

# Improving Adversarial Robustness of Masked Autoencoders via Test-time Frequency-domain Prompting

Qidong Huang<sup>1</sup> Xiaoyi Dong<sup>2</sup> Dongdong Chen<sup>3</sup> Yinpeng Chen<sup>3</sup> Lu Yuan<sup>3</sup>

Gang Hua<sup>4</sup> Weiming Zhang<sup>1,\*</sup> Nenghai Yu<sup>1</sup>

<sup>1</sup>University of Science and Technology of China <sup>2</sup>Shanghai AI Lab

<sup>3</sup>Microsoft Research <sup>4</sup>Wormpex AI Research

{hqd0037@mail., zhangwm@, ynh@}ustc.edu.cn dongxiaoyi@pjlab.org.cn

{cddlyf@, ganghua@}gmail.com {yiche@, luyuan@}microsoft.com

## Abstract

In this paper, we investigate the adversarial robustness of vision transformers that are equipped with BERT pretraining (e.g., BEiT, MAE). A surprising observation is that MAE has significantly worse adversarial robustness than other BERT pretraining methods. This observation drives us to rethink the basic differences between these BERT pretraining methods and how these differences affect the robustness against adversarial perturbations. Our empirical analysis reveals that the adversarial robustness of BERT pretraining is highly related to the reconstruction target, i.e., predicting the raw pixels of masked image patches will degrade more adversarial robustness of the model than predicting the semantic context, since it guides the model to concentrate more on medium-/high-frequency components of images. Based on our analysis, we provide a simple yet effective way to boost the adversarial robustness of MAE. The basic idea is using the dataset-extracted domain knowledge to occupy the medium-/high-frequency of images, thus narrowing the optimization space of adversarial perturbations. Specifically, we group the distribution of pretraining data and optimize a set of cluster-specific visual prompts on frequency domain. These prompts are incorporated with input images through prototype-based prompt selection during test period. Extensive evaluation shows that our method clearly boost MAE’s adversarial robustness while maintaining its clean performance on ImageNet-1k classification. Our code is available at: <https://github.com/shikiw/RobustMAE>.

## 1. Introduction

BERT pretraining for vision transformers [3, 25, 16, 60, 20] ignites rising interests recently and shows great po-

\*Corresponding author.

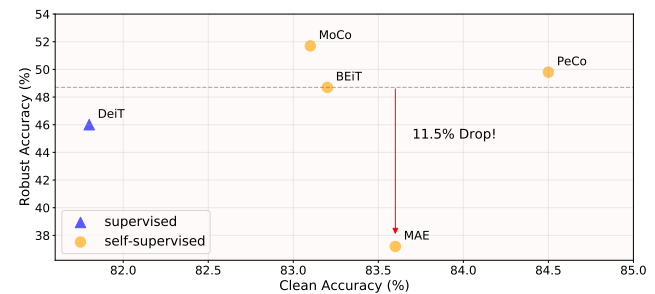


Figure 1. **MAE has much worse adversarial robustness than other pretraining counterparts.** We use ViT-Base as the backbone and evaluate on ImageNet-1k. The robust accuracy is tested under standard  $l_\infty$ -norm PGD attack.

tentials in improving downstream performance. Existing methods share the similar pretraining principle, i.e., randomly masking some patches on the input images and asking the transformer to reconstruct them based on unmasked patches, which encourages the transformer model to comprehend the semantic relationship among different patches for better reconstruction. Thus their main differences lie in the reconstruction targets.

Specifically, following the practices in natural language processing [14], BEiT [3] and PeCo [16] convert the image patches into discrete tokens with pretrained VQVAE [50] and regard such discrete tokens as the prediction targets. iBOT [60] integrates contrastive learning and uses online learned feature as prediction target. MAE [25] takes a step forward to directly use the raw pixel values as the target, achieving surprisingly good results. It can be regarded as a special form of denoising autoencoding, which provides a possibility to revitalize this classical direction. Although these methods have different reconstruction targets, they all achieve substantial improvement over multiple vision tasks.

However, we observe a surprising discrepancy between MAE and other pretraining counterparts (e.g., BEiT, PeCo)

on *adversarial robustness*. As shown in Figure 1, MAE is significantly more vulnerable to standard  $l_\infty$ -norm PGD [36] adversarial attack than BEiT and PeCo (drop 11.5% in ImageNet-1k top-1 accuracy). This paper is motivated by this observation to explore the causing factors, which naturally lead to our solution.

Our analysis manifests that MAE has the stronger dependence on medium-/high-frequency signals of images, thus it is more sensitive to the medium-/high-frequency noise introduced by regular adversarial attacks. This is exactly due to the difference between MAE and its counterparts (BEiT, PeCo) in reconstruction task of pretraining. Specifically, MAE reconstructs raw pixels of input images, while BEiT and PeCo reconstructs the semantic class for each image patch. Predicting raw pixels drives MAE to rely more on medium-/high-frequency information of input images. To validate this, we filter out both medium-/high-frequency components of clean images with a low pass filter and perform the evaluation. MAE suffers significantly more accuracy degradation than other pretraining methods (Please refer to Figure 4 and 5).

Based on the above observations, we propose a simple yet effective solution to improve the adversarial robustness of MAE. Our goal is to learn a set of frequency-domain visual prompts to enhance the medium-/high-frequency signals of input images during test time. In detail, we first analyze the representations learned by the pretrained model and then partition it into several clusters. On the samples distributed in each cluster, a visual prompt are optimized in frequency domain by minimizing the classification objective loss, of which the generation is constrained on medium-/high-frequency components of images. It naturally fills the medium-/high-frequency components with the dataset-extracted patterns, so that searching regular adversarial perturbations on these components can be more difficult. Experiments tells that the proposed solution not only significantly boost the adversarial robustness of MAE to surpass other pretraining counterparts, but also generally maintain the clean performance on the ImageNet-1k dataset.

## 2. Preliminaries

### 2.1. Vision BERT Pretraining

BERT pretraining [14] has become the dominant self-supervised learning mechanism in the field of natural language processing (NLP). The key idea is masked language modeling (MLM), *i.e.*, randomly masking some words in the sentence and asking the transformer to reconstruct them based on the remaining unmasked words. This reconstruction process requires the transformer model to understand underlying semantics and model the relationship between unmasked words and masked words, thus helping learn meaningful representations.

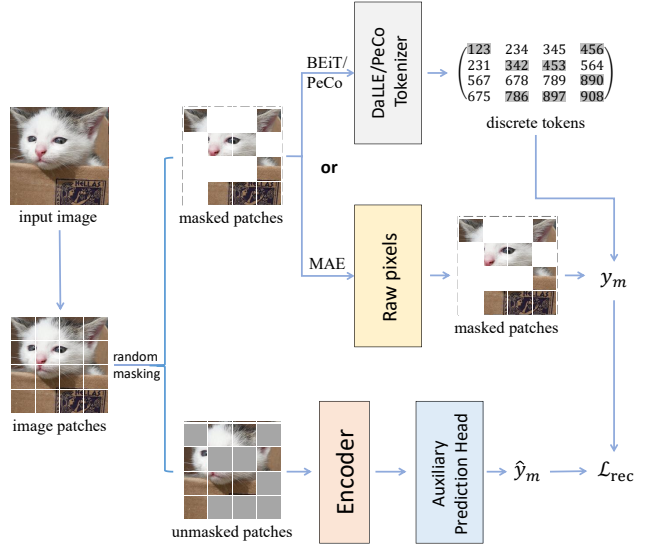


Figure 2. Key difference among BEiT, PeCo and MAE, where MAE directly use the original image patches for pixel-level reconstruction, BEiT and PeCo choose the discrete tokens as the target to reconstruct semantic context.

The recent surge of vision BERT pretraining follows the similar logic. As illustrated in Figure 2, each input image is converted into a series of patches (tokens) and some patches are randomly masked out. The unmasked patches  $\mathbf{x}_u$  are fed into the transformer encoder  $\mathbf{E}$ , on top of which an auxiliary prediction head  $\mathbf{h}$  is used to reconstruct masked patches  $\mathbf{x}_m$  by minimizing reconstruction loss  $\mathcal{L}_{rec}$ . Formally,

$$\begin{aligned} \hat{\mathbf{y}}_m &= \mathbf{h}(\mathbf{E}(\mathbf{x}_u, \mathbf{p}_{m,u}), \mathbf{p}_{m,u}), \\ \mathcal{L}_{rec} &= L(\hat{\mathbf{y}}_m, \mathbf{y}_m), \end{aligned} \quad (1)$$

where  $\mathbf{p}_{m,u}$  indicates the positional information of masked and unmasked tokens, thus the transformer model comprehends the spatial relationship of tokens.  $\hat{\mathbf{y}}_m, \mathbf{y}_m$  are the predicted and ground-truth reconstruction representations of masked tokens, respectively. Depending on the format of  $\mathbf{y}$ ,  $L$  is the corresponding reconstruction loss function (*e.g.*, cross-entropy loss in BEiT/PeCo and  $\ell_2$  loss in MAE). By analogy, the above vision BERT pretraining strategy is often called “masked image modeling (MIM)”.

The main difference among existing vision BERT pretraining methods lies in the design of  $\mathbf{y}$ . More specifically, BEiT [3] directly takes the pretrained DaLLE’s VQ-VAE tokenizer to convert each image patch into discrete tokens (indexed in a large codebook), where  $\mathbf{y}$  is the code index of each image patch. PeCo [16] follows the similar token design and designs a new perceptual VQVAE tokenizer, striving to make the discrete tokens more semantic. MaskFeat [54] extracts the low-level HOG features of image patches as  $\mathbf{y}$ . MAE[25] takes a big step forward, which directly uses the normalized raw pixel values as  $\mathbf{y}$ ,

*i.e.*,  $\mathbf{y}_m = \text{norm}(\mathbf{x}_m)$ . In the following sections, we will show that the design difference of  $\mathbf{y}$  is a key impact factor that influences the adversarial robustness of different BERT pretraining models.

After BERT pretraining, the pretrained transformers will be finetuned on downstream vision tasks, including but not limited to image classification, object detection and semantic segmentation. Considering it is difficult to evaluate the adversarial robustness of the pretrained model itself, we follow the common practice [30, 42, 37, 46, 38] and instead test the supervised finetuned model to evaluate the pretraining robustness.

## 2.2. Adversarial Attack and Defense

Deep neural networks (DNNs) are vulnerable to adversarial samples [47, 23, 55, 43, 48, 27, 19, 59, 17], which are generated by adding imperceptible perturbations on clean inputs to mislead the model prediction. Given a classification model  $\mathcal{C}(x)$ , the goal of adversarial attack is to search for a slight perturbation under the strict constraint until  $\mathcal{C}$  has been fooled. Usually, for an input  $x$  with its ground-truth label  $y$ , the perturbation added on the generated adversarial sample  $x'$  is constrained within  $l_p$ -norm ball, *i.e.*,  $\|x - x'\|_p \leq \epsilon$ . Typically, adversarial attacks can be divided into two categories: 1) Targeted attack that makes  $\mathcal{C}(x') = y'$  where  $y' \neq y$  is the label specified by the attacker. 2) Untargeted attack that just require  $x'$  to mislead the model  $\mathcal{C}$  to give wrong decision  $\mathcal{C}(x') \neq y$  without specifying the attack label. The adversarial robustness of one model is defined as whether this model can still correctly classify adversarial samples under the predefined attack strength, *i.e.*, classification accuracy on input adversarial samples.

Many defense methods have been proposed to improve model robustness against adversarial attacks, such as introducing stochasticity [15, 5], feature squeezing [40] or robust module designs [18]. Serving as the most effective way to defend against adversarial attack, adversarial training does not need any change on model architecture and directly engages adversarial samples during model training. Representative adversarial training methods include static ensemble models [49], adversarial logit pairing and universal first-order adversary [36, 58, 12] (*i.e.*, using PGD as the inner loop). Under the min-max optimization principle, adversarial training alternates between adversarial sample generation and model update, which follows the generic form

$$\min_{\theta} \mathbb{E}_{x \in \mathcal{D}} \left[ \max_{x'} \ell(x'; \theta) \right], \quad \text{s.t. } \|x - x'\|_p \leq \epsilon \quad (2)$$

where  $\theta$  denotes the learnable parameters of the model,  $\mathcal{D}$  denotes the training dataset,  $\ell$  is the training objective loss and  $\epsilon$  is the attack budget.

## 3. Adversarial Robustness of MAE

In this section, we prove that MAE has significant more degradation than other BERT pretraining methods when dealing with adversarial examples. *Our analysis shows that MAE’s sensitivity to adversarial noises is attributed to its strong dependence on medium-/high-frequency signals.*

### 3.1. MAE’s Degradation on Adversarial Robustness

We first verify the robust accuracy of different pretrained models on ImageNet-1k [13] classification. Five types of adversarial attacks, including PGD [36], BIM [33], MIM [21], C&W [6] and Auto-Attack (AA) [11], are selected here for the robustness evaluation, where PGD, BIM and MIM are implemented for 20 iterations, with different attack budget  $\epsilon$  under  $l_{\infty}$ -norm constraint (The  $l_2$ -norm results can be found in our supplementary). For fair comparison, all methods equally employ the same ViT-B/16 [22] backbone, and all of BERT pretraining methods are conducted for 800 epochs, sharing the same training recipe. The pretrained models are finetuned in a supervised manner on ImageNet-1k for 100 epochs. The adversarial robustness we study here concentrates on the finetuned models.

From the quantitative results listed in Table 1 and Figure 3, we can surprisingly find that, though MAE [25] pretrained model performs better or comparably on clean images, it encounters much larger performance drop on adversarial samples. For example, in Table 1, the robust accuracy of MAE is significantly lower than the supervised baseline, *i.e.*, 11.9% lower on PGD, 13.0% lower on MIM, 12.7% lower on C&W and 4.5% lower on Auto-Attack. But as for BEiT [3] and PeCo [16], both of them can obtain comparable performance with the supervised pretraining baseline by contrast. For the self-supervised contrastive pretraining MoCo v3 [8], it has slightly better robustness than all BERT pretraining methods but lower performance on clean samples. The better robustness should be relevant with the strong augmentation strategies used in contrastive-style pretraining, which are not required or considered in vision BERT pretraining. To summarize, though also adopting the BERT-style pre-text task during pretraining, MAE has much weaker ability to defend against adversarial attack compared with BEiT and PeCo. We give the further analysis in Sec. 3.2 to discuss the possible reason that causes MAE’s degradation on adversarial robustness.

### 3.2. Causing Factor Analysis

To clarify the reason, we start with the difference among these three methods in their implementation. As illustrated in Sec. 2.1, it is clear that the reconstruction target of MAE is significantly different from BEiT or PeCo. Specifically, MAE chooses the raw pixel values as the BERT pretraining prediction target for masked patches, where the simple

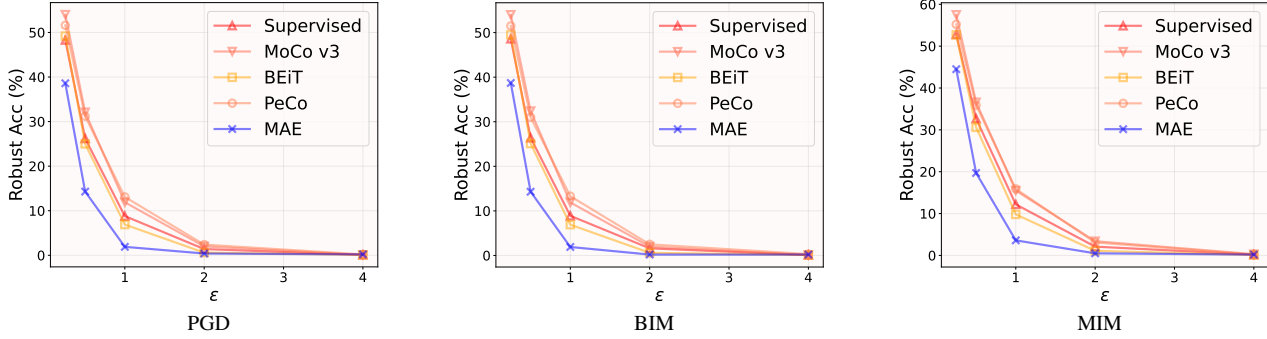


Figure 3. Comparison of adversarial robustness to various adversarial attacks ( $\epsilon = \epsilon/255$ ) among different BERT pretraining methods. All methods are equally implemented on ViT-B/16 backbone. **MAE shows significantly worse adversarial robustness than other methods.**

Method	Clean Acc (%)	Robust Acc (%) ( $\epsilon = 0.5/255$ )				Robust Acc (%) ( $\epsilon = 1.0/255$ )				
		PGD[36]	MIM[21]	BIM[33]	AA[11]	PGD[36]	MIM[21]	BIM[33]	AA[11]	C&W[6]
Supervised	81.8	26.2	32.7	26.4	6.1	8.8	12.2	8.9	0.1	18.6
MoCo v3 [8]	83.1	32.1	36.6	32.4	7.8	11.9	15.5	11.9	0.4	25.2
BEiT [3]	83.2	25.0	30.6	25.1	4.5	6.9	9.8	6.9	0.0	21.6
PeCo [16]	84.5	31.2	35.9	31.0	3.9	13.1	15.9	13.3	0.1	24.0
MAE [25]	83.6	<b>14.3</b>	<b>19.7</b>	<b>14.3</b>	<b>1.6</b>	<b>1.9</b>	<b>3.6</b>	<b>1.9</b>	<b>0.0</b>	<b>5.9</b>

Table 1. Comparison of adversarial robustness to various adversarial attacks among different BERT training methods. All methods are equally implemented on ViT-B/16 backbone, where MAE shows significantly worse robustness than other methods. Here “Clean Acc” and “Robust Acc” denote the top-1 classification accuracy on clean samples and adversarial samples, respectively.

mean square error loss on pixels is devised as the objective loss. Differently, the prediction target of BEiT or PeCo are discrete visual tokens that are encoded and quantized by a pretrained VQVAE, so these kind of targets are more semantic and information-denser than raw pixels. But how does the reconstruction target affects the adversarial robustness of pretrained models?

**Assumption 1.** *Predicting raw pixels of images makes MAE model more sensitive to medium-/high-frequency signals.*

It has been recently proved that Convolutional Neural Networks (CNNs) shows the strong ability [52] to perceive the high-frequency components of images that are almost imperceptible to human vision. This property, however, can be hardly observed in vision transformers (ViTs), since ViTs show less effective to capture the high-frequency components of images [41, 46, 2, 39]. In addition to the model architecture design, we hypothesis that the pretraining paradigm is also a crucial factor that affects the model sensitivity to the frequency-domain signals.

Following this clue, we conduct the experiment to study the classification performance of different BERT pretraining models on filtered images. These images are randomly sampled from the validation set of ImageNet-1k. We first try to remove the high-frequency components of images, where the images are converted into Fast Fourier Transform (FFT) space and only the central area (low frequency) is preserved with low pass filters of different sizes (similar to

the mask used in Figure 9). Figure 5 presents the relationship between top-1 classification accuracy and the radius used for low pass masking, from which we can observe that MAE’s performance is worse than other BERT pretraining methods and the supervised baseline. In other words, when the image discards its medium-/high-frequency information, it is more difficult for MAE to capture any discriminative representation from the image to support its recognition. It reveals that MAE relies relatively less on low-frequency components of images and behaves more sensitive to the removing of medium-/high-frequency signals.

We further visualize the frequency sensitivity of these BERT pretraining models. To this end, we convert images into the frequency spectrum with Discrete Cosine Transform (DCT) and directly compute the gradients on the spectrum through back-propagation (minimizing cross-entropy loss as the objective). Intuitively, the larger absolute value of the gradient on frequency domain indicates the greater sensitivity. We construct the gradient maps and visualize them with heatmaps in Figure 4, where it can be obviously observed that MAE requires more medium-/high-frequency signals (lower right triangle) for its prediction.

Therefore, we hypothesize that directly predicting pixels will make the MAE model more sensitive to medium-/high-frequency signals in the images. This claim is established on a simple fact that, raw pixels not only contain the basic information regarding the main content of images (usually presented as the low-frequency signals), but also includes

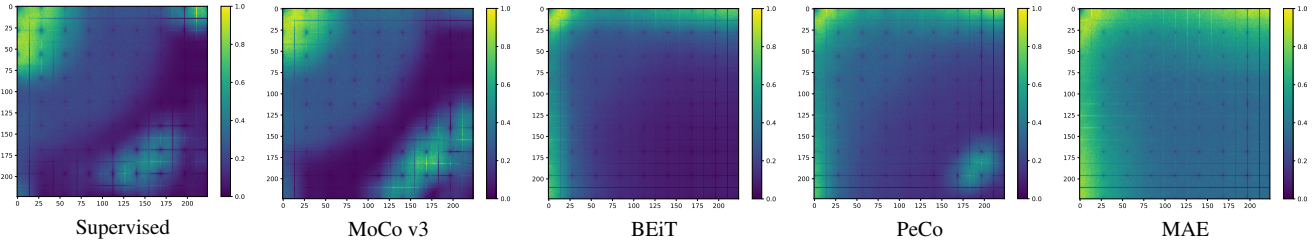


Figure 4. Frequency saliency maps of different pretrained models based on Discrete Cosine Transform (DCT). From the top left corner to the bottom right corner are low, medium and high frequencies. The brighter colors denote the greater sensitivity on the corresponding frequency component. **MAE shows the relatively stronger dependence on medium-/high-frequency components.**

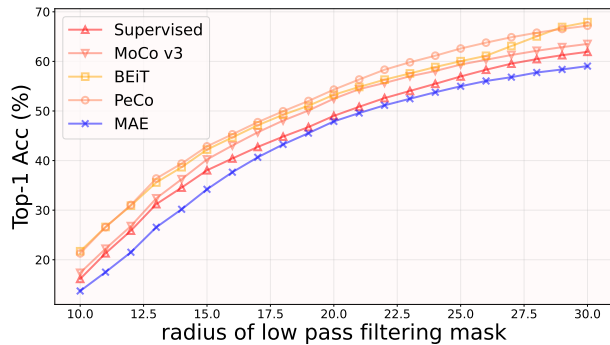


Figure 5. Top-1 accuracy comparison on images that are filtered by the low pass frequency mask among different pretraining methods. **MAE is apparently more sensitive to the removed medium-/high-frequency signal in input images.**

a lot of trivial details (*e.g.*, local textures or natural noises) that are not necessary for someone to recognize or describe the images. Accordingly, fitting on pixel values requires the neural learner to expend its capacity on precisely “remember” such details, *i.e.*, concentrating more on medium-/high-frequency signals that are relatively less helpful for image understanding.

On the other hand, though it is dataset-dependent, adversarial perturbation generated by regular attacks on natural images is generally a special kind of medium-/high-frequency signals. Through disrupting the medium-/high-frequency components of the input image, adversarial perturbation brings much more negative effects to MAE than other pretraining methods. Therefore, it is easier to fool MAE with adversarial samples, which is consistent with MAE’s degradation we observed in Table 1.

**Assumption 2.** *Predicting raw pixels of images helps MAE model become a more homogeneous vision learner.*

Some visualization results and similarity analysis tell us that, compared to other training methods, MAE is a more homogeneous vision learner. This unanticipated discovery hits us when we visualize the adversarial perturbations computed on vision transformers that are pretrained by different methods. Here we apply the  $\ell_\infty$ -norm PGD attack, where

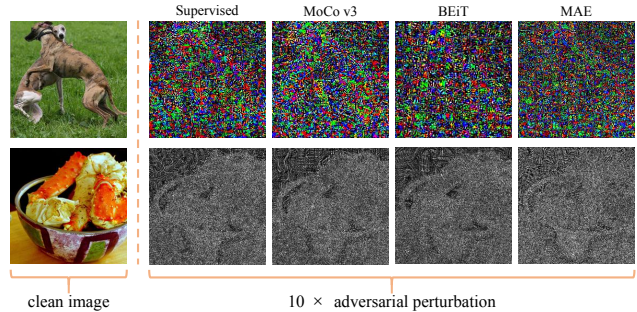


Figure 6. Visualization of adversarial perturbations computed on ViT-B/16 under different training methods, including both RGB images and gray images. **MAE’s perturbation is obviously more homogeneous across the spatial space than other methods.**

Method	Supervised	MoCo v3	BEiT	PeCo	MAE
Var ( $\times 10^{-4}$ )	1.77	1.84	1.69	1.66	<b>1.54</b>

Table 2. Average variance of adversarial perturbations on ViT-B/16 under different pretraining methods. **MAE’s adversarial perturbation has the lowest variance among pixels.**

the generation of adversarial examples heavily rely on the back-propagation of gradient information. Based on the spatial distribution and the local intensity of adversarial perturbations, it is reasonable to understand which patches are given more attention by the model. We provide two kinds of visualization in Figure 6: 1) the first row gives the raw RGB values of residuals; 2) the second row gives the gray images of residuals that are averaged over raw RGB values. The brighter part indicates the greater saliency.

From the visual results, we can easily distinguish the salient part for image recognition according to adversarial perturbations computed on the supervised baseline and MoCo v3. By contrast, we can hardly figure out the salient part on MAE’s perturbation. Furthermore, we calculate the average variance of these adversarial perturbations to quantitatively validate the observed vision homogeneity of MAE. From the results listed in Table 2, we find that MAE has the lowest perturbation variance among different pretraining methods. Additionally, we can see that vision

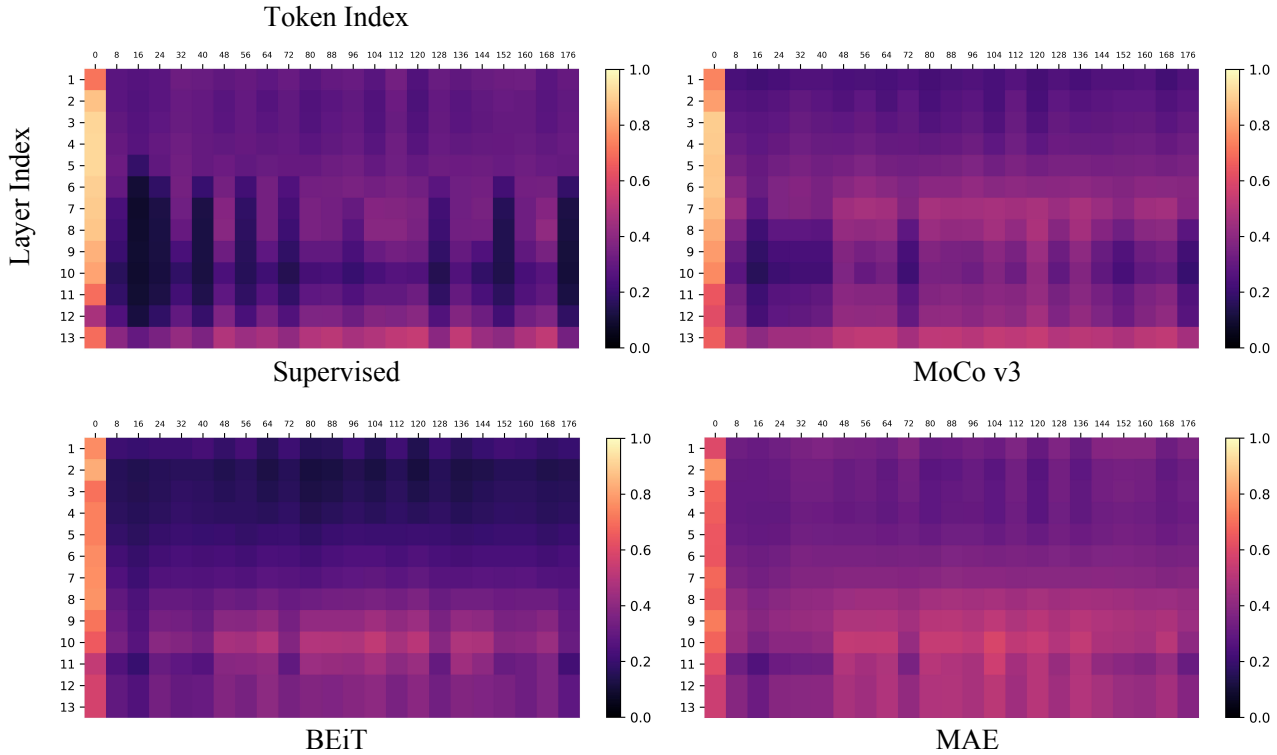


Figure 7. Centered kernel alignment (CKA) similarity calculated between spatial tokens and class tokens, where token index “0” means the class token and layer index “13” means the normalized activation. The brighter lattice denotes the larger CKA similarity, while the darker lattice denotes the weaker CKA similarity. **The tokens of both BEiT and MAE show the great correlation in each transformer layer.**

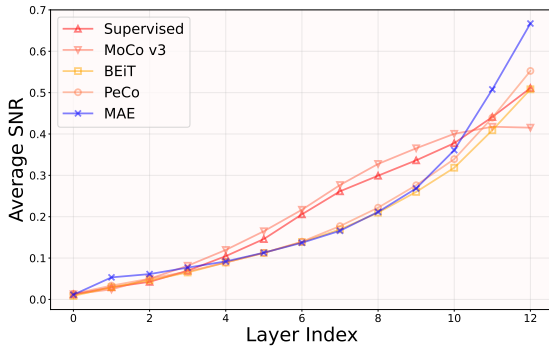


Figure 8. Average signal to noise (SNR) ratios for exploring which layers does adversarial perturbation affect more on different methods, where the larger SNR indicates the more feature deviation. **MAE shows much larger representation deviation on the last three layers when encountering adversarial perturbation.**

BERT pretraining methods are overall more homogeneous than previous methods in visual understanding.

Another interesting point is illustrated in Figure 8, where we intend to explore which layers are affected more by adversarial perturbations. To this end, we investigate the intermediate features in each layer of MAE encoder, including clean features that corresponds to clean input images and corrupted features from the corresponding adversarial samples. For each layer, we subsequently calculate the av-

erage signal to noise (SNR) value for each token between clean features and corrupted features, where the larger SNR value indicates the larger feature deviation. From Figure 8, it is clear that all of three BERT pretraining methods have similar SNR except for the last three layers, where MAE shows significantly larger deviation. The reason of this phenomenon is probably MAE pretrained model changes more on the last three layers during the finetuning stage, due to the gap between pretraining knowledge from raw pixels and finetuning knowledge from semantics. We also observe that MAE often adopts larger layer-wise learning rate decay rate than BEiT and PeCo, thus larger learning rates are used for pretrained transformer blocks during finetuning.

To verify our claim that vision BERT pretraining results in the more homogeneous attention on different regions of images, we further conduct the experiments to calculate the similarity of token representations. Formally, given two representations or activations  $\mathbf{X} \in \mathbb{R}^{k \times p_1}$  and  $\mathbf{Y} \in \mathbb{R}^{k \times p_2}$  computed on  $k$  images ( $p_1$  and  $p_2$  are the feature dimensions), we use Centered Kernel Alignment (CKA) [10, 32] of these two representations to measure their normalized similarity. Let  $\mathbf{K} = \mathbf{X}\mathbf{X}^\top$  and  $\mathbf{L} = \mathbf{Y}\mathbf{Y}^\top$ , we have

$$\text{CKA}(\mathbf{K}, \mathbf{L}) = \frac{\text{HSIC}(\mathbf{K}, \mathbf{L})}{\sqrt{\text{HSIC}(\mathbf{K}, \mathbf{K})\text{HSIC}(\mathbf{L}, \mathbf{L})}}, \quad (3)$$

here HSIC is Hilbert-Schmidt Independence Criterion [24].

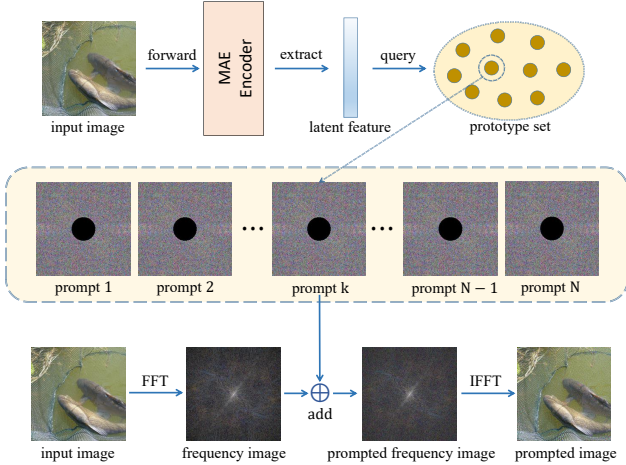


Figure 9. The pipeline of our frequency-domain prompting to boost the adversarial robustness of MAE during test time.

By calculating the CKA similarity between the class token and spatial tokens that are the output of each ViT self-attention block, we investigate the impacts contributed by different spatial tokens on the class token for the supervised baseline, MoCo v3, BEiT and MAE. We adopt the same ViT-B/16 backbone for these methods and test them on a randomly selected batch data (batch size is 256) from ImageNet-1k. With the class token as the reference, we can indirectly see the correlation between different spatial tokens. In Figure 7, we uniformly sample the spatial token indexes and compute the CKA similarity between the class token and the spatial tokens denoted by these indexes. Apparently, for both the supervised baseline and MoCo v3, the correlation among their spatial tokens is weak, since their similarities between each other are really disparate. By contrast, the spatial tokens of BEiT and MAE show more similarities with the class token, especially in 7<sup>th</sup>, 8<sup>th</sup> and 9<sup>th</sup> ViT self-attention block. This result also reveals that vision BERT pretraining, especially MAE, is a more homogeneous learner on image context. Note that, considering PeCo shares a similar CKA map as BEiT, we have not shown its result in Figure 7.

## 4. Boosting Adversarial Robustness of MAE

In this section, we propose a visual prompting-based solution to alleviate MAE’s weakness on adversarial robustness, *i.e.*, learning a set of cluster-specific visual prompts in frequency-domain to assist the image recognition of MAE. When the prompts are successfully tuned, we directly append an additional prompting stage before we input images to the MAE model, requiring no model change.

### 4.1. Frequency-Domain Visual Prompting

Visual prompting [29, 1, 26, 4] is an emerging topic that focuses to help large pretrained vision model [44, 57, 53, 56] adapt to downstream tasks through only changing the input space. Most of visual prompting methods consider their prompt design in the spatial domain of images. In our paper, it is the first time to explore the probability of prompting in the frequency domain of images. Formally, given a input image  $\mathbf{x}$ , we first convert it to frequency image  $\mathbf{x}_f$  using Fast Fourier Transform (FFT). The prompted image  $\mathbf{x}_p$  is derived by incorporating  $\mathbf{x}_f$  with a learnable patch  $\mathbf{p}$  and applying Inverse Fast Fourier Transform (IFFT), *i.e.*,

$$\mathbf{x}_p \triangleq \mathcal{T}_{ifft}(\mathbf{x}_f + \mathbf{m} \cdot \mathbf{p}(\mathbf{x})), \quad (4)$$

where  $\mathcal{T}_{ifft}$  is the IFFT transformation and  $\mathbf{p}(\mathbf{x})$  is the prompt selected for  $\mathbf{x}$ .  $\mathbf{m}$  denotes the mask that constrains the prompting in some particular frequency components (in our paper, we zero-out low-frequency areas with the radius of one-eighth the image size and preserve medium-/high-frequency areas by default, as shown in Figure 9).

### 4.2. Prompt Selection and Optimization

Instead of using a generic prompt across the dataset, we choose to cluster the training data of MAE model<sup>1</sup> into  $N$  clusters ( $N = 1,000$  in our default setting) by Mini-Batch K-means algorithm [45] and learn a set of cluster-specific prompts following “one prompt per cluster” principle. Formally, given the  $N$  clusters and their corresponding prototypes  $\{\mathbf{c}_i\}_{i=1}^N$ , we can assign input image  $\mathbf{x}$  with its frequency prompt  $\mathbf{p}(\mathbf{x})$  by extracting the latent feature of  $\mathbf{x}$  and searching the nearest prototype  $\mathbf{c}_k$ , *i.e.*,

$$\mathbf{p}(\mathbf{x}) = \mathbf{p}_k, \quad \text{s.t.} \quad k = \arg \min_i \|\mathbf{E}(\mathbf{x}) - \mathbf{c}_i\|_2^2. \quad (5)$$

For prompt optimization, our goal is to encourage each prompt  $\mathbf{p}_i$  to capture the common pattern of its corresponding cluster. Typically, we learn the prompt  $\mathbf{x}_p$  by minimizing the cross-entropy loss, *i.e.*,  $\mathcal{L}_{CE}(\mathbf{x}_p, y)$ , where  $y$  is the ground-truth label of  $\mathbf{x}_p$ . By default, we adopt AdamW optimizer [34] ( $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $eps = 10^{-8}$ ) with the base learning rate of  $10^{-3}$  and the total batch size of 1,024. The whole procedure is implemented on 4 RTX A6000 GPUs for 100 epochs, where we use the first 5 epochs for warmup and decay the learning rate with half-cycle cosine after warmup. Intuitively, the strong commonality among the data within a single cluster reduces the difficulty of prompt optimization and facilitates the quick convergence.

During test-time verification, we found that the engagement of frequency prompts will partially harm the model

<sup>1</sup>We use the checkpoint of pretrained ViT-B/16 provided in MAE official repository: <https://github.com/facebookresearch/mae>.

Method	CA (%)	Robust Acc (%)				
		PGD	MIM	BIM	C&W	AA
MAE	83.6	1.9	3.6	1.9	5.9	0.0
+ C-AVP [7]	50.2	31.3	34.0	30.7	OOM	OOM
+ Perceptual	83.9	9.3	15.1	9.1	14.6	2.3
+ ABP	83.7	8.4	11.6	8.8	7.2	0.0
<b>+ Ours*</b>	<b>82.7</b>	<b>53.8</b>	<b>53.3</b>	<b>53.2</b>	<b>68.6</b>	<b>14.0</b>
<b>+ Ours</b>	<b>82.7</b>	<b>69.3</b>	<b>69.2</b>	<b>69.3</b>	<b>56.8</b>	<b>15.9</b>
$\Delta$	<b>-0.9</b>	<b>+67.4</b>	<b>+65.6</b>	<b>+67.4</b>	<b>+50.9</b>	<b>+15.9</b>

Table 3. Comparison of different robust methods on MAE. Here “CA” reports Top-1 accuracy on clean images. “OOM” is “out of memory”. “Ours\*” means using differentiable prompt selection. We unify  $\epsilon = 1/255$  for PGD, MIM, BIM, AA attack.



Figure 10. Input images (top) and prompted images (bottom).

performance on clean images, though the robustness is improved. To address this, we adopt the voting result as the final prediction, which ensembles the logits of both the input image  $\mathbf{x}$  and the prompted image  $\mathbf{x}_p$  as  $\mathcal{M}(\mathbf{x}) + \lambda\mathcal{M}(\mathbf{x}_p)$ . To balance the trade-off, we simply set  $\lambda = 1$  by default.

### 4.3. Robustness Evaluation

We follow the same experimental setting in Section 3.1. Here we list several baseline methods for comparison:

**Class-wise Adversarial Visual Prompts (C-AVP) [7].** This solution devises class-wise visual prompts in the spatial domain of images, which are learned on adversarial examples to boost the robustness of classification models. Different from our method, it focuses on pixel-level prompts and learns one prompt per class in the dataset.

**Integrating Perceptual Loss.** This solution is to add an extra perceptual loss as the reconstruction supervision based on the analysis given in Section 3.2, thus encouraging MAE to reduce the dependency on high-frequency signals and comprehend more contextual information when learning the target semantics. This shares similar spirit to PeCo [16], but different from it, we directly use perceptual loss as BERT pretraining supervision rather than for VQ-VAE tokenizer training. To realize this, we leverage a pre-trained ViT-tiny to compute the multi-layer feature distance between MAE decoded tokens and original image tokens in

Dataset	Method	CA (%)	Robust Acc (%)				
			PGD	MIM	BIM	C&W	AA
DTD	MAE	<b>72.4</b>	39.4	39.8	39.8	43.4	5.2
	+ C-AVP	46.0	55.1	53.8	55.0	OOM	25.4
	+ Perceptual	75.9	42.6	43.1	42.7	49.4	11.2
	+ ABP	70.4	41.9	45.6	41.6	37.5	9.1
	<b>+ Ours</b>	72.2	<b>58.2</b>	<b>59.8</b>	<b>59.6</b>	<b>65.9</b>	<b>32.0</b>
CUB200	MAE	81.8	17.8	18.1	18.2	23.4	0.2
	+ C-AVP	43.5	39.1	38.8	39.0	OOM	12.7
	+ Perceptual	82.1	19.0	22.3	19.3	25.5	3.9
	+ ABP	78.0	24.3	26.1	24.5	35.8	5.4
	<b>+ Ours</b>	<b>81.9</b>	<b>62.8</b>	<b>66.8</b>	<b>64.8</b>	<b>64.6</b>	<b>23.0</b>

Table 4. Comparison of different robust methods on MAE. Here “CA” reports Top-1 accuracy on clean images. “OOM” is “out of memory”. We unify  $\epsilon = 4/255$  for PGD, MIM, BIM, AA attack.

Method	Clean Acc (%)	Robust Acc (%)				
		PGD	MIM	BIM	C&W	AA
Supervised	81.8	8.8	12.2	8.9	18.6	0.1
<b>+ Ours</b>	<b>81.1</b>	<b>45.0</b>	<b>45.6</b>	<b>44.9</b>	<b>60.1</b>	<b>11.9</b>
BEiT	83.2	6.9	9.8	6.9	21.6	0.0
<b>+ Ours</b>	<b>83.4</b>	<b>30.7</b>	<b>36.5</b>	<b>30.4</b>	<b>58.7</b>	<b>5.7</b>
PeCo	84.5	13.1	15.9	13.3	24.0	0.1
<b>+ Ours</b>	<b>83.9</b>	<b>35.4</b>	<b>37.5</b>	<b>35.3</b>	<b>56.3</b>	<b>8.9</b>

Table 5. Results of our method on other pretraining paradigms.

the masked region, so-called perceptual loss [31]. Here we empirically choose the features of  $3^{rd}$ ,  $6^{th}$ ,  $9^{th}$ ,  $12^{th}$  transformer block, and the final perceptual loss is the sum of feature loss of these four blocks. The final pretraining objective loss is the weighted sum of the original pixel-level mean square error loss  $L_{pix}$  and perceptual loss  $L_{perc}$ :

$$L_{total} = L_{pix} + \mu * L_{perc}. \quad (6)$$

where  $\mu = 1$  by default. More details are in supplementary.

**Adversarial BERT pretraining (ABP).** This solution is to naturally integrate adversarial training [35] into BERT pretraining, which we call “adversarial BERT pretrainin”. Specifically, we generate online adversarial perturbation upon unmasked patches by maximizing the reconstruction loss for masked patches, and conduct BERT pretraining with such online adversarial examples in a adversarial min-max manner. Similar to conventional adversarial learning, our intuition is to force the transformer model to learn on corrupted images that are adversarially generated by gradient-based attack in an online way. The difference is that we want to increase the robustness of pretraining rather than the supervised trained model. Additionally, we do not have any label information during BERT pretraining to guide the adversarial example generation. More details can be found in our supplementary.

Table 3 showcases the robustness of MAE models when



equipped with different training strategies. For fair comparison with baseline method C-AVP, our default setting uses arg min to select the prompt, so the gradient is only calculated within the classification branch. However, the gradient can not be back-propagated through arg min operation, thus we replace the non-differentiable arg min with a differentiable gumble-softmax [28] sampling and re-evaluate the performance. From Table 3, our method achieves the clear improvement on adversarial robustness while keeping limited clean accuracy sacrificed. Since our gumble-softmax version still achieves strong robustness (even better than the default version on C&W, it can also prove that the robustness of our method is not from the gradient obfuscated. Our visual prompts fulfill the medium-/high-frequency of images with the learned domain patterns, thus reducing the redundancy of these frequency components and narrowing the optimization space of adversarial perturbations.

#### 4.4. Additional Results

**More Datasets.** We further validate the proposed method on two classification datasets, *i.e.*, DTD [9] and CUB200 [51]. DTD is a dataset that collects 47 kinds of different textures and CUB200 contains images of various birds. On such datasets, visual prompts are expected to capture the more fine-grained representations than that are on ImageNet-1k. As shown in Table 4, our method keep the superiority on these fine-grained classification tasks, with its great improvement on robust accuracy.

**More Pretraining Methods.** We further apply our method on supervised pretraining, BEiT and PeCo. Table 5 shows that our method can also works for other pretraining paradigms but gains relatively less on MAE, *i.e.*, our medium-/high-frequency visual prompt better fits MAE.

#### 4.5. Ablation Study

Figure 10 showcases some instances of our prompted images. It indicates that though our experiments is focused on classification, our frequency-domain prompting has great potentials to well perform in other tasks like image segmentation or detection, since it does not use a visual-explicit way like spatial-domain patches used in [1] and rarely influences the visual quality. In Table 6, we also give the ablation studies to prove the validity about some of our designs: 1) we observe that prompting on the frequency domain obtains higher robust accuracy than that is on the spatial domain, frequency-domain prompting is more effective for boosting MAE’s robustness; 2) we show that learning cluster-specific prompts (“one cluster one prompt”) works better than learning a single prompt for all of images (“universal prompt”); 3) we show the validity of medium-/high-frequency prompting (*i.e.*, masking on low frequency in Figure 9) fits MAE better than low-/medium-frequency prompting (“masking high-freq.”); 4) we try to discard our logit ensemble and

Setting	CA (%)	Robust Acc (%)				
		PGD	MIM	BIM	C&W	AA
<i>spatial domain vs. frequency domain</i>						
full spatial.	78.3	10.5	11.6	10.4	16.3	0.9
full freq.	80.5	26.6	28.1	26.7	19.7	3.2
<i>universal prompt vs. cluster-specific prompt</i>						
universal	78.3	35.2	35.9	35.3	24.0	3.8
cluster-specific	82.7	69.3	69.2	69.3	56.8	15.9
<i>masking low frequency vs. masking high frequency</i>						
mask high-freq.	83.7	2.6	6.0	2.7	4.8	0.0
mask low-freq.	82.7	69.3	69.2	69.3	56.8	15.9
<i>whether using logit ensemble</i>						
w/o ensemble	80.5	70.3	70.4	70.4	51.8	16.1
w/ ensemble	82.7	69.3	69.2	69.3	56.8	15.9
<i>Number <math>N</math> of clusters</i>						
$N = 500$	81.3	68.6	68.8	68.8	53.1	12.2
$N = 1,500$	82.3	62.4	63.0	62.7	47.1	9.2
$N = 1,000$	82.7	69.3	69.2	69.3	56.8	15.9

Table 6. Ablation studies to verify the effect of different designs. The highlight results corresponds to our default setting.

find the minor gains on robust accuracy but non-negligible degradation on clean accuracy, proving that the ensemble is necessary to benefit the clean performance; 5) we ablate the number  $N$  of clusters and choose  $N = 1,000$  in our experiment considering both robustness and clean performance.

## 5. Conclusions

This paper analyzes the adversarial robustness of vision BERT pretraining. We observe that the main design difference of existing vision BERT pretraining methods, *i.e.*, BERT pretraining target designs, is the key impact factor. More specifically, adopting the raw pixels of masked patches as the BERT pretraining target (*e.g.*, MAE) degrades the adversarial robustness of the model than the counterparts predicting the semantic context (*e.g.*, BEiT, PeCo). Through the detailed analysis, we hypothesize that directly predicting raw pixels makes the MAE model more sensitive to medium-/high-frequency signals of images. To improve the robustness of MAE, a simple yet effective solution have been proposed based on our analysis, *i.e.*, learning frequency-domain visual prompts that contains dataset-extracted knowledge to assist the test-time model prediction. We surprisingly find that this solution can greatly boost the adversarial robustness of MAE while maintaining its performance on clean samples. We hope that the analysis of our paper can motivate more study in this direction.

**Acknowledgement.** This work was supported in part by the Natural Science Foundation of China under Grant U20B2047, 62072421, 62002334, 62102386 and 62121002, and by Fundamental Research Funds for the Central Universities under Grant WK210000011.

## References

- [1] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Exploring visual prompts for adapting large-scale models. *arXiv preprint arXiv:2203.17274*, 2022. 7, 9
- [2] Jiawang Bai, Li Yuan, Shu-Tao Xia, Shuicheng Yan, Zhifeng Li, and Wei Liu. Improving vision transformers by revisiting high-frequency components. In *ECCV*, 2022. 4
- [3] Hangbo Bao, Li Dong, and Furu Wei. Beit: BERT pre-training of image transformers. *CoRR*, abs/2106.08254, 2021. 1, 2, 3, 4
- [4] Amir Bar, Yossi Gandelsman, Trevor Darrell, Amir Globerson, and Alexei A Efros. Visual prompting via image inpainting. *arXiv preprint arXiv:2209.00647*, 2022. 7
- [5] Huanyu Bian, Dongdong Chen, Kui Zhang, Hang Zhou, Xiaoyi Dong, Wenbo Zhou, Weiming Zhang, and Nenghai Yu. Adversarial defense via self-orthogonal randomization super-network. *Neurocomputing*, 2021. 3
- [6] Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy (SP)*, 2017. 3, 4
- [7] Aochuan Chen, Peter Lorenz, Yuguang Yao, Pin-Yu Chen, and Sijia Liu. Visual prompting for adversarial robustness. *arXiv preprint arXiv:2210.06284*, 2022. 8
- [8] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021. 3, 4
- [9] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *CVPR*, 2014. 9
- [10] Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *J. Mach. Learn. Res.*, 2012. 6
- [11] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *ICML*, 2020. 3, 4
- [12] Jiequan Cui, Shu Liu, Liwei Wang, and Jiaya Jia. Learnable boundary guided adversarial training. In *ICCV*, 2021. 3
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 3
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. 1, 2
- [15] Guneet S. Dhillon, Kamyar Azizzadenesheli, Zachary C. Lipton, Jeremy Bernstein, Jean Kossaifi, Aran Khanna, and Animashree Anandkumar. Stochastic activation pruning for robust adversarial defense. In *ICLR*, 2018. 3
- [16] Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. Peco: Perceptual codebook for BERT pre-training of vision transformers. *CoRR*, abs/2111.12710, 2021. 1, 2, 3, 4, 8
- [17] Xiaoyi Dong, Dongdong Chen, Jianmin Bao, Chuan Qin, Lu Yuan, Weiming Zhang, Nenghai Yu, and Dong Chen. Greedyfool: Distortion-aware sparse adversarial attack. *Advances in Neural Information Processing Systems*, 33:11226–11236, 2020. 3
- [18] Xiaoyi Dong, Dongdong Chen, Hang Zhou, Gang Hua, Weiming Zhang, and Nenghai Yu. Self-robust 3d point recognition via gather-vector guidance. In *CVPR*, 2020. 3
- [19] Xiaoyi Dong, Jiangfan Han, Dongdong Chen, Jiayang Liu, Huanyu Bian, Zehua Ma, Hongsheng Li, Xiaogang Wang, Weiming Zhang, and Nenghai Yu. Robust superpixel-guided attentional adversarial attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12895–12904, 2020. 3
- [20] Xiaoyi Dong, Yinglin Zheng, Jianmin Bao, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, et al. Maskclip: Masked self-distillation advances contrastive language-image pretraining. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2023)*, 2023. 1
- [21] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *CVPR*, 2018. 3, 4
- [22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3
- [23] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 3
- [24] Arthur Gretton, Kenji Fukumizu, Choon Hui Teo, Le Song, Bernhard Schölkopf, and Alexander J. Smola. A kernel statistical test of independence. In *NeurIPS*, 2007. 6
- [25] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *CoRR*, abs/2111.06377, 2021. 1, 2, 3, 4
- [26] Qidong Huang, Xiaoyi Dong, Dongdong Chen, Weiming Zhang, Feifei Wang, Gang Hua, and Nenghai Yu. Diversity-aware meta visual prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 7
- [27] Qidong Huang, Xiaoyi Dong, Dongdong Chen, Weiming Zhang, and Nenghai Yu. Shape-invariant 3d adversarial point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3
- [28] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. 9
- [29] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, 2022. 7
- [30] Ziyu Jiang, Tianlong Chen, Ting Chen, and Zhangyang Wang. Robust pre-training by adversarial contrastive learning. In *NeurIPS*, 2020. 3
- [31] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 8

- [32] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey E. Hinton. Similarity of neural network representations revisited. In *ICML*, 2019. 6
- [33] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016. 3, 4
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 7
- [35] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. 8
- [36] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 2, 3, 4
- [37] Xiaofeng Mao, Gege Qi, Yuefeng Chen, Xiaodan Li, Ranjie Duan, Shaokai Ye, Yuan He, and Hui Xue. Towards robust vision transformer. In *CVPR*, 2022. 3
- [38] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Fahad Khan, and Fatih Porikli. On improving adversarial transferability of vision transformers. In *ICLR*, 2022. 3
- [39] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *NeurIPS*, 2021. 4
- [40] Nicolas Papernot and Patrick D. McDaniel. Extending defensive distillation. *CoRR*, abs/1705.05264, 2017. 3
- [41] Namuk Park and Songkuk Kim. How do vision transformers work? In *ICLR*, 2022. 4
- [42] Sayak Paul and Pin-Yu Chen. Vision transformers are robust learners. In *AAAI*, 2022. 3
- [43] Yao Qin, Chiyuan Zhang, Ting Chen, Balaji Lakshminarayanan, Alex Beutel, and Xuezhi Wang. Understanding and improving robustness of vision transformers through patch-based negative augmentation. In *NeurIPS*, 2021. 3
- [44] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 7
- [45] David Sculley. Web-scale k-means clustering. In *WWW*, 2010. 7
- [46] Rulin Shao, Zhouxing Shi, Jinfeng Yi, Pin-Yu Chen, and Cho-Jui Hsieh. On the adversarial robustness of vision transformers. *arXiv preprint arXiv:2103.15670*, 2021. 3, 4
- [47] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014. 3
- [48] Shiyu Tang, Ruihao Gong, Yan Wang, Aishan Liu, Jiakai Wang, Xinyun Chen, Fengwei Yu, Xianglong Liu, Dawn Song, Alan Yuille, et al. Robustart: Benchmarking robustness on architecture design and training techniques. *arXiv preprint arXiv:2109.05211*, 2021. 3
- [49] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian J. Goodfellow, Dan Boneh, and Patrick D. McDaniel. Ensemble adversarial training: Attacks and defenses. In *ICLR*, 2018. 3
- [50] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *NeurIPS*, 2017. 1
- [51] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 9
- [52] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *CVPR*, 2020. 4
- [53] Junke Wang, Dongdong Chen, Zuxuan Wu, Chong Luo, Lu-owei Zhou, Yucheng Zhao, Yujia Xie, Ce Liu, Yu-Gang Jiang, and Lu Yuan. Omnivl: One foundation model for image-language and video-language tasks. In *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS 2022)*, 2022. 7
- [54] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan L. Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. *CoRR*, abs/2112.09133, 2021. 2
- [55] Zhipeng Wei, Jingjing Chen, Micah Goldblum, Zuxuan Wu, Tom Goldstein, and Yu-Gang Jiang. Towards transferable adversarial attacks on vision transformers. In *AAAI*, 2022. 3
- [56] Zejia Weng, Xitong Yang, Ang Li, Zuxuan Wu, and Yu-Gang Jiang. Open-vclip: Transforming clip to an open-vocabulary video model via interpolated weight optimization. In *ICML*, 2023. 7
- [57] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. 7
- [58] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P. Xing, Laurent El Ghaoui, and Michael I. Jordan. Theoretically principled trade-off between robustness and accuracy. In *ICML*, 2019. 3
- [59] Hang Zhou, Dongdong Chen, Jing Liao, Kejiang Chen, Xiaoyi Dong, Kunlin Liu, Weiming Zhang, Gang Hua, and Nenghai Yu. Lg-gan: Label guided adversarial network for flexible targeted attack of point cloud based deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10356–10365, 2020. 3
- [60] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan L. Yuille, and Tao Kong. ibot: Image BERT pre-training with online tokenizer. *CoRR*, abs/2111.07832, 2021. 1