

# Investigating and Explaining the Frequency Bias in Image Classification

Zhiyu Lin<sup>1</sup>, Yifei Gao<sup>1</sup>, Jitao Sang<sup>1,2\*</sup>

<sup>1</sup>Beijing Jiaotong University, China

<sup>2</sup>Peng Cheng Lab, Shenzhen 518066, China

{zyllin, yf-gao, jtsang}@bjtu.edu.cn

## Abstract

CNNs exhibit many behaviors different from humans, one of which is the capability of employing high-frequency components. This paper discusses the frequency bias phenomenon in image classification tasks: the high-frequency components are actually much less exploited than the low- and mid-frequency components. We first investigate the frequency bias phenomenon by presenting two observations on feature discrimination and learning priority. Furthermore, we hypothesize that (i) the spectral density, (ii) class consistency directly affect the frequency bias. Specifically, our investigations verify that the spectral density of datasets mainly affects the learning priority, while the class consistency mainly affects the feature discrimination.

## 1 Introduction

Convolutional neural networks (CNNs) have now approached (and sometimes surpassed) "human-level" benchmarks on various tasks, especially those involving visual recognition. To understand this impressive performance, many recent findings show that DNNs differ in intriguing ways from human vision on processing the visual information [Hermann *et al.*, 2020]. One such human-model disparity is the texture bias in CNN-based image classification [Geirhos *et al.*, 2019]. It has been observed that, unlike humans, CNNs tend to classify images by texture rather than by shape. Another interesting finding is that CNNs can exploit high-frequency image components that are not perceivable to human [Wang *et al.*, 2020; Jo and Bengio, 2017]. The ability in capturing the high-frequency components of images partially explains the un-intuitive behaviors of CNNs like generalization advancement and adversarial vulnerability.

This paper follows the line of studies on analyzing CNNs' behavior in frequency domain. Our study is inspired by the observation that, although employed in classifying images, the high-frequency components are much less exploited than the low- and middle-frequency components. Fig.1 illustrates

the Kernel Density Estimation (KDE) curves of the 10 image classes in CIFAR-10 [Krizhevsky *et al.*, 2009] for low-, middle- and high-frequency components. It is shown that before CNNs feature extraction, HOG[Surasak *et al.*, 2018] feature for all frequency components manifest noticeable discrimination between classes. However, after CNNs feature extraction, while feature discrimination for the low- and middle-frequency components (left two sub-figures) are enhanced due to supervised learning, the high-frequency components (right two sub-figures) are considerably inhibited. Regarding the highest-frequency component in the rightmost sub-figure, the KDE curves of different classes almost collapse as one unique class. This demonstrates an obvious frequency bias phenomenon in image classification: CNNs prefer low- and middle-frequency components over high-frequency components. Examining the frequency bias phenomenon and understanding the reasons behind will help fully exploit the potential of high-frequency components and further contribute to model improvement.

Only a few studies have discussed the frequency bias phenomenon in image classification, with focus on exhibiting the biased accuracy by employing different frequency components. For example, [Wang *et al.*, 2020] discloses that low-frequency components are much more generalizable than high-frequency components, [Abello *et al.*, 2021] reports similar results by introducing a new way to divide frequency spectrum. In this paper, we investigate into more fundamental observations beyond the biased accuracy, and explain with novel hypothesis from data perspective what leads to the frequency bias<sup>1</sup>. The contributions are summarized as two-fold:

- We provide new observations along feature discrimination and learning priority to investigate the frequency bias phenomenon in terms of image classification tasks (Section 3). These two observations are correlated with each other, together offering supplementary perspectives to recognized biased accuracy in existing studies.
- We propose hypotheses to explain the frequency bias from perspective of spectral density and class consistency (Section 4). Experiment results and analyses verify our hypotheses, which sheds light on future solution to alleviate the frequency bias phenomenon.

\*Corresponding Author

<sup>1</sup>Our code is available at <https://github.com/zhiyugege/FreqBias>

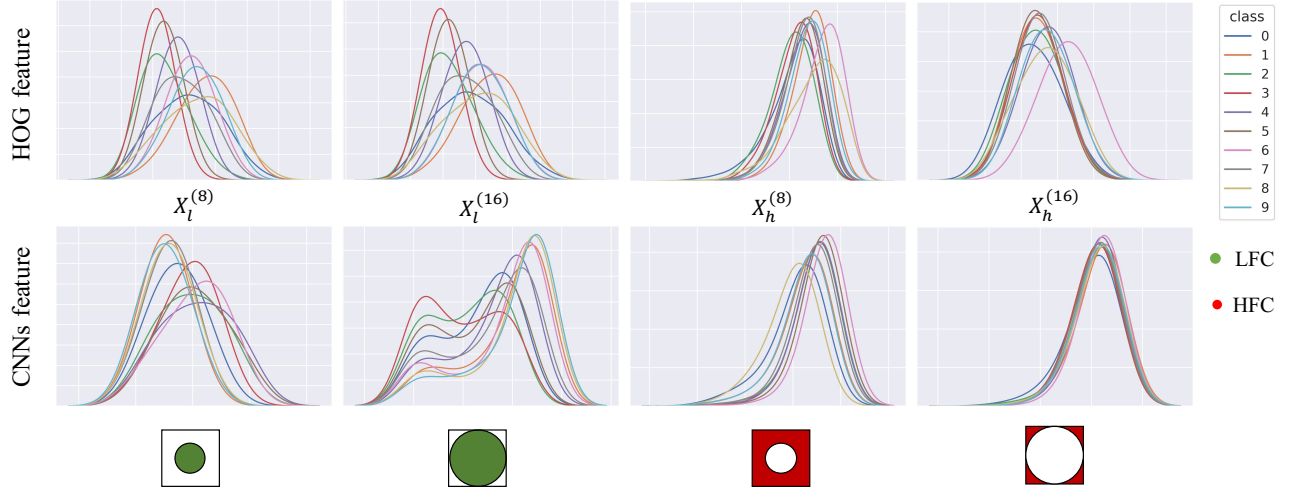


Figure 1: Visualization of KDE based inter-class variance for different frequency components (subfigures). We generate two types of frequency image following Eq.(2). The first row represents the class distribution of HOG features for all frequency components before CNNs extraction, it manifests noticeable discrimination between classes. The second row represents the features after CNNs extraction. While feature discrimination for the low- and middle- frequency components (left two sub-figures) are enhanced, the high-frequency components (right two sub-figures) are considerably inhibited.

## 2 Notations and Preliminaries

In this section, we set up the basic notations used in this paper. Given a data sample  $(X, y)$  where  $X$  represents the image with corresponding label  $y$ ,  $f(X; \theta) \in \mathbb{R}^D$  denotes the feature of a specific intermediate layer of the CNNs,  $D$  denotes the dimension of feature. Unless otherwise specified, we use the outputs from the penultimate layer as the feature representation for further analysis.  $\hat{y} = g(f(X; \theta), w) \in \mathbb{R}^C$  denotes the output of the CNNs, where  $C$  is the number of classes.  $\mathcal{L}(y, \hat{y})$  denotes a generic loss function (e.g., cross-entropy loss). In this study, we conduct frequency analysis with the 2D-discrete Fourier transform (shorten as DFT) computed as follows:<sup>2</sup>

$$\mathcal{F}(X)(u, v) = \frac{1}{HW} \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} X(h, w) \cdot e^{-j2\pi(\frac{uh}{H} + \frac{vw}{W})} \quad (1)$$

where  $\mathcal{F}(\cdot)$  denotes DFT and  $X_{\mathcal{F}} = \mathcal{F}(X)$  is the transformed Fourier image in frequency domain with the same dimension as the input image.  $\mathcal{F}^{-1}(\cdot)$  denote the 2D-inverse discrete Fourier transform (shorten as IDFT). We decompose the original image,  $X = \{X_l, X_h\}$ , where  $X_l$  and  $X_h$  denote the low-frequency components (shortened as LFC) and high-frequency components (shortened as HFC)<sup>3</sup>. Thus, we have the following two equations:

$$\begin{cases} X_l^{(r)} = \mathcal{F}^{-1}(\mathcal{F}(X) \odot \mathcal{M}_l^{(r)}) \\ X_h^{(r)} = \mathcal{F}^{-1}(\mathcal{F}(X) \odot \mathcal{M}_h^{(r)}) \end{cases} \quad (2)$$

where  $\mathcal{M}^{(r)}$  denotes the matrix of characteristic function with radius  $r$ ,  $\odot$  denotes the Hadamard product. We consider

<sup>2</sup> $H, W$  denotes the height and width of input image.

<sup>3</sup>Similarly, MFC indicates the mid-frequency components.

$d(\cdot, \cdot)$  as the Euclidean distance when dividing the spectrum in bands, and  $(c_i, c_j)$  as the center of an image. Therefore, we have:

$$\mathcal{M}_l^{(r)}(i, j) = \begin{cases} 1, & \text{if } d((i, j), (c_i, c_j)) \leq r \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

Obviously, there is  $\mathcal{M}_h^{(r)} = 1 - \mathcal{M}_l^{(r)}$ .

We follow the convention in [Durall *et al.*, 2020] and compute the azimuthally average of the magnitude of Fourier coefficients over radial frequencies to obtain the reduced spectrum and normalize it. Specifically, the spectral density is defined as the azimuthal integration over radial frequencies  $\phi$ .

$$AI_k(X) = \int_0^{2\pi} \|\mathcal{F}(X)(\omega_k \cdot \cos\phi, \sin\phi)\|^2 d\phi \quad (4)$$

where  $k = 0, \dots, H/2 - 1$ .

## 3 Observations of Frequency Bias

### 3.1 On Feature Discrimination

Following the results in Fig. 1 that the discriminative capability of HFC is inhibited after CNNs feature extraction, in this subsection, we investigate the frequency bias phenomenon by examining feature discrimination among different frequency components. Specifically, we propose to measure feature discrimination by computing the inter-class variance on KDE-based class feature distribution.

#### KDE-based Inter-Class Variance Calculation

We first introduce an estimation of class feature distribution based on KDE. Regarding the  $k^{th}$  feature dimension of class  $c$ , we sample  $N$  feature points  $\{x_i\}_{i=1}^N$  and estimate the kernel density on feature value  $x$  as follows:

$$S_c^k(x) = \frac{1}{N} \sum_{i=1}^N K(x - x_i) = \frac{1}{Nh} \sum_{i=1}^N K\left(\frac{x - x_i}{h}\right) \quad (5)$$




Figure 2: Illustration of KDE-based inter-class variance calculation for two feature dimensions (subfigure). For each feature dimension, the inter-class variance is estimated by the overlapped ratio (green area) of the two class curves (orange and blue).

where  $K(\cdot)$  denotes the Gaussian kernel function and  $h$  is the bandwidth selected by Silverman’s rule. Enumerating all possible feature values derives the feature distribution vector  $S_c^k$ , which is illustrated by the discrete histograms under each curve in Fig.2. The set of feature distributions  $\{S_c^k\}_{k=1}^D$  over all feature dimensions can be viewed as an estimation of the class feature distribution. Note that the curve of each class distribution in Fig.1 is estimated by all feature dimensions.

To evaluate the inter-class variance between two classes  $c_i$  and  $c_j$ , we propose to compute the overlapped area ratio between their feature distribution curves. Specifically, the overlapped area ratio is first computed for each feature dimension  $k$  and then averaged over all the  $D$  dimensions. Formally, it is defined as:

$$\text{Inter}(c_i, c_j) = \frac{1}{D} \sum_{k=1}^D \sum_x \frac{\min(S_{c_i}^k(x), S_{c_j}^k(x))}{\max(S_{c_i}^k(x), S_{c_j}^k(x))} \quad (6)$$

where  $\min(\cdot)$ ,  $\max(\cdot)$  are minimization and maximization functions respectively. Finally we measure the over inter-class variance by averaging all class pairs:

$$\text{Variance} = 1 - \frac{2}{C(C-1)} \sum_{i=1}^C \sum_{j=i}^C \text{Inter}(c_i, c_j) \quad (7)$$

### Observations and Discussions

We report results with ResNet-50 [He *et al.*, 2016] trained on CIFAR10 dataset<sup>4</sup>. To investigate the feature discrimination for different frequency components, for LFC and HFC, we respectively select frequency components with  $r = \{4, 8, 12, 16\}$  (defined as Eq.(2)). For each frequency component, the inter-class variance and test accuracy are calculated and shown in Fig.3.

The first observation is that feature discrimination capability exhibits significant bias among different frequency components. While features corresponding to the LMFC (e.g.,  $X_l^{(12)}$ ,  $X_l^{(16)}$ ,  $X_h^{(4)}$ ) capture adequate inter-class variance, the high- (e.g.,  $X_h^{(12)}$ ,  $X_h^{(16)}$ ) and very low- (e.g.,  $X_l^{(4)}$ ,  $X_l^{(8)}$ ) frequency components contribute to trivial discriminative features. With the fact that the inter-class variance reduces rapidly for the HFC as radius  $r$  increases, we demonstrate CNN’s frequency bias from the perspective of feature

<sup>4</sup>All experiments are repeated using ResNet-18 model and Restricted ImageNet dataset [Deng *et al.*, 2009]. The key observations are consistent. Detailed results are available in appendix.




Figure 3: Results of KDE-based inter-class variance and test accuracy between different frequency components. A positive correlation exhibits between variance and test accuracy.

discrimination that higher-frequency components are less exploited. A second observation is on the positive correlation between inter-class variance and test accuracy. It is easy to understand that discriminative features will contribute more to model’s generalization performance. This endorses the observed accuracy bias among frequency components in previous studies.

### 3.2 On Learning Priority

The previous analysis validates the phenomenon of feature discriminative bias led by frequency bias, which can be viewed as an explanation for generalization bias observed in previous works [Wang *et al.*, 2020]. This intuitively leads to an question: *How frequency bias affects the learning process?* While the affect of frequency bias can be effectively observed after training, understanding the evolution of bias phenomenon during training remains a problem. In this subsection, we further observe the learning priority of frequency components during training to investigate frequency bias phenomenon. We propose to evaluate the learning priority by analyzing gradient information.

#### Gradient Evaluation of Frequency Components

In the case of classification task, the measurement of gradients  $\frac{\partial \mathcal{L}}{\partial X}$  provides us valuable information about the contribution of spatial domain image  $X$  to loss function. In our work, we pay more attention on the gradient information of frequency domain image. Therefore, we first introduce a gradient map as an evaluation metric to measure the learning priority from frequency perspective. It is formally defined as:  $\frac{\partial \mathcal{L}}{\partial X^{(k)}}$ , where  $X_k$  represents a spatial image with only  $k^{th}$  frequency band information reserved from original image  $X$ . We expect to explore the priority performance of each frequency band of an image. One straight way to get the complete gradient information is as follows: we should first split the image spectrum into  $N$  bands, then traverse the value of  $k$  and use IDFT to obtain spatial images, finally compute the gradients for  $N$  times. However, directly calculating  $\frac{\partial \mathcal{L}}{\partial X^{(k)}}$  involves a high computational cost, we introduce the following proposition for approximated calculation. The proof of the proposition is available in appendix.

**Proposition**  $X$  denotes an input image.  $\mathcal{M}^{(k)}$  is the matrix of characteristic function which preserve the  $k^{th}$  band. Then the

---

**Algorithm 1** Visualizing the learning priority
 

---

**Input:**  $X$ : test image,  $N$ : training epochs

**Parameter:** The model parameters of each epoch

**Output:** The gradient spectrum

- 1: **for** epoch  $\leftarrow 0$  to  $N$  **do**
  - 2:    $\mathcal{L} \leftarrow$  Compute the loss
  - 3:    $\mathcal{G} \leftarrow$  Loss backward and obtain the gradient map
  - 4:    $\mathcal{S} \leftarrow AI(\mathcal{G})$  Compute the spectral density of gradient
  - 5: **end for**
  - 6: **return** spectral density of gradients  $\mathcal{S}$
- 

gradient map can be represented by the following equation:

$$\frac{\partial \mathcal{L}}{\partial X^{(k)}} = \mathcal{F}^{-1}(\mathcal{F}(\frac{\partial \mathcal{L}}{\partial X}) \odot \mathcal{M}^{(k)}), \text{ where } \mathcal{F}(\frac{\partial \mathcal{L}}{\partial X}) \text{ denotes the gradient spectrum.}$$

Following the proposition, the above computation process is equivalent to preserve the  $k^{th}$  band of the gradient spectrum in practice. It substantially reduces the computational cost, where we just need to compute the gradients once. In order to visualize the evolution of learning priority, we calculate the spectral density of gradients in each training epoch. The details are shown in Algo.1.

## Results and Discussions

The training setting is same as mentioned in section 3.1. To analyze the evolution of learning priority, we compute spectral density of gradients and average it over test set. We exhibit the evolution results during the first 50 epochs in Fig.4 (left). The spectral density of gradients at each row is notable: it mainly concentrates on some specific frequency bands (i.e., the colors on some frequency bands are much brighter than others) rather than uniformly distributes in all frequency bands, suggesting that the model exclusively pays attention to learning specific frequency information at each training stage. By comparing different rows, we find that the peak of gradient density (i.e., the brightest color) centralizes in the low-frequency bands at the early training stage, and then gradually shifts towards the middle-frequency bands and finally stop at the high-frequency bands. This significant trend, which reflects that models are strongly biased towards learning LFC first, reveals the learning priority bias on different components.

Interestingly, we find that the learning priority shifting from low frequency to high frequency occurs in a short time (about 20 epochs). Meanwhile, model focuses on learning HFC in the rest of training process (about 180 epochs). Recalling the discussion in Section 3.1, we illustrate accuracy curve of test set in Fig.4 (right). We mark the position of 15<sup>th</sup> epoch with a dash line. Note that test accuracy during the first 15 epochs increases much faster than the rest of training epochs. We argue that the bias of learning priority towards LFC contributes much to the generalizing capacity of model. Inversely, while the model manages to pick up HFC in a long training stage, it eventually achieves a poor performance on generalization improvement.




Figure 4: (Left): Evolution of the learned frequency components (x-axis for frequency bands) during the first 50 training epochs (y-axis). The colors show the measured amplitude of the gradients at the corresponding frequency bands (normalized to  $[0, 1]$ ). (Right): Evolution of the testing accuracy(x-axis for value of accuracy) during the first 50 training epochs (y-axis). The peaks of the gradient density is related to the learning priority, suggesting that models are strongly biased towards learning the LFC. Also, the LFC contribute more than HFC to accuracy.

## 4 Hypothesis behind Frequency Bias

### 4.1 On Spectral Density

To explain the frequency bias, we start with the investigation of frequency-domain characteristics on spectral density. Recent studies [Schwarz *et al.*, 2021] show that the architectures of GANs exhibit a frequency bias phenomenon towards generation task. They reports that an enhanced density of high frequency in spectrum is beneficial for the reconstruction of HFC. Motivated by this observation, we hypothesize that spectral density can serve as an explanation for frequency bias in terms of image classification. In this subsection, we propose a framework called *Convolutional Density Enhancement Strategy* (CDES) to modify spectral density of natural image and observe the performance changes in feature discrimination and learning priority.

#### Convolutional Density Enhancement Strategy (CDES)

Following the spectrum settings in [Schwarz *et al.*, 2021], the main idea of our framework is to create a density peak in high frequency. Since simply adding noise is difficult to modify the spectral density as we expect, we first propose to perform convolution operations on original images with a set of trainable convolution filters. Inspired by [Durall *et al.*, 2020; Dzanic *et al.*, 2020], the filters are optimized with loss function defined as:

$$\mathcal{L} = \sum_{k=0}^{H/2-1} \|AI_k(X^{(Conv)}) - AI_k^{(Target)}\|_2^2 + \|X^{(Conv)} - X\|_2^2 \quad (8)$$

where  $X^{(Conv)}$  denotes the images after convolution operation,  $AI^{(Target)}$  denotes the target spectral density. The first term of Eq.(8) is to match the spectral density of  $X^{(Conv)}$  with target spectral density  $AI^{(Target)}$  with L2-loss. To ensure the semantic information of the image, a regularization constraint at pixel level is illustrated in the second term of



Acc(%)	$X_h^{(10)}$	$X_h^{(12)}$	$X_h^{(14)}$	$X_h^{(16)}$	$X_h^{(18)}$
$f_{scr}$	<b>79.0</b>	<b>68.2</b>	<b>42.4</b>	<b>30.4</b>	<b>30.9</b>
$f_{wcr}$	20.4	15.5	12.8	10.0	10.0
$f_{baseline}$	43.2	27.2	14.3	10.5	8.3

Table 1: Here we report the test accuracy of the HFC among three models,  $f_{scr}$  significantly improves the generalization of HFC, while  $f_{wcr}$  has a poor performance that is even worse than the  $f_{baseline}$ .

Eq.(8). Since we are more concerned about the bias phenomenon towards HFC, we aim to maintain the immutability on LFC. Considering the disturbance on LFC after the above operations, we replace the LFC of  $X^{(Conv)}$  with LFC from original image  $X$  under a specified radius  $r$ . We select the radius of lowest point in expected spectral density curve. Then the recombin image  $X^{(Rec)}$  is defined as the following equation.

$$X^{(Rec)} = \mathcal{F}^{-1}(\mathcal{F}(X) \odot \mathcal{M}_l^{(r)} + \mathcal{F}(X^{(Conv)}) \odot \mathcal{M}_h^{(r)}) \quad (9)$$

The new dataset thus consists of images  $X^{(Rec)}$  and we split the train and test set on the new datasets.

### Experiment Setting

In our settings, we construct two variant datasets of CIFAR-10 based on CDES. (i)  $\mathcal{D}_{wcr}$ : We train one group of filters for the dataset. Note that this definition ensures that the filters is indistinguishable between classes. We argue that the added HFC is weakly class-related. (ii)  $\mathcal{D}_{scr}$ : We separately train a group of filters corresponding to each class, where the added HFC are strongly class-related<sup>5</sup>. Given the  $\mathcal{D}_{wcr}$ ,  $\mathcal{D}_{scr}$  and original datasets, we train the following classifiers:  $f_{wcr}, f_{scr}, f_{baseline}$ , and test them on the corresponding test set. The settings of the rest paper is to run the experiment with 150 epochs with SGD Momentum optimizer with learning rate set to be 1e-2 and batch size set to be 100. All experiments are repeated three times and averaged.

**Analyses of Feature Discrimination.** We report the test accuracy of HFC in Tab.1 with the following observations: (i)  $f_{scr}$  improves the feature discrimination on HFC with significantly increased accuracy. (ii)  $f_{wcr}$  exhibits a even worse performance of feature discrimination than  $f_{baseline}$ . We owe this huge disparity to the strategy of class-specific filters since the added HFC is weakly class-related. This also indicates that spectral density is not the only factor to explain the feature discrimination performance.

**Analyses of Learning Priority.** We visualize the gradient evolution during the first 10 epochs in the first row of Fig.5. An interesting finding is that the evolution tendency of  $f_{scr}$  and  $f_{wcr}$  are similar in the first few epochs. We hypothesize that the consistent spectral density of  $\mathcal{D}_{wcr}$  and  $\mathcal{D}_{scr}$  leads to the similar tendency. To verify this hypothesis, we average the spectral density of gradients in first 3 epochs and compare it with the spectral density of datasets. The second row

<sup>5</sup>It's worth noting that the expected spectrum of two datasets are the same and we respectively split the train and test set.

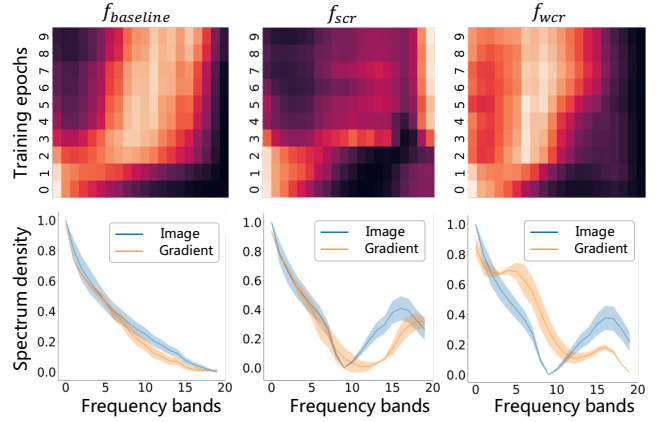


Figure 5: The first row shows the evolution of gradient spectrum in the first 10 training epochs among three models. The second row shows the Comparison of gradients and corresponding image on spectral density. The similar trends of spectral density at the early training stage suggests that the spectral density of dataset leads to the learning priority.

of Fig.5 confirms our hypothesis with the similar trends of two spectral density curves. This inspires us that the learning priority is led by the spectral density of the dataset at the early training stage.

The observations and analyses above partially verify our hypothesis that the spectral density of datasets explains well on learning priority phenomenon but is not a sufficient condition for feature discrimination.

## 4.2 On Class Consistency

In typical classification problem, the classifier is optimized to learn a consistent feature representation within the same class, so as to model the data-label relationship existed in the dataset. In this subsection, we aim to explain the frequency bias from the perspective of class consistency. Since data augmentation is an important training trick to ensure the class consistency at the data side, we first try to understand its validity and outstanding performance from frequency perspective. We choose the type of Mixed Sample Data Augmentation (shorten as MSDA) as our analyzing target. The method so far proposed can be broadly categorized as either combining samples with interpolation (e.g., Mixup[Zhang *et al.*, 2018] and CutMix[Yun *et al.*, 2019]).

### Frequency-based MSDA Analyses

To test the impact of MSDA strategy on the learning priority of different frequency bands, we use

$$\frac{AI_k(G^{(MSDA)}) - AI_k(G)}{AI_k(G)}$$

(for  $k = 1, 2, H/2 - 1$ , where  $G^{(MSDA)}$  and  $G$  denotes the gradient map of model trained with and without MSDA respectively) as a metric to represent the change rate of learning priority. We normalize the value to [-1,1] and results are shown in Fig.6.

An overall observation shows that MSDA models increase the utilization of MFC. We notice a significantly increasing

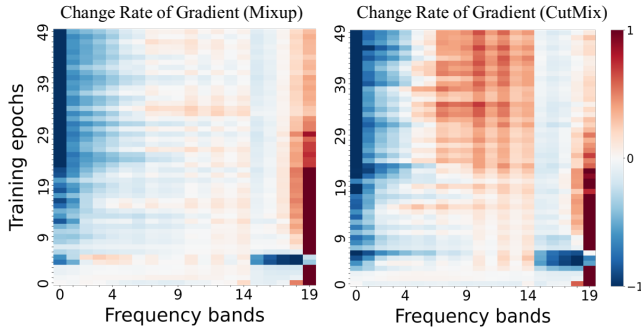


Figure 6: Visualization on the change rate of gradient evolution. Left reports the Mixup model, Right reports the CutMix model. Both of them show a tendency of utilizing MFC when the training process tends to be stable.

trend of the spectral density of gradient on MFC when the training process tends to be stable. The generalization results<sup>6</sup> of MFC also report a better performance on MSDA model than the baseline model (i.e., when  $r = 8$  or  $r = 12$ ), which align well with the phenomenon of learning priority. These findings imply that MSDA preserve the class consistency on MFC, which can effectively boost the performance of classification.

### Frequency Recombination-based Data Augmentation

Justified by the above analysis and inspired by the strategy of MSDA, we expect to introduce the idea of class consistency to the frequency domain. We hypothesize that the model only pick up the frequency components with the property of class consistency. For example, if we break up the class consistency on LFC, model will pay more attention to employing HFC.

In order to verify this hypothesis, we construct another variant dataset of CIFAR-10. We modify each image-label pair as follows: Given a clean image  $X$  with label  $y$ , we uniformly select another image  $X'$  with label  $y'$  from the dataset, (i.e.,  $y \neq y'$ ). Then we recombine the LFC of  $X'$  (shorten as  $x'_l$ ) and the HFC of  $X$  (shorten as  $X_h$ ) to create the new image  $X^*$ . For the selection strategy of radius  $r$  in Eq.(3), we consider the following two constraints: (i) For  $X'_l$ , it is indistinguishable from  $X'$  to human observation. (ii) For  $X_h$ , it preserves high frequency information from  $X$  as much as possible. In this case, we label  $X^*$  with label  $y$  so as to break the class consistency of LFC established by human annotation and preserve the class consistency of HFC. Then we create a new dataset named with HARS-dataset with image-label pair as  $(X^*, y)$ .

We report the generalization performance in Tab.2 and the results of learning priority are plotted in Fig.7. We have the following interesting findings:

**Analyses of Feature Discrimination.** The generalization performance of HFC on the clean test set is significantly enhanced. This indicates that training on HARS-dataset effectively increase the employment of HFC as well as the discrimination of high frequency features.

<sup>6</sup>The generalization results are available in appendix.




Figure 7: We visualize the evolution of gradient spectrum of HARS model in the first 50 epochs.

Acc(%)	LFC		HFC	
	Baseline	HARS	Baseline	HARS
r				
4	<b>15.91</b>	10.00	<b>82.30</b>	78.37
8	<b>32.12</b>	9.99	40.96	<b>76.25</b>
12	<b>81.54</b>	9.54	15.86	<b>64.20</b>
16	<b>90.68</b>	40.00	<b>11.42</b>	11.06

Table 2: We report the test accuracy on baseline model and HARS model. The HARS model has a good performance on HFC, while fail to generalize on LFC.

**Analyses of Learning Priority.** (i) Although we have broken up the class consistency of LFC, model still picked up LFC at the early training stage. The reason for the remaining bias towards LFC is due to the dominate power of LFC’s density in spectrum, which is consistent with original dataset. This observation further justify our hypothesis that the spectral density affects learning priority of frequency information. (ii) Compared with the gradient spectrum evolution of vanilla set-up as shown in Fig.7, we find that model starts paying attention on HFC much earlier.

## 5 Conclusion

This paper investigated and explained the frequency bias in image classification. We have extended the frequency bias phenomenon with observations on feature discrimination and learning priority. Possible reasons leading to the frequency bias phenomenon are also discussed with validated hypothesis. There remains an along way to go before solving the frequency bias problem and exploiting the biased frequency towards practical model improvement. Some future works include: (1) More reasons need to be analyzed from the model perspective, e.g., whether local receptive field and layered representation benefit high-frequency feature extraction. (2) It is necessary to reaching a balance between HFC and LMFC, i.e., enhancing the biased HFC without devastating the LMFC. (3) We found in preliminary observations that the inhibited high-frequency feature is closely related to the adversarial vulnerability. It will be interesting to examine whether addressing frequency bias provides possible solution for the generalization-robustness trade-off.

## Acknowledgments

This work is supported by the National Key R&D Program of China (Grant No.2018AAA0100604), the National Natural Science Foundation of China (Grant No.61832002, 62172094), and Beijing Natural Science Foundation (No.JQ20023).

## References

- [Abello *et al.*, 2021] Antonio A Abello, Roberto Hirata, and Zhangyang Wang. Dissecting the high-frequency bias in convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 863–871, June 2021.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [Durall *et al.*, 2020] Ricard Durall, Margret Keuper, and Janis Keuper. Watch your up-convolution: Cnn based generative deep neural networks are failing to reproduce spectral distributions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7890–7899, 2020.
- [Dzanic *et al.*, 2020] Tarik Dzanic, Karan Shah, and Freddie Witherden. Fourier spectrum discrepancies in deep network generated images. *Advances in neural information processing systems*, 33:3022–3032, 2020.
- [Geirhos *et al.*, 2019] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2019.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [Hermann *et al.*, 2020] Katherine Hermann, Ting Chen, and Simon Kornblith. The origins and prevalence of texture bias in convolutional neural networks. *Advances in Neural Information Processing Systems*, 33:19000–19015, 2020.
- [Jo and Bengio, 2017] Jason Jo and Yoshua Bengio. Measuring the tendency of cnns to learn surface statistical regularities. *arXiv preprint arXiv:1711.11561*, 2017.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Handbook of Systemic Autoimmune Diseases*, 1(4), 2009.
- [Schwarz *et al.*, 2021] Katja Schwarz, Yiyi Liao, and Andreas Geiger. On the frequency bias of generative models. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Surasak *et al.*, 2018] Thattapon Surasak, Ito Takahiro, Cheng-hsuan Cheng, Chi-en Wang, and Pao-you Sheng. Histogram of oriented gradients for human detection in video. In *2018 5th International conference on business and industrial research (ICBIR)*, pages 172–176. IEEE, 2018.
- [Wang *et al.*, 2020] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8684–8694, 2020.
- [Yun *et al.*, 2019] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019.
- [Zhang *et al.*, 2018] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018.