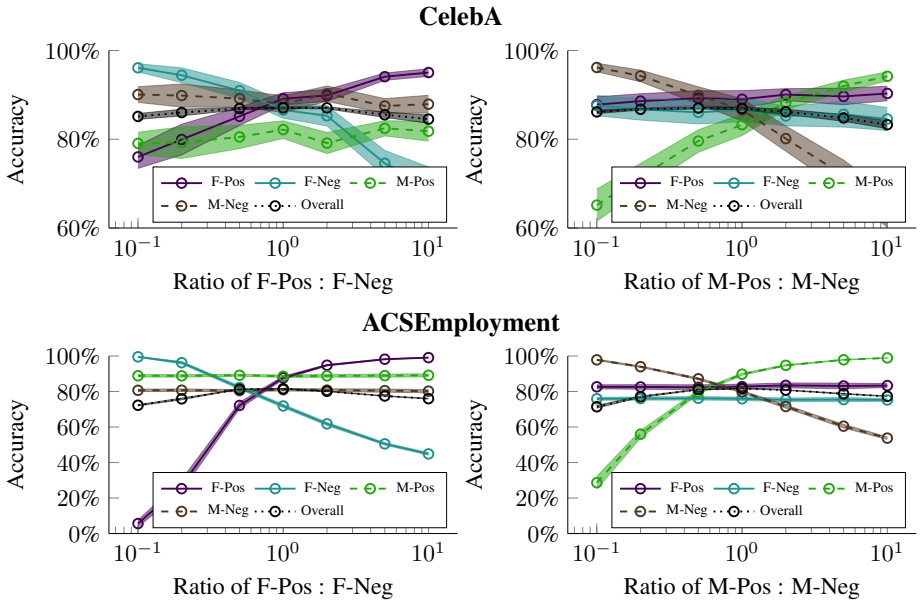

Figure 15: Additional experiments on CelebA and ACSEmployment datasets reveal similar trends as seen in Figure 8. In only a single epoch of training, we are able to manipulate the group level accuracy trade-off, with relatively small impact on overall accuracy.

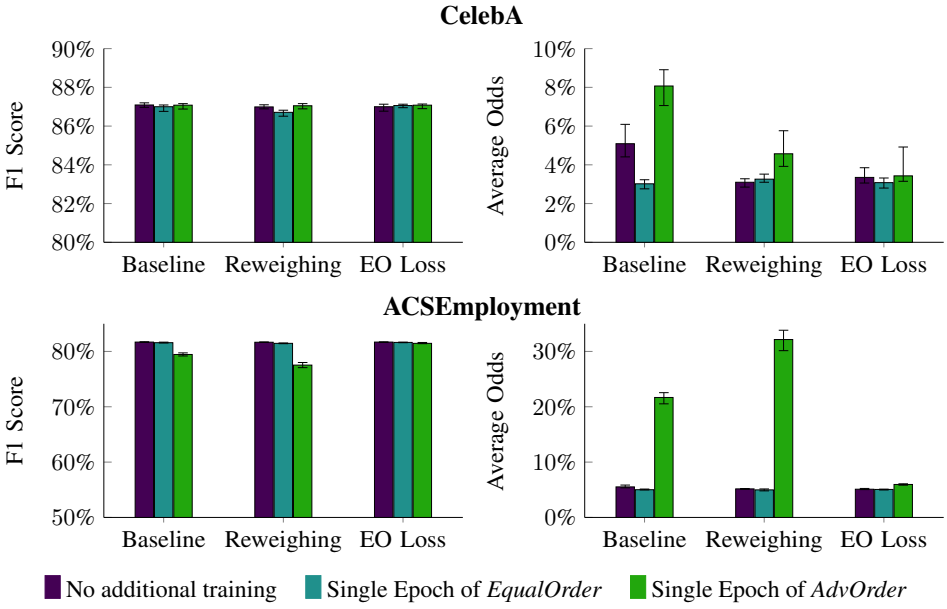

Figure 16: Additional experiments on CelebA and ACSEmployment datasets reveal similar trends as seen in Figure 9. *EqualOrder* gets competitive performance to commonly used bias mitigation methods. Similarly, *AdvOrder* gets significantly worse fairness than even the baseline.