




VSA: Learning Varied-Size Window Attention in Vision Transformers

Qiming Zhang^{1*}, Yufei Xu^{1*}, Jing Zhang¹, and Dacheng Tao^{2,1}

¹ University of Sydney, Australia

² JD Explore Academy, China

{yuxu7116,qzha2506}@uni.sydney.edu.au,
jing.zhang1@sydney.edu.au, dacheng.tao@gmail.com

Abstract. Attention within windows has been widely explored in vision transformers to balance the performance, computation complexity, and memory footprint. However, current models adopt a hand-crafted fixed-size window design, which restricts their capacity of modeling long-term dependencies and adapting to objects of different sizes. To address this drawback, we propose **Varied-Size Window Attention (VSA)** to learn adaptive window configurations from data. Specifically, based on the tokens within each default window, VSA employs a window regression module to predict the size and location of the target window, *i.e.*, the attention area where the key and value tokens are sampled. By adopting VSA independently for each attention head, it can model long-term dependencies, capture rich context from diverse windows, and promote information exchange among overlapped windows. VSA is an easy-to-implement module that can replace the window attention in state-of-the-art representative models with minor modifications and negligible extra computational cost while improving their performance by a large margin, e.g., 1.1% for Swin-T on ImageNet classification. In addition, the performance gain increases when using larger images for training and test. Experimental results on more downstream tasks, including object detection, instance segmentation, and semantic segmentation, further demonstrate the superiority of VSA over the vanilla window attention in dealing with objects of different sizes. The code is available at <https://github.com/ViTAE-Transformer/ViTAE-VSA>.

1 Introduction

Recent Vision transformers have shown great potential in various vision tasks. By stacking multiple transformer blocks with vanilla attention, ViT [14] processes non-overlapping image patches and obtain superior classification performance. However, vanilla attention with quadratic complexity over the input length is hard to adapt to vision tasks with high-resolution images as input due to the expensive computational cost. To alleviate such issues, window-based attention [29] is proposed to partition the images into local windows and conduct attention

* Equal contribution.




Fig. 1. The comparison of the current works (hand-crafted windows) and the proposed VSA (varied-size windows).




Fig. 2. The performance with different image sizes.

within each window to balance the performance, computation complexity, as well as memory footprint. This mechanism enables vision transformers to make a great success in many downstream visual tasks [29,43,13,50,44,41,40,32]. However, it also enforces a spatial constraint on transformers’ attention distance, i.e., within the predefined window at each layer, thereby limiting the transformer’s ability to deal with objects at different scales.

Recent works have explored heuristic designs of attending to more tokens to alleviate such a spatial constraint. For example, Swin transformer [29] enlarges the window sizes from 7×7 to 12×12 when varying the image size from 224×224 to 384×384 , and sets the window size as 32×32 to deal with image size 640×640 in SwinV2 [28]. Some other methods try to find a good trade-off between attending to more tokens and increasing attention distance, *e.g.*, multiple window mechanisms have been explored in Focal attention [43], where coarse granularity tokens are involved in capturing long-distance information. Cross-shaped window attention [13] relaxes the spatial constraint of the window in vertical and horizontal directions and allows the transformer to attend to far-away relevant tokens along with the two directions while keeping the constraint along the diagonal direction. Pale [36] further increases the diagonal-direction attention distance by attending to tokens in the dilated vertical/horizontal directions. These methods have achieved superior performance in image classification tasks by enlarging the attention distance. However, they sacrifice computational efficiency and consume more memory, especially when training large models with high-resolution images. Besides, all these methods determine the window sizes heuristically. Intuitively, using a fixed-size window may be sub-optimal for dealing with objects of different sizes, although stacking more layers could mitigate this issue to some extent, which may also result in more parameters and optimization difficulty. In this paper, we argue that if the window can be relaxed to a varied-size rectangular one, whose size and position are learned directly from data, the transformer can capture rich context from diverse windows and learn more powerful object feature representation.

To this end, we propose a novel **V**aried-**S**ize **W**indow **A**ttention (VSA) mechanism to learn adaptive window configurations from data. Different from the previous window-based transformers where query, key, and value tokens are all sampled from the same window as shown in Figure 1(a), VSA employs a window regression module to predict the size and location of the target window based on the tokens within each default window. Then, the key and values tokens are sampled from the target window. By adopting VSA independently for each attention head, it enables the attention layers to model long-term dependencies, capture rich context from diverse windows, and promote information exchange among overlapped windows, as illustrated in Figure 1(b). VSA is an easy-to-implementation module that can replace the window attention in state-of-the-art representative models with minor modifications and negligible extra computational cost while improving their performance by a large margin, e.g., 1.1% for Swin-T on ImageNet classification. In addition, the performance gain increases when using larger images for training and test, as shown in Figure 2. With the larger images as input, Swin-T with predefined window sizes cannot adapt to large objects well, and the improvement brought by enlarging image sizes is marginal, i.e., a gain of 0.3% from 224×224 to 480×480 . In contrast, the performance gain of VSA over Swin-T increases significantly from 1.1% to 1.9%, owing to the varied-size window attention. Besides, as VSA can effectively promote information exchange across overlapped windows via token sampling, it does not need the shifted windows mechanism in Swin.

In conclusion, the contribution of this study is threefold. (1) We introduce a novel VSA mechanism that can directly learn adaptive window size and location from data. It breaks the spatial constraint of the fixed-size window in existing works and makes it easier for window-based transformers to adapt to objects at different scales. (2) VSA can serve as an easy-to-implement module to improve various window-based transformers, including but not limited to Swin [29,28] and ViTAEv2 [41,50], with minor modifications and negligible extra computational cost. (3) Extensive experimental results on public benchmarks demonstrate the superiority of VSA over the vanilla window attention on various visual tasks, including image classification, object detection, and semantic segmentation.

2 Related Work

2.1 Window-based vision transformers

Vision transformers [14] have demonstrated superior performance in many vision tasks by modeling long-term dependencies among local image patches (a.k.a. tokens) [40,24,23]. However, vanilla full attention performs poorly in training efficiency due to the shortage of inductive bias. To improve the efficiency, the following works either implicitly or explicitly introduce inductive bias into vision transformers [31,41,11,42] and obtain superior classification performance. After that, multi-stage design has been explored in [34,33,29,35,50] to better adapt vision transformers to downstream vision tasks. Among them, Swin [29]

is a representative work. By partitioning the tokens into non-overlapping windows and conducting attention within each window, Swin alleviates the huge computational cost caused by attention when dealing with larger input images. Although it balances the performance, computational cost, and memory footprints well, window-based attentions bring a spatial constraint on the attention distance due to the constant maximum size of windows. To alleviate such issues, different techniques have been explored to recover the transformer’s ability to model long-term dependency gradually, *e.g.*, using additional tokens for efficient cross-window feature exchange or designing delicate windows to allow the transformer layers to attend to far-away tokens in specific directions [16,13,36,22]. However, they still 1) rely on heuristic-designed windows for attention computation and 2) need to stack the transformers layers sequentially to enable feature exchange across all windows and model long-term dependencies. Thus, they lack the flexibility to adapt well to inputs of various sizes since their maximum attention distances are restricted by the constant and data-agnostic window size and model depth.

Unlike them, the proposed VSA estimates window sizes and locations adaptively based on input features and calculates attention within such windows. Therefore, VSA allows transformer layers to model long-term dependencies, capture rich context, and promote cross-window information exchange from diverse varied-size windows. As VSA learns the window sizes in a data-driven manner, it can benefit window-based vision transformers to adapt to objects at various scales and thus helps boost their performance on image classification, object detection, and semantic segmentation.

2.2 Deformable sampling

Deformable sampling has been widely explored previously to help the convolution networks [10,52] to focus on regions of interest and extract better features. Similar mechanisms have been exploited in deformable-DETR [53] to help the transformer detector to find and utilize the most valuable token features for object detection in a sparse manner. Recently, DPT [6] designs deformable patch merging layers based on PVT [34] to help the transformer to preserve better features after downsampling. VSA, from another perspective, introduces learnable varied-size window attention into transformers. By flexibly estimating the window sizes and locations for attention calculation, VSA breaks the spatial constraint of fixed-size windows and makes it easier for window-based transformers to better adapt to the objects at various scales.

3 Method

In this section, we will take Swin transformer [29] as an example and give a detailed description of applying VSA in Swin. The details of incorporating VSA into ViTAE [50] will be presented in the supplementary.

3.1 Preliminary

We will first briefly review the window attention operation in the baseline method Swin transformer. Given the input features $X \in \mathcal{R}^{H \times W \times C}$ as input, Swin transformer employs several window-based attention layers for feature extraction. In each window-based attention layer, the input features are firstly partitioned into several non-overlapping windows, *i.e.*, $\{X_w^i \in \mathcal{R}^{w \times w \times C} | i \in [1, \dots, \frac{H \times W}{w^2}]\}$, where w is the predefined window size. After that, the partitioned tokens are flatten along the spatial dimension and projected to query, key, and value tokens, *i.e.*, $\{Q_{w,f}^i, K_{w,f}^i, V_{w,f}^i \in \mathcal{R}^{w^2 \times N \times C'} | i \in [1, \dots, \frac{H \times W}{w^2}]\}$, where Q, K, V represent the query, key, and value tokens, respectively, N denotes the head number and C' is the channel dimension along each head. It is noted that $N \times C'$ equals the channel dimension C of the given feature. Given the flattened query, key, and value tokens from the same default window, the window-based attention layers conduct full attention within the window, *i.e.*,

$$F_{w,f}^i = MHSA(Q_{w,f}^i, K_{w,f}^i, V_{w,f}^i). \quad (1)$$

The $F_{w,f}^i \in \mathcal{R}^{w^2 \times N \times C'}$ is the features after attention and $MHSA$ represents the vanilla multi-head self-attention operation [14]. The relative position embeddings are utilized during the attention calculation to encode spatial information into the features. The extracted features F are reshaped back to the window shape, *i.e.*, $F_w^i \in \mathcal{R}^{w \times w \times C}$, and added with the input feature X_w^i . The same operation is individually repeated for each window and the generated features from all windows are then concatenated to recover the shape of input features. After that, an FFN module is employed to refine the extracted features, which contains two linear layers with hidden dimension αC , where α is the expansion ratio. For notation simplification, we dismiss the window index notation i in the following since each window's operations are the same.

With the usage of window-based attention, the computational complexity decreases to linear to the input size, *i.e.*, each window attention's complexity is $\mathcal{O}(w^4 C)$ and the computation complexity of window attention for each image is $\mathcal{O}(w^2 HWC)$. To bridge connections between different windows, shifted operations are used between two adjacent transformer layers in Swin [29]. As a result, the receptive field of the model is gradually enlarged with layers stacking in sequence. However, current window-based attentions restrict the attention area of the tokens within the corresponding hand-crafted window at each transformer layer. It limits the model's ability to capture far-away contextual information and learn better feature representations for objects at different scales.

3.2 Varied-size window attention

Base window generation. Rather than stacking layers with hand-crafted windows to gradually enlarge the receptive field, our VSA allows the query tokens to attend to far-away regions and empower the network with the flexibility to determine the target window size, *i.e.*, attention area, given specific input data

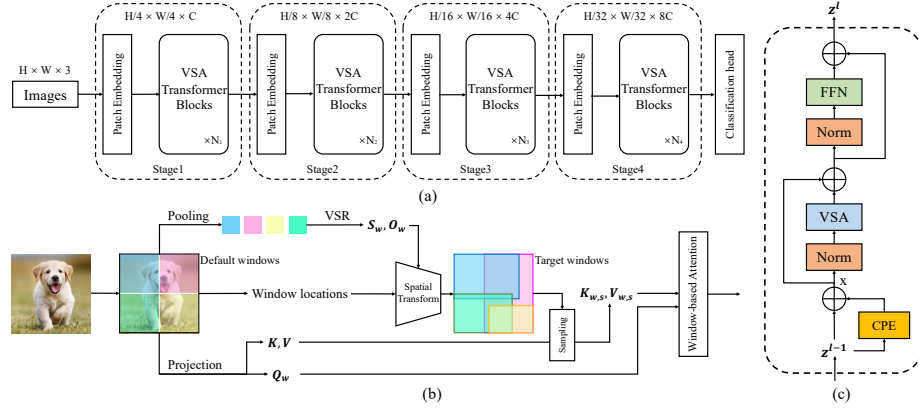


Fig. 3. The pipeline of the transformer with our proposed varied-size window attention. (a) The overall structure of stacking VSA transformers blocks; (b) The details of the proposed VSA module; (c) The pipeline of the VSA transformer block.

at each layer. VSA only needs to make minor modifications to the basic structure of backbone networks and serves as an easy-to-implement module to replace the vanilla window attention in window-based transformers as in Figure 3 (a). Technically, given the input features X , VSA first partitions these tokens into several windows X_w with the predefined window size w , following the baseline methods' routine. We refer to these windows as default windows and get the query features from the default windows, *i.e.*,

$$Q_w = \text{Linear}(X_w). \quad (2)$$

Varied-size window regression module. To estimate the size and location of the target window for each default window, VSA considers the size and location of the default window as a reference and adopts a varied-size window regression (VSR) module to predict the scale and offset upon the references as shown in Figure 3(b). The VSR module consists of an average pooling layer, a LeakyReLU [39] activation layer, and a 1×1 convolutional layer with stride 1 in sequence. The kernel size and stride of the pooling layer follow the default window size, *i.e.*,

$$S_w, O_w = \text{Conv} \circ \text{LeakyReLU} \circ \text{AveragePool}(X_w), \quad (3)$$

where S_w and $O_w \in \mathcal{R}^{2 \times N}$ represent the estimated scales and offsets in the horizontal and vertical directions w.r.t. the default window locations, independently for N attention heads. The generated windows are referred to as target windows.

Varied-size window-based attention. We first get the key and value tokens $K, V \in \mathcal{R}^{H \times W \times C}$ from the feature map X , *i.e.*,

$$K, V = \text{Reshape} \circ \text{Linear}(X). \quad (4)$$

Then the VSA module uniformly samples M features from each varied-size window over K, V respectively, and obtains $K_{w,v}, V_{w,v} \in \mathcal{R}^{M \times N \times C'}$ to serve as the key/value tokens for the query tokens Q_w . To keep the computational cost as window attention, we set M equal to $w \times w$. The sampled tokens $K_{w,v}, V_{w,v}$ are then fed into $MHSA$ with queries Q_w for attention calculation. However, as the key/value tokens are sampled from different locations with the query tokens, the relative position embeddings between the query and key tokens may not describe the spatial relationship well. Following the spirit in CPVT [8], we adopt conditional position embedding (CPE) before the MHSA layers to supply the spatial relationships into the model as shown in Figure 3 (c), *i.e.*,

$$X = Z^{l-1} + CPE(Z^{l-1}), \quad (5)$$

where Z^{l-1} is the feature from the previous transformer block and CPE is implemented by a depth-wise convolution layer with kernel size equal the window size, *i.e.*, 7×7 by default, and stride 1.

3.3 Computation complexity analysis

The extra computations caused by VSA come from the CPE and VSR module, while the other parts, including the window-based multi-head self-attention and FFN network, are exactly the same as the baseline models. Given the input features $X \in \mathcal{R}^{H \times W \times C}$, VSA firstly uses a depth-wise convolutional layer with 7×7 kernels to generate CPE, which brings extra $\mathcal{O}(49 \cdot HWC)$ computations. In the VSR module, we first employ an average pooling layer with kernel size and stride equal to the window size to aggregate features from the default windows, whose complexity is $\mathcal{O}(HWC)$. The following activation function does not introduce extra computations, and the last convolutional layer with kernel size 1×1 takes $X_{pool} \in \mathcal{R}^{\frac{H}{w} \times \frac{W}{w} \times C}$ as the input and estimates the scales S_w and offsets O_w . Both the scales and offsets belong to $\mathcal{R}^{2 \times N}$. Thus, the computational complexity of the convolutional layer is $\mathcal{O}(\frac{4N}{w^2} HWC)$, where N is the number of the attention heads in the transformer layers, and w is the window size. After obtaining the scales and offsets, we transform the default windows to the varied-size windows and uniformly sample $w \times w$ tokens within each target window. The computational complexity for each window is $w^2 \times 4 \times C$, and the total computational complexity for the sampling operation is $\mathcal{O}(4 \cdot HWC)$. Thus, the total extra computations brought by VSA is $\mathcal{O}\{(54 + \frac{4N}{w^2})HWC\}$, which is far less ($\leq 5\%$) than the total computational cost of the baseline models, regarding the complexity of FFN is $\mathcal{O}(2\alpha HWC^2)$ and C is always larger than 96.

4 Experiments

4.1 Implementation details

We evaluate the performance of the proposed VSA based on Swin [29] and ViTAEv2 [50]. The former is a pure transformer model with shifted windows

between two adjacent layers, while the latter is an improved transformer model by introducing convolution inductive bias, which models long- and short-term dependencies jointly. In this paper, we adopt the full-window version of ViTAEv2 as the baseline. All the models are trained for 300 epochs from scratch on the standard ImageNet-1k [12] dataset with input resolution 224×224 . We follow the hyper-parameters setting in the baseline methods to train the variants with VSA, *e.g.*, we use AdamW [30] optimizer with cosine learning rate schedulers during training. A 20-epochs linear warm-up is utilized following Swin [29] to stabilize training. The initial learning rate is set to 0.001 for 1024 batch size during training. The data augmentation is the same as [29] and [50], *i.e.*, random cropping, auto-augmentation [25], CutMix [47], MixUp [48], and random erasing are used to augment the input images. Besides, label smoothing with a weight of 0.1 is adopted. It is also noteworthy that there is no shifted window mechanism in the models with VSA, since VSA enables cross-window information exchange among overlapped varied-size windows.

4.2 Image Classification on ImageNet

We evaluate the classification performance of different models on the ImageNet [12] validation set. As shown in Table 1, the proposed VSA helps boost the classification accuracy of Swin transformer by 1.1% absolute Top-1 accuracy, *i.e.* from 81.2% to 82.3%, even without the shifted window mechanisms. It indicates that VSA can flexibly determine the appropriate window sizes and locations given the input features, allow the tokens to effectively attend far-away but relevant tokens outside the default windows to extract rich context, and learn better feature representations. Besides, Swin-T with VSA obtains comparable performance with MSG-T [16], which adopts extra messenger tokens for feature exchange across windows, *i.e.*, 82.3% v.s. 82.4%, demonstrating that our varied-size window mechanism can enable sufficient feature exchange across windows without the need of using extra tokens. For ViTAEv2 [50], ViTAEv2-S with VSA obtains 82.7% (+0.5%) classification accuracy with only 20M parameters, demonstrating that the proposed varied-size window attention is compatible with not only the transformers with vanilla window attentions but also those with convolutions for feature exchange across windows.

When scaling the input images to higher resolutions, *i.e.*, from 224×224 to 384×384 and 480×480 , the performance gains from VSA become larger owing to its ability to learn adaptive target window sizes from data. Specifically, the performance gain brought by VSA increases from 1.1% to 1.8% absolute accuracy over Swin-T when scaling the input size from 224 to 384, respectively. For the 480×480 input resolution, the performance gain of VSA further increases to 1.9%, while the Swin transformers only benefit from the higher resolution marginally (*i.e.*, 0.2%). The reason is that the fixed-size window attention in Swin limits the attention region at each transformer layer, which brings difficulty in handling objects at different scales. In contrast, VSA can learn to vary the window size to adapt to the objects and capture rich contextual information from different

Model	Params	FLOPs	Input	ImageNet [12]		Real [1]
	(M)	(G)	Size	Top-1	Top-5	Top-1
DeiT-S [31]	22	4.6	224	81.2	95.4	86.8
PVT-S [34]	25	3.8	224	79.8	-	-
ViL-S [49]	25	4.9	224	82.4	-	-
PiT-S [21]	24	4.8	224	80.9	-	-
TNT-S [18]	24	5.2	224	81.3	95.6	-
MSG-T [16]	25	3.8	224	82.4	-	-
Twins-PCPVT-S [7]	24	3.8	224	81.2	-	-
Twins-SVT-S [7]	24	2.9	224	81.7	-	-
T2T-ViT-14 [46]	22	5.2	224	81.5	95.7	86.8
Swin-T [29]	29	4.5	224	81.2	-	-
Swin-T+VSA	29	4.6	224	82.3	96.1	87.5
ViTAEv2-S ¹ [50]	20	5.4	224	82.2	96.1	87.5
ViTAEv2-S¹+VSA	20	5.6	224	82.7	96.3	87.8
Swin-T [29]	29	14.2	384	81.4	95.4	86.4
Swin-T+VSA	29	14.9	384	83.2	96.5	88.0
Swin-T [29]	29	23.2	480	81.5	95.7	86.3
Swin-T+VSA	29	24.0	480	83.4	96.7	88.0
PiT-B [21]	74	12.5	224	82.0	-	-
TNT-B [18]	66	14.1	224	82.8	96.3	-
Focal-B [43]	90	16.0	224	83.8	-	-
ViL-B [49]	56	13.4	224	83.7	-	-
MSG-S [16]	56	8.4	224	83.4	-	-
PVTv2-B5 [33]	82	11.8	224	83.8	-	-
Swin-S [29]	50	8.7	224	83.0	-	-
Swin-S+VSA	50	8.9	224	83.8	96.8	88.54
Swin-B [29]	88	15.4	224	83.3	-	88.0
Swin-B+VSA	88	16.0	224	83.9	96.7	88.6

¹ The full window version.

Table 1. Image classification results on ImageNet. ‘Input Size’ denotes the image size used for training and test.

attention heads at each layer, which is beneficial for learning powerful object feature representations.

4.3 Object detection and instance segmentation on MS COCO

Settings. We evaluate the backbone models for the object detection and instance segmentation tasks on the MS COCO [27] dataset, which contains 118K training, 5K validation, and 20K test images with full annotations. We adopt the models trained on ImageNet with 224×224 input resolutions as backbones and use three typical object detection frameworks, *i.e.*, the two-stage frameworks Mask RCNN [19] and Cascade RCNN [2,3], and the one-stage framework RetinaNet [26]. We follow the common practice in mmdetection [5], *i.e.*, multi-scale training with an AdamW optimizer and a batch size of 16. The initial learning rate is 0.0001 and the weight decay is 0.05. We adopt both $1 \times$ (12 epochs) and $3 \times$ (36 epochs) training schedules for the Mask RCNN framework to evaluate the object detection performance w.r.t. different backbones. For RetinaNet and Cascade RCNN, the models are trained with $1 \times$ and $3 \times$ schedules, respectively. The results on other settings are reported in the supplementary.

Results. The results of baseline models and those with VSA on the MS COCO dataset with Mask RCNN, RetinaNet, and Cascade RCNN are reported in Tables 2, 3, and 4, respectively. Compared to the baseline method Swin-T [29] and

	Params (M)	Mask RCNN 1x						Mask RCNN 3x					
		AP ^{bb}	AP ^{bb} ₅₀	AP ^{bb} ₇₅	AP ^{mk}	AP ^{mk} ₅₀	AP ^{mk} ₇₅	AP ^{bb}	AP ^{bb} ₅₀	AP ^{bb} ₇₅	AP ^{mk}	AP ^{mk} ₅₀	AP ^{mk} ₇₅
ResNet50 [20]	44	38.6	59.5	42.1	35.2	56.3	37.5	40.8	61.2	44.4	37.0	58.4	39.3
ViL-S [45]	45	44.9	67.1	49.3	41.0	64.2	44.1	47.1	68.7	51.5	42.7	65.9	46.2
PVT-M [34]	64	42.0	64.4	45.6	39.0	61.6	42.1	-	-	-	-	-	-
PVT-L [34]	81	42.9	65.0	46.6	39.5	61.9	42.5	-	-	-	-	-	-
PVTv2-B2 [33]	45	45.3	67.1	49.6	41.2	64.2	44.4	-	-	-	-	-	-
CMT-S [17]	45	44.6	66.8	48.9	40.7	63.9	43.4	-	-	-	-	-	-
RegionViT-S [4]	50	42.5	-	-	39.5	-	-	46.3	-	-	42.3	-	-
XCiT-S12/16 [15]	44	-	-	-	-	-	-	45.3	67.0	49.5	40.8	64.0	43.8
DPT-M [6]	66	43.8	66.2	48.3	40.3	63.1	43.4	44.3	65.6	48.8	40.7	63.1	44.1
ResT-Base [51]	50	41.6	64.9	45.1	38.7	61.6	41.4	-	-	-	-	-	-
Shuffle-T [22]	48	-	-	-	-	-	-	46.8	68.9	51.5	42.3	66.0	45.6
Focal-T [43]	49	44.8	-	-	41.0	-	-	47.2	69.4	51.9	42.7	66.5	45.9
Swin-T [29]	48	43.7	66.6	47.7	39.8	63.3	42.7	46.0	68.1	50.3	41.6	65.1	44.9
Swin-T+VSA	48	45.6	68.4	50.1	41.4	65.2	44.4	47.5	69.4	52.3	42.8	66.3	46.0
ViTAEv2-S ¹ [50]	39	43.5	65.8	47.4	39.4	62.6	41.8	44.7	65.8	49.1	40.0	62.6	42.8
ViTAEv2-S¹+VSA	39	45.9	68.2	50.4	41.4	65.1	44.5	48.1	69.8	52.9	42.9	66.9	46.2

¹ The full window version.**Table 2.** Object detection results on MS COCO with Mask RCNN.

ViTAEv2 [50], their VSA variants obtain better performance on both object detection and instance segmentation tasks with all detection frameworks, *e.g.*, VSA brings a gain of 1.9 and 2.4 mAP^{bb} for Swin-T and ViTAEv2-S with Mask RCNN 1× training schedule, confirming that VSA learns better object features than the vanilla window attention via the varied-size window attention that can better deal with objects at different scales for object detection. Besides, a longer training schedule (3×) also sees a significant performance gain from VSA over the vanilla window attention. For example, the performance gain of VSA on Swin-T and ViTAEv2 reaches 1.5 mAP^{bb} and 3.4 mAP^{bb}, respectively. We attribute this to the better attention regions learned by the VSR module in our VSA under longer training epochs. Similar conclusions can also be drawn when using RetinaNet [26] and Cascade RCNN [2] as detection frameworks, where VSA brings a gain of at least 2.0 and 1.2 mAP^{bb}, respectively. It is also noteworthy that the performance gains on ViTAEv2 are more significant than those on Swin-T. This is because there is no shifted window mechanism existing in ViTAEv2, and thus the ability to model long-range dependencies via attention is constrained within each window. In contrast, the varied size window attention in VSA empowers ViTAEv2 models to have such an ability and efficiently exchange rich contextual information across windows.

4.4 Semantic segmentation on Cityscapes

Settings. The Cityscapes [9] dataset is adopted to evaluate the performance of different backbones for semantic segmentation. The dataset contains over 5K well-annotated images of street scenes from 50 different cities. UperNet [38] is adopted as the segmentation framework. The training and evaluation of the models follow the common practice, *i.e.*, using the Adam optimizer with polynomial learning rate schedulers. The models are trained for 40k iterations and 80k iterations separately with both 512×1024 and 769×769 input resolutions.

	Params (M)	RetinaNet		
		AP ^{bb}	AP ^{bb} ₅₀	AP ^{bb} ₇₅
ResNet50 [20]	38	36.3	55.3	38.6
PVTv2-B1 [33]	24	41.2	61.9	43.9
ResT-Base [51]	41	42.0	63.2	44.8
DAT-T [37]	39	42.8	64.4	45.2
Twins-SVT-S [7]	34	42.3	63.4	45.2
Swin-T [29]	39	41.6	62.1	44.2
Swin-T+VSA	39	43.6	64.8	46.6
ViTAEv2-S ¹ [50]	30	42.1	62.7	44.8
ViTAEv2-S¹+VSA	30	44.3	65.2	47.6

¹ The full window version.

	Params (M)	Cascade RCNN		
		AP ^{bb}	AP ^{bb} ₅₀	AP ^{bb} ₇₅
ResNet50 [20]	82	44.3	62.4	48.5
PVTv2-B2 [33]	83	51.1	69.8	55.3
PVTv2-B2-Li [33]	80	50.9	69.5	55.2
MSG-T [16]	83	51.4	70.1	56.0
Swin-T [29]	86	50.2	68.8	54.7
Swin-T+VSA	86	51.4	70.4	55.9
ViTAEv2-S ¹ [50]	77	48.0	65.7	52.5
ViTAEv2-S¹+VSA	77	51.9	70.6	56.2

¹ The full window version.**Table 3.** Object detection results on MS COCO [27] with RetinaNet [26].**Table 4.** Object detection results on MS COCO [27] with Cascade RCNN [2].

	512×1024						769×769					
	40k			80k			40k			80k		
	mIoU	mAcc	mIoU*	mIoU	mAcc	mIoU*	mIoU	mAcc	mIoU*	mIoU	mAcc	mIoU*
ResNet50 [20]	77.1	84.0	78.4	78.2	84.6	79.2	78.0	86.7	79.7	79.4	87.2	80.9
Swin-T [29]	78.9	85.3	79.9	79.3	85.7	80.2	79.3	86.7	79.8	79.6	86.6	80.1
Swin-T+VSA	80.8	87.6	81.7	81.5	87.8	82.4	81.0	88.0	81.9	81.6	88.3	82.5
ViTAEv2-S [50]	80.1	86.5	80.9	80.8	87.0	81.0	79.6	86.1	80.6	80.5	86.8	81.2
ViTAEv2-S+VSA	81.4	87.9	82.3	82.2	88.6	83.0	80.6	87.1	81.4	81.5	88.0	82.4
Swin-S [29]	80.7	87.3	82.0	81.2	87.4	82.2	80.9	87.8	81.6	81.5	88.1	82.3
Swin-S+VSA	82.1	88.5	83.2	82.8	88.9	83.6	82.0	88.7	83.0	82.8	89.5	83.6

Table 5. Semantic segmentation results on Cityscapes [9] with UperNet [38]. * denotes results are obtained with multi-scale test.

Results. The results are available in Table 5. With 512×1024 input size, VSA brings over 1.3 mIoU and 1.4 mAcc gains for both Swin-T [29] and ViTAEv2-S [50], no matter with 40k or 80k training schedules. This observations hold with 769×769 resolution images as input, where VSA brings over 1.0 mIoU and 1.0 mAcc gains for both models. Such phenomena validates the effectiveness of the proposed VSA in improving the baseline models’ performance on semantic segmentation tasks. With more training iterations (80k), the performance gains of VSA over Swin-T increases from 1.9 to 2.2 mIoU with 512×1024 and from 1.7 to 2.0 mIoU with 769×769, owing to the better attention regions learned by the VSR module. Besides, with multi-scale testing, the performance of using VSA further improves, indicating that VSA can implicit capture multi-scale features as the target windows have different scales and locations for each head.

4.5 Ablation Study

We adopt Swin-T [29] with VSA for ablation studies. The models are trained for 300 epochs with AdamW optimizer. To find the optimal configuration of VSA, we gradually substitute the window attention in different stages of Swin with VSA. The results are shown in Table 6, where ✓ indicates that VSA replaces the vanilla window attention. We can see that the performance gradually improves with more VSA used and reaches the best when using VSA in all four stages. Meanwhile, it only takes a few extra parameters and FLOPs. Therefore, we choose to use VSA at all stages as the default setting in this paper.

VSA at stages				FLOPs Param		Acc.
stage 1	stage 2	stage 3	stage 4	(G)	(M)	
			✓	4.5	28.2	81.2
		✓	✓	4.5	28.3	81.4
	✓	✓	✓	4.6	28.7	81.9
	✓	✓	✓	4.6	28.7	82.1
✓	✓	✓	✓	4.6	28.7	82.3

CPE	VSR	Shift	Acc.
		✓	
	✓		81.2
✓	✓		81.6
✓		✓	81.6
✓	✓		82.3
✓	✓	✓	82.3

Table 6. The ablation study of using VSA in each stage of Swin-T [29].

Table 7. The ablation study of each component in VSA based on Swin-T [29].

We take Swin-T as the baseline and further validate the contribution of each component in VSA. The results are available in Table 7, where ✓ denotes using the specific component. ‘Shift’ is short for the shifted window mechanism. With only ‘Shift’ marked, the model becomes the baseline Swin-T. As can be seen, the model with ‘VSR’ alone outperforms Swin-T by 0.3% absolute accuracy, implying (1) the effectiveness of varied-size windows in cross-window information exchange and (2) the advantage of adapting the window sizes and locations, *i.e.*, attention regions, to the objects at different scales. Besides, using CPE and VSR in VSA further boosts to 82.3%, which outperforms the variant of ‘CPE’ + ‘Shift’ by 0.6% accuracy. It indicates that CPE is better compatible with varied size windows by providing local positional information. It is also noteworthy that there is no need to use the shifted-window mechanism in VSA according to the results in the last two rows, confirming that varied-size windows can guarantee the feature exchange across overlapped windows.

4.6 Throughputs & GPU memory comparison

	Throughputs on A100	Throughputs on V100	Memory
	(fps)	(fps)	(G)
Swin-T	1557	679	15.8
Swin-T+VSA	1297	595	16.1
Swin-S	961	401	23.0
Swin-S+VSA	769	352	23.5

Table 8. Throughput & GPU memory comparison with VSA.

We also evaluate the model’s throughputs during inference and GPU memory consumption during training, with batch size 128 and input resolution 224×224 . We run each model 20 times firstly as warmup and count the average throughputs of the subsequent 30 runs as the throughputs of the models. All of the experiments are conducted on the NVIDIA A100 and V100 GPUs. As shown in Table 8, VSA slows down the Swin model by about 12%~17% on different hardware platforms and consumes 2% more GPU memory, with much better performance on both classification and downstream dense prediction tasks. Such slow-down and extra memory consumption is mainly due to the sub-optimal optimization of sampling operations compared with the matrix multiply operations in the PyTorch framework, where the latter is sufficiently optimized with

cuBLAS. Integrating the sampling operation with following linear projection operations with CUDA optimization can help alleviate the speed concerns, which we leave as our future work to implement the proposed VSR module better.

4.7 Visual inspection and analysis

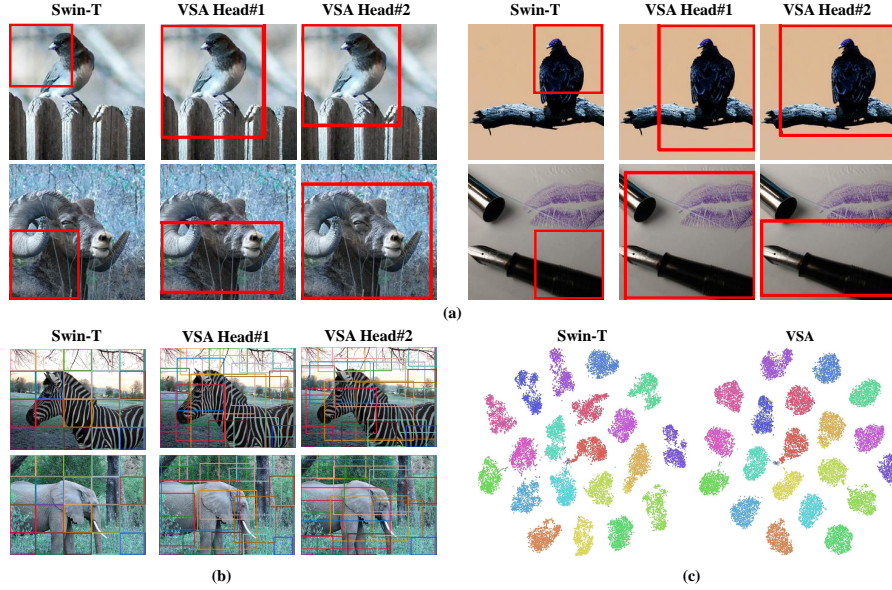


Fig. 4. Visualization of the varied-size windows generated by VSA from ImageNet (a) and MS COCO (b). The t-SNE analysis is also provided in (c).

Visualization of target windows. We visualize the default windows used in Swin-T [29] and the varied-size windows generated by VSA on images from the ImageNet [12] and MS COCO [27] datasets to see where VSA learns to attend for different images. The results are visualized in Figure 4. As shown in Figure 4(a), the generated windows from VSA can better cover the target objects in the images while the fixed-size windows adopted in Swin can only capture part of the targets. It can also be inferred from Figure 4(b) that the windows generated by different heads in VSA have different sizes and locations to focus on different parts of the targets, which helps to capture rich contextual information and learn better object feature representations. Besides, the windows that cover the target objects have more variance in size and location compared with those covering background as shown in (b), *e.g.*, the windows on the zebra and elephant vary (the blue, red, orange, pink, *etc.*) significantly while others in the background are less varied. In addition, the target windows are overlapped with each other, thus enabling abundant cross-window feature exchange and making it possible to drop the shifted window mechanism in VSA.

t-SNE analysis. We further use t-SNE to analyze the features generated by Swin-T models with and without VSA. We randomly select 20 categories from the ImageNet dataset and use t-SNE to visualize the extracted features. As shown in Figure 4(c), the features generated by Swin-T with VSA are better clustered, demonstrating that VSA can help the models deal with objects of different sizes and learn more discriminative features.

5 Limitation and Discussion

Although VSA has been proven efficient in dealing with images of varied resolutions and has shown its effectiveness on various vision tasks, including classification, detection, instance segmentation, and semantic segmentation, we only evaluate VSA with Swin [29] and ViTAEv2 [50] in this paper. It will be our future work to explore the usage of VSA on other transformers with window-based attentions, *e.g.*, CSwin [13] and Pale [36], which use cross-shaped attentions. Besides, to keep the computational cost as the vanilla window attention, we only sample sparse tokens from each target window, *i.e.*, the number of sampled tokens equals the default window size, which may ignore some details when the window becomes extremely large. Although the missed details may be complemented from other windows via feature exchange, a more efficient sampling strategy can be explored in the future study.

6 Conclusion

This paper presents a novel varied-size window attention (VSA), *i.e.*, an easy-to-implement module that can help boost the performance of representative window-based vision transformers such as Swin in various vision tasks, including image classification, object detection, instance segmentation, and semantic segmentation. By estimating the appropriate window size and location for each image in a data-driven manner, VSA enables the transformers to attend to far-away yet relevant tokens with negligible extra computational cost, thereby modeling long-term dependencies among tokens, capturing rich context from diverse windows, and promoting information exchange among overlapped window. In the future, we will investigate the usage of VSA in more attentions types including cross-shaped windows, axial attentions, and others as long as they can be parameterized w.r.t. size (*e.g.*, height, width, or radius), rotation angle, and position. We hope that this study can provide useful insight to the community in developing more advanced attention mechanisms as well as vision transformers.

Acknowledgement Mr. Qiming Zhang, Mr. Yufei Xu, and Dr. Jing Zhang are supported by ARC FL-170100117.

References

1. Beyer, L., Hénaff, O.J., Kolesnikov, A., Zhai, X., Oord, A.v.d.: Are we done with imagenet? arXiv preprint arXiv:2006.07159 (2020)
2. Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6154–6162 (2018)
3. Cai, Z., Vasconcelos, N.: Cascade r-cnn: High quality object detection and instance segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2019)
4. Chen, C.F., Panda, R., Fan, Q.: Regionvit: Regional-to-local attention for vision transformers. In: International Conference on Learning Representations (2022)
5. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 (2019)
6. Chen, Z., Zhu, Y., Zhao, C., Hu, G., Zeng, W., Wang, J., Tang, M.: Dpt: Deformable patch-based transformer for visual recognition. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 2899–2907 (2021)
7. 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: Advances in Neural Information Processing Systems (2021)
8. Chu, X., Tian, Z., Zhang, B., Wang, X., Wei, X., Xia, H., Shen, C.: Conditional positional encodings for vision transformers. arXiv preprint arXiv:2102.10882 (2021)
9. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
10. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 764–773 (2017)
11. Dai, Z., Liu, H., Le, Q.V., Tan, M.: Coatnet: Marrying convolution and attention for all data sizes. In: Advances in Neural Information Processing Systems (2021)
12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 248–255. Ieee (2009)
13. Dong, X., Bao, J., Chen, D., Zhang, W., Yu, N., Yuan, L., Chen, D., Guo, B.: Cswin transformer: A general vision transformer backbone with cross-shaped windows. arXiv preprint arXiv:2107.00652 (2021)
14. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021)
15. El-Nouby, A., Touvron, H., Caron, M., Bojanowski, P., Douze, M., Joulin, A., Laptev, I., Neverova, N., Synnaeve, G., Verbeek, J., et al.: Xcit: Cross-covariance image transformers. arXiv preprint arXiv:2106.09681 (2021)
16. Fang, J., Xie, L., Wang, X., Zhang, X., Liu, W., Tian, Q.: Msg-transformer: Exchanging local spatial information by manipulating messenger tokens. arXiv preprint arXiv:2105.15168 (2021)
17. Guo, J., Han, K., Wu, H., Xu, C., Tang, Y., Xu, C., Wang, Y.: Cmt: Convolutional neural networks meet vision transformers. arXiv preprint arXiv:2107.06263 (2021)

18. Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. *Advances in Neural Information Processing Systems* **34** (2021)
19. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 2961–2969 (2017)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)
21. Heo, B., Yun, S., Han, D., Chun, S., Choe, J., Oh, S.J.: Rethinking spatial dimensions of vision transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2021)
22. Huang, Z., Ben, Y., Luo, G., Cheng, P., Yu, G., Fu, B.: Shuffle transformer: Rethinking spatial shuffle for vision transformer. *arXiv preprint arXiv:2106.03650* (2021)
23. Jing, Y., Liu, X., Ding, Y., Wang, X., Ding, E., Song, M., Wen, S.: Dynamic instance normalization for arbitrary style transfer. In: *AAAI* (2020)
24. Li, Y., Mao, H., Girshick, R., He, K.: Exploring plain vision transformer backbones for object detection. *arXiv preprint arXiv:2203.16527* (2022)
25. Lin, C., Guo, M., Li, C., Yuan, X., Wu, W., Yan, J., Lin, D., Ouyang, W.: Online hyper-parameter learning for auto-augmentation strategy. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 6579–6588 (2019)
26. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 2980–2988 (2017)
27. 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: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 740–755. Springer (2014)
28. Liu, Z., Hu, H., Lin, Y., Yao, Z., Xie, Z., Wei, Y., Ning, J., Cao, Y., Zhang, Z., Dong, L., et al.: Swin transformer v2: Scaling up capacity and resolution. *arXiv preprint arXiv:2111.09883* (2021)
29. 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: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 10012–10022 (2021)
30. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: *International Conference on Learning Representations* (2018)
31. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H.: Training data-efficient image transformers; distillation through attention. In: *International Conference on Machine Learning*. PMLR (2021)
32. Wang, P., Wang, X., Wang, F., Lin, M., Chang, S., Xie, W., Li, H., Jin, R.: Kvt: k-nn attention for boosting vision transformers. *arXiv preprint arXiv:2106.00515* (2021)
33. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pvtv2: Improved baselines with pyramid vision transformer (2021)
34. 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: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 568–578 (2021)

35. 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: International Conference on Learning Representations (2022)
36. Wu, S., Wu, T., Tan, H., Guo, G.: Pale transformer: A general vision transformer backbone with pale-shaped attention. In: Proceedings of the AAAI Conference on Artificial Intelligence (2022)
37. Xia, Z., Pan, X., Song, S., Li, L.E., Huang, G.: Vision transformer with deformable attention (2022)
38. 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)
39. Xu, B., Wang, N., Chen, T., Li, M.: Empirical evaluation of rectified activations in convolutional network. arXiv preprint arXiv:1505.00853 (2015)
40. Xu, Y., Zhang, J., Zhang, Q., Tao, D.: Vitpose: Simple vision transformer baselines for human pose estimation. arXiv preprint arXiv:2204.12484 (2022)
41. Xu, Y., ZHANG, Q., Zhang, J., Tao, D.: ViTAE: Vision transformer advanced by exploring intrinsic inductive bias. In: Advances in Neural Information Processing Systems (2021)
42. Yan, H., Li, Z., Li, W., Wang, C., Wu, M., Zhang, C.: Contnet: Why not use convolution and transformer at the same time? arXiv preprint arXiv:2104.13497 (2021)
43. Yang, J., Li, C., Zhang, P., Dai, X., Xiao, B., Yuan, L., Gao, J.: Focal attention for long-range interactions in vision transformers. In: Advances in Neural Information Processing Systems (2021)
44. Yang, Z., Liu, D., Wang, C., Yang, J., Tao, D.: Modeling image composition for complex scene generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7764–7773 (2022)
45. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: International Conference on Learning Representations (2016)
46. 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: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 558–567 (2021)
47. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6023–6032 (2019)
48. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)
49. Zhang, P., Dai, X., Yang, J., Xiao, B., Yuan, L., Zhang, L., Gao, J.: Multi-scale vision longformer: A new vision transformer for high-resolution image encoding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2998–3008 (October 2021)
50. Zhang, Q., Xu, Y., Zhang, J., Tao, D.: Vitaev2: Vision transformer advanced by exploring inductive bias for image recognition and beyond. arXiv preprint arXiv:2202.10108 (2022)
51. Zhang, Q., Yang, Y.B.: Rest: An efficient transformer for visual recognition. Advances in Neural Information Processing Systems **34** (2021)
52. Zhu, X., Hu, H., Lin, S., Dai, J.: Deformable convnets v2: More deformable, better results. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9308–9316 (2019)

53. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: International Conference on Learning Representations (2021)