

Navigate Beyond Shortcuts: Debiased Learning through the Lens of Neural Collapse

Yining Wang, Junjie Sun, Chenyue Wang, Mi Zhang[†], Min Yang
 School of Computer Science, Fudan University, China

{ynwang22@m., jjsun22@m., wangcy23@m., mi_zhang@, m_yang@}fudan.edu.cn

Abstract

Recent studies have noted an intriguing phenomenon termed *Neural Collapse*, that is, when the neural networks establish the right correlation between feature spaces and the training targets, their last-layer features, together with the classifier weights, will collapse into a stable and symmetric structure. In this paper, we extend the investigation of Neural Collapse to the biased datasets with imbalanced attributes. We observe that models will easily fall into the pitfall of shortcut learning and form a biased, non-collapsed feature space at the early period of training, which is hard to reverse and limits the generalization capability. To tackle the root cause of biased classification, we follow the recent inspiration of prime training, and propose an avoid-shortcut learning framework without additional training complexity. With well-designed shortcut primes based on Neural Collapse structure, the models are encouraged to skip the pursuit of simple shortcuts and naturally capture the intrinsic correlations. Experimental results demonstrate that our method induces better convergence properties during training, and achieves state-of-the-art generalization performance on both synthetic and real-world biased datasets.

1. Introduction

When the input-output correlation learned by a neural network is consistent with its training target, the last-layer features and classifier weights will attract and reinforce each other, forming a stable, symmetric and robust structure. Just as the Neural Collapse phenomenon discovered by Papayan et al. [24], at the terminal phase of training on balanced datasets, a model will witness its last-layer features of the same class converge towards the class centers, and the classifier weights align to these class centers correspondingly. The convergence will ultimately lead to the collapse of feature space into a simplex equiangular tight




Figure 1. Illustration of (a) **Neural Collapse** phenomenon on balanced datasets, where the simplex ETF structure maximizes the class-wise angles, and (b) **Biased Classification** on datasets with imbalanced attributes, where the model takes the shortcut of attributes to make predictions and fails to collapse into the simplex ETF. The color of points represents different class labels and the shape of points represents different attributes.

frame (ETF) structure, as illustrated in Fig. 1(a). The elegant structure has demonstrated its efficacy in enhancing the generalization, robustness, and interpretability of the trained models [5, 24]. Therefore, a wave of empirical and theoretical analysis of Neural Collapse has been proposed [2, 6, 8, 25, 26, 38, 40], and a series of studies have adopted the simplex ETF as the optimal geometric structure of the classifier, to guide the maximized class-wise separation in class-imbalanced training [20, 35–37, 43].

However, in practical visual recognition tasks, besides the challenge of inter-class imbalance, we also encounter intra-class imbalance, where the majority of samples are dominated by the bias attributes (e.g., some misleading contents such as background, color, texture, etc.). For example, the widely used LFW dataset [12] for facial recognition has been demonstrated severely imbalanced in gender, age and ethnicity [3]. A biased dataset often contains a majority of *bias-aligned* samples and a minority of *bias-conflicting* ones. The prevalent bias-aligned samples exhibit a strong correlation between the ground-truth labels

[†]Corresponding Author.

and bias attributes, while the scarce bias-conflicting samples have no such correlation. Once a model relies on the simple but spurious *shortcut* of bias attributes for prediction, it will ignore the intrinsic relations and struggle to generalize on out-of-distribution test samples. The potential impact of biased classification may range from political and economic disparities to social inequalities within AI systems, as emphasized in EDRI’s latest report [1].

Therefore, the fundamental solution to biased classification lies in deferring, or ideally, preventing the learning of shortcut correlations. However, previous *debiased learning* methods rely heavily on additional training expenses. For example, a bias-amplified auxiliary model is often adopted to identify and up-weight the bias-conflicting samples [16, 19, 23], or employed to guide the input-level and feature-level augmentations [13, 18, 21]. Some disentangle-based debiasing methods, from the perspective of causal intervention [32, 42] or Information Bottleneck theory [30], also require large amounts of contrastive samples or pre-training process to disentangle the biased features, significantly increasing the burden of debiased learning.

In this paper, we extend the investigation of Neural Collapse to the biased visual datasets with imbalanced attributes. Through the lens of Neural Collapse, we observe that models prioritize the period of *shortcut learning*, and quickly form the biased feature space based on misleading attributes at the early stage of training. After the bias-aligned samples reach zero training error, the intrinsic correlation within bias-conflicting samples will then be discovered. However, due to i) the scarcity of bias-conflicting samples and ii) the stability of the established feature space, the learned shortcut correlation is challenging to reverse and eliminate. The mismatch between bias feature space and the training target induces inferior generalizability, and hinders the convergence of Neural Collapse, as shown in Fig. 1(b).

To achieve efficient model debiasing, we follow the inspiration of *prime training*, and encourage the model to skip the active learning of shortcut correlations. The *primes* are often provided as additional supervisory signals to redirect the model’s reliance on shortcuts, which helps improve generalization in image classification and CARLA autonomous driving [33]. To rectify models’ attention on the intrinsic correlations, we define the primes with a training-free simplex ETF structure, which approximates the “optimal” shortcut features and guides the model to pursue unbiased classification from the beginning of training. Our method is free of auxiliary models or additional optimization of prime features. Experimental results also substantiate its state-of-the-art debiasing performance on both synthetic and real-world biased datasets.

Our contributions are summarized as follows:

- For the first time, we investigate the Neural Collapse phenomenon on biased datasets with imbalanced attributes.

Through the empirical results of feature convergence, we analyze the shortcut learning stage of training, as well as the fundamental issues of biased classification.

- We propose an efficient avoid-shortcut training paradigm, which introduces the simplex ETF structure as prime features, to rectify models’ attention on the intrinsic correlations.
- We demonstrate the state-of-the-art debiasing performance of our method on 2 synthetic and 3 real-world biased datasets, as well as the better convergence properties of debiased models.

2. Related Works

Debiased Learning. Extensive efforts have been dedicated to model debiasing, but they are significantly limited by additional training costs. Recent advances can be divided into three categories: reweight-based, augmentation-based, and disentangle-based. Based on the easy-to-learn heuristic of biased features, reweight-based approaches require pre-trained bias-amplified models to identify and emphasize the bias-conflicting training samples [16, 19, 23]. Augmentation-based approaches, with the guidance of explicit bias annotations, conduct image-level and feature-level augmentations to enhance the diversity of training datasets [13, 18, 21]. Other disentangle-based approaches attempt to remove the bias-related part of features, from the perspective of Information Bottleneck theory [30] or causal intervention [32, 42], but at the cost of substantial contrastive samples. Additionally, model debiasing is also well studied in graph neural networks [4, 41], language models [7, 22] and multi-modal tasks [11, 34].

Neural Collapse. Discovered by Pappayan et al. [24], the Neural Collapse phenomenon reveals the convergence of the last-layer feature space to an elegant geometry. At the terminal phase of training on balanced datasets, the feature centers and classifier weights will collapse together into the structure of a simplex ETF, which is illustrated in Section 3.1. Recent works have dug deeper into the phenomenon and provided theoretical supports under different constraints or regularizations [2, 8, 40], as well as empirical studies of intermediate features and transfer learning [6, 25, 26, 38]. Considering the class-imbalanced datasets, Fang et al. [5] point out the Minority Collapse phenomenon, where features of long-tailed classes will merge together and be hard to classify. As a remedy, they fix the classifier as an ETF structure during training, which guarantees the optimal geometric property in imbalanced learning [36], semantic segmentation [43], and federated learning [37]. To take a step further, our work fills the gap of Neural Collapse analysis on biased datasets with shortcut correlations.

Avoid-shortcut Learning. The recent inspiration of avoid-shortcut learning aims to postpone, or even prevent the learning of shortcut relations in model training. With well-

crafted contrastive samples [27–29] or artificial shortcut signals [33, 42], avoid-shortcut learning has demonstrated its efficacy in image classification, autonomous driving and question answering models. One of the representative methods is named prime training, which provides richer supervisory signals of key input features (i.e., *primes*) to guide the establishment of correct correlations, therefore improving generalization on OOD samples [33]. In this work, we leverage the approximated “optimal” shortcuts as primes to encourage the models to bypass shortcut learning.

3. Preliminaries

3.1. Neural Collapse Phenomenon

Consider a biased dataset \mathcal{D} with K classes of training samples, we denote $\mathbf{x}_{k,i}$ as the i -th sample of the k -th class and $\mathbf{z}_{k,i} \in \mathbb{R}^d$ as its corresponding last-layer feature. A linear classifier with weights $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_K] \in \mathbb{R}^{d \times K}$ is trained upon the last-layer features to make predictions.

The Neural Collapse (NC) phenomenon discovered that, when neural networks are trained on balanced datasets, the correctly learned correlations will naturally lead to the convergence of feature spaces. Given enough training steps after the zero classification error, the last-layer features and classifier weights will collapse to the vertices of a simplex equiangular tight frame (ETF), which is defined as below.

Definition 1 (Simplex Equiangular Tight Frame) A collection of vectors $\mathbf{m}_k \in \mathbb{R}^d$, $k = 1, 2, \dots, K$, $d \geq K - 1$ is said to be a k -simplex equiangular tight frame if:

$$\mathbf{M} = \sqrt{\frac{K}{K-1}} \mathbf{P} (\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^T) \quad (1)$$

where $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_K] \in \mathbb{R}^{d \times K}$, and $\mathbf{P} \in \mathbb{R}^{d \times K}$ is an orthogonal matrix which satisfies $\mathbf{P}^T \mathbf{P} = \mathbf{I}_K$, with \mathbf{I}_K denotes the identity matrix and $\mathbf{1}_K$ denotes the all-ones vector. Within the ETF structure, all vectors have the maximal pair-wise angle of $-\frac{1}{K-1}$, namely the maximal equiangular separation.

Besides the convergence to simplex ETF structure, the Neural Collapse phenomenon could be concluded as the following properties during the terminal phase of training:

NC1: Variability collapse. The last-layer features $\mathbf{z}_{k,i}$ of the same class k will collapse to their class means $\bar{\mathbf{z}}_k = \text{Avg}_i \{\mathbf{z}_{k,i}\}$, and the within-class variation of the last-layer features will approach 0.

NC2: Convergence to simplex ETF. The normalized class means will collapse to the vertices of a simplex ETF. We denote the global mean of all last-layer features as $\mathbf{z}_G = \text{Avg}_{i,k} \{\mathbf{z}_{k,i}\}$, $k \in [1, \dots, K]$ and the normalized class means as $\tilde{\mathbf{z}}_k = (\mathbf{z}_k - \mathbf{z}_G) / \|\mathbf{z}_k - \mathbf{z}_G\|$, which satisfies Eq. 1.

NC3: Self duality. The classifier weights \mathbf{w}_k will align with the corresponding normalized class means $\tilde{\mathbf{z}}_k$, which satisfies $\tilde{\mathbf{z}}_k = \mathbf{w}_k / \|\mathbf{w}_k\|$.

NC4: Simplification to nearest class center. After convergence, the model’s prediction will collapse to simply choosing the nearest class mean to the input feature (in standard Euclidean distance). The prediction of \mathbf{z} could be denoted as $\arg \max_k \langle \mathbf{z}, \mathbf{w}_k \rangle = \arg \min_k \|\mathbf{z} - \bar{\mathbf{z}}_k\|$.

3.2. Neural Collapse Observation on Biased Dataset

Besides the findings on balanced datasets, some studies have explored Neural Collapse under the class-imbalanced situation [5, 36]. Taking a step further, we investigate the phenomenon on biased datasets with imbalanced attributes, to advance the understanding of biased classification. To examine the convergence of last-layer features and classifier weights, we compare the metrics of Neural Collapse on both unbiased and synthetic biased datasets. As shown in Fig. 2, we report the result of NC1-NC3, which corresponds to the first three convergence properties in Section 3.1 and respectively evaluates the convergence of same-class features, the structure of feature space and self-duality. The details of NC metrics are concluded in Tab. 1.

When trained on unbiased datasets (**black lines** in Fig. 2), the model displays the expected convergence properties, with metrics NC1-NC3 all converge to zero. We owe the elegant collapse phenomenon to the right correlation between the feature space and training objective, which is also supported by the analysis of benign global landscapes [45, 46].

However, when trained on biased datasets, the training process exhibits two stages: first the *shortcut learning* period and then the *intrinsic learning* period, as divided by the vertical dashed line. During the shortcut learning period, the accuracy of bias-aligned samples increases quickly, and the NC1-NC3 metrics show a rapid decline (**green lines with ▲** in Fig. 2). It indicates that when simple shortcuts exist in the training distribution, the model will quickly establish its feature space based on the bias attributes, and exhibit a converging trend towards the simplex ETF structure.

After the bias-aligned samples approach zero error, the model turns to the period of *intrinsic learning*, which focuses on the intrinsic correlations within bias-conflicting samples to further reduce the empirical loss. However, although their final loss reduces to zero, the bias-conflicting samples still display low accuracy and poor convergence results (**green lines with ▼**). It implies that the intrinsic learning period merely induces the over-fitting of bias-conflicting samples and does not benefit in generalization. We attribute the failure of collapse to the early establishment of shortcut correlations. Once the biased feature space is established based on misleading attributes, rectifying it becomes challenging, particularly with scarce bias-conflicting samples. In the subsequent training steps, the misled features of bias-conflicting samples will hinder the convergence of same-class features, thereby halting the converging trend towards the simplex ETF structure and lead-




Figure 2. Comparison of (a) testset accuracy and (b-d) Neural Collapse metrics on unbiased (CIFAR-10) and synthetic biased (Corrupted CIFAR-10 with the bias ratio of 5.0%) datasets. All vanilla models are trained with standard cross-entropy loss for 500 epochs. The postfix *-Aligned* and *-Conflicting* indicate the results of bias-aligned and bias-conflicting samples respectively. The NC1 metric evaluates the convergence of same-class features, NC2 evaluates the difference between the feature space and a simplex ETF, and NC3 measures the duality between feature centers and classifier weights. The vertical dashed line at the epoch of 60 divides two stages of training.

Metrics	Computational details
NC1	$\mathcal{NC}_1 = \frac{1}{K} \text{Tr}(\Sigma_W \Sigma_B^\dagger)$, where Tr is the trace of matrix and Σ_B^\dagger denotes the pseudo-inverse of Σ_B $\Sigma_B = \frac{1}{K} \sum_{k \in [K]} (\bar{\mathbf{z}}_k - \mathbf{z}_G)(\bar{\mathbf{z}}_k - \mathbf{z}_G)^\top$, $\Sigma_W = \frac{1}{K} \sum_{k \in [K]} \frac{1}{n_k} \sum_{i=1}^{n_k} (\mathbf{z}_{k,i} - \bar{\mathbf{z}}_k)(\mathbf{z}_{k,i} - \bar{\mathbf{z}}_k)^\top$
NC2	$\mathcal{NC}_2 = \left\ \frac{\mathbf{W}\mathbf{W}^\top}{\ \mathbf{W}\mathbf{W}^\top\ _{\mathcal{F}}} - \frac{1}{\sqrt{K-1}}(\mathbf{I}_K - \frac{1}{K}\mathbf{1}_K\mathbf{1}_K^\top) \right\ _{\mathcal{F}}$
NC3	$\mathcal{NC}_3 = \left\ \frac{\mathbf{W}\bar{\mathbf{Z}}}{\ \mathbf{W}\bar{\mathbf{Z}}\ _{\mathcal{F}}} - \frac{1}{\sqrt{K-1}}(\mathbf{I}_K - \frac{1}{K}\mathbf{1}_K\mathbf{1}_K^\top) \right\ _{\mathcal{F}}$, where $\bar{\mathbf{Z}} = [\bar{\mathbf{z}}_1 - \mathbf{z}_G, \dots, \bar{\mathbf{z}}_K - \mathbf{z}_G]$

Table 1. The metrics of evaluating the Neural Collapse phenomenon, which are generally adopted in previous studies [24, 38, 46]. $\|\cdot\|_{\mathcal{F}}$ denotes the Frobenius norm, \mathbf{I}_K is the identity matrix and $\mathbf{1}_K$ is the all-ones vector.

ing to a non-collapsed, sub-optimal feature space.

To break the curse of shortcut learning, we turn the tricky shortcut into a training prime, which effectively guides the models to focus on intrinsic correlations and form a naturally collapsed feature space (blue lines in Fig. 2). Our method is presented in the following sections.

4. Methodology

4.1. Motivation

Following the previous analysis, we highlight the importance of redirecting the model’s emphasis from simple shortcuts to intrinsic relations. Since the models can be easily misled by shortcuts in the training distribution, is it feasible to supply a “perfectly learned” shortcut feature, to deceive the models into skipping the active learning of shortcuts, and directly focusing on the intrinsic correlations?

We observe in Fig. 2 that the NC1-NC3 metrics show a rapid decrease during shortcut learning, but remain stable in the subsequent training epochs. However, if the training distribution does follow the shortcut correlation (with no obstacle from bias-conflicting samples), the convergence will end up with the optimal structure of simplex ETF, just as the results on unbiased datasets. This inspires us to approximate the “perfectly learned” shortcut features with a simplex ETF structure, which requires no additional training and represents the optimal geometry of feature space.

Therefore, following the outstanding performance of prime training in OOD generalization [33], we introduce the approximated “perfect” shortcuts as the primes for debiased learning. The provided shortcut primes are constructed with a training-free simplex ETF structure, which encourages the models to directly capture the intrinsic correlations, therefore exhibit superior generalizability and convergence properties in our experiments.

4.2. Avoid-shortcut Learning with Neural Collapse

Building upon our motivation of avoid-shortcut learning, the illustration of the proposed **ETF-Debias** is shown in Fig. 3. The debiased learning framework can be divided into three stages: *prime construction*, *prime training*, and *unbiased classification*. Firstly, a prime ETF will be constructed to approximate the “perfect” shortcut features. Then during the prime training, the model will be guided to directly capture the intrinsic correlations with the prime training and the prime reinforcement regularization. In evaluation, we rely on the intrinsic correlations to perform unbiased classification. The details are as follows.

Prime construction. When constructing the prime ETF, we first randomly initialize a simplex ETF as $\mathbf{M} \in \mathbb{R}^{d \times B}$, which satisfies the definition in Eq. 1. The dimension d is the same as the learnable features, and the number of vectors in \mathbf{M} is determined by the categories of bias attributes $\mathbf{b}_i \in \{1, \dots, B\}$, which are pre-defined in the training dis-

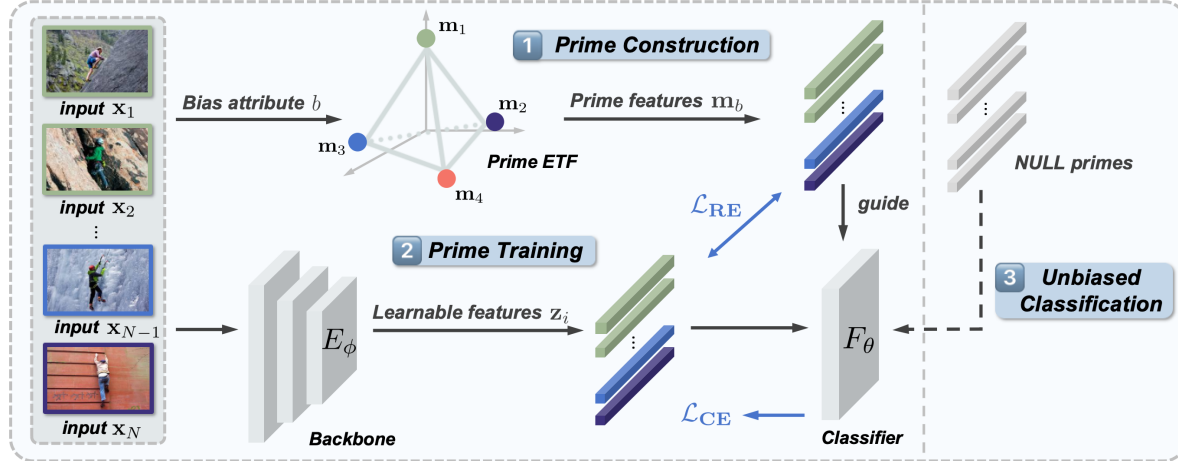


Figure 3. The illustration of our method. We take the class *climbing* from BAR [23] as an example, which contains samples of human climbing but with the bias attribute of different backgrounds (as indicated with the color of image frames). The framework contains: 1) **Prime Construction**: Before training, a randomly initialized ETF structure is constructed as the shortcut primes, and 2) During **Prime Training**, the prime features \mathbf{m}_b are retrieved based on the bias attribute b of the input samples, to guide the optimization of learnable features \mathbf{z} towards the intrinsic correlations. The classifier F_θ will take both the learnable features and fixed prime features to make predictions. 3) In **Unbiased Classification**, the prime features are assigned as null vectors to evaluate the debiased model on test distributions.

tribution. After initialization, the vertices of prime ETF $[\mathbf{m}_1, \dots, \mathbf{m}_B]$ are considered as the approximation of the “perfect” shortcut features for each attribute, which serve as the prime features for avoid-shortcut training. During training, the prime features will be retrieved based on the bias attribute b of each input sample.

Prime training. During the prime training, we take the end-to-end model architecture with a backbone E_ϕ and a classifier F_θ . For the i -th input $\mathbf{x}_{i,b}$ with the bias attribute of b , we first extract its learnable feature $\mathbf{z}_{i,b} = E_\phi(\mathbf{x}_{i,b})$ with the backbone model, and retrieve its prime feature \mathbf{m}_b based on the bias attribute b . The classifier F_θ will take both the learnable feature $\mathbf{z}_{i,b}$ and the prime feature \mathbf{m}_b to make softmaxed predictions $\hat{\mathbf{y}} = F_\theta(\mathbf{z}_{i,b}, \mathbf{m}_b)$. The standard classification objective is defined as:

$$\min_{\phi, \theta} \mathcal{L}_{\text{CE}}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^N \mathcal{L}(F_\theta(\mathbf{z}_{i,b}, \mathbf{m}_b), \mathbf{y}_{i,b}) \quad (2)$$

where \mathbf{y} is the ground-truth label. In our implementation, we use the standard cross-entropy loss as \mathcal{L} , and concatenate the prime features after the learnable features to perform predictions.

In essence, we provide a pre-defined prime feature for each training sample based on its bias attribute. The prime features, with a strong correlation with the bias attributes, can be viewed as the optimal solution to shortcut learning. By leveraging the already “perfect” representation of shortcut correlations, the model will be forced to explore the intrinsic correlations within the training distribution. The prime-guided mechanism targets at the fundamental issue of

biased classification, without inducing extra training costs. **Prime reinforcement regularization.** Given the prime features, the model is encouraged to grasp the intrinsic correlation of the training distributions. However, we raise another potential risk that, despite the provided “perfectly learned” shortcut features, the model may still pursue the easy-to-follow shortcuts, leading to the redundancy between the learnable feature $\mathbf{z}_{i,b}$ and the fixed \mathbf{m}_b . We point out that the model may not establish a strong correlation between the prime features and the bias attributes, and continues to optimize the learnable features for the missing connections.

Therefore, we introduce a *prime reinforcement regularization* mechanism to enhance the model’s dependency on prime features. We encourage the model to classify the bias attributes with only the prime features, and the regularization loss is defined as:

$$\mathcal{L}_{\text{RE}}(\mathbf{x}, \mathbf{b}) = \sum_{i=1}^N \mathcal{L}(F_\theta(\mathbf{z}_{i,b}, \mathbf{m}_b) - F_\theta(\mathbf{z}_{i,b}, \mathbf{m}_{\text{null}}), \mathbf{b}) \quad (3)$$

where \mathbf{m}_{null} is implemented as all-zero vectors with the same dimension as \mathbf{m}_b , and \mathcal{L} is the standard cross-entropy loss. In the ablation studies in Section 5.3, we observe an improved generalization capability across test distributions, as the result of the strengthened reliance on prime features. Regarding the entire framework, we define the overall training objective as:

$$\min_{\phi, \theta} \mathcal{L}_{\text{CE}}(\mathbf{x}, \mathbf{y}) + \alpha \mathcal{L}_{\text{RE}}(\mathbf{x}, \mathbf{b}) \quad (4)$$

where α is the hyper-parameter to adjust the regularization. **Unbiased classification.** In evaluation, we rely on the intrinsic correlations to perform unbiased classification.

Given a test sample $\mathbf{x}_{i,b}$, we extract its learnable feature $\mathbf{z}_{i,b}$ and set its prime feature as \mathbf{m}_{null} to obtain the final output $\hat{y} = F_{\theta}(\mathbf{z}_{i,b}, \mathbf{m}_{\text{null}})$.

4.3. Theoretical Justification

Based on the analysis of Neural Collapse from the perspective of gradients [36, 43], we provide a brief theoretical justification for our method.

With the priming mechanism, we denote the i -th feature of the k -th class as $\tilde{\mathbf{z}}_{k,i} = [\mathbf{z}_{k,i}, \mathbf{m}_{i,b}] \in \mathbb{R}^{2 \times d}$, which represents the concatenation of learnable feature $\mathbf{z}_{k,i}$ and prime feature $\mathbf{m}_{i,b}$ based on its bias attribute b . To keep the same form, we also denote the classifier weights as $\tilde{\mathbf{w}}_k = [\mathbf{w}_k, \mathbf{a}_k] \in \mathbb{R}^{2 \times d}$, where \mathbf{w}_k represents the weight for intrinsic correlations and \mathbf{a}_k represents the one for shortcut correlations. We observe that, due to the fixed prime features during training, \mathbf{a}_k will quickly collapse to the bias-correlated prime features of class k , and can be viewed as constant after just a few steps of training. With the definition, the cross-entropy (CE) loss can be written as:

$$\mathcal{L}_{\text{CE}}(\tilde{\mathbf{z}}_{k,i}, \tilde{\mathbf{w}}_k) = -\log \left(\frac{\exp([\mathbf{z}_{k,i}, \mathbf{m}_{i,b}]^T [\mathbf{w}_k, \mathbf{a}_k])}{\sum_{k'=1}^K \exp([\mathbf{z}_{k,i}, \mathbf{m}_{i,b}]^T [\mathbf{w}_{k'}, \mathbf{a}_{k'}])} \right) \quad (5)$$

We follow the analysis of previous works and compute the gradients of \mathcal{L}_{CE} w.r.t both classifier weights and features.

Gradient w.r.t classifier weights. We first compute the gradient of \mathcal{L}_{CE} w.r.t classifier weights $\tilde{\mathbf{W}} = [\tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_K]$:

$$\begin{aligned} \frac{\partial \mathcal{L}_{\text{CE}}}{\partial \tilde{\mathbf{w}}_k} &= \sum_{i=1}^{n_k} -(1 - p_k(\tilde{\mathbf{z}}_{k,i})) \tilde{\mathbf{z}}_{k,i} + \sum_{k' \neq k} \sum_{j=1}^{n_{k'}} p_k(\tilde{\mathbf{z}}_{k',j}) \tilde{\mathbf{z}}_{k',j} \\ &\leq \underbrace{\sum_{i=1}^{n_k} -(1 - p_k^{(b)}(\mathbf{m}_{i,b}) - p_k^{(l)}(\mathbf{z}_{k,i})) \tilde{\mathbf{z}}_{k,i}}_{\text{pulling part}} \\ &\quad + \underbrace{\sum_{k' \neq k} \sum_{j=1}^{n_{k'}} (p_k^{(b)}(\mathbf{m}_{j,b'}) + p_k^{(l)}(\mathbf{z}_{k',j})) \tilde{\mathbf{z}}_{k',j}}_{\text{forcing part}} \end{aligned} \quad (6)$$

where $p_k^{(l)}$ and $p_k^{(b)}$ are the predicted probabilities for class labels and bias attributes, calculated with softmax:

$$p_k^{(l)}(\mathbf{z}_{k,i}) = \frac{\exp(\mathbf{z}_{k,i}^T \mathbf{w}_k)}{\sum_{k'=1}^K \exp([\mathbf{z}_{k,i}, \mathbf{m}_{i,b}]^T [\mathbf{w}_{k'}, \mathbf{a}_{k'}])} \quad (7)$$

$$p_k^{(b)}(\mathbf{m}_{i,b}) = \frac{\exp(\mathbf{m}_{i,b}^T \mathbf{a}_k)}{\sum_{k'=1}^K \exp([\mathbf{z}_{k,i}, \mathbf{m}_{i,b}]^T [\mathbf{w}_{k'}, \mathbf{a}_{k'}])} \quad (8)$$

In Eq. 6, the gradient w.r.t classifier weights are divided into two parts. The *pulling part* is composed of features from the same class that pulls \mathbf{w}_k towards the direction of the k -th feature cluster, while the *forcing part* contains the features of other classes and pushes \mathbf{w}_k away from their clusters. The weight factor of each feature $\tilde{\mathbf{z}}_{k,i}$ represents

its influence on the optimization of $\tilde{\mathbf{w}}_k$, which implicitly plays the role of re-weighting in our method.

We assume that class k is strongly correlated with bias attribute b . As the weight \mathbf{a}_k is observed to collapse quickly to the bias-correlated prime feature \mathbf{m}_b , the probability $p_k^{(b)} \propto \exp(\mathbf{m}_b^T \mathbf{a}_k)$ of bias-aligned samples (with prime features \mathbf{m}_b) are much greater than that of bias-conflicting samples (with prime features $\mathbf{m}_{b'}$). Thus, with the weight factors in Eq. 6, the pulling and forcing effects of bias-aligned samples will be relatively down-weighted, and the impact of bias-conflicting samples will be up-weighted. The re-weighting mechanism of gradient mitigates the tendency of pulling $\tilde{\mathbf{w}}_k$ towards the center of bias-aligned samples, which alleviates the misdirection of bias attributes.

Gradient w.r.t features. Similarly, we compute the gradient of \mathcal{L}_{CE} w.r.t the feature $\tilde{\mathbf{z}}_{k,i}$:

$$\begin{aligned} \frac{\partial \mathcal{L}_{\text{CE}}}{\partial \tilde{\mathbf{z}}_{k,i}} &= -(1 - p_k(\tilde{\mathbf{z}}_{k,i})) \mathbf{w}_k + \sum_{k' \neq k} p_{k'}(\tilde{\mathbf{z}}_{k,i}) \mathbf{w}_{k'} \\ &\leq \underbrace{-(1 - p_k^{(b)}(\mathbf{m}_{i,b}) - p_k^{(l)}(\mathbf{z}_{k,i})) \mathbf{w}_k}_{\text{pulling part}} \\ &\quad + \underbrace{\sum_{k' \neq k} (p_{k'}^{(b)}(\mathbf{m}_{i,b}) + p_{k'}^{(l)}(\mathbf{z}_{k,i})) \mathbf{w}_{k'}}_{\text{forcing part}} \end{aligned} \quad (9)$$

In the gradient w.r.t features, the *pulling part* directs the feature $\tilde{\mathbf{z}}_{k,i}$ towards the weight of its class \mathbf{w}_k , and the *forcing part* repels it from wrong classes. Regarding the weight factors, the probability of bias attribute $p_k^{(b)}$ also re-weights the influence of classifier weights. Bias-aligned samples, with high $p_k^{(b)}$ probability, will have smaller pulling effects towards the classifier weight \mathbf{w}_k , which avoids the dominance of bias-aligned features around the weight centers and hinders the tendency of shortcut learning. In comparison, the bias-conflicting samples are granted stronger pulling and pushing effects, which strengthens their convergence toward the right class. The detailed theoretical justification of our method, along with the comparison with vanilla training, are available in Appendix A.

5. Experiments

5.1. Experimental Settings

Datasets and models. We validate the effectiveness of ETF-Debias on general debiasing benchmarks, which cover various types of bias attributes including color, corruption, gender, and background. We adopt 2 synthetic biased datasets, Colored MNIST [15] and Corrupted CIFAR-10 [10] with the ratio of bias-conflicting training samples $\{0.5\%, 1.0\%, 2.0\%, 5.0\%\}$, and 3 real-world biased datasets, Biased FFHQ (BFFHQ) [18] with bias ratio

Table 2. Comparison of debiasing performance on synthetic datasets. We report the accuracy on the unbiased test sets of Colored MNIST and Corrupted CIFAR-10. Best performances are marked in bold, and the number in brackets indicates the improvement compared to the best result in baselines. (*) and (°) denote methods with/without bias supervision respectively.

Dataset	Ratio(%)	Vanilla	LfF [°] [23]	LfF+BE [°] [19]	EnD [*] [30]	SD [*] [42]	DisEnt [*] [18]	Selecmix [°] [13]	ETF-Debias
Colored MNIST	0.5	32.22±0.13	57.78±0.81	69.69±1.99	35.93±0.40	56.96±0.37	68.83±1.62	70.53±0.46	71.63 ±0.28 (+1.10)
	1.0	48.45±0.06	72.29±1.69	80.90±1.40	49.32±0.58	72.46±0.18	79.49±1.44	83.34 ±0.37	81.97±0.26 (-1.37)
	2.0	58.90±0.12	79.51±1.82	84.90±1.14	65.58±0.46	79.37±0.46	84.56±1.19	85.90±0.23	86.00 ±0.03 (+0.10)
	5.0	74.19±0.04	83.96±1.44	90.28±0.18	80.70±0.17	88.89±0.21	88.83±0.15	91.27±0.31	91.36 ±0.21 (+0.09)
Corrupted CIFAR-10	0.5	17.06±0.12	31.00±2.67	23.68±0.50	14.30±0.10	36.66±0.74	30.12±1.60	33.30±0.26	40.06 ±0.03 (+3.40)
	1.0	21.48±0.55	34.33±1.76	30.72±0.12	20.17±0.19	45.66±1.05	35.28±1.39	38.72±0.27	47.52 ±0.26 (+1.86)
	2.0	27.15±0.46	39.68±1.15	42.22±0.60	30.10±0.54	50.11±0.69	40.34±1.41	47.09±0.17	54.64 ±0.42 (+4.53)
	5.0	39.46±0.58	53.04±0.76	57.93±0.58	45.85±0.21	62.43±0.57	49.99±0.84	54.69±0.29	65.34 ±0.60 (+2.91)

Table 3. Comparison of debiasing performance on real-world datasets. We report the accuracy on the unbiased test sets of Biased FFHQ, Dogs & Cats, and BAR. The class-wise accuracy on BAR is reported in Appendix D. Best performances are marked in bold, and the number in brackets indicates the improvement compared to the best result in baselines. (*) and (°) denote methods with/without bias supervision respectively.

Dataset	Ratio(%)	Vanilla	LfF [°] [23]	LfF+BE [°] [19]	EnD [*] [30]	SD [*] [42]	DisEnt [*] [18]	Selecmix [°] [13]	ETF-Debias
Biased FFHQ	0.5	53.27±0.61	65.60±2.27	67.07±2.37	55.93±1.62	65.60±0.20	63.07±1.14	65.00±0.82	73.60 ±1.22 (+6.53)
	1.0	57.13±0.64	72.33±2.19	73.53±1.62	61.13±0.50	69.20±0.20	68.53±2.32	67.50±0.30	76.53 ±1.10 (+3.00)
	2.0	67.67±0.81	74.80±2.03	80.20±2.78	66.87±0.64	78.40±0.20	72.00±2.51	69.80±0.87	85.20 ±0.61 (+5.00)
	5.0	78.87±0.83	80.27±2.02	87.40±2.00	80.87±0.42	84.80±0.20	80.60±0.53	83.47±0.61	94.00 ±0.72 (+6.60)
Dogs & Cats	1.0	51.96±0.90	71.17±5.24	78.87±2.40	51.91±0.24	78.13±1.06	65.13±2.07	54.19±1.61	80.07 ±0.90 (+1.20)
	5.0	76.59±1.27	85.83±1.62	88.60±1.21	79.07±0.28	89.12±0.18	82.47±2.86	81.50±1.06	92.18 ±0.62 (+3.06)
BAR	1.0	68.00±0.43	68.30±0.97	71.70±1.33	68.25±0.19	67.33±0.35	69.30±1.27	69.83±1.02	72.79 ±0.21 (+1.09)
	5.0	79.34±0.19	80.25±1.27	82.00±1.24	78.86±0.36	79.10±0.42	81.19±0.70	78.79±0.52	83.66 ±0.21 (+1.66)

{0.5%, 1.0%, 2.0%, 5.0%}, BAR [23], and Dogs & Cats [15] with bias ratio {1.0%, 5.0%}.

As for the model architecture, we adopt a three-layer MLP for Colored MNIST and ResNet-20 [9] for other datasets. Since BAR has a tiny training set, we follow the previous work [19] and initialize the parameters with pre-trained models on corresponding datasets. All results are averaged over three independent trials. More details about datasets and implementation are available in Appendix B.

Baselines. According to the three categories of debiased learning in Section 2, we compare the performance of ETF-Debias with six recent methods. For reweight-based debiasing, we consider LfF [23] with auxiliary bias models, and its improved version LfF+BE [19]. For disentangle-based debiasing, we consider EnD [30] and SD [42], which stem from the Information Bottleneck theory and causal intervention respectively. For augmentation-based debiasing, we consider DisEnt [18] and Selecmix [13], to include both the feature-level and image-level augmentations.

5.2. Main Results

Comparison on synthetic datasets. To display the debiasing performance, we report the accuracy on the unbiased

test set of 2 synthetic datasets in Tab. 2. It’s notable that ETF-Debias consistently outperforms baselines in the generalization capability towards test samples, on almost all levels of bias ratio. We observe that some baseline methods (e.g., EnD) do not display a satisfactory debiasing effect on synthetic datasets, as they rely heavily on diverse contrastive samples to identify and mitigate the bias features. In contrast, our approach directly provides the approximated shortcut features as training primes, which achieves superior performance on synthetic bias attributes.

Comparison on real-world datasets. To verify the scalability of ETF-Debias in real-world scenarios with more diverse bias attributes, we test our method on 3 real-world biased datasets in Tab. 3. We observe that ETF-Debias shows an even greater performance gain on real-world datasets than on synthetic ones, which may be attributed to the semantically meaningful prime features constructed with the simplex ETF structure. On the large-scale BFFHQ dataset, our method achieves up to 6.6% accuracy improvements compared to baseline methods, demonstrating its potential in real-world applications.

Convergence of Neural Collapse. In Fig. 2, we display the trajectory of NC metrics during training on the Corrupted

CIFAR-10 dataset. Guided by the prime features, the model establishes a right correlation and shows a much better convergence property on biased datasets, contributing to the superior generalization capability. More convergence results are available in Appendix C.

5.3. Ablation Study

Ablation on the influence of regularization. To measure the sensitivity of our method to different levels of prime reinforcement regularization, we compare the accuracy on the unbiased test set with α range from 0.0 to 1.0 in Fig. 4(a). It’s been shown that the debiasing performance remains significant with different strengths of regularization, and achieves extra performance gain with the proper level of prime reinforcement.




Figure 4. Ablation studies on hyper-parameter and prime features. We report (a) test set accuracy on different datasets, with hyper-parameter α ranging from 0.0 to 1.0, and (b) test set accuracy on 5 datasets with different prime features. The shaded areas represent the standard deviation, and the bias ratio of all datasets is 5%.

Ablation on the influence of ETF prime features. As illustrated before, we choose the vertices of ETF as the “perfectly learned” shortcut features, thus redirecting the model’s attention to intrinsic correlations. To demonstrate the efficacy of ETF prime features, we compare the results of randomly initialized prime features with the same dimension as the ETF-based ones. As shown in Fig. 4(b), the randomly initialized primes suffer a severe performance degradation, underscoring the advantages of ETF-based prime features in approximating the optimal structure.

5.4. Visualization

To intuitively reveal the effectiveness of our method, we compare the CAM [44] visualization results on vanilla models and debiased models trained with ETF-Debias. As shown in Fig. 5, the model’s attention is significantly rectified with ETF-Debias. For example, on Corrupted CIFAR-10 dataset, vanilla models are easily misled by the corruptions on the entire images, but with the guide of prime fea-




Figure 5. The comparison of visualization results on bias-conflicting samples (first row) between vanilla models (second row) and ETF-Debias models (third row). We display the result of CAM [44] on (a) Corrupted CIFAR-10 dataset, with the bias attribute as different types of corruption on the entire image, (b) Dogs & Cats dataset, with the bias attribute as the color of animals and (c) BFFHQ dataset, with the bias attribute as gender.

tures in our method, the debiased models shift their attention to the objects themselves. It’s also notable that our method circumvents the wrong attention area in classification and encourages the focus on more discriminative and finer-grained regions. On datasets of facial recognition, our method also breaks the spurious correlation on specific visual attributes [19] and considers more facial features, as shown in Fig. 5(c). More visualization results are available in Appendix E.

6. Conclusion

In this paper, we propose an avoid-shortcut learning framework with the insights of the Neural Collapse phenomenon. By extending the analysis of Neural Collapse to biased datasets, we introduce the simplex ETF as the prime features to redirect the model’s attention to intrinsic correlations. With the state-of-the-art debiasing performance on various benchmarks, we hope our work may advance the understanding of Neural Collapse and shed light on the fundamental solutions to model debiasing.

Acknowledgement

We appreciate the valuable comments from the anonymous reviewers that improves the paper’s quality. This work was supported in part by the National Key Research and Development Program (2021YFB3101200), National Natural Science Foundation of China (U1736208, U1836210, U1836213, 62172104, 62172105, 61902374, 62102093, 62102091). Min Yang is a faculty of Shanghai Institute of Intelligent Electronics & Systems, Shanghai Institute for Advanced Communication and Data Science, and Engineering Research Center of Cyber Security Auditing and Monitoring, Ministry of Education, China.

References

- [1] Agathe Balayn and Seda Gürses. Beyond debiasing: Regulating ai and its inequalities. *EDRi Report*. https://edri.org/wp-content/uploads/2021/09/EDRi_Beyond-Debiasing-Report-Online.pdf, 2021. 2
- [2] Hien Dang, Tan Nguyen, Tho Tran, Hung Tran, and Nhat Ho. Neural collapse in deep linear network: From balanced to imbalanced data. *ICML*, 2023. 1, 2
- [3] Athiya Deviyani. Assessing dataset bias in computer vision. *arXiv preprint arXiv:2205.01811*, 2022. 1
- [4] Shaohua Fan, Xiao Wang, Yanhu Mo, Chuan Shi, and Jian Tang. Debiasing graph neural networks via learning disentangled causal substructure. *NeurIPS*, 35:24934–24946, 2022. 2
- [5] Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43):e2103091118, 2021. 1, 2, 3
- [6] Tomer Galanti, András György, and Marcus Hutter. On the role of neural collapse in transfer learning. In *ICLR*, 2022. 1, 2
- [7] Yue Guo, Yi Yang, and Ahmed Abbasi. Auto-debias: Debiasing masked language models with automated biased prompts. In *ACL*, pages 1012–1023, 2022. 2
- [8] XY Han, Vardan Papyan, and David L Donoho. Neural collapse under mse loss: Proximity to and dynamics on the central path. In *ICLR*, 2022. 1, 2
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 7
- [10] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2018. 6, 13
- [11] Yusuke Hirota, Yuta Nakashima, and Noa Garcia. Model-agnostic gender debiased image captioning. In *CVPR*, pages 15191–15200, 2023. 2
- [12] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008. 1
- [13] Inwoo Hwang, Sangjun Lee, Yunhyeok Kwak, Seong Joon Oh, Damien Teney, Jin-Hwa Kim, and Byoung-Tak Zhang. SelecMix: Debiasing learning by contradicting-pair sampling. *NeurIPS*, 35:14345–14357, 2022. 2, 7, 13, 14, 16
- [14] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019. 14
- [15] Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *CVPR*, pages 9012–9020, 2019. 6, 7, 13, 14
- [16] Nayeong Kim, Sehyun Hwang, Sungsoo Ahn, Jaesik Park, and Suha Kwak. Learning debiased classifier with biased committee. *NeurIPS*, 35:18403–18415, 2022. 2
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 13
- [18] Jungsoo Lee, Eungyeup Kim, Juyoung Lee, Jihyeon Lee, and Jaegul Choo. Learning debiased representation via disentangled feature augmentation. *NeurIPS*, 34:25123–25133, 2021. 2, 6, 7, 13, 14, 16
- [19] Jungsoo Lee, Jeonghoon Park, Daeyoung Kim, Juyoung Lee, Edward Choi, and Jaegul Choo. Revisiting the importance of amplifying bias for debiasing. In *AAAI*, pages 14974–14981, 2023. 2, 7, 8, 13, 16
- [20] Zexi Li, Xinyi Shang, Rui He, Tao Lin, and Chao Wu. No fear of classifier biases: Neural collapse inspired federated learning with synthetic and fixed classifier. *ICCV*, 2023. 1
- [21] Jongin Lim, Youngdong Kim, Byungjai Kim, Chanho Ahn, Jinwoo Shin, Eunho Yang, and Seungju Han. Biasadv: Bias-adversarial augmentation for model debiasing. In *CVPR*, pages 3832–3841, 2023. 2
- [22] Yougang Lyu, Piji Li, Yechang Yang, Maarten de Rijke, Pengjie Ren, Yukun Zhao, Dawei Yin, and Zhaochun Ren. Feature-level debiased natural language understanding. In *AAAI*, pages 13353–13361, 2023. 2
- [23] Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure: De-biasing classifier from biased classifier. *NeurIPS*, 33:20673–20684, 2020. 2, 5, 7, 13, 14, 16
- [24] Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020. 1, 2, 4
- [25] Gao Peifeng, Qianqian Xu, Peisong Wen, Zhiyong Yang, Huiyang Shao, and Qingming Huang. Feature directions matter: Long-tailed learning via rotated balanced representation. *ICML*, 2023. 1, 2
- [26] Akshay Rangamani, Marius Lindegaard, Tomer Galanti, and Tomaso A Poggio. Feature learning in deep classifiers through intermediate neural collapse. In *ICML*, pages 28729–28745. PMLR, 2023. 1, 2
- [27] Joshua Robinson, Li Sun, Ke Yu, Kayhan Batmanghelich, Stefanie Jegelka, and Suvrit Sra. Can contrastive learning avoid shortcut solutions? *NeurIPS*, 34:4974–4986, 2021. 3
- [28] Piyapat Saranrittichai, Chaithanya Kumar Mummadi, Claudia Blaiotta, Mauricio Munoz, and Volker Fischer. Overcoming shortcut learning in a target domain by generalizing basic visual factors from a source domain. In *ECCV*, pages 294–309. Springer, 2022.
- [29] Kazutoshi Shinoda, Saku Sugawara, and Akiko Aizawa. Which shortcut solution do question answering models prefer to learn? In *AAAI*, pages 13564–13572, 2023. 3
- [30] Enzo Tartaglione, Carlo Alberto Barbano, and Marco Grangetto. End: Entangling and disentangling deep representations for bias correction. In *CVPR*, pages 13508–13517, 2021. 2, 7, 14, 16
- [31] Christos Thrampoulidis, Ganesh Ramachandra Kini, Vala Vakilian, and Tina Behnia. Imbalance trouble: Revisiting neural-collapse geometry. *Advances in Neural Information Processing Systems*, 35:27225–27238, 2022. 15
- [32] Tan Wang, Chang Zhou, Qianru Sun, and Hanwang Zhang. Causal attention for unbiased visual recognition. In *CVPR*, pages 3091–3100, 2021. 2

- [33] Chuan Wen, Jianing Qian, Jierui Lin, Jiaye Teng, Dinesh Jayaraman, and Yang Gao. Fighting fire with fire: Avoiding dnn shortcuts through priming. In *ICML*, pages 23723–23750. PMLR, 2022. 2, 3, 4
- [34] Zhiqian Wen, Guanghai Xu, Mingkui Tan, Qingyao Wu, and Qi Wu. Debaised visual question answering from feature and sample perspectives. *NeurIPS*, 34:3784–3796, 2021. 2
- [35] Liang Xie, Yibo Yang, Deng Cai, and Xiaofei He. Neural collapse inspired attraction–repulsion-balanced loss for imbalanced learning. *Neurocomputing*, 527:60–70, 2023. 1
- [36] Yibo Yang, Shixiang Chen, Xiangtai Li, Liang Xie, Zhouchen Lin, and Dacheng Tao. Inducing neural collapse in imbalanced learning: Do we really need a learnable classifier at the end of deep neural network? *NeurIPS*, 35:37991–38002, 2022. 2, 3, 6, 11
- [37] Yibo Yang, Haobo Yuan, Xiangtai Li, Zhouchen Lin, Philip Torr, and Dacheng Tao. Neural collapse inspired feature-classifier alignment for few-shot class-incremental learning. In *ICLR*, 2022. 1, 2
- [38] Yongyi Yang, Jacob Steinhardt, and Wei Hu. Are neurons actually collapsed? on the fine-grained structure in neural representations. *ICML*, 2023. 1, 2, 4
- [39] Corinna Cortes and Yann LeCun. Mnist handwritten digit database. Available at <http://yann.lecun.com/exdb/mnist/>, 2010. 13
- [40] Can Yaras, Peng Wang, Zhihui Zhu, Laura Balzano, and Qing Qu. Neural collapse with normalized features: A geometric analysis over the riemannian manifold. *NeurIPS*, 35:11547–11560, 2022. 1, 2
- [41] Qing Zhang, Xiaoying Zhang, Yang Liu, Hongning Wang, Min Gao, Jiheng Zhang, and Ruocheng Guo. Debiasing recommendation by learning identifiable latent confounders. *KDD*, 2023. 2
- [42] Yi Zhang, Jitao Sang, Junyang Wang, Dongmei Jiang, and Yaowei Wang. Benign shortcut for debiasing: Fair visual recognition via intervention with shortcut features. *ACM MM*, 2023. 2, 3, 7, 16
- [43] Zhisheng Zhong, Jiequan Cui, Yibo Yang, Xiaoyang Wu, Xiaojuan Qi, Xiangyu Zhang, and Jiaya Jia. Understanding imbalanced semantic segmentation through neural collapse. In *CVPR*, pages 19550–19560, 2023. 1, 2, 6, 11
- [44] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 8, 17
- [45] Jinxin Zhou, Chong You, Xiao Li, Kangning Liu, Sheng Liu, Qing Qu, and Zhihui Zhu. Are all losses created equal: A neural collapse perspective. *NeurIPS*, 35:31697–31710, 2022. 3
- [46] Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. *NeurIPS*, 34:29820–29834, 2021. 3, 4