

Robustifying Token Attention for Vision Transformers

Yong Guo, David Stutz, Bernt Schiele

Max Planck Institute for Informatics, Saarland Informatics Campus

guoyongcs@gmail.com, {david.stutz, schiele}@mpi-inf.mpg.de

Abstract

Despite the success of vision transformers (ViTs), they still suffer from significant drops in accuracy in the presence of common corruptions, such as noise or blur. Interestingly, we observe that the attention mechanism of ViTs tends to rely on few important tokens, a phenomenon we call *token overfocusing*. More critically, these tokens are not robust to corruptions, often leading to highly diverging attention patterns. In this paper, we intend to alleviate this overfocusing issue and make attention more stable through two general techniques: First, our **Token-aware Average Pooling (TAP)** module encourages the local neighborhood of each token to take part in the attention mechanism. Specifically, TAP learns average pooling schemes for each token such that the information of potentially important tokens in the neighborhood can adaptively be taken into account. Second, we force the output tokens to aggregate information from a diverse set of input tokens rather than focusing on just a few by using our **Attention Diversification Loss (ADL)**. We achieve this by penalizing high cosine similarity between the attention vectors of different tokens. In experiments, we apply our methods to a wide range of transformer architectures and improve robustness significantly. For example, we improve corruption robustness on ImageNet-C by 2.4% while improving accuracy by 0.4% based on state-of-the-art robust architecture FAN. Also, when fine-tuning on semantic segmentation tasks, we improve robustness on CityScapes-C by 2.4% and ACDC by 3.0%. Our code is available at <https://github.com/guoyongcs/TAPADL>.

1. Introduction

Despite the success of vision transformers (ViTs), their performance still drops significantly on common image corruptions such as ImageNet-C [24, 52, 19], adversarial examples [18, 16, 40, 56], and out-of-distribution examples as benchmarked in ImageNet-A/R/P [64, 24]. In this paper, we examine a key component of ViTs, i.e., the self-attention mechanism, to understand these performance drops. Interestingly, we discover a phenomenon we call *token overfo-*




Figure 1. Stability against image corruptions in terms of attention visualization (left, a matrix of 196×196) and cosine similarity of attention between clean and corrupted examples (right). *Left:* We average the attention maps across different heads for visualization and show the results of the last layer. We observe that ViTs put too much focus on very few tokens, a phenomenon we call *token overfocusing*. More critically, the attention of the baseline FAN model [16] is fragile to image corruptions, e.g., with Gaussian noise. Our approach, in contrast, alleviates token overfocusing and thereby improves stability of the attention against corruptions. *Right:* On ImageNet, we plot the distribution of cosine similarities across all layers (without averaging heads) between clean and corrupted examples. We show that our model yields a significantly higher similarity score than the baseline model.

using, where only few important tokens are relied upon by the attention mechanism across all heads and layers. We hypothesize that this overfocusing is particularly fragile to the corruptions on input images and greatly hampers the robustness of the attention mechanism.

Starting from the state-of-the-art robust architecture FAN [67], we exemplarily investigate the last attention layer (see attention of other layers in supplementary). Specifically, Figure 1 shows a clean and a corrupted input image as well as the corresponding attention maps. These are matrices of $N \times N$, with N being the number of input/output tokens. Here, the i -th row indicates which input tokens (columns) the i -th output token “attends” to – darker red indicates higher attention scores. **Token overfocusing** can then, informally, be defined by observing pronounced vertical lines in the attention map: First, each output token (row) focuses on only few important input tokens, ignoring most

of the other information. Second, all output tokens seem to focus on the same input tokens, leading to a very low diversity among the attention vectors in different rows. We highlight that the overfocusing issue is present throughout the entire ImageNet dataset, and also across diverse architectures (see more examples in Figure 3 and supplementary). Moreover, this overfocusing issue has also been observed in existing works. For example, Figure 5 of [10] observes very few important tokens (in deep red color) in layers 3~6; Figure 1 of [16] also reports very few important tokens (in yellow color). More critically, we find that these important tokens are extremely fragile in the presence of common corruptions. To be specific, when applying Gaussian noise on the input image, the tokens recognized as important change entirely, see Figure 1 (left, second column). Quantitatively, this can be captured by computing the cosine similarity between the clean and corrupted attention maps. Unsurprisingly, as shown by the blue box in Figure 1 (right), the cosine similarity is indeed extremely low, confirming our initial hypothesis. This motivates us to robustify the attention by alleviating the token overfocusing issue.

To this end, we encourage the attention module to focus on diverse input tokens. This can be achieved by *changing the patterns in both columns and rows* of an attention map, which motivates the two key components of our method. First, when comparing attention columns in Figure 1, very few tokens are important, leading to a fragile attention. To address this, we encourage output tokens to not only focus on individual input tokens but take into account the local neighborhood around these tokens, in order to make more tokens (columns) contribute meaningful information. Intuitively, an individual token itself may not be important but can be enhanced by aggregating the information from potentially important tokens located within its neighborhood. We achieve this using a learnable average pooling mechanism applied to each input token to aggregate information before computing self-attention. Second, when comparing attention rows in Figure 1, most output tokens focus on the same input tokens. Once these tokens are distorted, the whole attention may fail. To address this issue, we seek to diversify the set of input tokens (columns) that the output tokens (rows) rely on. We achieve this using a loss that explicitly penalizes high cosine similarity across rows. In Figure 1, the combination of these techniques leads to more balanced attention across columns and more diverse attention across rows. More critically, these attention maps are more stable in the light of image corruptions. Again, we quantitatively confirm this on ImageNet using the cosine similarity which is significantly higher.

Overall, we make three key **contributions**: 1) we propose a *Token-aware Average Pooling (TAP)* module that encourages the local neighborhood of tokens to participate in the self-attention mechanism. To be specific, we conduct

averaging pooling to aggregate information and control it by adjusting the pooling area for each token. 2) We further develop an *Attention Diversification Loss (ADL)* to improve the diversity of attention received by different tokens, i.e., rows in the attention map. To this end, we compute the attention similarity between rows in each layer and minimize it during training. 3) We highlight that both TAP and ADL can be applied on top of diverse transformer architectures. In the experiments, the proposed methods consistently improve the model robustness on out-of-distribution benchmarks by a large margin while preserving a competitive improvement of clean accuracy at the same time. For example, we improve robustness against corruptions and distribution shifts on ImageNet-A/C/P/R by at least 1.5%, as shown in Table 2. In addition, the improvement also generalizes well to other downstream tasks, e.g., semantic segmentation, with an improvement of 2.4% in Table 4.

2. Related Work

The remarkable performance of ViTs on various learning tasks is largely attributed to the use of self-attention [26]. Naturally, the self-attention mechanism has been extended in various aspects, addressing shortcomings such as the heavy reliance on pre-training [42, 62, 35] or the high computational complexity [29, 28, 13, 1, 8, 3, 47] and adapting the architecture to vision tasks by considering multi-scale transformers [25, 15, 46, 48, 49, 9, 57, 53, 25, 50, 6, 58]. Similar to us, some related works [66, 10, 44, 34, 45] also investigate and visualize the self-attention in order to improve transformer architectures, e.g., to learn deeper transformers [66] or prune attention heads [45]. However, to the best of our knowledge, we are the first to report a token overfocusing issue and link it to poor robustness of ViTs.

There is also a lot of interests in understanding and improving the robustness of ViTs [5, 2, 41, 37, 4, 21, 31]. For example, [32] develops a robust vision transformer (RVT) by combining various components to boost robustness, including an improved attention scaling. FAN [67] combines token attention and channel attention [1] and can be considered state-of-the-art. Both approaches rely on modified self-attention mechanisms and can be shown to suffer from token overfocusing. Thus, our techniques can be shown to improve robustness on top of RVT and FAN, respectively. Our TAP approach also shares similarities to recent ideas of trying to introduce locality into the self-attention mechanism [54, 53, 17, 38, 61, 54, 51]. For example, CvT [53] introduces convolutions into the self-attention architecture to enhance the local information inside tokens. MViTv2 [27] exploits average pooling to extract local features and improve clean performance. However, all of these strategies use the same aggregation across all tokens, while our approach adaptively chooses an aggregation neighborhood for each token to improve robustness.

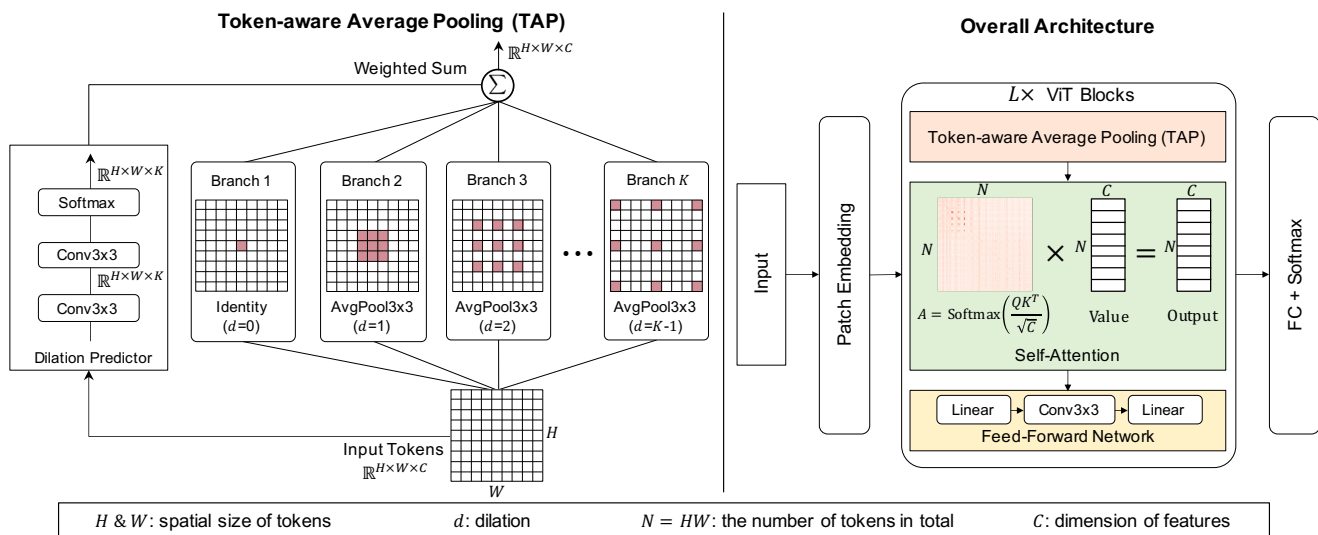


Figure 2. The proposed Token-aware Average Pooling (TAP) module (left) and the overall architecture (right). *Left*: In TAP, we introduce K branches to enable tokens to consider different pooling areas and compute the weighted sum over them. Besides, we build a lightweight dilation predictor to learn the weights for different branches. *Right*: We introduce a TAP layer into every basic block to encourage more tokens to be actively involved in the following self-attention mechanism.

3. Robust Token Self-Attention

In the following, we focus on the attention mechanism in ViTs, aiming to improve their overall robustness. To get started, we first describe the observed *token overfocusing issue* where ViTs focus on very few but unstable tokens in detail in Section 3.1. Then, we propose two general techniques for alleviating this issue: First, in Section 3.2, we propose a **Token-aware Average Pooling (TAP)** module that encourages the neighborhood of tokens to participate in the attention mechanism. This is achieved using a learnable pooling area as illustrated in Figure 2. Second, in Section 3.3, we develop a new **Attention Diversification Loss (ADL)** to improve the diversity of attention patterns across tokens. Both methods can be applied to most transformer architectures and we will show they greatly improve robustness with negligible training overhead.

3.1. Token Overfocusing

As illustrated in Figure 1, we can visualize the self-attention mechanism in each layer as $N \times N$ attention matrix. Here, N is the number of input and output tokens and each entry (i, j) denotes the attention that the i -th output token (i -th row) puts on the j -th input token (j -th column) – deeper red denoting higher attention scores. To handle multi-head self-attention, we visualize this matrix by averaging across attention heads.

As baseline, Figure 1 highlights the recent FAN [67] architecture, showing the attention map of the last layer for

example. We observe that the attention is generally very sparse in columns, meaning that most input tokens are not attended to and very few tokens are overfocused. More importantly, these “important” tokens are often similar across output tokens (rows). We refer to this phenomenon as *token overfocusing*. This leads to problems when facing corruptions such as Gaussian noise: the set of important tokens might change entirely (Figure 1, second column). This can be understood as the original tokens not capturing robust information. We can also quantitatively capture this instability by computing the cosine similarity between clean and corrupted attention maps across all ImageNet images. As shown by the blue box in Figure 1 (right), the baseline model obtains very low cosine similarities, indicating the poor robustness of standard self-attention.

We found that this phenomenon exists across diverse architectures, e.g., DeiT [43] and RVT [32], and diverse tasks, including both image classification and semantic segmentation (see visual examples in supplementary). Moreover, we highlight that our observation is consistent with existing works [10, 36]. Specifically, these works observe that the middle and deep layers (often with overfocusing issue) tend to capture global information but put focus on very few tokens, e.g., very few deep red tokens in Figure 5 of [10] and Figure 3 (Stage 4) in [36]. This further justifies the existence and importance of the token overfocusing issue. To alleviate this issue, we propose two general techniques to robustify self-attention in the remainder of this section.

3.2. Token-aware Average Pooling (TAP)

In the first part of our method, we seek to encourage more input tokens to participate in the self-attention mechanism, i.e., obtaining more columns with high scores in the attention map. To this end, we encourage each input token to explicitly aggregate useful information from its local neighborhood, in case the token itself does not contain important information. This approach is justified by existing works [53, 38, 27] and the observation that introducing any local aggregation before self-attention consistently improves robustness, see Table 1 (last column). In fact, these methods apply a fixed convolutional kernel or pooling area to all the tokens. Nevertheless, tokens often differ from each other and each token should require a specific local aggregation strategy. This motivates us to adaptively choose the right neighborhood size and aggregation strategy.

Inspired by this, we enable each token to select an appropriate area/strategy to conduct local aggregation. Specifically, we develop a **Token-aware Average Pooling (TAP)** that conducts average pooling and adaptively adjusts the pooling area for each token. As shown in Figure 2, we exploit a multi-branch structure that computes the weighted sum over multiple branches, each with a specific pooling area. Instead of simply changing the kernel size similar to [59], TAP changes the dilation to adjust the pooling area. The main observation behind this is that average pooling with a large kernel, without dilation, leads to extremely large overlaps between adjacent pooling regions and thereby severe redundancy in output tokens. This can, for example, be seen in Table 1 where AvgPool5x5 incurs a large drop in clean accuracy of around 1.2%. Similar to [54, 53, 17, 38], we also investigated using learnable convolutions, but observed only marginal improvement alongside a significant increase in computational cost.

Based on these observations, we build TAP based on average pooling with diverse dilations: Without loss of generality, given K branches, we consider the dilation within the range $d \in [0, K-1]$. Here, $d = 0$ implies identity mapping without any computation, i.e., no local aggregation. The maximum dilation determined by K is a hyper-parameter and we find that performance and robustness improvements saturate at around $K = 5$ (see Figure 7 top). Within the allowed dilation range, our method includes a lightweight dilation predictor to predict which dilation (i.e., which branch in Figure 2) to utilize. Note that this can also be a weighted combination of multiple d . We emphasize that this predictor is very efficient since it reduces the feature dimension (from C to K in Figure 2) such that it adds minimal computational overhead and model parameters. The same approach can also be applied to non-dilated average-pooling where d controls the kernel size, named TAP (multi-kernel). From Table 1, our TAP greatly outperforms this variant, indicating the effectiveness of using dilations for pooling.

Model	#FLOPs (G)	#Params (M)	ImageNet	ImageNet-C ↓
Baseline (FAN-B-Hybrid)	11.7	50.4	83.9	46.1
+ AvgPool3x3	11.7	50.4	83.6	45.6 (-0.5)
+ AvgPool5x5	11.7	50.4	82.7	45.5 (-0.6)
+ Conv3x3	17.3	79.4	84.0	45.9 (-0.2)
+ Conv5x5	27.4	130.7	84.4	45.8 (-0.3)
+ TAP (multi-kernel)	11.8	50.7	84.1	45.5 (-0.6)
+ TAP (Ours)	11.8	50.7	84.3	44.9 (-1.2)

Table 1. Comparisons of local aggregation approaches based on FAN-B-Hybrid. We show that conducting average pooling for all the tokens improves the robustness but impedes clean accuracy. Introducing a convolution into each block greatly increases model complexity. In addition, we also compare a variant of our TAP, namely TAP (multi-kernel), that considers multiple kernel sizes for pooling and learns weights for each branch. Our Tap greatly outperforms this variant, indicating the effectiveness of using dilation. Moreover, TAP yields the best tradeoff between accuracy and robustness along with negligible computational overhead.

3.3. Attention Diversification Loss (ADL)

In the second part of our method, we seek to improve the diversity of attention across output tokens, i.e., encourage different rows in Figure 1 to attend to different input tokens. Based on this objective, we propose an **Attention Diversification Loss (ADL)** that explicitly reduces the cosine similarity of attention among different output tokens (rows). However, for this approach to work, there are several challenges to overcome. First, computing the cosine similarity between attentions is numerically tricky. For example, if two rows (i.e., output tokens) have very disjoint attention patterns, we expect a low cosine similarity close to 0. However, even for tokens that are not attended to, the attention scores will not be zero. For a large N , computing dot product and adding these values up tend to result in a cosine similarity significantly above zero. To alleviate this issue, we exploit a thresholding trick to filter out those very small values and only focus on the most important ones. Let $\mathbb{1}(\cdot)$ be the indication function, and $A_i^{(l)}$ be the attention vector of the i -th token (row) in the l -th layer. We introduce a threshold τ (see ablation in Table 7) that depends on the number of tokens N , i.e., τ/N . Thus, the attention after thresholding becomes

$$\hat{A}_i^{(l)} = \mathbb{1}(A_i^{(l)} \geq \tau/N) \cdot A_i^{(l)}. \quad (1)$$

Second, to avoid the quadratic complexity of computing similarities between pairs of N rows, we approximate it by computing the cosine similarity between each individual attention vector $\hat{A}_i^{(l)}$ with the average attention $\bar{A}^{(l)} := \frac{1}{N} \sum_{i=1}^N \hat{A}_i^{(l)}$. When considering a model with L layers, we average the ADL loss across all the layers by:

$$\mathcal{L}_{ADL} = \frac{1}{L} \sum_{l=1}^L \mathcal{L}_{ADL}^{(l)}, \quad \mathcal{L}_{ADL}^{(l)} = \frac{1}{N} \sum_{i=1}^N \frac{\hat{A}_i^{(l)} \bar{A}^{(l)}}{\|\hat{A}_i^{(l)}\| \|\bar{A}^{(l)}\|}. \quad (2)$$

Method	#Params (M)	#FLOPs (G)	ImageNet \uparrow	ImageNet-C \downarrow	ImageNet-P \downarrow	ImageNet-A \uparrow	ImageNet-R \uparrow
ConvNeXt-B [30]	88.6	15.4	83.8	46.8	-	36.7	51.3
ConViT-B [14]	86.5	17.7	82.4	46.9	32.2	29.0	48.4
Swin-B [29]	87.8	15.4	83.4	54.4	32.7	35.8	46.6
T2T-ViT_t-24 [62]	64.1	15.0	82.6	48.0	31.8	28.9	47.9
RSPC (FAN-B-Hybrid) [20]	50.5	11.7	84.2	44.5	30.0	41.1	-
RVT-B [32]	91.8	17.7	82.6	46.8	31.9	28.5	48.7
+ TAP	92.1	17.9	83.0 (+0.4)	45.5 (-1.3)	30.6 (-1.3)	30.0 (+1.5)	49.4 (+0.7)
+ ADL	91.8	17.7	82.6 (+0.0)	45.2 (-1.6)	30.2 (-1.7)	30.8 (+2.3)	49.8 (+1.1)
+ TAP & ADL	92.1	17.9	83.1 (+0.5)	44.7 (-2.1)	29.6 (-2.3)	32.7 (+4.2)	50.2 (+1.5)
FAN-B-Hybrid [67]	50.4	11.7	83.9	46.1	31.3	39.6	52.7
+ TAP	50.7	11.8	84.3 (+0.4)	44.9 (-1.2)	30.3 (-1.0)	41.0 (+1.4)	53.9 (+1.2)
+ ADL	50.4	11.7	84.0 (+0.1)	44.4 (-1.7)	29.8 (-1.5)	41.4 (+1.8)	54.2 (+1.5)
+ TAP & ADL	50.7	11.8	84.3 (+0.4)	43.7 (-2.4)	29.2 (-2.1)	42.3 (+2.7)	54.6 (+1.9)

Table 2. Comparisons on ImageNet and diverse robustness benchmarks. We report the mean corruption error (mCE) on ImageNet-C and mean flip rate (mFR) on ImageNet-P. For these metrics, lower is better. Moreover, we directly report the accuracy on ImageNet-A and ImageNet-R. Based on the considered two baselines, our models consistently improve the accuracy and robustness on diverse benchmarks.

In practice, we combine our ADL with the standard cross-entropy (CE) loss and introduce a hyper-parameter λ (see ablation in Figure 7) to control the importance of ADL:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda \mathcal{L}_{ADL}. \quad (3)$$

We highlight that our ADL can be applied to boost the robustness on diverse tasks, including image classification and semantic segmentation (see Table 2 and Table 4).

4. Experiments

We conduct extensive experiments to verify our method on both image classification and semantic segmentation tasks. In Section 4.1, we first train classification models on ImageNet [12] and demonstrate that our models obtain significant improvement on various robustness benchmarks, including ImageNet-A [64], ImageNet-C [24], ImageNet-R [23], and ImageNet-P [24]. Then, in Section 4.2, we take our best pre-trained model and further finetune it on Cityscapes [11] for semantic segmentation. In practice, our models greatly improve mIoU on two popular robustness benchmarks, including Cityscapes-C [33] and ACDC [39], along with competitive performance on clean data. Both the code and pretrained models will be available soon.

4.1. Results on Image Classification

In this experiment, we build our method on top of two state-of-the-art robust architectures: RVT [32] and FAN [67] with the “Base” model size, i.e., RVT-B and FAN-B-Hybrid. We train the models on ImageNet and evaluate them on several robustness benchmarks. We closely follow the settings of RVT and FAN for training. Specifically, we train the models using the same augmentation schemes and adopt the batch size of 2048. We set the learning rate to 2×10^{-3} and train all the models for 300 epochs. In all the experiments, by default, we set $K = 4$ and $\lambda = 1$ to

train our models. To evaluate robustness, we consider several robustness benchmarks, including ImageNet-A [64], ImageNet-C [24], ImageNet-R [23], and ImageNet-P [24]. Note that, we report the mean corruption error (mCE) on ImageNet-C and mean flip rate (mFR) on ImageNet-P. For both metrics, lower is better. Empirically, we demonstrate that using either TAP or ADL individually is able to improve the robustness. When combining them together, our models outperform the baselines by a larger margin and the performance improvement generalizes well to diverse architectures (see Table 5).

4.1.1 Comparisons on ImageNet

As shown in Table 2, compared with the strong baselines RVT and FAN, our models consistently improve the robustness on ImageNet-C by $>2.1\%$ and also yield comparable improvement on other robustness benchmarks, including ImageNet-A/R/SK. Moreover, our models also obtain a competitive improvement in terms of clean accuracy on ImageNet. For example, we improve the accuracy by $>0.4\%$ on both considered baseline architectures. More critically, we highlight that these improvements only come with negligible computational cost in terms of both the number of parameters and the number of floating-point operations (FLOPs). In addition, we also report the detailed corruption error on individual corruption types of ImageNet-C based on FAN-B-Hybrid. From Table 3, our best model (combining TAP and ADL together) obtains the best results on most of the corruption types. It is worth noting that our model is particularly effective against noise corruptions, e.g., yielding a large improvement of 6.25% on Gaussian noise corruption. Overall, these experiments indicate that robustifying attention consistently improves robustness across different architectures and benchmarks.

Method	mCE	Noise			Blur				Weather				Digital			
		Gaussian	Shot	Impulse	Defocus	Glass	Motion	Zoom	Snow	Frost	Fog	Brightness	Contrast	Elastic	Pixelate	JPEG
FAN-B-Hybrid	46.2	40.12	39.27	36.80	51.58	63.96	47.53	54.98	40.24	43.96	36.98	36.68	34.17	61.59	53.25	51.86
+ TAP	44.9	36.02	36.26	34.16	52.72	65.07	45.73	54.90	39.75	42.39	35.68	37.38	33.21	62.75	48.85	49.46
+ ADL	44.3	35.61	35.55	33.51	50.80	64.27	45.47	54.47	38.06	40.46	37.92	36.99	32.70	61.45	47.78	49.00
+ TAP & ADL	43.7	33.87	34.24	32.04	51.29	61.51	44.76	54.39	38.14	40.12	35.27	36.43	32.25	62.12	46.78	49.55

Table 3. Comparisons of corruption error (lower is better) on individual corruption type of ImageNet-C based on FAN-B-Hybrid. Combining TAP and ADL together yields the best results on most of the corruption types.

4.1.2 Attention Stability and Visualization Results

We demonstrate that our models greatly improve attention stability against corruptions both qualitatively and quantitatively. Interestingly, in each layer, our models obtain a higher attention diversity among different attention heads.

Attention stability. In Figure 3 we first visualize how much the attention would be changed when facing image corruptions, e.g., Gaussian noise. We take FAN-B-Hybrid as the baseline and compare our best model with two variants that only contain TAP and ADL individually. Across diverse examples, the baseline model incurs a severe token overfocusing issue that it puts too much focus on very few tokens and comes with a significant attention shift when facing corruptions. With the help of our local aggregation module TAP, our TAP model assigns the attention to more tokens surrounding some important ones, alleviating the token overfocusing issue to some extent. Nevertheless, we still observe an attention shift between clean and corrupted examples. When training models with our ADL, the attention follows a diagonal pattern such that tokens aggregate information from the others while retaining most of the information from itself. We highlight that this diagonal pattern is somehow similar to the residual learning [22] that additionally learns a residual branch while keeping the features unchanged using an identity shortcut. Clearly, the attention becomes much more stable against corruptions. When combining TAP and ADL together, we further encourage the diagonal pattern to expand within a local region. In this way, each token puts more focus on its neighborhood beyond itself and thus obtains stronger features. Quantitatively, we also evaluate the attention stability by computing cosine similarity of attention between clean and corrupted examples across the whole ImageNet. From Figure 5, TAP yields higher attention stability than the baseline model. Thanks to the diagonal pattern in attention, the model with ADL greatly improves the similarity score by a large margin, indicating that the attention is very stable against corruptions. When combining both TAP and ADL, we can further improve attention stability.

Attention of each head and attention diversity. In this part, we further investigate the attention map in each individual head. As shown in Figure 4, for the baseline model, the most important tokens are always the most important ones in almost all the heads, resulting in a very low di-

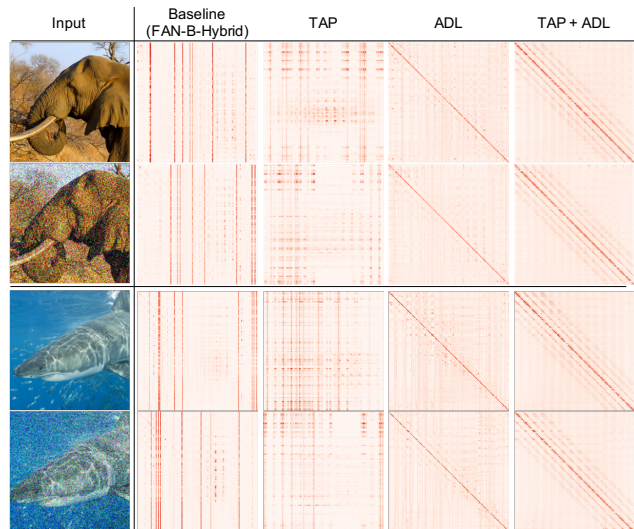


Figure 3. Comparisons of attention maps among different models. Compared with the baseline model, our TAP alleviates the overfocusing issue by encouraging the tokens surrounding the most important ones to have higher attention scores. When only using ADL, we obtain an attention map that follows the diagonal pattern, i.e., preserving the token itself while aggregating information from other tokens. Apparently, the attention rows are different from each other, fulfilling our goal of reducing similarity between rows. When combining TAP and ADL together, the diagonal pattern is further expanded to the nearby areas thanks to the TAP module.

versity of attention among different heads. By contrast, in our model, only two heads follow the diagonal pattern and the other six heads have different attention patterns from each other. We highlight that our attention is a combination of both local and global filters w.r.t. different heads. Specifically, these two diagonal heads can be regarded as local filters to extract local information since the diagonal pattern encourages tokens to aggregate information within their local neighborhood. As for the other six heads, the attention is distributed across the whole map and thus they can be regarded as global filters. From the visualization result in Figure 4, our model has a higher diversity of attention among different heads. To quantify this, following the previously discussed approach, we directly compute the cosine similarity of attention maps between any two heads in each layer (see detailed computation method in supple-




Figure 4. Attention maps of different attention heads in the last layer. For the baseline model, the most important tokens are often shared across different heads. In our model, two heads have the attention with the diagonal pattern and the other heads have specific patterns to extract different features, yielding higher attention diversity among heads (see quantitative results in Section 4.1.2).




Figure 5. Distributions of cosine similarity of intermediate attention maps between clean and corrupted examples (e.g., with Gaussian noise) on ImageNet. We demonstrate that either TAP or ADL is able to improve the stability of attention independently. When combining them together, we further improve the stability/similarity of attention against corruptions.

mentary). In other words, the lower the similarity score is, the higher the attention diversity will be. In practice, the baseline model obtains a high similarity score of 0.63 when averaging across the whole ImageNet, indicating a very low attention diversity among different heads. By contrast, our model yields a lower similarity of 0.27, which is consistent with the previous observation that our model produces diverse attention across different heads.

4.2. Results on Semantic Segmentation

In this part, we further apply our method to semantic segmentation tasks. We train the models on Cityscapes [11] and evaluate the robustness on two popular benchmarks, including Cityscapes-C [33] and ACDC [39]. Specifically, Cityscapes-C contains 16 corruption types which can be divided into 4 categories: noise, blur, weather, and digital. ACDC collects the images with adverse conditions, including night, fog, rain, and snow. In this paper, we report mIoU across diverse datasets. During training, we follow the same settings of SegFormer [55] to train our models. We demonstrate that our TAP and ADL also generalize well to segmentation tasks and significantly improve robustness.

Model	Cityscapes \uparrow	Cityscapes-C \uparrow	ACDC \uparrow
DeepLabv3+ (R101) [7]	77.1	39.4	41.6
ICNet [63]	65.9	28.0	-
DilatedNet [60]	68.6	30.3	-
Swin-T [29]	78.1	47.3	56.3
SETR [65]	79.5	63.1	60.2
Segformer-B5 [55]	82.4	65.8	62.0
FAN-B-Hybrid [67]	82.2	67.3	60.6
+ TAP	82.7 (+0.5)	69.2 (+1.9)	62.7 (+2.1)
+ ADL	82.4 (+0.2)	69.4 (+2.1)	63.1 (+2.5)
+ TAP & ADL	82.9 (+0.7)	69.7 (+2.4)	63.6 (+3.0)

Table 4. Comparisons of semantic segmentation models on Cityscapes validation set, Cityscapes-C, and ACDC test set. Both our TAP and ADL greatly improve the robustness. We can further improve the results when combining TAP and ADL together.

4.2.1 Quantitative Comparisons

As shown in Table 4, compared with the considered baseline model, using either our TAP or ADL individually can greatly improve the robustness on Cityscapes-C and ACDC. With the help of our effective local aggregation module TAP, we highlight that we also obtain a promising improvement of 0.5% mIoU on clean data. When combining TAP and ADL together, we further improve the performance and obtain a larger improvement of 2.4% and 3.0% on Cityscapes-C and ACDC, respectively. Moreover, our best model also significantly outperforms several popular segmentation models with comparable model size. Overall, these results indicate that the proposed two techniques not only work for image classification but also generalize well to semantic segmentation tasks.

4.2.2 Visual Comparisons

In this part, we compare the visualization results of the predicted segmentation masks based on examples of diverse robustness benchmarks. As shown in Figure 6, for the first example with snow corruptions, the baseline model cannot detect a large region of road (highlighted by the red




Figure 6. Visual comparisons of segmentation results. When facing image corruptions or adverse conditions, the baseline FAN-B-Hybrid model fails to detect some important objects (e.g., road in the first example) or mistakenly recognizes a part of car as a rider (in the second example). By contrast, our model is much more robust against these corruptions and adverse conditions.

box), which poses potential risks when applied in some real-world applications, e.g., autonomous driving. Moreover, in the second example with night conditions, the baseline model recognizes a part of the car as a rider and introduces a lot of artifacts in the predicted mask. By contrast, our model is much more robust and is able to accurately detect most parts of the road and the car in both cases. We highlight that our superiority of robustness can be observed on most examples of the considered benchmarks. Please refer to more visual comparisons in supplementary.

5. Analysis and Discussions

In the following, we present further ablation experiments and discussions. In Section 5.1, we demonstrate that the proposed two methods are general techniques and can be applied on top of diverse transformer architectures. In Section 5.2, we study the impact of the number of branches K in TAP and the weight of ADL in the training loss. In Section 5.3, we investigate the strategy of how to allocate TAP layers. In practice, uniformly allocating a TAP into every block yields the best results. In Section 5.4, we study the effect of the attention threshold τ used in Eqn. (1).

5.1. Effectiveness on Diverse Architectures

Besides RVT and FAN, we additionally apply our methods on top of more transformer architectures, including DeiT [43] and Swin [29]. In this experiment, we report the accuracy on ImageNet and robustness in terms of mCE (the lower the better) on ImageNet-C. As shown in Table 5, based on DeiT-B, we greatly improve the corruption robustness by reducing the mCE by 1.9% and yield a promising improvement of 0.4% on clean data. As for Swin-B, we obtain a similar observation that our methods are particularly effective in improving corruption robustness, reducing mCE from 54.4% to 51.9%. These results indicate that our methods can generalize well across diverse architectures.

Method	ImageNet	ImageNet-C (mCE) ↓
DeiT-B [43]	82.0	48.5
+ TAP & ADL	82.4 (+0.4)	46.6 (-1.9)
Swin-B [29]	83.4	54.4
+ TAP & ADL	84.0 (+0.6)	51.9 (-2.5)
RVT-B [32]	82.6	46.8
+ TAP & ADL	83.1 (+0.5)	44.7 (-2.1)
FAN-B-Hybrid [67]	83.9	46.1
+ TAP & ADL	84.3 (+0.4)	43.7 (-2.4)

Table 5. Results on top of diverse architectures. We report accuracy and mean corruption error (mCE) on ImageNet and ImageNet-C, respectively. Our method consistently improves robustness and accuracy across different architectures.

5.2. Impact of Hyperparameters K and λ

We conduct ablations on ImageNet-C to study the impact of two hyperparameters of our methods, including the number of branches K in TAP and the weight of ADL.

The number of branches K . As detailed in Figure 2, we build our TAP with K branches to enable tokens to consider diverse pooling areas. In fact, the value of K is an important factor for the performance of our method. It is worth noting that $K = 1$ is essentially equivalent to the baseline model without TAP. As shown in Figure 7 (top), we greatly improve the robustness with a lower mCE score on ImageNet-C when gradually increasing K from 1 to 4. If we further increase K , learning weights for too many candidate pooling areas (dilations) for each token becomes increasingly difficult and we cannot observe a significant improvement. Although additional branches only introduce minimal overhead in terms of model size and computational complexity, a large K would inevitably require a larger memory footprint. Thus, we choose $K = 4$ to obtain the best results at the minimal cost of extra memory footprint.

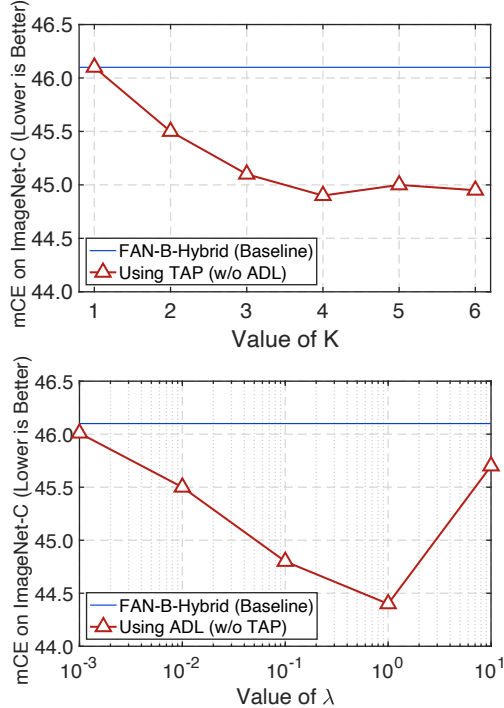


Figure 7. Robustness in terms of mean corruption error (mCE, lower is better) on ImageNet-C against the number of branches K (top) and the importance of our ADL loss λ (bottom). *Top*: When only introducing TAP without ADL, our model consistently outperforms the baseline model when increasing the value of K and yields the best result with $K = 4$. *Bottom*: When using ADL to train the model (without TAP), we observe that a too small or too large λ reduces the benefit of our method. In practice, $\lambda = 1$ performs best in most cases.

Weight of ADL λ . In Figure 7 (bottom), we change the value of λ in Eqn. (3). In practice, a larger λ encourages models to diversify the attention among different rows in the attention map more aggressively. Given a set of values $\lambda \in \{0.001, 0.01, 0.1, 1, 10\}$, we gradually improve the robustness (reduce mCE score) until $\lambda = 1$. When considering an even larger $\lambda = 10$, we observe a significant performance drop since a too large λ for ADL may hamper the standard cross-entropy loss during training. In this paper, we set $\lambda = 1$ and it generalizes well across all the considered architectures and learning tasks.

5.3. Strategies of Allocating TAP Layers

In this part, we explicitly investigate how to allocate TAP layers into a transformer model. As discussed in Section 3.2, TAP is a learnable module that can adjust itself to fit different layers. Specifically, when overfocusing is not severe, TAP can be reduced to identity mapping by setting the weight of the first branch to 1 in Figure 2. Empirically, in Table 6 (without ADL), when introducing TAP into

Method	Baseline	TAP (Shallow)	TAP (Middle+Deep)	TAP (All)
ImageNet \uparrow	83.9	84.1 (+0.2)	84.3 (+0.4)	84.3 (+0.4)
ImageNet-C \downarrow	46.1	45.7 (-0.4)	45.0 (-1.1)	44.9 (-1.2)

Table 6. Comparisons of accuracy on ImageNet and mCE (lower is better) on ImageNet-C between uniformly and non-uniformly allocating TAP. Taking FAN-B-Hybrid as the baseline, allocating TAP into every block yields better results than non-uniform allocation strategies in terms of both accuracy and robustness.

τ	0 (Baseline)	1	2	3	5
ImageNet \uparrow	83.9	83.9	84.0	84.0	83.9
ImageNet-C (mCE) \downarrow	46.1	45.1 (-1.0)	44.4 (-1.7)	45.5 (-0.6)	45.8 (-0.3)

Table 7. Comparisons of accuracy on ImageNet and mCE (lower is better) on ImageNet-C across diverse τ . We take FAN-B-Hybrid as the baseline and observe that a too small or too large τ reduces the benefit of ADL. In practice, $\tau = 2$ performs best in most cases.

shallow layers (first 30% layers) where overfocusing is not severe, we observe marginal improvement on ImageNet-C. However, for middle and deep layers (the rest 70% layers with overfocusing), we obtain large improvement, similar to uniformly allocating TAP. To avoid manually selecting layers, we suggest allocating TAP into every block.

5.4. Effect of Attention Threshold τ

According to Eqn. (1), we use a threshold τ/N with N being the number of tokens to filter out very small values and only focus on the most important ones in the attention map via $\hat{A}_i^{(l)} = \mathbb{1}(A_i^{(l)} \geq \tau/N) \cdot A_i^{(l)}$. Here, we explicitly study the effect of the attention threshold τ for computing our ADL. As detailed in Table 7, when using ADL to train the model (without TAP), our ADL only brings marginal improvements in terms of clean performance on ImageNet. However, our ADL becomes particularly effective in improving robustness, e.g., greatly reducing mCE on ImageNet-C. In practice, a too small or too large τ reduces the benefit of our ADL. In our experiments, we set $\tau = 2$ performs to obtain the best results.

6. Conclusion

In this paper, we address the token overfocusing issue in vision transformers (ViTs) such that ViTs tend to rely on very few important tokens in the attention mechanism. In fact, the attention is not robust and often obtains highly diverging attention patterns in the presence of corruptions. To alleviate this, we propose two general techniques. First, our Token-aware Average Pooling (TAP) module encourages the local neighborhood of tokens to take part in the self-attention by learning an adaptive average pooling scheme for each token. Second, our Attention Diversification Loss (ADL) explicitly reduces the cosine similarity of attention among tokens. In practice, we apply our methods to diverse architectures and obtain a significant improvement of robustness on different benchmarks and learning tasks.

References

- [1] Alaaeldin Ali, Hugo Touvron, Mathilde Caron, Piotr Bojanowski, Matthijs Douze, Armand Joulin, Ivan Laptev, Natalia Neverova, Gabriel Synnaeve, Jakob Verbeek, et al. Xcit: Cross-covariance image transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 2021. 2
- [2] Yutong Bai, Jieru Mei, Alan L Yuille, and Cihang Xie. Are transformers more robust than cnns? In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 2021. 2
- [3] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020. 2
- [4] Philipp Benz, Chaoning Zhang, Soomin Ham, Adil Karjauv, and I Kweon. Robustness comparison of vision transformer and mlp-mixer to cnns. In *Proceedings of the CVPR 2021 Workshop on Adversarial Machine Learning in Real-World Computer Vision Systems and Online Challenges (AML-CV)*, pages 21–24, 2021. 2
- [5] Srinadh Bhojanapalli, Ayan Chakrabarti, Daniel Glasner, Daliang Li, Thomas Unterthiner, and Andreas Veit. Understanding robustness of transformers for image classification. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, pages 10231–10241, 2021. 2
- [6] Chun-Fu Chen, Rameswar Panda, and Quanfu Fan. Regionvit: Regional-to-local attention for vision transformers. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2021. 2
- [7] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 7
- [8] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv:1904.10509*, 2019. 2
- [9] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. *Advances in Neural Information Processing Systems*, 34:9355–9366, 2021. 2
- [10] Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. On the relationship between self-attention and convolutional layers. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2020. 2, 3
- [11] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5, 7
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 5
- [13] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12124–12134, 2022. 2
- [14] Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *Proc. of the International Conference on Machine Learning (ICML)*, pages 2286–2296. PMLR, 2021. 5
- [15] Jiemin Fang, Lingxi Xie, Xinggang Wang, Xiaopeng Zhang, Wenyu Liu, and Qi Tian. Msg-transformer: Exchanging local spatial information by manipulating messenger tokens. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12063–12072, 2022. 2
- [16] Yonggan Fu, Shun Yao Zhang, Shang Wu, Cheng Wan, and Yingyan Lin. Patch-fool: Are vision transformers always robust against adversarial perturbations? In *Proc. of the International Conference on Learning Representations (ICLR)*, 2022. 1, 2
- [17] Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, pages 12259–12269, 2021. 2, 4
- [18] Jindong Gu, Volker Tresp, and Yao Qin. Are vision transformers robust to patch perturbations? In *Proc. of the European Conference on Computer Vision (ECCV)*, pages 404–421. Springer, 2022. 1
- [19] Yong Guo, David Stutz, and Bernt Schiele. Improving robustness by enhancing weak subnets. In *European Conference on Computer Vision*, pages 320–338. Springer, 2022. 1
- [20] Yong Guo, David Stutz, and Bernt Schiele. Improving robustness of vision transformers by reducing sensitivity to patch corruptions. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4108–4118, 2023. 5
- [21] Xing Han, Tongzheng Ren, Tan Minh Nguyen, Khai Nguyen, Joydeep Ghosh, and Nhat Ho. Robustify transformers with robust kernel density estimation. *arXiv.org*, 2210.05794, 2022. 2
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [23] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. 5
- [24] Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2019. 1, 5
- [25] Zilong Huang, Youcheng Ben, Guozhong Luo, Pei Cheng, Gang Yu, and Bin Fu. Shuffle transformer: Rethink-

- ing spatial shuffle for vision transformer. *arXiv preprint arXiv:2106.03650*, 2021. [2](#)
- [26] Salman H. Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM Comput. Surv.*, 54(10s):200:1–200:41, 2022. [2](#)
- [27] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Mvity2: Improved multiscale vision transformers for classification and detection. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4804–4814, 2022. [2](#), [4](#)
- [28] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12009–12019, 2022. [2](#)
- [29] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, pages 10012–10022, 2021. [2](#), [5](#), [7](#), [8](#)
- [30] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11976–11986, 2022. [5](#)
- [31] Kaleel Mahmood, Rigel Mahmood, and Marten van Dijk. On the robustness of vision transformers to adversarial examples. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [32] Xiaofeng Mao, Gege Qi, Yuefeng Chen, Xiaodan Li, Ranjie Duan, Shaokai Ye, Yuan He, and Hui Xue. Towards robust vision transformer. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#), [3](#), [5](#), [8](#)
- [33] Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander S Ecker, Matthias Bethge, and Wieland Brendel. Benchmarking robustness in object detection: Autonomous driving when winter is coming. *arXiv preprint arXiv:1907.07484*, 2019. [5](#), [7](#)
- [34] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:23296–23308, 2021. [2](#)
- [35] Haolin Pan, Yong Guo, Qinyi Deng, Haomin Yang, Jian Chen, and Yiqun Chen. Improving fine-tuning of self-supervised models with contrastive initialization. *Neural Networks*, 159:198–207, 2023. [2](#)
- [36] Zizheng Pan, Bohan Zhuang, Haoyu He, Jing Liu, and Jianfei Cai. Less is more: Pay less attention in vision transformers. In *Proc. of the Conference on Artificial Intelligence (AAAI)*, volume 36, pages 2035–2043, 2022. [3](#)
- [37] Sayak Paul and Pin-Yu Chen. Vision transformers are robust learners. In *Proc. of the Conference on Artificial Intelligence (AAAI)*, volume 36, pages 2071–2081, 2022. [2](#)
- [38] Zhiliang Peng, Wei Huang, Shanzhi Gu, Lingxi Xie, Yaowei Wang, Jianbin Jiao, and Qixiang Ye. Conformer: Local features coupling global representations for visual recognition. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, pages 367–376, 2021. [2](#), [4](#)
- [39] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Acdd: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10765–10775, 2021. [5](#), [7](#)
- [40] Yucheng Shi, Yahong Han, Yu-an Tan, and Xiaohui Kuang. Decision-based black-box attack against vision transformers via patch-wise adversarial removal. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:12921–12933, 2022. [1](#)
- [41] Zhouxing Shi, Huan Zhang, Kai-Wei Chang, Minlie Huang, and Cho-Jui Hsieh. Robustness verification for transformers. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2020. [2](#)
- [42] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *Proc. of the International Conference on Machine Learning (ICML)*, pages 10347–10357. PMLR, 2021. [2](#)
- [43] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *Proc. of the International Conference on Machine Learning (ICML)*, 2021. [3](#), [8](#)
- [44] Jesse Vig. A multiscale visualization of attention in the transformer model. In Marta R. Costa-jussà and Enrique Alfonseca, editors, *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 37–42, 2019. [2](#)
- [45] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019. [2](#)
- [46] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In *Proc. of the European Conference on Computer Vision (ECCV)*, pages 108–126. Springer, 2020. [2](#)
- [47] Sinong Wang, Belinda Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. [2](#)
- [48] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 568–578, 2021. [2](#)
- [49] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3):415–424, 2022. [2](#)
- [50] Wenxiao Wang, Lu Yao, Long Chen, Deng Cai, Xiaofei He, and Wei Liu. Crossformer: A versatile vision transformer

- based on cross-scale attention. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2021. [2](#)
- [51] Yujing Wang, Yaming Yang, Jiangang Bai, Mingliang Zhang, Jing Bai, Jing Yu, Ce Zhang, Gao Huang, and Yunhai Tong. Evolving attention with residual convolutions. In *Proc. of the International Conference on Machine Learning (ICML)*, pages 10971–10980. PMLR, 2021. [2](#)
- [52] Florian Wenzel, Andrea Dittadi, Peter Vincent Gehler, Carl-Johann Simon-Gabriel, Max Horn, Dominik Zietlow, David Kernert, Chris Russell, Thomas Brox, Bernt Schiele, et al. Assaying out-of-distribution generalization in transfer learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. [1](#)
- [53] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, pages 22–31, 2021. [2](#), [4](#)
- [54] Tete Xiao, Mannat Singh, Eric Mintun, Trevor Darrell, Piotr Dollár, and Ross Girshick. Early convolutions help transformers see better. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 30392–30400, 2021. [2](#), [4](#)
- [55] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:12077–12090, 2021. [7](#)
- [56] Bingna Xu, Yong Guo, Luoqian Jiang, Mianjie Yu, and Jian Chen. Downscaled representation matters: Improving image rescaling with collaborative downscaled images. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, 2023. [1](#)
- [57] Weijian Xu, Yifan Xu, Tyler Chang, and Zhuowen Tu. Co-scale conv-attentional image transformers. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, pages 9981–9990, 2021. [2](#)
- [58] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers. *arXiv preprint arXiv:2107.00641*, 2021. [2](#)
- [59] Donggeun Yoo, Sunggyun Park, Joon-Young Lee, and In So Kweon. Multi-scale pyramid pooling for deep convolutional representation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 71–80, 2015. [4](#)
- [60] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2016. [7](#)
- [61] Kun Yuan, Shaopeng Guo, Ziwei Liu, Aojun Zhou, Fengwei Yu, and Wei Wu. Incorporating convolution designs into visual transformers. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, pages 579–588, 2021. [2](#)
- [62] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, pages 558–567, 2021. [2](#), [5](#)
- [63] Hengshuang Zhao, Xiaojuan Qi, Xiaoyong Shen, Jianping Shi, and Jiaya Jia. Icnets for real-time semantic segmentation on high-resolution images. In *Proc. of the European Conference on Computer Vision (ECCV)*, pages 405–420, 2018. [7](#)
- [64] Zhengli Zhao, Dheeru Dua, and Sameer Singh. Generating natural adversarial examples. *Proc. of the International Conference on Learning Representations (ICLR)*, 2018. [1](#), [5](#)
- [65] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6881–6890, 2021. [7](#)
- [66] Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*, 2021. [2](#)
- [67] Daquan Zhou, Zhiding Yu, Enze Xie, Chaowei Xiao, Animesh Anandkumar, Jiashi Feng, and Jose M Alvarez. Understanding the robustness in vision transformers. In *Proc. of the International Conference on Machine Learning (ICML)*, pages 27378–27394. PMLR, 2022. [1](#), [2](#), [3](#), [5](#), [7](#), [8](#)