

FairViT: Fair Vision Transformer via Adaptive Masking

Bowei Tian¹, Ruijie Du², and Yanning Shen²

¹ Wuhan University, Wuhan, Hubei 430072, China
boweitian@whu.edu.cn

² University of California, Irvine, CA 92697, USA
{ruijied, yannings}@uci.edu

Abstract. Vision Transformer (ViT) has achieved excellent performance and demonstrated its promising potential in various computer vision tasks. The wide deployment of ViT in real-world tasks requires a thorough understanding of the societal impact of the model. However, most ViT-based works do not take fairness into account and it is unclear whether directly applying CNN-oriented debiased algorithm to ViT is feasible. Moreover, previous works typically sacrifice accuracy for fairness. Therefore, we aim to develop an algorithm that improves accuracy without sacrificing fairness. In this paper, we propose FairViT, a novel accurate and fair ViT framework. To this end, we introduce a novel distance loss and deploy adaptive fairness-aware masks on attention layers updating with model parameters. Experimental results show FairViT can achieve accuracy better than other alternatives, even with competitive computational efficiency. Furthermore, FairViT achieves appreciable fairness results.

Keywords: Vision Transformer · Accuracy · Fairness · Adaptive Masking

1 Introduction

Vision transformer (ViT) [7, 16] has been widely adopted in various computer vision (CV) tasks, and is considered a viable alternative to the Convolutional Neural Network (CNN) [14]. Unlike CNN, ViT has a specialized structure that can extract global relationships via a self-attention mechanism, leading to improved performance in various CV tasks, including image classification [4, 16, 19], object detection [3, 6, 10] and instance segmentation [12, 26, 30]. Due to its excellent performance, the structure has formed the architectural backbone of many CV algorithms for real-world applications. However, wide deployment of CV algorithms highly depends on how trustworthy they are [22, 27]. This prompts an investigation of the fairness aspects on ViT models.

Despite the abundance of debiased algorithms targeting Convolutional Neural Networks (CNNs) [5, 21, 31], there is a lack of literature concerning debiased algorithms on vision transformers. Different from CNNs that capture pixel-wise

local features through convolutions, vision transformers extract global contextual information through image patches. Vision transformers interpolate these patches through the attention mechanism with a stronger shape recognition capacity [33]. It is unclear whether directly applying CNN-oriented debiased algorithm to vision transformers is feasible [5]. Besides, vision transformers are shown to be more robust to input perturbations and latent features than CNNs [7, 25], which may be challenging for a specific fair ViT design.

Fairness in ViT is investigated in several recent works [22, 24, 27], yet a majority of them either sacrifice accuracy for fairness, or require a huge amount of computational cost. TADeT, a targeted alignment technique proposed in [27], seeks to identify and eliminate bias from the query matrix in ViT. Their results demonstrate an effective debiased performance, and it is easy to implement in real scenarios. Whereas their directly manipulating parameters in the model will sacrifice accuracy for fairness. A bilevel optimization is designed in [24], which finds the optimal data sampling ratios between real and generated data and leads to an improving tradeoff between fairness and accuracy, yet this method require relatively high computing power. Debiased Self-Attention (DSA) proposed in [22] is a fairness-target approach that enforces ViT to eliminate spurious features correlated with the sensitive label. DSA uses adversarial machine learning to enhance fairness-accuracy balance. However, it requires costly two-stage training, which is hard to deploy in real scenarios.

To address the aforementioned challenges, we propose FairViT, including adaptive masking and distance loss, are innovative and effective frameworks aimed at addressing fairness and accuracy concerns. Rather than deploying a high computing mechanism, the distance loss is a regularizer that is convenient to deploy, and the adaptive masking is convenient to calculate as well. With the assistance of adaptive masking, the model can reach a better performance on fairness and accuracy matrices. At the same time, distance loss is an extendable, convenient approach that can be used in other applications, not limited in ViT. The code is available at <https://github.com/abdd68/Fair-Vision-Transformer>.

Our main contributions can be summarized as follows:

- We introduce an adaptive masking framework wherein group-specific masks and weights are learned to enhance fairness. We equip the adaptive masking with a backward algorithm that optimizes the masks and weights.
- We incorporate an extendable distance loss function manipulating the output scores to augment accuracy.
- We conduct extensive experiments on real datasets and demonstrate FairViT achieves accuracy better than alternatives, even with competitive computational efficiency. Furthermore, FairViT attains appreciable fairness results.

2 Related Work

2.1 Vision Transformer

The transformer architecture was initially designed for natural language processing (NLP) [29] tasks. Unlike convolutional neural networks, the transformer

network relies on the attention mechanism to process sequences of input tokens in parallel.

Recently, the transformer architecture has been adapted to computer vision tasks [7], utilizing the self-attention mechanism to model relationships between different parts of an image. ViT’s advantages include flexibility in handling various resolutions, capturing global information, parameter efficiency, and potential for better generalization. In many scenarios, ViT outperforms CNN and achieves considerable robustness [23].

In ViT, sample \mathbf{x} contains p input image patches. ViT first employs an embedding layer to each patch to convert it into a embedding vector. Subsequently, ViT applies a series of transformer encoder layers to the embeddings, and each encoder layer consists of two parts: a Multi-Head Attention mechanism (MHA) and a position-wise FeedForward Network (FFN). The MHA layer models the interactions between the patch embeddings using self-attention, while the FFN layer implements a non-linear transformation on each patch embedding individually. The self-attention mechanism [7] can be illustrated as:

$$\text{Attn}(\mathbf{x}) = \text{S} \left(\frac{QK^T}{\sqrt{d}} \right) V, \quad (1)$$

where Q , K , and V are the query, key, and value matrices, respectively, $\text{S}(\cdot)$ is the softmax function, d is the dimension of the key vectors. Self-attention is an important building block for transformers and raises a huge amount of interest in the CV domain, since it is shown that the reliance on CNNs is not necessary and a pure transformer applied directly to sequences of image patches can perform very well on image classification tasks [7]. There are abundant works that aim to explore the transformers’ attention mechanism, such as Gradient Attention Rollout [1] in explainability and Swin Transformer [16] in efficiency.

2.2 Fairness in Neural Networks

Most of the existing debiasing methods for image classification tasks are specified for CNN or deep neural network (DNN) models [21, 31], and can not be directly applied to ViTs. However, several studies show that CV models make predictions by mixing sensitive features with input features [21, 35], and the sensitive features may capture biased relationships between the input features and the target labels. For example, the sensitive feature “gender” usually influences the accuracy of a face recognition task. In this case, it may lead to discriminatory results towards underrepresented groups, which causes serious social and ethical problems.

Fairness of ViT has been investigated in several recent works [22, 24, 27]. A targeted alignment technique TADeT was proposed in [27], which seeks to identify and eliminate bias from the query matrix in ViT. However, their directly manipulating Q sacrifices accuracy for fairness. Dr-Fairness [24] proposes a bilevel optimization that finds the optimal data sampling ratios between real and generated data, and it leads to an improving tradeoff between fairness and accuracy, yet they require relatively high computing power. Debaised Self-Attention

(DSA) [22] is a fairness-target approach that enforces ViT to eliminate spurious features correlated with the sensitive label, and DSA uses adversarial machine learning to enhance fairness-accuracy balance.

In this paper, we present a novel fair and accurate training framework designed for vision transformers. FairViT outperforms existing works by demonstrating superior accuracy and appreciable fairness. Furthermore, our time cost experiment and multi-task testing show that FairViT is applicable in real deployments, and maintains reasonable computational efficiency.

3 Problem formulation

We formulate the fairness-accuracy issue as a supervised classification problem, where the goal is to train a model f using training samples $\{\mathbf{x}, s, y\}$ and learn patterns from the data in order to make predictions, where \mathbf{x} is the input feature, y is the target label, and s is a sensitive label. Let y belongs to the space \mathbf{Y} , and s belongs to the space \mathbf{S} , some examples of \mathbf{S} include gender, race, or other attributes that can determine a sensitive group. We assume that s can only be accessed in the training phase, and are not accessible in the validation or testing phase. The classification framework in training embodies the following form:

$$\min_{\theta} L(f(\mathbf{x}; \theta), s, y) \quad (2)$$

where $f(\mathbf{x}; \theta)$ is the learned model parameterized by θ , L is the loss function characterizing the discrepancy between the estimated label and the target label. One common selection for L is the cross-entropy loss [18]. However, the cross-entropy loss does not take s into account [8]. Therefore, our objective is to devise a novel framework $f(\mathbf{x}; \theta)$ to alleviate bias, and we use s in adaptive masking during the back propagation. During the validation and testing stage, since the sensitive attribute s is not available, the model treats s as \emptyset and computes the weighted sum within the adaptive masking.

4 Fairness-aware Vision Transformer Design

Our design comprises two pivotal parts, i.e. the adaptive masking and distance loss. First, we introduce adaptive masking, which is an assistance of the attention mechanism, concentrating on manipulating the model structure to enhance accuracy and maintain fairness. We optimize the adaptive masking by updating the masks and weights iteratively. Then, the distance loss is introduced to further enhance the accuracy. Figure 1 illustrates the overall procedure of FairViT. Algorithm 1 outlines the entire fairness-aware process of FairViT.

4.1 Adaptive Masking

Forward Propagation Design Existing ViTs have demonstrated remarkable capabilities in various image recognition tasks. However, ViTs rely heavily on

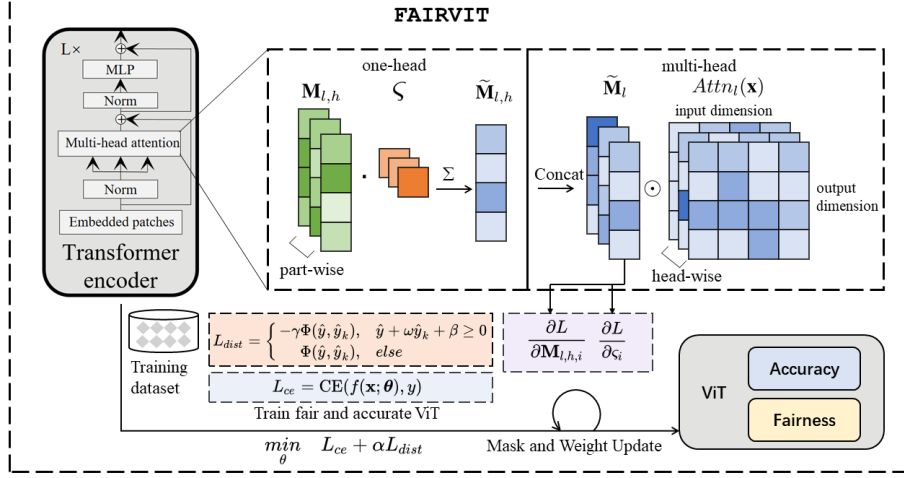


Fig. 1: An illustration of FairViT . For the forward propagation, we first apply weight ς to $\mathbf{M}_{l,h}$, calculate the weighted sum $\tilde{\mathbf{M}}_{l,h}$, which is utilized to assist attention mechanism to control the information flow. For the backward propagation, we optimize $\mathbf{M}_{l,h,i}$ and ς_i . Additionally, we introduce a novel distance loss L_{dist} .

vast datasets for training, and its attention mechanism extracts information from every processed patches of an image, potentially perpetuating biases inherent in those datasets. Consequently, the accuracy of ViTs might vary across different groups, leading to disparities in performance. Therefore, we are seeking a solution to address this issue. Motivated by multi-channel convolution [28], where the convolution kernel channels align with the input channels, we seek to integrate analogous concepts into the ViT structure. This integration aims to enhance accuracy while upholding fairness. Our approach, named adaptive masking, first splits the training dataset into G distinct parts, where each sensitive group has $\lfloor G/2 \rfloor$ parts. For each sensitive group, every part contains the same number of images. Since the number of samples in each sensitive group is different, the number of images in each part does not have to be equal between different sensitive groups. The splitting process is illustrated in Figure 2. Subsequently, we associate each part with a corresponding mask and weight. Each part i has a corresponding mask $\mathbf{M}_{l,h,i}$ and weight ς_i as parameters. We then introduce one-Head Attention (HA) as the example:

$$\text{Attn}_{l,h}(\mathbf{x}) = \text{S} \left(\frac{QK^T}{\sqrt{d}} \right) V \quad (3)$$

$$\tilde{\mathbf{M}}_{l,h} = \sum_{i=1}^G (\varsigma_i \mathbf{M}_{l,h,i}) \quad (4)$$

Algorithm 1 Pseudo-code of FairViT

Require: Transformer model parameters θ , training data set T_t , validation data set T_v , threshold t , epoch E and learning rate lr .

Ensure: The accurate and fair model parameters θ^* .

- 1: Initialize $\mathbf{M}_{l,h,i} = \mathbf{0}$, $L = INF$, $h = 0$
- 2: **while** $h < E$ **and** $L > t$ **do**
- 3: // Training Stage
- 4: **for all** $\mathbf{x} \in T_t$ **do**
- 5: **if** $h = 0$ **then**
- 6: $L = L_{ce}$.
- 7: Obtain $\frac{\partial L}{\partial \theta}$ and back-propagate.
- 8: **else**
- 9: $L = L_{ce} + \alpha \cdot L_{dist}$.
- 10: Obtain $\frac{\partial L}{\partial \theta}$, $\frac{\partial L}{\partial \mathbf{M}_{l,h,i}}$, $\frac{\partial L}{\partial \varsigma_i}$ by Equation (7-8) and back-propagate.
- 11: **end if**
- 12: **end for**
- 13: // Validation stage
- 14: **for all** $\mathbf{x} \in T_v$ **do**
- 15: $\hat{y} = f_y(\mathbf{x}; \theta)$.
- 16: $\hat{y}_k = \sum_{i \in \{topk\}/\{y\}} f_i(\mathbf{x}; \theta)$.
- 17: **end for**
- 18: Update w and β in Equation (9) by classifying $(\hat{y}, \hat{y}_k) \rightarrow z$.
- 19: $h = h + 1$.
- 20: **end while**
- 21: $\theta^* = \theta$.
- 22: **return** θ^* .

$$\text{HA}(\mathbf{x}, \mathbf{M}_{l,h}) = \widetilde{\mathbf{M}}_{l,h} \odot \text{Attn}_{l,h}(\mathbf{x}) \quad (5)$$

where \odot is the element-wise product, $\mathbf{M}_{l,h,i}$ represents the i_{th} mask ($i \in \{1, \dots, G\}$) within layer l and head h , ς_i is the weight of $\mathbf{M}_{l,h,i}$, and $\widetilde{\mathbf{M}}_{l,h}$ is the weighted sum of $\mathbf{M}_{l,h,i}$. We follow the Multi-Head Attention (MHA) as [7], formulated as:

$$\text{MHA}(\mathbf{x}, \mathbf{M}_l) = \Lambda(\text{HA}(\mathbf{x}, \mathbf{M}_{l,1}), \dots, \text{HA}(\mathbf{x}, \mathbf{M}_{l,H})) \quad (6)$$

where Λ denotes the concatenate operation, $\text{HA}(\mathbf{x}, \mathbf{M}_{l,h}) \in \mathbb{R}^{p \times d}$, $\text{MHA}(\mathbf{x}, \mathbf{M}_l) \in \mathbb{R}^{p \times (Hd)}$, and H is the number of heads in a transformer encoder layer. Adaptive masking can regulate information flow during forward propagation: if a particular group exhibits lower accuracy, the model trainer can potentially adapt by adjusting the weight assigned to this group. This adjustment framework establishes a criterion wherein adequate information is acquired for effective group classification, while simultaneously ensuring a fair balance of attention across the groups. Our subsequent objective is to guarantee that each mask and weight maintain an applicable distribution to attain the global optimum. Nonetheless, we observe that static values of ς demonstrate a sub-optimal performance in




Fig. 2: The split of the dataset in our design. Each part only contains samples from one sensitive group, and each part in one sensitive group contains the same number of images, but the number of images in one part between different sensitive groups does not have to be equal.

specific scenarios, as shown in Table 2. This observation motivates us to develop a gradient-based method that automatically optimizes the mask and weight.

Updating the Adaptive Masking Instead of manually setting static masks and weights, we propose to update them iteratively during training. Specifically, given a sample belonging to part $g \in \{1, \dots, G\}$, the $\mathbf{M}_{l,h}$ and ς are updated through gradient descent. The gradient of $\mathbf{M}_{l,h,i}$ can be obtained as:

$$\frac{\partial L}{\partial \mathbf{M}_{l,h,i}} = \begin{cases} \frac{\partial L}{\partial \text{HA}} \text{Attn}_{l,h}(\mathbf{x}) \cdot \varsigma_i, & \text{if } i = g \\ \mathbf{0}, & \text{otherwise,} \end{cases} \quad (7)$$

To update ς_i , we first obtain the computing map of ς_i towards $\widetilde{\mathbf{M}}_{l,h}$. Based on Equation 4, the gradient of ς_i can be obtained as

$$\frac{\partial L}{\partial \varsigma_i} = \begin{cases} \sum_p \sum_d \left(\frac{\partial L}{\partial \text{HA}} \text{Attn}_{l,h}(\mathbf{x}) \cdot \left(\sum_d \mathbf{M}_{l,h,i} \right) \right), & \text{if } i = g \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

An illustration of this process is shown as Figure 3. Our method maintains reasonable computational efficiency, as $\frac{\partial L}{\partial \text{HA}}$ has been calculated during the backward propagation, requiring limited matrix multiplications to compute the gradients of $\mathbf{M}_{l,h,i}$ and ς_i . Moreover, the experimental results in Table 2 show that our iteratively updating the masks significantly improves the accuracy and upholds the fairness of the model compared with static masks. This improvement can be attributed to the parameterization of the $\mathbf{M}_{l,h}$ and ς , i.e. they can be viewed as trainable parameters of the model, creating generalization to the sensitive groups.

4.2 Distance Loss

Cross-entropy loss uses the sign function that activates the output from the target label, deactivates the outputs from other labels. Solely minimizing cross-entropy loss actually omits the information from other labels. Therefore, we design the distance loss, considering not only maximizing the target label’s score




Fig. 3: An illustration of the update process. We ascertain the specific part i to which the training sample belongs, and ∇ refers to the gradient calculation, specified in Equation (7-8). The gray blocks signify that the gradients are zero during the backward pass of this training sample.

but also minimizing other labels' scores. Following some regularization techniques [8, 34], we formulate a regularizer that improves accuracy. During the validation phase, we use logistic regression, a binary classifier to extract a hyperplane that underlines data distribution. Subsequently, in the training stage, we leverage this hyperplane to guide the processing of each sample. This innovative strategy allows us to align the training process more closely with the actual distribution of the data, enhancing adaptability and performance.

In details, we define $\hat{y} = f_y(\mathbf{x}; \boldsymbol{\theta})$ as the predicted score corresponding to the target label y . Additionally, we denote $\hat{y}_k = \sum_{i \in \{\text{top}k\} / \{y\}} f_i(\mathbf{x}; \boldsymbol{\theta})$ as the cumulative score derived from labels in the top k set, excluding the target label. In the validation stage, we train a linear classifier utilizing logistic regression, i.e. $\zeta(\hat{y}, \hat{y}_k) = \mathcal{S}(\hat{y} + \omega \hat{y}_k + \beta)$, where $\mathcal{S}(x) = \frac{1}{1+e^{-x}}$ is the sigmoid function, ω and β are trainable parameters. The samples are labeled by $z = \mathbb{1}(\hat{y} = \max_i(f_i(\mathbf{x}; \boldsymbol{\theta}))) \in \{0, 1\}$, indicating if the sample is classified correctly. The decision boundary of this linear classifier can be shown as

$$\hat{y} + \omega \hat{y}_k + \beta = 0. \quad (9)$$

Since ω and β are updated during the validation stage, they remain constant in the training stage. Then we introduce the distance term as [2] for the training stage

$$\Phi(\hat{y}, \hat{y}_k) = \frac{|\hat{y} + \omega \hat{y}_k + \beta|}{\sqrt{1 + \omega^2}} \quad (10)$$

which measures the distance between point (\hat{y}, \hat{y}_k) and the hyperplane in Equation (9). We train ζ in the validation stage to obtain ω and β , which remain constant in the next training stage. These fixed values facilitate the computation of our distance loss. To elaborate, the distance loss is as follows:

$$L_{dist} = \begin{cases} -\gamma\Phi(\hat{y}, \hat{y}_k) & \text{if } \hat{y} + \omega\hat{y}_k + \beta \geq 0, \\ \Phi(\hat{y}, \hat{y}_k) & \text{otherwise,} \end{cases} \quad (11)$$

where γ is a non-negative hyperparameter. Minimizing L_{dist} makes the points (\hat{y}, \hat{y}_k) that $\hat{y} + \omega\hat{y}_k + \beta \geq 0$ retain, and the points that $\hat{y} + \omega\hat{y}_k + \beta < 0$ are moved to the decision boundary, as our objective is to shift all the points to $\hat{y} + \omega\hat{y}_k + \beta \geq 0$. The overall loss function consists of two parts, shown as

$$L = L_{ce} + \alpha L_{dist}, \quad (12)$$

where $L_{ce} = \text{CE}(f(\mathbf{x}; \boldsymbol{\theta}), y)$, and it guides the initial training phase when the hyperplane lacks meaningful definition. During the first epoch, $\boldsymbol{\theta}$ undergoes training solely based on L_{ce} . In each subsequent of the epochs, we update and refine the hyperplane, training the model with $L = L_{ce} + \alpha L_{dist}$.

Introducing the distance loss contributes to accuracy enhancement. Recall the definitions of \hat{y} and \hat{y}_k , since the model selects the highest score’s label as the predicted label, a higher \hat{y} and a lower \hat{y}_k suggest a greater likelihood of correct classification. Because we determine ω by the logistic regression on z , ω should assume a negative value, resulting in a positive slope for the hyperplane in Equation (9). Equation (10) is the distance between (\hat{y}, \hat{y}_k) and the hyperplane. In Equation (11), L_{dist} encourages both points that satisfy $\hat{y} + \omega\hat{y}_k + \beta < 0$ and $\hat{y} + \omega\hat{y}_k + \beta \geq 0$ to move towards $\hat{y} + \omega\hat{y}_k + \beta \geq 0$. In both scenarios, the distance loss encourages an increase in \hat{y} and a decrease in \hat{y}_k , ultimately improving accuracy. Furthermore, we can extend the distance loss directly to other models, such as deep neural networks (DNNs) and convolutional neural networks (CNNs), as it solely requires the information from the model’s output.

5 Experiments

5.1 Experimental Setup

We conduct experiments in three distinct scenarios on the CelebA dataset [17], a large-scale dataset of facial attributes, containing over 200,000 images of celebrities’ faces. We attached our code in the supplementary materials and we will open-source it upon publication. In standard deployments, we initialize ς as **2** and clamp the weight to the range $(\epsilon, 4 - \epsilon)$, where ϵ is a small value, set as $1e - 8$ in our experiment. We initialize $\mathbf{M}_{l,h}$ with $\mathbf{1}/2\mathbf{G}$ and clamp it to the range $(-1, 1)$, therefore $\widetilde{\mathbf{M}}_{l,h}$ is initialized with **1**. We set the training-validation split ratio as 0.9 : 0.1, $\gamma = 0.5$, $G = 10$ and $\alpha = 0.01$ in our experiments. For measuring fairness, the evaluation metric is as follows:

Balanced Accuracy (BA) [20] measures the performance of a classification model, particularly when dealing with imbalanced datasets. Specifically, The formulation is shown as

$$\text{BA} = \frac{1}{4}(\text{TPR}_{s=0} + \text{TNR}_{s=0} + \text{TPR}_{s=1} + \text{TNR}_{s=1}). \quad (13)$$

It takes into account the imbalance in the dataset by calculating the average accuracy of each sensitive group and the target label.

Demographic Parity (DP) [9] measures how algorithms make predictions or decisions fairly among different demographic groups, or how much algorithms introduce biases or unfairness based on individual characteristics such as race, gender, and age, particularly between the sensitive groups ($s = 0$ and $s = 1$). Formally,

$$DP = |P(\hat{y} = 1|s = 1) - P(\hat{y} = 1|s = 0)|, \quad (14)$$

where P denotes the probability calculated over the test set. A smaller DP typically means fewer differences between various groups in the outcomes of the algorithm.

Equalized Opportunity (EO) [13] is a simple and interpretable notion of nondiscrimination with respect to a specified protected attribute. Specifically,

$$EO = |P(\hat{y} = 1|s = 1, y = 1) - P(\hat{y} = 1|s = 0, y = 1)|. \quad (15)$$

As two of the most popular fairness matrices, DP focuses on the probability of being assigned to *positive prediction* among different sensitive groups; EO focuses on the *true positive rate* among different sensitive groups.

5.2 Comparison with Baselines

We select five SOTA fairness-aware baselines to compare with our work, i.e., Vanilla [7], TADeT-MMD [27], TADeT [27], FSCL [21] and FSCL+ [21]. As the source code of DSA [22] is not publicly available, we did not include DSA in our experiment. We implement Vanilla, FSCL and FSCL+ based on their published source codes, TADeT-MMD, and TADeT using illustrations in the paper. Table 1 demonstrates that FairViT exhibits superior fairness performance alongside appreciable accuracy. In comparison to FSCL+, which is our main competitor, FairViT achieves a significantly higher accuracy of at least 4.5%. In terms of fairness metrics, FairViT showcases excellent unbiased effects.

Table 1: The performance of image classification on CelebA dataset [17] with Vanilla [7], TADeT-MMD [27], TADeT [27], FSCL [21], FSCL+ [21] and FairViT. Shown is the mean of 3 independent runs. Highlighted is the best result.

method	Y: Attraction, S: Gender				Y: Expression, S: Gender				Y: Attraction, S: Hair color			
	ACC%	BA%	EO _{e-2}	DP _{e-1}	ACC%	BA%	EO _{e-2}	DP _{e-1}	ACC%	BA%	EO _{e-2}	DP _{e-1}
Vanilla	74.01	72.36	14.43	3.245	88.42	88.85	4.91	1.489	76.48	74.55	3.61	1.896
TADeT-MMD	79.89	73.85	7.10	3.693	92.51	93.03	2.48	1.290	77.97	75.64	2.27	1.491
TADeT	78.73	74.52	3.11	3.116	90.05	90.68	4.86	1.443	78.49	77.42	3.78	1.057
FSCL	79.09	74.76	1.78	3.004	89.37	90.08	1.76	1.344	78.85	78.06	2.65	0.989
FSCL+	77.26	73.42	0.79	2.604	88.83	89.02	1.20	1.263	78.02	77.37	1.79	0.834
FairViT	84.01	79.96	1.15	2.837	94.27	94.12	1.52	1.205	82.52	81.56	2.10	0.701

5.3 Ablation Study

Method ablation study: We conducted an ablation study to evaluate the effectiveness of adaptive masking and the distance loss. The results are presented in Table 2. Here, Θ denotes the adaptive masking method without updating masks and weights, while $\Delta\Theta$ signifies the deployment of adaptive masking with updating masks and weights. The experiment results show that both L_{dist} and $\Delta\Theta$ contribute to accuracy improvement, with $\Delta\Theta$ playing a much more crucial role in enhancing accuracy and fairness. Continuously updating adaptive masks performs much better in both accuracy and fairness than static masks.

Table 2: Ablation study of FairViT. Shown is the mean of 3 independent runs. Highlighted is the best result.

method	Y: Attraction, S: Gender				Y: Expression, S: Gender				Y: Attraction, S: Hair color			
	ACC%	BA%	EO _{e-2}	DP _{e-1}	ACC%	BA%	EO _{e-2}	DP _{e-1}	ACC%	BA%	EO _{e-2}	DP _{e-1}
L_{ce}	74.01	72.36	14.43	3.245	88.42	88.85	4.91	1.489	76.48	74.55	3.61	1.886
$L_{ce} + L_{dist}$	77.01	72.68	12.54	3.166	89.49	89.98	3.85	1.426	77.08	74.89	2.10	1.741
$L_{ce} + L_{dist} + \Theta$	79.96	74.06	8.99	3.929	92.04	92.77	6.86	1.484	80.10	78.79	7.98	2.051
$L_{ce} + L_{dist} + \Delta\Theta$	84.01	79.96	1.15	2.837	94.27	94.12	1.52	1.268	82.52	81.56	2.10	0.701

Impact of α in the loss function: In Table 3, we observe that as α increases, the accuracy gradually improves while EO and BA are not hurt at all. we can empirically observe that α is even beneficial to fairness at around 0.1. When α reaches a certain threshold (e.g. between 0.1 and 1 in the second case from Table 3), the model experiences a decrease in both accuracy and fairness. This phenomenon could be attributed to the model’s need to strike a balance between the distance loss and the cross-entropy loss. An inappropriate α might disrupt the optimization process, keeping the model from focusing on the optimization objective. For subsequent experiments, we set $\alpha = 0.01$.

Impact of γ in Distance Loss: In Table 4, we note that as the values of γ increase, both EO and BA typically exhibit an initial rise until reaching 0.5, after which they decline. Meanwhile, accuracy follows an increasing trend till 0.5, then it begins to decrease. We can observe γ positively influences the impact of the distance loss, benefiting accuracy. However, excessively high values of γ might create an imbalance between the two losses, potentially leading to decreased accuracy. Given these observations, we use $\gamma = 0.5$ as it represents an optimal balance between fairness and accuracy based on our analysis.


Impact of G in the adaptive masking: In Figure 4, we demonstrate the impact of G , the number of distinct parts, in the adaptive masking. When G is small, the accuracy and fairness matrices do not bring huge benefits, however, as G reaches a threshold, the accuracy and fairness tend to be stabilized and perform well. A possible explanation is that G has a positive correlation with the adjustment capabilities of the model, as the model can consider more different parts at the same time and judge them in a more personalized way. However,

Table 3: Impact of α , the weight of distance loss. Shown is the mean of 3 independent runs. Highlighted is the best result.

α	Y: Attraction, S: Gender				Y: Expression, S: Gender				Y: Attraction, S: Hair color			
	ACC%	BA%	EO _{e-2}	DP _{e-1}	ACC%	BA%	EO _{e-2}	DP _{e-1}	ACC%	BA%	EO _{e-2}	DP _{e-1}
0	81.08	77.19	4.89	3.027	91.49	89.67	3.43	1.373	80.25	77.97	5.25	0.788
0.001	81.15	79.19	4.03	2.989	92.18	89.81	3.60	1.320	82.25	80.07	3.92	1.009
0.01	82.51	79.41	4.48	3.337	93.00	92.87	2.40	1.424	82.49	79.62	5.08	0.942
0.1	81.73	76.51	3.33	3.156	93.54	93.70	1.23	1.204	80.35	77.43	4.24	0.589
1	75.53	70.38	5.92	3.597	82.63	83.80	3.34	1.228	75.18	73.49	5.87	1.053

Table 4: Impact of γ . Shown is the mean of 3 independent runs. Highlighted is the best result.

γ	Y: Attraction, S: Gender				Y: Expression, S: Gender				Y: Attraction, S: Hair color			
	ACC%	BA%	EO _{e-2}	DP _{e-1}	ACC%	BA%	EO _{e-2}	DP _{e-1}	ACC%	BA%	EO _{e-2}	DP _{e-1}
0.1	81.65	77.82	5.11	3.386	92.17	91.31	2.94	1.556	81.59	78.83	3.67	0.807
0.3	82.89	78.96	2.29	3.205	92.58	91.53	1.75	1.296	82.17	80.21	3.51	0.742
0.5	82.75	79.12	3.06	2.685	93.89	94.12	1.52	1.205	82.51	80.91	3.07	0.535
0.7	82.46	77.23	3.22	2.909	93.05	93.86	1.57	1.016	80.67	77.69	4.59	0.627
0.9	81.67	75.85	3.31	3.237	92.25	92.73	2.02	1.338	80.83	76.84	4.87	0.939

**Fig. 4:** Impact of G . Shown is the mean \pm standard deviation of 3 independent runs.




Fig. 5: The interpretability study of FairViT .

an exceeding G may result in too few images in one part, where each part will not have enough training data to obtain adequate performance, resulting in a slight decrease in performance. Therefore, the best choice of G may vary across different problem scenarios and datasets.

5.4 Interpretability Study

We conduct an interpretability study on FairViT to elucidate the reasons behind its superior performance across diverse scenarios, shown in Figure 5. We use Gradient Attention Rollout (GAR) [1] to generate heat maps that accentuate crucial decision-making zones in ViT, and details about GAR are in Appendix A. We observe there are noteworthy distinctions of focused areas between Vanilla and FairViT. The Vanilla method appears to capture information relevant to sensitive attributes like Gender in the first scenario and Hair color in the second scenario. In contrast, FairViT, leveraging adaptive masking, demonstrates a tendency to extract information relevant to the target attributes, such as Expression in the first scenario and Attraction in the second scenario. This phenomenon illustrates the effectiveness of FairViT in fairness and accuracy. Furthermore, FairViT generates heat maps that are distributed more distinctly and densely in space, potentially indicating enhanced model learning.

5.5 Time Cost

To evaluate the efficiency of FairViT , we conduct a comparative analysis of computational costs between FairViT and the baselines, as illustrated in Table 5.

Our findings reveal that FairViT exhibits comparable computational costs to the baselines. Moreover, FairViT attains a superior balance between accuracy and fairness while maintaining reasonable computational efficiency. Compared with FSCL+, FairViT is 6 times faster while achieving better accuracy and competitive fairness results. The core incremental consumption in FairViT is the adaptive masking, requiring $O(p * d)$, where p is the patch number and d is the dimension of the key. The time complexity is better than FSCL, which requires cubic time complexity.

Table 5: The time cost of FairViT compared with other baselines. Shown is the mean of 3 independent runs. Highlighted is the best result.

method	AG ³ _{min}	EG ⁴ _{min}	AH ⁵ _{min}
Vallina	4.62	4.38	4.51
TADeT-MMD	4.68	4.50	4.61
TADeT	7.74	7.08	7.38
FSCL	34.52	34.84	35.34
FSCL+	34.73	34.98	35.79
FairViT	5.09	5.16	5.27

6 Conclusion, Limitation and Discussion

In this paper, we proposed FairViT, addressing the fairness-accuracy issue in vision transformers. FairViT employs adaptive masks to alleviate bias without compromising accuracy and crafts a versatile distance loss to enhance overall accuracy. Extensive experiments validate that FairViT can enhance fairness while upholding comparable levels of accuracy.

In the future, it would be interesting to extend the proposed techniques to a broader range of neural networks. In upcoming research endeavors, we aim to further explore the intrinsic mechanisms driving the effectiveness of distance loss and adaptive masking. Furthermore, more experimental evaluations on other learning tasks other than classification are also of great interest. For example, we plan to further explore fair generation tasks such as text-to-image generation [11] and graph generation [15, 32].

³ AG is short of **Y**: Attraction, **S**: Gender.

⁴ EG is short of **Y**: Expression, **S**: Gender.

⁵ AH is short of **Y**: Attraction, **S**: Hair color.

References

1. Abnar, S., Zuidema, W.: Quantifying attention flow in transformers. arXiv preprint arXiv:2005.00928 (2020)
2. Ballantine, J.P., Jerbert, A.R.: Distance from a line, or plane, to a poin. *The American Mathematical Monthly* **59**(4), 242–243 (1952)
3. Beal, J., Kim, E., Tzeng, E., Park, D.H., Zhai, A., Kislyuk, D.: Toward transformer-based object detection. arXiv preprint arXiv:2012.09958 (2020)
4. Bhojanapalli, S., Chakrabarti, A., Glasner, D., Li, D., Unterthiner, T., Veit, A.: Understanding robustness of transformers for image classification. In: *IEEE/CVF International Conference on Computer Vision*. pp. 10231–10241 (2021)
5. Chen, Q., Syrgkanis, V., Austern, M.: Debiased machine learning without sample-splitting for stable estimators. *Advances in Neural Information Processing Systems* **35**, 3096–3109 (2022)
6. Dai, Z., Cai, B., Lin, Y., Chen, J.: Up-detr: Unsupervised pre-training for object detection with transformers. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1601–1610 (2021)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
8. Du, R., Shen, Y.: Fairness-aware user classification in power grids. In: *2022 30th European Signal Processing Conference (EUSIPCO)*. pp. 1671–1675. IEEE (2022)
9. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: *Proceedings of the 3rd innovations in theoretical computer science conference*. pp. 214–226 (2012)
10. Fang, Y., Liao, B., Wang, X., Fang, J., Qi, J., Wu, R., Niu, J., Liu, W.: You only look at one sequence: Rethinking transformer in vision through object detection. In: *Advances in Neural Information Processing Systems*. pp. 26183–26197 (2021)
11. Friedrich, F., Brack, M., Struppek, L., Hintersdorf, D., Schramowski, P., Luccioni, S., Kersting, K.: Fair diffusion: Instructing text-to-image generation models on fairness. arXiv preprint arXiv:2302.10893 (2023)
12. Gu, J., Kwon, H., Wang, D., Ye, W., Li, M., Chen, Y.H., Lai, L., Chandra, V., Pan, D.Z.: Multi-scale high-resolution vision transformer for semantic segmentation. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12094–12103 (2022)
13. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. *Advances in neural information processing systems* **29** (2016)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
15. Kose, O.D., Shen, Y.: Fast&fair: Training acceleration and bias mitigation for gnns. *Transactions on Machine Learning Research* (2023)
16. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 10012–10022 (2021)
17. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: *Proceedings of International Conference on Computer Vision (ICCV)* (December 2015)

18. Mao, A., Mohri, M., Zhong, Y.: Cross-entropy loss functions: Theoretical analysis and applications. arXiv preprint arXiv:2304.07288 (2023)
19. Moayeri, M., Pope, P., Balaji, Y., Feizi, S.: A comprehensive study of image classification model sensitivity to foregrounds, backgrounds, and visual attributes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19087–19097 (2022)
20. Park, S., Kim, D., Hwang, S., Byun, H.: Readme: Representation learning by fairness-aware disentangling method. arXiv preprint arXiv:2007.03775 (2020)
21. Park, S., Lee, J., Lee, P., Hwang, S., Kim, D., Byun, H.: Fair contrastive learning for facial attribute classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10389–10398 (2022)
22. Qiang, Y., Li, C., Khanduri, P., Zhu, D.: Fairness-aware vision transformer via debiased self-attention. arXiv preprint arXiv:2301.13803 (2023)
23. Raghu, M., Unterthiner, T., Kornblith, S., Zhang, C., Dosovitskiy, A.: Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems* **34**, 12116–12128 (2021)
24. Roh, Y., Nie, W., Huang, D.A., Whang, S.E., Vahdat, A., Anandkumar, A.: Dr-fairness: Dynamic data ratio adjustment for fair training on real and generated data. *Transactions on Machine Learning Research* (2023)
25. Shao, R., Shi, Z., Yi, J., Chen, P.Y., Hsieh, C.J.: On the adversarial robustness of visual transformers. arXiv preprint arXiv:2103.15670 **1**(2) (2021)
26. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: IEEE/CVF International Conference on Computer Vision. pp. 7262–7272 (2021)
27. Sudhakar, S., Prabhu, V., Krishnakumar, A., Hoffman, J.: Mitigating bias in visual transformers via targeted alignment. arXiv preprint arXiv:2302.04358 (2023)
28. Vasudevan, A., Anderson, A., Gregg, D.: Parallel multi channel convolution using general matrix multiplication. In: 2017 IEEE 28th international conference on application-specific systems, architectures and processors (ASAP). pp. 19–24. IEEE (2017)
29. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
30. Wang, Y., Xu, Z., Wang, X., Shen, C., Cheng, B., Shen, H., Xia, H.: End-to-end video instance segmentation with transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8741–8750 (2021)
31. Wang, Z., Dong, X., Xue, H., Zhang, Z., Chiu, W., Wei, T., Ren, K.: 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)
32. Wang, Z., Wallace, C., Bifet, A., Yao, X., Zhang, W.: FG²AN: Fairness-aware graph generative adversarial networks. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases. pp. 259–275. Springer (2023)
33. Xie, W., Li, X.H., Cao, C.C., Zhang, N.L.: Vit-cx: Causal explanation of vision transformers. arXiv preprint arXiv:2211.03064 (2022)
34. Zafar, M.B., Valera, I., Rogniguez, M.G., Gummadi, K.P.: Fairness constraints: Mechanisms for fair classification. In: Artificial intelligence and statistics. pp. 962–970. PMLR (2017)
35. Zhao, J., Wang, T., Yatskar, M., Ordonez, V., Chang, K.W.: Men also like shopping: Reducing gender bias amplification using corpus-level constraints. arXiv preprint arXiv:1707.09457 (2017)