

FAIR MULTI-EXIT FRAMEWORK FOR FACIAL ATTRIBUTE CLASSIFICATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Fairness has become increasingly pivotal in facial recognition. Without bias mitigation, deploying unfair AI would harm the interest of the underprivileged population. In this paper, we observe that though the higher accuracy that features from the deeper layer of a neural networks generally offer, fairness conditions deteriorate as we extract features from deeper layers. This phenomenon motivates us to extend the concept of multi-exit framework. Unlike existing works mainly focusing on accuracy, our multi-exit framework is fairness-oriented, where the internal classifiers are trained to be more accurate and fairer. During inference, any instance with high confidence from an internal classifier is allowed to exit early. Moreover, our framework can be applied to most existing fairness-aware frameworks. Experiment results show that the proposed framework can largely improve the fairness condition over the state-of-the-art in CelebA and UTK Face datasets.

1 INTRODUCTION

Machine learning has been applied in various fields and has impacted our daily life in recent years. Many institutions have introduced machine learning-based systems to help them decide on administrative operations. Although the machine learning model achieves accurate prediction, there exists some bias in such an AI system (Mehrabi et al., 2021; Dressel & Farid, 2018). The discriminative nature of the machine learning model will harm the opportunity of different races, religions, and genders and thus tear society apart.

Several methods are proposed to ameliorate the bias in machine learning models. Many of them (Wang et al., 2022; Zhang et al., 2018; Kim et al., 2019; Ngxande et al., 2020) adopted adversarial training to eliminate bias by training the network to learn a classifier while disabling the adversary’s ability to categorize the sensitive attribute. Disentanglement representation (Creager et al., 2019) is another mainstream to achieve fairness. It forces the latent vector of the sensitive group to be independent of that of the target group and thus reaches fairness.

In this paper, we observe that although features from a deep layer of a neural network bring high accuracy in classification, they cause fairness conditions to deteriorate, and we will demonstrate this observation in Section 2. This finding reminds us of the “overthinking” phenomenon in deep neural networks (Kaya et al., 2019), where accuracy decreases as the features come from deeper in a neural network. This problem is successfully addressed through multi-exit neural networks by introducing multiple internal classifiers and treating high-confidence results from these internal classifiers as the final result. We conjecture that a similar approach can be used to address the issue concerning fairness.

In the proposed multi-exit framework for fairness, both accuracy and fairness constraints are included when training every internal classifier to keep the fairness and accuracy from shallow to deep. Specifically, with early exits at the inference stage, a sufficient discriminative basis can be obtained based on low-level features when classifying easier samples. This contributes to selecting the optimal prediction for each test instance regarding the trade-off between accuracy and fairness.

To validate the effectiveness of our framework, we perform the facial attribute classification on CelebA (Liu et al., 2018) and UTK Face (Zhang et al., 2017) datasets. Experiments show that the proposed method significantly improves the results from the baseline and different state-of-the-

arts in terms of the trade-off between classification accuracy and fairness. Our code is released at <https://anonymous.4open.science/r/ME-Fair-ED26/>.

The main contributions of the proposed method are as follows:

- Extensive experiments show that although the features from a deep layer of a neural network are highly discriminative and thus bring high accuracy, they cause fairness to deteriorate.
- We explore the use of multi-exit training framework to deal with the fairness issue.
- We introduce a simple debias framework with high extensibility that could apply to different baseline and state-of-the-art, and achieve further improvement.
- Through extensive experiments on different settings and comparisons, our framework achieves the best trade-off performances between top-1 accuracy and fairness on CelebA and UTK Face datasets.

2 MOTIVATION

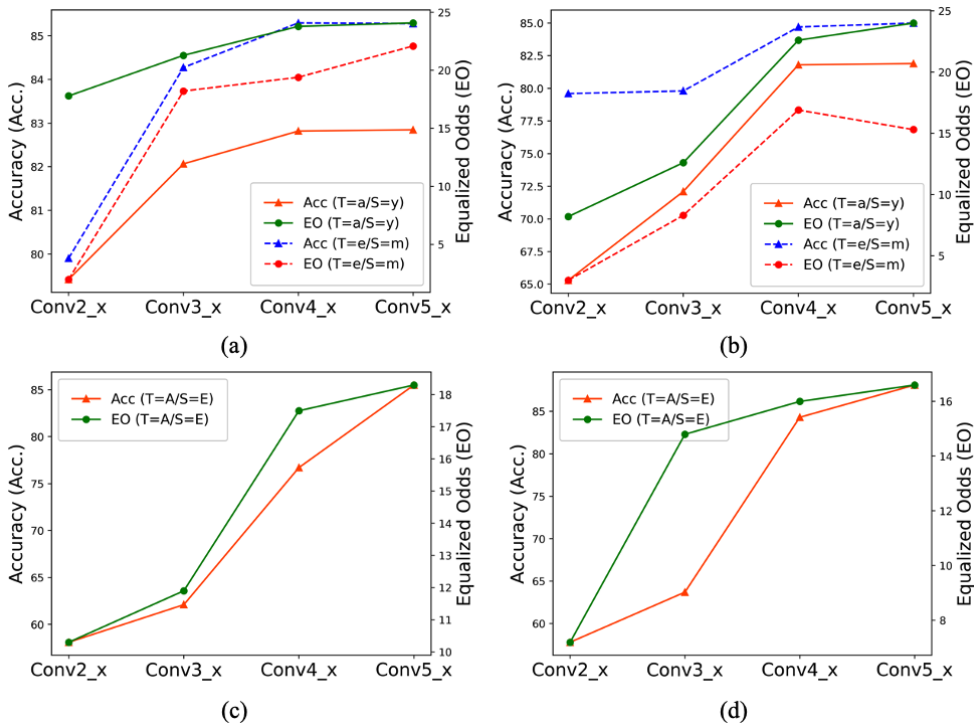


Figure 1: Equalized odds (EO) and accuracy (Acc.) for the internal layers of the conventional ResNet on CelebA (a-b) and UTK Face (c-d) dataset. Note that lower EO represents fairer. (a) and (c) are the results using ResNet-18, while (b) and (d) use the ResNet-50. T and S stand for target and sensitive attributes, respectively. In the CelebA dataset, *a*, *e*, *m*, and *y* represent *attractiveness*, *bags-under-eyes*, *male*, and *young*, respectively. For the UTK Face dataset, *A* and *E* are the *Age* and *Ethnicity*, respectively.

As shown in Fig. 1, we observe that despite the features from a deep layer of a neural network bring high accuracy in classification, they cause fairness condition to deteriorate. We report the equalized odds (EO) and accuracy (Acc.) of ResNet-18 and ResNet-50 on the CelebA and the UTK Face dataset. We first trained a vanilla CNN, and froze the backbone. Then, we trained 3 MLP classifiers with the features from each residual module (i.e., Conv2_x, Conv3_x, Conv4_x, Conv5_x).

For CelebA dataset, in the experiments of the ResNet-18 (Fig. 1(a)) and the ResNet-50 (Fig. 1(b)), both the equalized odds (EO) and accuracy (Acc.) increase when the features are extracted from

deeper layers. When comparing the accuracy between conv2_x and the final layer in ResNet-18, the final layer has improved 6% on average. However, the EO also increases over six times larger on average. As for the ResNet-50, the accuracy increased by about 16%, and the EO increased more than eight times larger on average.

In Fig. 1(c) and (d), results on UTK Face dataset again show a similar trend on the increasing EO and accuracy from shallow to deep layers. The ResNet-18 (Fig. 1(c)) shows a 47% improvement in accuracy with 1.7 times higher EO from the shallow layer (conv2_x) to the deeper layer (Final layer). As for the ResNet-50 (Fig. 1(d)), accuracy increases about 52% and EO increases over 2.3 times.

As higher EO stands for larger bias (unfair) of the predicted result, intuitively, choosing the result at a shallow layer for final prediction could ameliorate the bias condition of the model. Our extensive experiments demonstrate that our observation could be applied to different network architectures and datasets.

3 RELATED WORK

3.1 MULTI-EXIT NETWORKS AND EARLY EXIT POLICY

The multi-exit network is designed by putting additional loss constraints at internal exit branches (internal classifier) to increase the accuracy at shallow layer. With the early exit scheme, early exit branches reduce the computational resource during the model inference time while enhancing the model’s accuracy. In the inference phase, the instances will stop inferencing and leave the model from different branches following the pre-defined early exit rules or criteria. BranchyNet Teerapitayanon et al. (2016) calculates entropy at each branch after obtaining the result and exits if the entropy of the predicted result is less than the threshold value. Shallow-Deep Networks Kaya et al. (2019) calculates the softmax score of each internal classifier’s prediction and takes the maximum probability value as the confidence score. Once the score exceeds the threshold during the forward passing, the instance will exit from the branch prematurely. This model further mitigates the “over-thinking” problem of deep neural networks. Schwartz et al. (2020) leveraged the early exit inference scheme of Kaya et al. (2019) and applied a confidence-based strategy to the natural language processing task. Zhou et al. (2020) makes predictions using adjacent layers and stops inference when the predicted value of the internal classifier remains constant in a given inference unit times. In this paper, we borrow the confidence-based early exit strategy proposed by Kaya et al. (2019) to make sure the the early exit instance is confident enough to be correct. To our best knowledge, we are the first work that leverages the early exit technique to improve fairness.

3.2 BIAS MITIGATION METHODS

Bias mitigation methods are designed to reduce the native bias in the dataset to reduce the chance of unfair prediction. There are largely three avenues for current debiased strategy, including pre-processing, in-processing, and post-processing. (1) Pre-processing strategies usually remove the information which may cause “discrimination” from training data before training. Kamiran & Calders (2012) use different weight to neutralize the effect of the sensitive information in training phase and present the experiments result on real-life data. Ngxande et al. (2020) and Lu et al. (2020) achieve fairness via data pre-process methods including data generation and data augmentation.

In-processing method usually modify on-the-shelf model architecture, loss function, and model regularization to achieve fairness goal. Adversarial training mitigates the bias through adversarially trains an encoder and classifier to learn a fair representation. Zhang et al. (2018) adversarially cooperate a predictor and an adversary to remove the sensitive attributes from the representation. (Kim et al., 2019) eliminated the correlations between extracted feature and sensitive attribute to achieve fairness by unlearn the bias in the feature domain. Wang et al. (2022) adversarially learned a perturb to mask out the sensitive information of the input images, and the proposed framework do not need to alter the deep models’ parameters and structure. Some regularization-based methods, such as Quadrianto et al. (2019) used Hilbert-Schmidt norm to learn a fair representation that retain the features’ semantics from input domain. Jung et al. (2021) learned a fair representation by distilling the fair information of the teacher model to the student model with the Maximum Mean Discrepancy

(MMD) (Gretton et al., 2012) loss. Park et al. (2022) introduced a group-wise normalization and penalizing the inclusion of sensitive attribute to mitigate the intrinsic unbiased condition of supervised contrastive learning.

Post-processing method aims to calibrate the model’s output to enhance fairness. They need to use sensitive attribute and prediction distribution to modify the previous distribution result. Hardt et al. (2016) reveals the limit of demographic parity and give a new metric to fairness, equalized odds, and show how to adjust the learned prediction. Zhao et al. (2017) used Lagrangian relaxation to designed an inference algorithm which reduces bias but maintain accuracy in the meantime.

In this paper, we focus on improving the existed in-processing methods by introducing a general multi-exit (ME) training framework. We compare the regularization-based (Kim et al., 2019; Jung et al., 2021) fairness methods w and w/o the ME framework. Moreover, we show that our framework could apply to complex training structures; for instance, the adversarial debias method Kim et al. (2019) and the fair contrastive learning Park et al. (2022).

4 METHOD

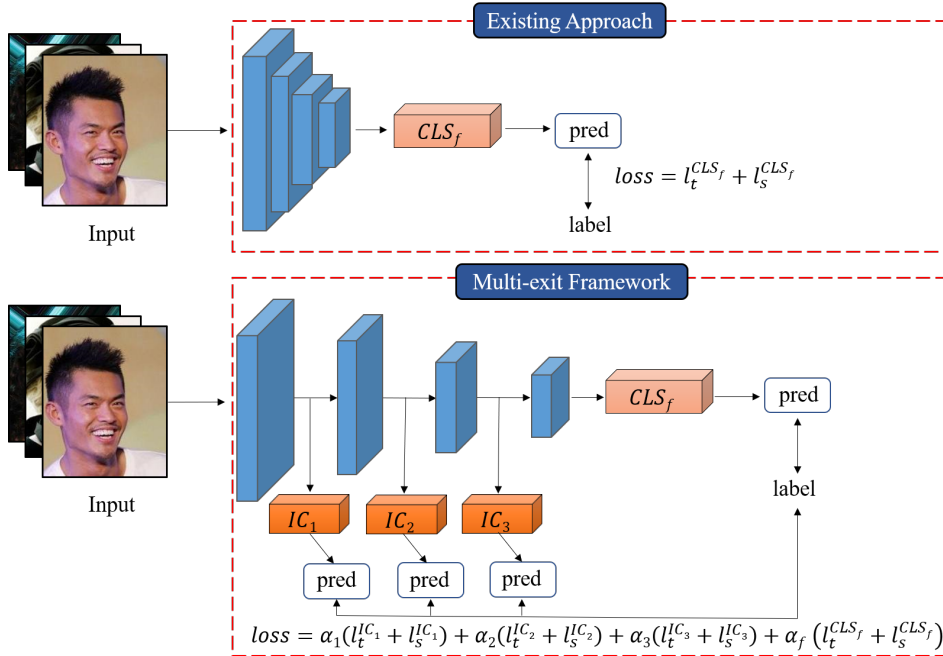


Figure 2: Illustration of the multi-exit training framework. l_t and l_s are the loss function related to target and sensitive attributes, respectively.

In this section, we provide a clear definition of our goal: overcoming the prediction bias of the deep neural network, and we define our problem formulation in 4.1. Afterward, in Section 4.2, we introduce our main approach, multi-exit (ME) training framework and the early exit policy, which allow us to improve the fairness of state-of-the-arts.

4.1 PROBLEM FORMULATION

In the classification task, define input features $X \in x \in \mathbb{R}^d$, target class, $Y \in y = \{1, 2, \dots, N\}$, and sensitive attributes $A \in a = \{1, 2, \dots, M\}$. The goal is to learn a classifier $f : x \rightarrow y$ that predicts the target class Y to achieve high accuracy while being unbiased to the sensitive attributes A . Several criteria are proposed to evaluate the bias against sensitive attributes A , and we will discuss the fairness criteria in our experiments in Section 5.2.

4.2 MULTI-EXIT (ME) TRAINING FRAMEWORK

Fundamental to our approach is that although deep neural networks usually achieve high accuracy in the deeper layer, the prediction would be unfair to the different sensitive groups, e.g., race, gender, age, etc. This phenomenon allows the possibility to select the result at a shallow layer with high confidence to solve the unfair issue and maintain the predicted accuracy. Our method is based on a multi-exit training framework and an early exit policy similar to previous works (Kaya et al., 2019). The main contribution lies in introducing the use of multi-exit to improve the fairness of most state-of-the-arts.

As shown in Fig. 2, existing fairness approaches usually contain two loss term, target classification loss l_t and fairness regularization loss l_s . As most fairness research is done under classification tasks, l_t could be either cross-entropy loss or multi-label soft margin loss, optimizing the training data’s accuracy. As for the fairness regularization loss l_s , it is designed to remove the bias between two sensitive groups, such as Maximum Mean Discrepancy (MMD) (Gretton et al., 2012), and Hilbert-Schmidt Independence Criterion (HSIC) (Gretton et al., 2005).

In our multi-exit training framework, we duplicate the loss function, $loss = (l_t + \lambda l_s)$, used in previous work into every internal classifier (IC), $loss^{IC} = (l_t^{IC} + \lambda l_s^{IC})$, and the final loss is defined by the weighted summation of them, $loss = \alpha_1 \cdot loss^{IC1} + \alpha_2 \cdot loss^{IC2} + \alpha_3 \cdot loss^{IC3} + \alpha_f \cdot loss^{CLS}$, where λ is a hyperparameter that controls the trade-off between fairness and accuracy. As both features at shallow and deep layer are included into the loss function, the model will optimize to increase the accuracy and the fairness from shallow to deep naturally.

In addition, as introduced in Section 1, since we observe that the fairness would drop at the deeper layer, it is recommended to replace the prediction of the final layer with the internal layer. Based on the heuristic that confidence indicates the correctness of a prediction, we preserve both fairness and accuracy during inference by allowing any instance with high confidence from an internal classifier to exit early. We pre-define a confidence threshold θ , and select the result from the earliest internal classifier in which the confidence is above the threshold. This early exit policy successfully selects the optimal prediction in terms of the trade-off between accuracy and fairness.

5 EXPERIMENTAL SETTINGS

5.1 DATASETS

In this work, we evaluate our framework on two facial attribute datasets, CelebA (Liu et al., 2018) and UTK Face. The CelebA dataset consists of over 200k images and for each one, it has 40 binary attributes. Similar to Park et al. (2022), we abbreviate the target attribute (T) and the sensitive attributes (S), *Attractive*, *Big Nose*, *Bags Under Eyes*, *Male*, and *Young* as a , b , e , m , and y , respectively.

UTK Face dataset consists of over 20k face images with 3 annotations, *Ethnicity*, *Age*, and *Gender*. We follow the setting in Jung et al. (2021) to set *Ethnicity* as the sensitive attribute (including *White*, *Black*, *Asian*, and *Indian*.) and divided the *Age* into 3 ranges (ages between 0 to 19, 20 to 40, and larger than 40.) for the target attribute.

To show a fair comparison, we follow the recommended setting in previous works to divide both datasets into training/val/test. Results of the test set are reported and discussed in Section 6.

5.2 EVALUATION METRICS

Several fairness metrics are proposed to evaluate the degree of fairness in classification task. Demographic parity Dwork et al. (2012) and equalized odds (EO) (Hardt et al., 2016; Dwork et al., 2012) are two well-known fairness criterion. We first define an input feature $X \in \mathbb{R}^d$ with sensitive attribute A , ground truth target class Y , and the predicted target class \hat{Y} . Demographic parity is satisfied if

$$P(\hat{Y} = 1|A = 0) = P(\hat{Y} = 1|A = 1). \quad (1)$$

The drawback of demographic parity is that the classifier could achieve the fairness condition by adjusting the proportion of correct rate of two sensitive attributes through misclassifying some in-

stances. On the other hand, EO forces the true positive rate and the false positive rate along different groups to be equal, that is,

$$P(\hat{Y} = 1|A = a, Y = y) = P(\hat{Y} = 1|A = b, Y = y) \quad (2)$$

where $y \in \{0, 1\}$, and $a, b \in A$. This metric addresses unfair wrong prediction of the model, and therefore EO will be the suitable fairness metric to evaluate our model. We calculate the degree of EO by calculating the disparity of the TPR and FPR alone different sensitive attributes as follows:

$$\sum_{k=1}^K |TPR_k^1 - TPR_k^0 + FPR_k^1 - FPR_k^0| \quad (3)$$

where TPR_k^a and FPR_k^a are the True Positive Rate and the False Positive Rate respectively of target class k and sensitive attribute a , this equation can also extend to multi-attribute cases. In conclusion, the optimal EO score becomes 0 when the True Positive Rate and the False Positive of each target and sensitive class are the same. It indicates that a lower EO represents the prediction is fairer.

5.3 IMPLEMENTATION DETAILS

We utilize Resnet-18 as the backbone of every state-of-the-art and the baseline CNN used in our experiments. In the training phase, the data is augmented by random flipping, rotation, and scaling before being fed into the model. The network is trained for 200 epochs using SGD optimizer (Ruder, 2016) with an initial learning rate set as 0.01. We attach three internal classifiers at the end of each residual block. The internal classifier contains one adaptive average pooling layer and a two layer MLP. As introduced in Section 4.2, the modified loss function in the multi-exit training framework is a weighted summation of the loss from each internal classifier. The coefficients, α_1 , α_2 , α_3 , and α_f , are set to 0.3, 0.45, 0.6, and 0.9, respectively. We set the confidence threshold $\theta = 0.85$ for the early exit policy at the inference phase.

To show that our scheme can be widely applied to most existing frameworks, we conduct experiments on four baselines selected from state-of-art approaches for fairness-aware learning: LNL (Kim et al., 2019), HSIC (Quadrianto et al., 2019), MFD (Jung et al., 2021), and FSCL (Park et al., 2022). All the baselines are reproduced by following the recommended hyperparameter settings of the original papers or resources.

6 RESULTS

6.1 COMPARISON WITH STATE-OF-THE-ART


In this section, we show the feasibility of ME training framework on four state-of-the-art, LNL (Kim et al., 2019), HSIC (Quadrianto et al., 2019), FSCL (Park et al., 2022), and MFD (Jung et al., 2021), by comparing the result with and without the ME training framework on CelebA and UTK Face dataset in Table 1 and Fig. 3, respectively.

For the CelebA dataset, in Table 1 we follow the sensitive and target groups setting in Park et al. (2022) and compare our results with their reproduce results accordingly. The ME-CNN is the baseline ME framework training without any fairness constraint. That is, the loss function in each internal classifier is $loss^{IC} = l_t^{IC}$. The ME-CNN framework improves the EO in 20.3% while losing only 0.13% of accuracy in average. In the comparison with different state-of-the-art, our results achieve a 38.5% EO improvement in average while keeping the competitive accuracy. It is noteworthy that in ME-MFD and ME-LNL, our framework also improves the accuracy by an average of 3.8% and 1.3%, respectively.

For UTK Face dataset in Fig. 3, we follow Jung et al. (2021) to define *Ethnicity* as the sensitive attribute and *Age* as the target groups. Compared with the CNN baseline, the EO decreased from 17.8 to 16, which is a 10% improvement. As for the comparison with different state-of-the-art, our results achieve a 11.3% EO improvement in average. In addition, in ME-LNL and ME-MFD, the accuracy also shows a 3% improvement in average. Since the FSCL (Park et al., 2022) selected the *Gender* as the target attribute, we follow their data imbalance setting to product the experiment. Performance comparison between FSCL and ME-FSCL shows that ME-FSCL achieves a 13.5%

Table 1: Classification results of the fairness (EO) and accuracy (Acc.) evaluation on the test set of the CelebA dataset. * denotes our own implementation.

Methods	T=a / S=m		T=a / S=y		T=b / S=m		T=b / S=y		T=e / S=m		T=e / S=y	
	EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.
CNN	27.8	79.6	16.8	79.8	17.6	84.0	14.7	84.5	15.0	83.9	12.7	83.8
ME-CNN	23.7	82.3	16.1	76.8	12.9	84.8	12.9	84.8	10.8	82.5	8.1	83.6
LNL	21.8	79.9	13.7	74.3	10.7	82.3	6.8	82.3	5.0	81.6	3.3	80.3
ME-LNL	14.4	82.2	13.1	72.7	7.3	82.8	5.5	83.1	2.7	84.0	1.0	82.6
HSIC*	19.4	81.7	16.5	80.3	11.2	80.8	10.5	82.6	12.5	84.0	7.4	84.2
ME-HSIC	12.9	78.8	15.9	78.7	3.6	80.9	3.2	82.2	7.8	83.5	4.1	82.0
FSCL+	6.5	79.1	12.4	79.1	4.7	82.9	4.8	84.1	3.0	83.4	1.6	83.5
ME-FSCL+	5.8	78.2	7.6	76.4	3.6	81.2	2.8	84.4	2.0	83.5	1.4	80.8
MFD	7.4	78.0	14.9	80.0	7.3	78.0	5.4	78.0	8.7	79.0	5.2	78.0
ME-MFD	5.8	78.3	11.4	79.5	2.6	82.1	3.3	82.6	1.4	81.9	1.5	84.2

Figure 3: Classification results of the fairness (EO) and accuracy (Acc.) evaluation on the test set of the UTK Face dataset. We follow previous work to define *Ethnicity* as the sensitive attribute and *Age* as the target groups for experiments on CNN, LNL, HSIC, and MFD. Since FSCL selected the *Gender* as the target attribute, we follow their data imbalance setting to product the experiment.

improvement at EO than FSCL, where the accuracy is almost the same. Comparisons in Fig. 3 successfully demonstrate that our framework outperforms all the baseline on the fairness score with a competitive accuracy in UTK Face dataset.

6.2 ABLATION STUDY OF MULTI-EXIT TRAINING FRAMEWORK

In this section, we establish the ablation study of applying a multi-exit training framework to the MFD (Jung et al., 2021). In Fig. 4(a), we visualized the representation of the test instance from each internal classifier. Since the multi-exit training framework optimized each internal classifier (IC_x

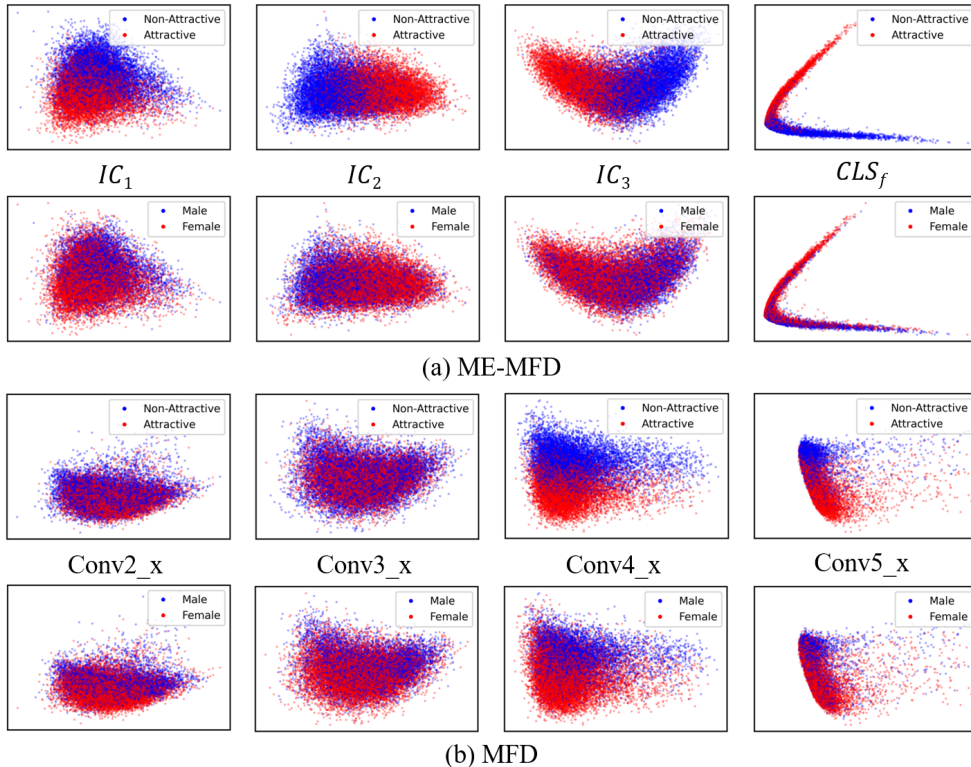


Figure 4: Visualization of the features of the CelebA dataset (T=a / S=m). (a) demonstrates the feature of the proposed multi-exit training framework on MFD, while (b) are the features without using the multi-exit training. IC_x represents the x^{th} internal classifier. CLS_f stands for the classifier following behind the last layer. Since the MFD did not contain internal classifier, we extract the feature from the same layers of ME-MFD, which are Conv2_x, Conv3_x, Conv4_x, and Conv5_x.

and CLS_f) for high accuracy and fairness, we can observe a clear decision boundary between the attractive and non-attractive target class. However, the features of the sensitive attribute (gender) are mixed and uniformly distributed, which indicates that the model remains fair.

In addition, we also visualize the representation of the test instance from the same layer of MFD, which is the output of each residual module of the network in Fig. 4(b). Obviously, since the intermediate feature does not pass through any direct target optimization, both the target and the sensitive attributes could not be easily recognized in shallow layers (Conv2_x and Conv3_x). At the deeper layer, the loss function maximizes the accuracy and minimizes the bias of the sensitive group, the feature of the target group can be clearly separated, whereas the sensitive attribute remains the same. We also report the quantitative results of the experiment mentioned above to show the feasibility of the multi-exit training framework.

In Table 2, we compare the classification results of ME-MFD and MFD selected from each exit branch. We report the features of MFD at different residual modules (i.e., Conv2_x, Conv3_x, Conv4_x, and Conv5_x) and each IC’s features of ME-MFD. The multi-exit framework improves the EO by an average of 27% and the accuracy by 1%. Besides, the EO of MFD increases from Conv2_x to Conv4 but drops at the last layer. This phenomenon is due to the Maximum Mean Discrepancy (MMD) constraint being imposed at the last layer while there is no fairness constraint on the features of shallow layers. ME-MFD regularizes the fairness condition of each IC; as a result, the trend of EO score still holds our observation in Section 2.

Experiments demonstrate the significance of using the multi-exit training framework, which allows us to select prediction at a shallow layer to achieve a fair and high accuracy result.

Table 2: Classification results of the proposed ME-MFD and MFD selected from each exit. IC_x represents the x^{th} internal classifier. CLS_f stands for the classifier following behind the last layer. Since the MFD did not contain internal classifier, we extract the feature from the same layers of ME-MFD, which are Conv2_x, Conv3_x, Conv4_x, and Conv5_x. The fairness (EO) and accuracy (Acc.) evaluation on the test set of the CelebA dataset are reported.

Methods		T=a / S=m		T=a / S=y		T=b / S=m		T=b / S=y		T=e / S=m		T=e / S=y	
		EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.
ME-MFD	IC_1	1.4	70.5	12.9	76.9	1.4	80.1	3.4	78.2	0.8	79.6	1.8	80.0
	IC_2	2.0	75.7	13.4	80.0	1.5	81.7	7.1	78.2	2.4	83.2	3.2	84.0
	IC_3	2.2	77.0	16.4	81.7	4.6	83.7	12.0	78.9	4.4	84.6	3.5	85.1
	CLS_f	13.3	79.9	23.4	83.0	17.7	84.6	18.0	79.2	12.2	85.3	9.6	85.3
MFD	Conv2_x	12.2	68.9	13.3	70.5	4.5	78.5	3.7	78.7	0.1	79.7	1.9	79.4
	Conv3_x	16.0	73.2	15.3	75.3	19.1	80.0	10.9	79.7	8.4	80.8	3.7	80.3
	Conv4_x	23.1	80.5	19.2	79.3	27.6	80.0	16.9	81.4	18.7	83.1	7.8	82.8
	Conv5_x	23.1	81.3	15.6	79.0	12.5	80.7	8.0	83.0	13.1	83.3	6.5	83.1

6.3 ABLATION STUDY OF EARLY EXIT POLICY

In this section, we study the impact of using the proposed early exit policy. Table 3 shows the fairness and accuracy of the proposed ME-MFD selected from a different exit. From the table, we observe that the IC_1 obtains the smallest EO and the accuracy, while the CLS_f achieves the largest. The increasing trend in both metrics indicates that there is a trade-off between fairness and accuracy. Thus, the proposed early exit policy independently selects the exit of each test instance to preserve a large fairness improvement with only a slight drop in accuracy. In the comparison without using the proposed early exit (CLS_f), our method achieves 74.5% lower EO, but the accuracy drop is less than 1.7% on average. The significant improvement demonstrates the importance of deciding the exit for each instance, which also shows that the proposed early exit algorithm is essential.

Table 3: Classification results of the proposed ME-MFD selected from each exit. IC_x represents the x^{th} internal classifier. CLS_f stands for the classifier following behind the last layer. ME-MFD (EE) denotes the result of using the confidence-based early exit algorithm. The fairness (EO) and accuracy (Acc.) evaluation on the test set of the CelebA dataset are reported.

Methods	T=a / S=m		T=a / S=y		T=b / S=m		T=b / S=y		T=e / S=m		T=e / S=y	
	EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.	EO	Acc.
IC_1	1.4	70.5	12.9	76.9	1.4	80.1	3.4	78.2	0.8	79.6	1.8	80.0
IC_2	2.0	75.7	13.4	80.0	1.5	81.7	7.1	78.2	2.4	83.2	3.2	84.0
IC_3	2.2	77.0	16.4	81.7	4.6	83.7	12.0	78.9	4.4	84.6	3.5	85.1
CLS_f	13.3	79.9	23.4	83.0	17.7	84.6	18.0	79.2	12.2	85.3	9.6	85.3
ME-MFD (EE)	5.8	78.3	11.4	79.5	2.6	82.1	3.3	82.6	1.4	81.9	1.5	84.2

7 CONCLUSION

In this paper, we first explored the problem that the fairness condition deteriorates as we classify the features in deeper layers. Then, we introduced the multi-exit training framework, with high extensibility that could be applied to many bias mitigation methods. With the confidence-based exit strategy, we select the optimal exit for each test instance to achieve both high accuracy and fairness. The extensive results and the ablation studies have shown that our framework can achieve the best trade-off of accuracy and fairness conditions compared to the state-of-the-art on two well-known facial datasets.

REFERENCES

- Elliot Creager, David Madras, Jörn-Henrik Jacobsen, Marissa Weis, Kevin Swersky, Toniann Pitassi, and Richard Zemel. Flexibly fair representation learning by disentanglement. In *International conference on machine learning*, pp. 1436–1445. PMLR, 2019.
- Julia Dressel and Hany Farid. The accuracy, fairness, and limits of predicting recidivism. *Science advances*, 4(1):eaao5580, 2018.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pp. 214–226, 2012.
- Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *International conference on algorithmic learning theory*, pp. 63–77. Springer, 2005.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.
- Sangwon Jung, Donggyu Lee, Taeon Park, and Taesup Moon. Fair feature distillation for visual recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12115–12124, 2021.
- Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems*, 33(1):1–33, 2012.
- Yigitcan Kaya, Sanghyun Hong, and Tudor Dumitras. Shallow-deep networks: Understanding and mitigating network overthinking. In *International conference on machine learning*, pp. 3301–3310. PMLR, 2019.
- Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9012–9020, 2019.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Large-scale celebfaces attributes (celeba) dataset. *Retrieved August*, 15(2018):11, 2018.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. Gender bias in neural natural language processing. In *Logic, Language, and Security*, pp. 189–202. Springer, 2020.
- 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.
- Mkhuseli Ngxande, Jules-Raymond Tapamo, and Michael Burke. Bias remediation in driver drowsiness detection systems using generative adversarial networks. *IEEE Access*, 8:55592–55601, 2020.
- Sungho Park, Jewook Lee, Pilhyeon Lee, Sunhee Hwang, Dohyung Kim, and Hyeran Byun. Fair contrastive learning for facial attribute classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10389–10398, 2022.
- Novi Quadrianto, Viktoriia Sharmanska, and Oliver Thomas. Discovering fair representations in the data domain. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8227–8236, 2019.
- Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.

- Roy Schwartz, Gabriel Stanovsky, Swabha Swayamdipta, Jesse Dodge, and Noah A Smith. The right tool for the job: Matching model and instance complexities. *arXiv preprint arXiv:2004.07453*, 2020.
- Surat Teerapittayanon, Bradley McDanel, and Hsiang-Tsung Kung. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 2464–2469. IEEE, 2016.
- Zhibo Wang, Xiaowei Dong, Henry Xue, Zhifei Zhang, Weifeng Chiu, Tao Wei, and Kui Ren. Fairness-aware adversarial perturbation towards bias mitigation for deployed deep models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10379–10388, 2022.
- Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 335–340, 2018.
- Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5810–5818, 2017.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *arXiv preprint arXiv:1707.09457*, 2017.
- Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian McAuley, Ke Xu, and Furu Wei. Bert loses patience: Fast and robust inference with early exit. *Advances in Neural Information Processing Systems*, 33:18330–18341, 2020.