Improving Vision Transformers by Revisiting High-frequency Components

Jiawang Bai¹, Li Yuan^{2,5, \boxtimes}, Shu-Tao Xia^{1,5, \boxtimes}, Shuicheng Yan⁴, Zhifeng Li^{3, \boxtimes}, and Wei Liu^{3, \boxtimes}

¹Tsinghua Shenzhen International Graduate School, Tsinghua University ²School of ECE at Peking University ³Data Platform, Tencent ⁴ Sea AI Lab ⁵Peng Cheng Laboratory bjw19@mails.tsinghua.edu.cn; yuanli-ece@pku.edu.cn; xiast@sz.tsinghua.edu.cn; yansc@sea.com; michaelzfli@tencent.com; wl2223@columbia.edu

Abstract. The transformer models have shown promising effectiveness in dealing with various vision tasks. However, compared with training Convolutional Neural Network (CNN) models, training Vision Transformer (ViT) models is more difficult and relies on the large-scale training set. To explain this observation we make a hypothesis that ViT models are less effective in capturing the high-frequency components of images than CNN models, and verify it by a frequency analysis. Inspired by this finding, we first investigate the effects of existing techniques for improving ViT models from a new frequency perspective, and find that the success of some techniques (e.g., RandAugment) can be attributed to the better usage of the high-frequency components. Then, to compensate for this insufficient ability of ViT models, we propose HAT, which directly augments high-frequency components of images via adversarial training. We show that HAT can consistently boost the performance of various ViT models (e.g., +1.2% for ViT-B, +0.5% for Swin-B), and especially enhance the advanced model VOLO-D5 to 87.3% that only uses ImageNet-1K data, and the superiority can also be maintained on out-of-distribution data and transferred to downstream tasks. The code is available at: https://github.com/jiawangbai/HAT.

1 Introduction

Recently, transformer models have shown high effectiveness in various vision tasks and attracted growing attention. The pioneering work is Vision Transformer (ViT) [20], which is a full-transformer architecture directly inherited from natural language processing [55] but taking raw image patches as input. After that, many ViT variants [21,40,13,32,59,69,60] have been proposed and achieved competitive performance with Convolutional Neural Network (CNN) models. Though promising in vision tasks, ViT models suffer training difficulty and require significantly more training samples [20] compared with CNN models.

 $[\]boxtimes$ Corresponding authors.

 $\mathbf{2}$



(b) High-pass Filtering

Fig. 1: Comparison of ViT-B, ResNet-50, ResNet-101, and ViT-B(+KD) on lowand high-pass filtered validation set with different filter sizes. ViT-B is the base ViT model taking as input a sequence of 16×16 patches. KD denotes knowledge distillation, where the teacher model is a RegNetY-16GF [46] following [54]. The top-1 accuracy of ViT-B, ResNet-101, ResNet-50, and ViT-B (+KD) on the ImageNet validation set is 82.0%, 79.8%, 81.6%, and 83.6%, respectively.

One reason for this difficulty may be that ViT models can not effectively exploit the local structures as they split an image to a sequence of patches and model their dependencies with the self-attention mechanism [68,44]. In contrast, CNN models can effectively extract local features within the receptive fields [3,23] with convolution operation. From some previous studies [6,17,53], the local structures (e.g., edges and lines) are more related to the high-frequency components of the images. We then naturally make such a hypothesis: ViT models are less effective in capturing the high-frequency components of images than CNN models.

To verify our hypothesis, we use the discrete Fourier transform (DFT) to decompose the original images into the low- and high-frequency components and evaluate the model performance on them respectively [67,57]. Figure 1 shows a comparison between ViT-B [20] and ResNet-50 [28], where a larger filter size for the low- and high-pass filtering means more low- and high-frequency components, respectively. In our experiments, ViT-B has a higher top-1 accuracy on the original ImageNet validation set (82.0% vs. 79.8%) and a larger model size (86.6M vs. 25.6M). We can see that ViT-B performs better than ResNet-50 on the low-frequency components (Figure 1 (a)), but worse than ResNet-50 on the high-frequency components (Figure 1 (b)), which supports our hypothesis. Motivated by the above observation, we study existing techniques for ViT models from a frequency perspective, including knowledge distillation [37,54], architecture design [21,40,13], and data augmentations [54,44], and provide some useful insights which are beneficial to improving performance of ViT models. Through extensive experiments, we find that i) knowledge distillation is helpful to a ViT model using a CNN teacher in capturing high-frequency components of the images; ii) compared to the original ViT [20], some advanced architectures utilizing convolutional-like operation [21] or multi-scale feature maps [40,13] can more effectively exploit high-frequency components of the images; iii) RandAugment [14] is more helpful for catching high-frequency components of the images than CutMix [70] and Mixup [73].

Furthermore, we propose to compensate for the insufficient capacity of ViT models in capturing the high-frequency components of the images by directly augmenting the high-frequency components via adversarial training [43,72]. Specifically, we craft adversarial examples by altering clean images with high-frequency perturbations, and jointly train ViT models over clean images and adversarial examples. Our results indicate that this training strategy improves the performance of the ViT model by compensating for its ability to capture the high-frequency components of the images. Moreover, since adversarial perturbations can naturally influence the high-frequency components in our case, we directly use adversarial perturbations without high-frequency limitation, resulting in a simple but effective method, named HAT, standing for improving ViT models on the high-frequency components via adversarial training. Note that HAT does not bring extra complexity during the inference stage or alter the model architecture.

Our main contributions are summarized as follows:

- Based on our frequency analysis, we validate that compared to CNN models, ViT models are less effective in capturing the high-frequency components of images, which may lead to the difficulty of training ViT models.
- We analyze the effects of existing techniques for improving the performance of ViT models from a frequency perspective.
- We propose HAT, which improves the performance of ViT models by influencing the high-frequency components of images directly.
- Our results on ImageNet classification and out-of-distribution data demonstrate the superiority of HAT. We also find that pre-trained models with HAT are beneficial to downstream tasks.

2 Related Work

Transformer Models in Vision Tasks. Transformer models [55] entirely rely on the self-attention mechanism to build long-distance dependencies, which have achieved great success in almost all natural language processing tasks [19,39,4]. Vision Transformer (ViT) [20] is one of the earlier attempts to introduce transformer models into vision tasks, which applies a pure transformer architecture on non-overlapping image patches for image classification and has achieved state-of-the-art accuracy. Since ViT models excel at capturing spatial information, they

have also been extended to more challenging tasks, including object detection [7,77,15], segmentation [52,50], image enhancement [9,66], and video processing [75,71,61]. Besides, many efforts have been devoted to designing new ViT architectures [21,40,13,32,59,69,60]. For example, Liu *et al.* [40] presented a hierarchical architecture with shifted window based attention that can efficiently extract multi-scale features; Yuan *et al.* [69] introduced outlook attention to efficiently encode finer-level features and contexts into tokens.

Training Strategies for ViT Models. It is shown that training ViT models is more challenging than training CNN models, and requires large-scale datasets (*e.g.*, ImageNet-22K [18] and JFT-300M [51]) to perform pre-training [20]. To enable ViT to be effective on the smaller ImageNet-1K dataset [18], many training strategies have been explored. In [54,49], applying strong data augmentation and model regularization makes a quick solution to this problem. Among them, CutMix [70], Mixup [73], and RandAugment [14] are proven to be particularly helpful [54]. Besides, some customized augmentations for training ViT models are presented [10,59]. Utilizing a trained CNN teacher, knowledge distillation (KD) [54,37] can significantly boost the performance of ViT models. There are also some works solving this problem by using a better optimization strategy, such as promoting patch diversification [24] and sharpness-aware minimizer [11]. Unlike these works, we focus on directly compensating for the ability of ViT models in capturing the high-frequency components for better performance.

3 Revisiting ViT Models from a Frequency Perspective

To investigate ViT models from a frequency perspective, we use the discrete Fourier transform (DFT) to evaluate the model performance on certain frequency components of test samples [67]. Let $\boldsymbol{x} \in \mathbb{R}^{H \times W}$ (omitting the dimension of image channels) and $\boldsymbol{y} \in \mathbb{R}^C$ represent an image in the spatial domain and its label vector, where C is the number of classes. We transform \boldsymbol{x} to the frequency spectrum by the DFT $\mathcal{F} : \mathbb{R}^{H \times W} \to \mathbb{C}^{H \times W}$ and transform signals of the image from frequency back to the spatial domain by the inverse DFT $\mathcal{F}^{-1} : \mathbb{C}^{H \times W} \to \mathbb{R}^{H \times W}$. In this work, the low-frequency components are shifted to the center of the frequency spectrum.

For a mask $\boldsymbol{m} \in \{0, 1\}^{H \times W}$, the low-pass filtering \mathcal{M}_l^S and high-pass filtering \mathcal{M}_h^S with the filter size S are formally defined as:

$$\mathcal{M}_{l}^{S}(\boldsymbol{x}) = \mathcal{F}^{-1}(\boldsymbol{m} \odot \mathcal{F}(\boldsymbol{x})), \text{ where } \boldsymbol{m}_{i,j} = \begin{cases} 1, \text{ if } \min(|i - \frac{H}{2}|, |j - \frac{W}{2}|) \leq \frac{S}{2} \\ 0, \text{ otherwise} \end{cases}, \quad (1)$$

$$\mathcal{M}_{h}^{S}(\boldsymbol{x}) = \mathcal{F}^{-1}(\boldsymbol{m} \odot \mathcal{F}(\boldsymbol{x})), \text{ where } \boldsymbol{m}_{i,j} = \begin{cases} 0, \text{ if } \min(|i - \frac{H}{2}|, |j - \frac{W}{2}|) \leqslant \frac{\min(H, W) - S}{2} \\ 1, \text{ otherwise} \end{cases}$$
(2)

where \odot is element-wise multiplication and $m_{i,j}$ denotes the value of m at position (i, j). For images containing multiple color channels, the filtering operates

4



Fig. 2: Comparison of ViT-B, ConViT-B, Twins-SVT-L, and Swin-B on low- and high-pass filtered validation set with different filter sizes. The top-1 accuracy of ViT-B, ConViT-B, Twins-SVT-L, and Swin-B on the ImageNet validation set is 82.0%, 82.3%, 83.7%, and 83.5%, respectively.

on each channel independently. To make a comprehensive analysis, we evaluate various ViT architectures and training strategies with different filter sizes based on the ImageNet validation set. We provide the visualized examples in Figure 1. Comparison of ViT and CNN Models. Firstly, we compare the performance of ViT-B with ResNet-50 and ResNet-101, which are trained with the same data augmentations. The plots in Figure 1(a) show that ViT-B surpasses ResNet on the low-frequency components of images. However, although ViT-B achieves a higher accuracy (82.0%) than ResNet-50 (79.8%) and ResNet-101 (81.6%) on the original ImageNet validation set, its performance is lower than CNN models on the high-frequency components of images, as shown in Figure 1(b). This observation indicates that ViT models can capture the global contexts effectively, but fails to well leverage local details compared to CNN models. It may be because cascading self-attention blocks in ViT models is equivalent to repeatedly applying a low-pass filter, corresponding to the theoretical justification in [58], while CNN models utilizing convolution operations behave like a series of highpass filters [44] to catch more high-frequency components [57].

We further study the distillation method introduced in [54], *i.e.*, transferring the learned knowledge in a ViT model using a CNN teacher. We use a RegNetY-16GF model [46] as a teacher with the hard-label distillation, and adopt all settings in [54]. The results in Figure 1 show that the improvement of KD (from 82.0% to 83.6%) is primarily attributed to the stronger ability to exploit the high-frequency components of images. It also confirms that there is a gap between the abilities of ViT and CNN models in capturing the high-frequency components.

Various ViT Architectures. Recently, various ViT architectures are proposed and show excellent results [21,40,13]. We compare ViT-B with three advanced architectures, including ConViT-B [21], Twins-SVT-L [40], and Swin-B [13], with a similar model size, and present a reason for the success of these architectures from the frequency perspective. As shown in Figure 2, all architectures perform similarly on the low-frequency components, while three advanced architectures achieve higher accuracy than ViT-B on the high-frequency components. Our results also provide evidence for the effects of the proposed components of these $\mathbf{6}$



Fig. 3: Comparison of vanilla training and three data augmentations on low- and high-pass filtered validation set with different filter sizes. The top-1 accuracy of Vanilla, +CutMix, +Mixup, and +RandAugment on the ImageNet validation set is 76.7%, 80.8%, 79.9%, and 78.8%, respectively.

recent architectures. Specifically, the convolutional-like operation in ConViT and multi-scale feature maps in Swin Transformer and Twins-SVT can help the vision transformer capture the high-frequency components.

Data Augmentations. As demonstrated in recent works [54,44], training ViT models relies heavily on strong data augmentation. Compared to the vanilla training, the improvements of the commonly used augmentations are significant, including CutMix [70] (+4.1%), Mixup [73] (+3.2%), and RandAugment [14] (+2.1%). We make a comparison between the effects of these three augmentations from the frequency perspective. The results are shown in Figure 3. We can see that the ranking of these three augmentations w.r.t. improvements they bring is CutMix > Mixup > RandAugment on the low-frequency components. However, on the high-frequency components, the case is opposite: RandAugment > Mixup > CutMix. Our observation reveals that CutMix can help ViT models leverage the global context information of an image by removing a random region and replacing it with a patch from another image. Moreover, it also indicates that the transformations used in RandAugment can force the trained model to pay more attention to high-frequency information.

4 The Proposed Method

In this section, we firstly describe the proposed HAT, and then demonstrate its effects on ViT models via a case study.

4.1 Adversarial Training with High-frequency Perturbations

As demonstrated by the analysis in Section 3, the ability of ViT models to capture the high-frequency components is limited, and compensating for this limitation is a key to boosting their performance. Therefore, different from previous data augmentation methods, we propose to directly augment the high-frequency components during the training stage. We alter the high-frequency components of training images by adding adversarial perturbations and training ViT models on these altered images. It corresponds to adversarial training [43,72] with high-frequency perturbations and is stated formally below.

Given a ViT model f with the weights θ , $f_{\theta}(x)$ denotes its softmax output of the input sample x. Inspired by the min-max formulation of adversarial training [43], the objective function of adversarial training with high-frequency perturbations is as follows:

$$\mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}}\left[L(\boldsymbol{\theta},\boldsymbol{x},\boldsymbol{y}) + \max_{||\boldsymbol{\delta}||_{\infty}\leqslant\epsilon} \left(\alpha L(\boldsymbol{\theta},\boldsymbol{x}+\mathcal{M}_{h}^{S}(\boldsymbol{\delta}),\boldsymbol{y}) + \beta L_{kl}(\boldsymbol{\theta},\boldsymbol{x}+\mathcal{M}_{h}^{S}(\boldsymbol{\delta}),\boldsymbol{x})\right)\right],\tag{3}$$

where ϵ denotes the maximum perturbation strength. $L(\boldsymbol{\theta}, \boldsymbol{x}, \boldsymbol{y}) = \text{CE}(f_{\boldsymbol{\theta}}(\boldsymbol{x}), \boldsymbol{y})$ and $L_{kl}(\boldsymbol{\theta}, \boldsymbol{x}_1, \boldsymbol{x}_2) = \frac{1}{2}[\text{KL}(f_{\boldsymbol{\theta}}(\boldsymbol{x}_1), f_{\boldsymbol{\theta}}(\boldsymbol{x}_2)) + \text{KL}(f_{\boldsymbol{\theta}}(\boldsymbol{x}_2), f_{\boldsymbol{\theta}}(\boldsymbol{x}_1))]$, where $\text{CE}(\cdot)$

and KL(·) calculate the cross-entropy and the Kullback-Leibler divergence, respectively. α and β are two hyperparameters. We use the high-pass filtering \mathcal{M}_h^S with a given filter size to limit the perturbations in the high-frequency domain. Our experiments in Section 4.2 demonstrate that optimizing Eq. (3) can compensate for the ability of the ViT model to capture the high-frequency components of images and thus improve its performance.



Fig. 4: Heat maps of Fourier spectrum for natural images and adversarial perturbations. They are obtained by averaging over a batch of data.

Then, we notice that adversarial perturbations are naturally imposed on the high-frequency components in our case [16 41] It

high-frequency components in our case [16,41]. It is validated in Figure 4 that compared to natural images, adversarial perturbations show higher concentrations in the high-frequency domain. Therefore, we directly use full-frequency adversarial perturbations in our HAT with the below objective:

$$\mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}}\left[L(\boldsymbol{\theta},\boldsymbol{x},\boldsymbol{y}) + \max_{||\boldsymbol{\delta}||_{\infty}\leqslant\epsilon} \left(\alpha L(\boldsymbol{\theta},\boldsymbol{x}+\boldsymbol{\delta},\boldsymbol{y}) + \beta L_{kl}(\boldsymbol{\theta},\boldsymbol{x}+\boldsymbol{\delta},\boldsymbol{x})\right)\right].$$
(4)

The inner maximization in Eq. (4) can be solved by project gradient descent (PGD) for K steps [43]. Different from the standard PGD, following [76,22], we accumulate the gradients of the model weights in each PGD step, and update the parameters at once with the accumulated gradients. In this way, the perturbations in each PGD step can be used for training. This procedure is detailed in Algorithm 1 in Appendix C. Besides, to address the mismatched distribution between clean images and adversarial examples [65], we perform adversarial training in some initial epochs (200 epochs in our setting) and train normally in the rest epochs.

Table 1: Comparison of the baseline and adversarial training (AT) with three types of perturbations, where the case of using the full-frequency perturbations corresponds to the proposed HAT.

8

Training Strategy	Top-1 ACC (%)
Baseline	82.0
+AT (Low-freq. Pert.)	81.9
+AT (High-freq. Pert.)	83.0
+AT (Full-freq. Pert.)	83.2



Fig. 5: Comparison of the baseline and adversarial training (AT) with three types of perturbations on low- and high-pass filtered validation set with different filter sizes.

4.2 A Case Study using ViT-B

To illustrate how the proposed method influences the ViT models, we conduct a case study using ViT-B on ImageNet. For training ViT-B, we adopt the hyperparameters in [25] in all cases. Without considering our adversarial training, these hyper-parameters result in a strong baseline with a 82.0% top-1 accuracy. For our method, we set $\epsilon = 2/255$, K = 3, and $\eta = 1/255$. The parameters α and β are fixed at 3 and 0.01, respectively. We compare three types of adversarial perturbations: low-frequency perturbations with the filter size 10, high-frequency perturbations with the filter size 10, and full-frequency perturbations.

The results on the ImageNet validation set are shown in Table 1. We can see that adversarial training with the high-frequency perturbations brings 1.0% gains over the baseline. Figure 5 reveals the reason for that: the ability of ViT-B to capture the high-frequency components of images is stronger than the baseline, which exactly confirms our expectation. In contrast, there is no improvement for the case of using low-frequency perturbations. For adversarial training with the full-frequency perturbations, without the high-filter operation and setting the filter size, it is more simple, but can also improve the performance of ViT-B. This case study illustrates that using the full-frequency perturbations is a reasonable choice for our HAT, which will be further verified in our below experiments.

5 Experiments

5.1 Experimental Setup

We evaluate the proposed method on ImageNet [18]. Our code is implemented based on PyTorch [45] and the timm library [62]. We conduct experiments on various model architectures: three variants of ViT [20] (ViT-T, ViT-S, and ViT-B with 16×16 input patch size) following [54], three variants of Swin Transformer [40] (Swin-T, Swin-S, and Swin-B), and two variants of VOLO [69] (VOLO-D1 and VOLO-D5). "T", "S", and "B" denote tiny, small, and base model sizes, respectively. Following the standard training schedule, we train all models on the ImageNet-1K training set for 300 epochs with strong data augmentation

Model	Params	FLOPs	Test Size	Top-1	Real Top-1	V2 Top-1
ViT-T	5 7M	1.60	994	72.2	80.0	60.1
+HAT	5.7 M	1.0G	224	73.3	81.1	61.0
ViT-S	22.1M	4 7C	224	80.1	85.7	68.2
+HAT	22.111	4.70	224	80.9	86.6	70.0
ViT-B	86 GM	17.6C	224	82.0	87.1	71.0
+HAT	00.0101	17.00	224	83.2	87.9	72.6
Swin-T	96 SM	4.5C	224	81.2	86.8	70.5
+HAT	20.31/1	4.0G	224	82.0	87.3	71.5
Swin-S	40.6M	8 7 C	224	83.0	87.8	72.4
+HAT	49.0101	0.7G	224	83.3	87.7	72.8
Swin-B	87 8M	15.40	224	83.5	87.9	72.9
+HAT	07.011	10.40	224	84.0	88.2	73.8
VOLO-D1	26 GM	C o C	22.4	84.2	89.0	74.0
+HAT	20.01/1	0.6G	224	84.5	89.2	74.9
VOLO-D1	26 GM	22.00	291	85.2	89.6	75.6
+HAT	20.01/1	22.0G	304	85.5	89.8	76.6
VOLO-D5	205.5M	60.00	224	86.1	89.9	76.3
+HAT	290.0M	09.00	224	86.3	90.2	76.8
VOLO-D5	205 FM	2040	119	87.0	90.6	77.8
+HAT	290.0M	504G	440	87.2	90.6	78.6
VOLO-D5	205 FM	419C	F19	87.1	90.6	78.0
+HAT	230.0M	412G	512	87.3	90.7	78.7

Table 2: Performance of various ViT architectures trained without and with the proposed HAT on the ImageNet, ImageNet-Real, and ImageNet-V2 validation set.

(e.g., CutMix [70], Mixup [73], and RandAugment [14]) and model regularization (e.g., stochastic depth [35] and weight decay [42]). Specifically, we use the hyperparameters in [25] for training ViT-B, and train ViT-T and ViT-S with the same hyper-parameters except for throwing away EMA, resulting in strong baselines. For training variants of Swin Transformer and VOLO, we follow the training setup of the original paper [40,69] (including token labeling [37] for VOLO). The default image resolution for these models is 224×224 . We also finetune variants of VOLO on larger image resolutions (384×384 , 448×448 , and 512×512).

For the proposed HAT, in all cases, the PGD learning rate η is 1/255, and the parameters α and β are set to 3 and 0.01, respectively. We set the maximum perturbation strength ϵ as 2/255 and the number of PGD steps K as 3 by default. For VOLO-D5, the largest model in our experiments, we set K = 2with $\epsilon = 1/255$ to reduce the training time. For all ViT models, we adopt the training strategy in Algorithm 1 in the first 200 epochs and perform normal training in the rest 100 epochs. In our HAT, each PGD step requires one forward and backward pass. Accordingly, for the whole training, HAT leads to about $1.7 \times$ and $2.3 \times$ computation cost for K = 1 and K = 2, respectively.

5.2 Results on ImageNet Classification

Results of Various ViT Architectures. We present the results of variants of ViT, Swin Transformer, and VOLO trained without and with our HAT in Table 2.

"Top-1", "Real Top-1", and "V2 Top-1" refer to the top-1 accuracy evaluated on the ImageNet [18], ImageNet-Real [2], and ImageNet-V2 [48] validation set, respectively, where ImageNet-Real is built by relabeling the validation set of the original ImageNet for correcting labeling errors and ImageNet-V2 is a newly collected version of the ImageNet validation set.

Note that these models in Table 2 are with different architectures and sizes, and the baselines are all carefully tuned with various data augmentation and model regularization techniques. As can be seen, HAT can steadily improve the performance of all models. To be specific, we boost top-1 accuracy on the ImageNet validation set by 1.1%, 0.8%, and 1.2% for ViT-T, ViT-S, and ViT-B, respectively. Even for Swin Transformer and VOLO, more advanced architectures, we can still consistently improve the performance of their variants.

Table 3: Comparison of HAT with other training strategies. All results are based on the ImageNet-1K training set and ViT-B. Results of other methods are drawn from original papers. The methods with blue color mean self-supervised learning.

Ton 1	Real	V2	
10p-1	Top-1	Top-1	
76.7	82.3	64.1	
81.8	86.7	71.5	
83.4	88.3	73.2	
81.7	86.8	70.8	
82.4	-	-	
79.9	85.2	67.5	
82.8	-	-	
83.2	-	-	
83.2	-	-	
83.6	-	-	
83.2	87.9	72.6	
84.3	88.8	73.9	
	Top-1 76.7 81.8 83.4 81.7 82.4 79.9 82.8 83.2 83.2 83.2 83.2 83.2 83.2 83.2 83.2 83.2 83.2 83.2 83.2	Real Top-1 Real Top-1 76.7 82.3 81.8 86.7 83.4 88.3 81.7 86.8 82.4 - 79.9 85.2 82.8 - 83.2 - 83.2 - 83.2 - 83.2 - 83.2 - 83.2 - 83.2 - 83.2 - 83.4 87.9 84.3 88.8	

The performance gains of our method are preserved through finetuning at higher resolutions. In particular, when the image resolution is 512×512 , VOLO-D5 with our HAT reaches a top-1 accuracy of 87.3% on the ImageNet validation set.

Comparison to Other Methods with ViT-B. We compare our proposed HAT with other state-of-the-art training strategies in Table 3. We conduct these experiments using ViT-B. Compared to DeiT [54] and TransMix [10], which utilize data augmentations to empower ViT models, we achieve a significantly higher top-1 accuracy. Most closely related to our work is pyramid adversarial training (PyramidAT) [33], which leverages structured adversarial perturbations. However, due to a bigger number of PGD steps, its training cost is twice as high as ours, but resulting in a lower top-1 accuracy than our HAT. Besides the supervised training methods, we also compare with the methods that pre-train on the ImageNet-1K training set in a self-supervised manner and then perform supervised finetuning, including DINO [8], MoCo v3 [12], BEiT [1], and MAE [25]. Our HAT shows very competitive performance with them. Furthermore, combining the proposed HAT with knowledge distillation [54], we obtain the performance of 84.3%, which is the highest top-1 accuracy among these methods.

5.3 Results on Out-of-distribution Data

We evaluate the proposed HAT on five out-of-distribution datasets: ImageNet-A which contains 7,500 examples that are harder and may cause mistakes across various models [31]; ImageNet-C [30] which applies a set of common visual corruptions to the ImageNet validation set; ImageNet-Sketch [56] which contains

11

(Ince) , lower is better. The test resolution of an below models is 224×224 .											
Model	ImageNet-A		ImageNet-C		Sketch		Rendition		Stylized		
Model	Top 1	+HAT	mCEL	+HAT	Top 1	+HAT	Top 1	+HAT	Top 1	+HAT	
	10p-1	Top-1	$mCE\downarrow$ mCE↓		Top-1 Top-1		Top-1 Top-1		10p-1	Top-1	
ViT-T	7.7	7.3	70.0	66.8	19.8	22.9	31.9	35.8	9.5	12.5	
ViT-S	18.5	23.1	53.3	49.7	29.3	32.3	41.6	45.1	15.8	18.1	
ViT-B	25.3	30.6	46.4	42.2	36.1	38.5	49.6	51.3	21.8	24.7	
Swin-T	22.1	25.7	58.0	53.9	28.5	31.0	41.4	43.8	13.1	13.8	
Swin-S	32.7	34.6	51.8	48.6	32.7	33.9	45.2	46.5	14.2	15.1	
Swin-B	35.8	40.0	51.7	46.9	32.2	36.4	45.8	49.0	15.7	16.4	
VOLO-D1	39.0	42.9	46.8	43.7	38.5	39.5	50.3	51.9	19.2	21.2	
VOLO-D5	50.9	54.5	41.8	38.4	44.3	45.7	57.9	59.7	24.6	25.9	

Table 4: Performance of various ViT architectures trained without and with the proposed HAT on five out-of distribution datasets. Note that for mean Corruption Error (mCE), lower is better. The test resolution of all below models is 224×224 .

sketch-like images and matches the ImageNet validation set in categories and scale; ImageNet-Rendition [29], a 30,000 image test set containing various renditions (*e.g.*, paintings, embroidery); Stylized ImageNet [23] which is a stylized version of ImageNet created by applying AdaIN style transfer [36] to ImageNet images. We report the top-1 accuracy on all datasets, except ImageNet-C where we report the normalized mean Corruption Error (mCE) (lower is better) following the original paper [30].

The results are shown in Table 4. Note that all models are trained on the ImageNet-1K training set and tested on these five out-of-distribution datasets. As can be seen, our method can bring performance gains for all architectures, demonstrating that HAT can enhance the robustness of ViT models to out-of-distribution data. Accordingly, HAT breaks the trade-off between in-distribution and out-of-distribution generalization [47,72], or in other words, it can achieve better performance in these two cases simultaneously.

5.4 Transfer Learning to Downstream Tasks

ImageNet pre-training is widely used in various vision tasks [26]. For the downstream tasks, the backbones can be initialized by the model weights pre-trained on ImageNet. In this section, we demonstrate that the advantages of HAT can be transferred to the downstream tasks, including object detection, instance segmentation, and semantic segmentation. More implementation details can be found in Appendix D.

Object Detection and Instance Segmentation. We take three variants of Swin Transformer trained without and with our HAT as pre-trained models to evaluate the performance in object detection and instance segmentation. The experiments are conducted on COCO 2017 [38] with the Cascade Mask R-CNN object detection framework [5,27]. We present the results in Table 5. As can be seen, HAT helps three variants of Swin Transformer achieve higher detection performance. These results show that the superiority of HAT can be transferred

Table 5: Object detection and instance segmentation performance of the models pre-trained without and with HAT on COCO val 2017. We adopt the Cascade Mask R-CNN object detection framework. AP^{box} and AP^{mask} are box average precision and mask average precision, respectively.

Backbone	Params	FLOPs	AP^{box}	$\mathrm{AP}^{\mathrm{box}}_{50}$	AP_{75}^{box}	AP^{mask}	AP_{50}^{mask}	AP_{75}^{mask}
Swin-T	96M	745 C	50.5	69.3	54.9	43.7	66.6	47.1
+HAT	801VI	740G	50.9	69.4	55.6	43.9	66.8	47.3
Swin-S	1071	0200	51.8	70.4	56.3	44.7	67.9	48.5
+HAT	107101	030G	52.5	71.2	57.1	45.4	68.8	49.4
Swin-B	145M	089C	51.9	70.9	56.5	45.0	68.4	48.7
+HAT	140101	962G	52.8	71.5	57.5	45.6	69.1	49.6

Table 6: Semantic segmentation performance of the models pre-trained without and with HAT on the ADE20K validation set. We adopt the UperNet segmentation framework. MS denotes testing with variable input size.

Table 7: Performance of ViT-B trained with the proposed HAT under different maximum perturbation strength ϵ .

$ \begin{array}{c} \hline ckbone \ Params \ FLOPs \ mIoU \ mIoU(MS) \ mAcc} \\ \hline in-T & 60M & 945G & 44.5 & 46.1 & 55.6 \\ \hline in-S & 81M & 1038G & 47.6 & 49.5 & 58.8 \\ \hline in-B & 121G & 1088G & 48.1 & 49.7 & 59.1 \\ \hline in-B & 121G & 1088G & 48.9 & 50.3 & 60.2 \\ \hline in-B & 5/25 \\ \hline in-B & 5/25 \\ \hline in-B & 121G & 1088G & 48.9 & 50.3 & 60.2 \\ \hline in-B & 5/25 \\ \hline in-B & 121G & 1088G & 48.9 & 50.3 & 60.2 \\ \hline in-B & 5/25 \\ \hline in-$	1					
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	Backbone	Params	FLOPs	mIoU	mIoU(MS	b) mAcc
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	Swin-T	COM	0450	44.5	46.1	55.6
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	+HAT	60M	945G	45.6	46.7	57.4
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	Swin-S	<u>ө</u> 1М	10290	47.6	49.5	58.8
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	+HAT	01111	1056G	48.1	49.7	59.5
AT 48.9 50.3 60.2 5/255	Swin-B	191C	10880	48.1	49.7	59.1
	+HAT	1210	1000G	48.9	50.3	60.2

to downstream tasks. Moreover, we would like to emphasize that HAT does not introduce extra parameters or computation cost in inference.

Semantic Segmentation. We also use Swin Transformer to evaluate the performance in semantic segmentation. We report results on the widely-used segmentation benchmark ADE20K [74] with the UperNet [64] segmentation framework. The results are shown in Table 6. We can see that HAT brings significant gains for these three variants. Especially for Swin-T and Swin-S, the improvements of HAT are more than 1.0% mIOU. These results further show the benefits of our proposed HAT to downstream tasks.

5.5 Ablation Studies

Ablation on Maximum Perturbation Strength. We ablate the effects of the maximum perturbation strength ϵ , where a larger ϵ indicates stronger adversarial examples for the adversarial training. We test HAT with $\epsilon \in \{1/255, 2/255, 3/255, 4/255, 5/255\}$. Correspondingly, we set the number of PGD steps as $K \in \{2, 3, 4, 5, 6\}$. The PGD learning rate η is fixed at 1/255. The results are shown in Table 7. We can see that HAT can achieve superior performance when





Fig. 6: Performance of ViT-B trained without and with the proposed HAT on various sizes of training sets.

Fig. 7: Fourier heat maps of ViT-B trained without and with HAT. The Fourier heat map reflects the sensitivity of a model to high- and low-frequency corruptions. Error rates are averaged over the entire ImageNet validation set.

 ϵ is set as 2/255 or 3/255. It demonstrates that medium strength adversarial examples are more helpful for training ViT models than weaker or stronger adversarial examples. This ablation study and other experiments in this paper verify that $\epsilon = 2/255$ and K = 3 are reasonable choices in most cases.

Scaling the Training Set Size. HAT In this part, we investigate the effects of with various training set sizes. Specifically, we experiment on training sets with different sizes by randomly sampling 1/8, 1/4, and 1/2 images from the original ImageNet-1K training set, resulting in training sets with 16K, 32K, and 64K samples. Then, we train ViT-B on these datasets without and with the proposed HAT and evaluate on the original ImageNet validation set. The comparison is presented in Figure 6. It is shown that HAT can improve the performance of ViT-B under all sizes of the training set. Especially for the small-scale training set, HAT brings more gains, *e.g.*, +7.1 for the smallest training set. It may be because the insufficient ability of ViT models to capture high-frequency components is amplified on the small-scale dataset and the effectiveness of HAT on compensating for this ability is more significant in this case. In short, our results illustrate that HAT enables ViT models to handle the small-scale training set better.

5.6 Discussions

Fourier Heat Maps of ViT Models. We investigate the effects of HAT on the model sensitivity to low- and high-frequency corruptions. We adopt the Fourier heat map [67], which visualizes the error rates of a model tested on perturbed images with additive Fourier basis noise. We fix ℓ_2 -norm of the additive noise as 15.7 and average the error rates over the entire ImageNet validation set. Following [67], we present the 63×63 square centered at the lowest frequency in the Fourier domain. The Fourier heat maps of ViT-B trained without and with HAT are shown in Figure 7. As we can see, the baseline model is highly sensitive to additive noise in the high-frequency. In contrast, the model trained with HAT is more robust to the noise, especially in the high-frequency.

14 Jiawang Bai, Li Yuan, Shu-Tao Xia et al.

Table 8: Comparison of HAT and longer Table 9: Performance of CNN (ResNet-50) normal training with ViT-B under vari- and MLP (ViP-Small/7) models trained ous computational costs.

Jus computational costs.						thout and	WITTI TH	e proposed	IAI.
Cost	Normal Training		+HAT			Model	Top-1	Real Top-1	V2 Top-1
	Setting	Top-1	Setting	Top-1		ResNet-50	79.8	86.9	70.7
$1 \times$	300 epochs	82.0	-	-		+HAT	80.2	87.2	71.3
$1.7 \times$	500 epochs	82.5	K = 2	83.1		ViP-Small/	7 81.6	85.5	67.8
$2.3 \times$	690 epochs	82.4	K = 3	83.2		+HAT	82.2	86.1	69.6

Longer Normal Training. To further verify the effectiveness of HAT, we increase the number of epochs of normal training to match the computational cost of HAT. The results are presented in Table 8. We can see that 500 epochs are enough for the normal training to converge, and HAT surpasses normal training under $1.7 \times$ and $2.3 \times$ cost.

Beyond ViT Models. We explore the performance of HAT on the CNN and MLP models. We conduct experiments on ResNet-50 [28] and ViP-Small-7 [34]. We train ResNet-50 for 800 epochs following the setup in [63] and ViP-Small-7 for 300 epochs following the setup in [34]. For our HAT, we perform adversarial training in the first 600 epochs for ResNet-50 and in the first 200 epochs for ViP-Small-7, and keep other settings unchanged. The results in Table 9 show that HAT brings improvements of 0.4% and 0.6% for ResNet-50 and ViP-Small-7, respectively. Therefore, the proposed training strategy is promising to be extended to other models.

6 Conclusions and Future Work

In this paper, we study ViT models from a frequency perspective. We find that compared to CNN models, ViT models can not well exploit the high-frequency components of images. We also present a new frequency analysis of existing techniques for improving the performance of ViT models. To compensate for this insufficient ability of ViT models, we propose HAT, a simple but effective training strategy based on adversarial training. Extensive experiments verify its effectiveness on diverse benchmarks.

Despite achieving higher performance, HAT has an increased training time compared to the normal training. Therefore, a future study is to improve the efficiency of the proposed HAT. Also, the insights provided in this paper further prompt us to explore other techniques to compensate for the ability of ViT models to capture the high-frequency components of images.

Acknowledgments. This work is supported in part by the National Natural Science Foundation of China under Grant 62171248, and the PCNL KEY project (PCL2021A07). Li Yuan was supported in part by PKU-Shenzhen Start-Up Research Fund (1270110283) and PengCheng Laboratory.

References

- Bao, H., Dong, L., Wei, F.: Beit: Bert pre-training of image transformers. arXiv preprint arXiv:2106.08254 (2021)
- Beyer, L., Hénaff, O.J., Kolesnikov, A., Zhai, X., Oord, A.v.d.: Are we done with imagenet? arXiv preprint arXiv:2006.07159 (2020)
- 3. Brendel, W., Bethge, M.: Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. In: ICLR (2019)
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. In: NeurIPS (2020)
- Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. In: CVPR (2018)
- Campbell, F.W., Robson, J.G.: Application of fourier analysis to the visibility of gratings. The Journal of physiology 197(3), 551 (1968)
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: Endto-end object detection with transformers. In: ECCV (2020)
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021)
- Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. In: CVPR (2021)
- Chen, J.N., Sun, S., He, J., Torr, P., Yuille, A., Bai, S.: Transmix: Attend to mix for vision transformers. arXiv preprint arXiv:2111.09833 (2021)
- 11. Chen, X., Hsieh, C.J., Gong, B.: When vision transformers outperform resnets without pre-training or strong data augmentations. In: ICLR (2022)
- Chen, X., Xie, S., He, K.: An empirical study of training self-supervised vision transformers. In: ICCV (2021)
- 13. Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C.: Twins: Revisiting the design of spatial attention in vision transformers. In: NeurIPS (2021)
- 14. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: CVPR Workshops (2020)
- 15. Dai, Z., Cai, B., Lin, Y., Chen, J.: Up-detr: Unsupervised pre-training for object detection with transformers. In: CVPR (2021)
- Das, N., Shanbhogue, M., Chen, S.T., Hohman, F., Li, S., Chen, L., Kounavis, M.E., Chau, D.H.: Shield: Fast, practical defense and vaccination for deep learning using jpeg compression. In: ACM SIGKDD (2018)
- De Valois, R.L., De Valois, K.K.: Spatial vision. Annual review of psychology **31**(1), 309–341 (1980)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
- 19. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAACL (2019)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2020)
- d'Ascoli, S., Touvron, H., Leavitt, M.L., Morcos, A.S., Biroli, G., Sagun, L.: Convit: Improving vision transformers with soft convolutional inductive biases. In: ICML (2021)
- 22. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. In: NeurIPS (2020)

- 16 Jiawang Bai, Li Yuan, Shu-Tao Xia et al.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In: ICLR (2018)
- Gong, C., Wang, D., Li, M., Chandra, V., Liu, Q.: Vision transformers with patch diversification. arXiv preprint arXiv:2104.12753 (2021)
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. arXiv preprint arXiv:2111.06377 (2021)
- 26. He, K., Girshick, R., Dollár, P.: Rethinking imagenet pre-training. In: CVPR (2019)
- 27. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
- Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: ICCV (2021)
- Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: ICLR (2018)
- 31. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. In: CVPR (2021)
- 32. Heo, B., Yun, S., Han, D., Chun, S., Choe, J., Oh, S.J.: Rethinking spatial dimensions of vision transformers. In: ICCV (2021)
- Herrmann, C., Sargent, K., Jiang, L., Zabih, R., Chang, H., Liu, C., Krishnan, D., Sun, D.: Pyramid adversarial training improves vit performance. arXiv preprint arXiv:2111.15121 (2021)
- 34. Hou, Q., Jiang, Z., Yuan, L., Cheng, M.M., Yan, S., Feng, J.: Vision permutator: A permutable mlp-like architecture for visual recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence (2022)
- Huang, G., Sun, Y., Liu, Z., Sedra, D., Weinberger, K.Q.: Deep networks with stochastic depth. In: ECCV (2016)
- 36. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV (2017)
- Jiang, Z.H., Hou, Q., Yuan, L., Zhou, D., Shi, Y., Jin, X., Wang, A., Feng, J.: All tokens matter: Token labeling for training better vision transformers. In: NeurIPS (2021)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
- 40. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: CVPR (2021)
- Liu, Z., Liu, Q., Liu, T., Xu, N., Lin, X., Wang, Y., Wen, W.: Feature distillation: Dnn-oriented jpeg compression against adversarial examples. In: CVPR (2019)
- 42. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (2018)
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: ICLR (2018)
- 44. Park, N., Kim, S.: How do vision transformers work? In: ICLR (2022)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: an imperative style, highperformance deep learning library. In: NeurIPS (2019)

Improving Vision Transformers by Revisiting High-frequency Components

- Radosavovic, I., Kosaraju, R.P., Girshick, R., He, K., Dollár, P.: Designing network design spaces. In: CVPR (2020)
- 47. Raghunathan, A., Xie, S.M., Yang, F., Duchi, J., Liang, P.: Understanding and mitigating the tradeoff between robustness and accuracy. In: ICML (2020)
- Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imagenet classifiers generalize to imagenet? In: ICML (2019)
- 49. Steiner, A., Kolesnikov, A., Zhai, X., Wightman, R., Uszkoreit, J., Beyer, L.: How to train your vit? data, augmentation, and regularization in vision transformers. arXiv preprint arXiv:2106.10270 (2021)
- 50. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: ICCV (2021)
- 51. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: ICCV (2017)
- 52. Sun, Z., Cao, S., Yang, Y., Kitani, K.M.: Rethinking transformer-based set prediction for object detection. In: ICCV (2021)
- Sweldens, W.: The lifting scheme: A construction of second generation wavelets. SIAM journal on mathematical analysis 29(2), 511–546 (1998)
- 54. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: ICML (2021)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
- 56. Wang, H., Ge, S., Lipton, Z.C., Xing, E.P.: Learning robust global representations by penalizing local predictive power. In: NeurIPS (2019)
- 57. Wang, H., Wu, X., Huang, Z., Xing, E.P.: High-frequency component helps explain the generalization of convolutional neural networks. In: CVPR (2020)
- Wang, P., Zheng, W., Chen, T., Wang, Z.: Anti-oversmoothing in deep vision transformers via the fourier domain analysis: From theory to practice. In: ICLR (2022)
- Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: ICCV (2021)
- 60. Wang, W., Yao, L., Chen, L., Lin, B., Cai, D., He, X., Liu, W.: Crossformer: A versatile vision transformer hinging on cross-scale attention. In: ICLR (2022)
- Wang, Y., Xu, Z., Wang, X., Shen, C., Cheng, B., Shen, H., Xia, H.: End-to-end video instance segmentation with transformers. In: CVPR (2021)
- Wightman, R.: Pytorch image models. https://github.com/rwightman/ pytorch-image-models (2019). https://doi.org/10.5281/zenodo.4414861
- 63. Wightman, R., Touvron, H., Jégou, H.: Resnet strikes back: An improved training procedure in timm. arXiv preprint arXiv:2110.00476 (2021)
- Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 418–434 (2018)
- Xie, C., Tan, M., Gong, B., Wang, J., Yuille, A.L., Le, Q.V.: Adversarial examples improve image recognition. In: CVPR (2020)
- Yang, F., Yang, H., Fu, J., Lu, H., Guo, B.: Learning texture transformer network for image super-resolution. In: CVPR (2020)
- 67. Yin, D., Gontijo Lopes, R., Shlens, J., Cubuk, E.D., Gilmer, J.: A fourier perspective on model robustness in computer vision. In: NeurIPS (2019)
- Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.H., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token vit: Training vision transformers from scratch on imagenet. In: ICCV (2021)

- 18 Jiawang Bai, Li Yuan, Shu-Tao Xia et al.
- Yuan, L., Hou, Q., Jiang, Z., Feng, J., Yan, S.: Volo: Vision outlooker for visual recognition. arXiv preprint arXiv:2106.13112 (2021)
- 70. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: ICCV (2019)
- Zeng, Y., Fu, J., Chao, H.: Learning joint spatial-temporal transformations for video inpainting. In: ECCV (2020)
- 72. Zhang, H., Yu, Y., Jiao, J., Xing, E., El Ghaoui, L., Jordan, M.: Theoretically principled trade-off between robustness and accuracy. In: ICML (2019)
- Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: ICLR (2018)
- Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. International Journal of Computer Vision 127(3), 302–321 (2019)
- Zhou, L., Zhou, Y., Corso, J.J., Socher, R., Xiong, C.: End-to-end dense video captioning with masked transformer. In: CVPR (2018)
- Zhu, C., Cheng, Y., Gan, Z., Sun, S., Goldstein, T., Liu, J.: Freelb: Enhanced adversarial training for natural language understanding. In: ICLR (2020)
- 77. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: ICLR (2020)