# ScalableViT: Rethinking the Context-oriented Generalization of Vision Transformer

Rui Yang<sup>1†\*</sup>, Hailong Ma<sup>2†</sup>, Jie Wu<sup>2‡</sup>, Yansong Tang<sup>1</sup>, Xuefeng Xiao<sup>2</sup>, Min Zheng<sup>2</sup>, Xiu Li<sup>1‡</sup>

<sup>1</sup> Tsinghua Shenzhen International Graduate School, Tsinghua University {r-yang20@mails,tang.yansong@sz,li.xiu@sz}.tsinghua.edu.cn <sup>2</sup> ByteDance Inc.

{mahailong.1206,wujie.10,xiaoxuefeng.ailab,zhengmin.666}@bytedance.com

Abstract. The vanilla self-attention mechanism inherently relies on predefined and steadfast computational dimensions. Such inflexibility restricts it from possessing context-oriented generalization that can bring more contextual cues and global representations. To mitigate this issue, we propose a Scalable Self-Attention (SSA) mechanism that leverages two scaling factors to release dimensions of query, key, and value matrices while unbinding them with the input. This scalability fetches context-oriented generalization and enhances object sensitivity, which pushes the whole network into a more effective trade-off state between accuracy and cost. Furthermore, we propose an Interactive Windowbased Self-Attention (IWSA), which establishes interaction between nonoverlapping regions by re-merging independent value tokens and aggregating spatial information from adjacent windows. By stacking the SSA and IWSA alternately, the Scalable Vision Transformer (ScalableViT) achieves state-of-the-art performance on general-purpose vision tasks. For example, ScalableViT-S outperforms Twins-SVT-S by 1.4% and Swin-T by 1.8% on ImageNet-1K classification.

**Keywords:** Vision Transformer, Self-Attention Mechanism, Classification, Detection, Semantic Segmentation

## 1 Introduction

Convolutional Neural Networks (CNNs) dominated the computer vision field last few years, which attributes to their capacity in modeling realistic images from a local to global perception. Although they have been widely applied in various vision tasks, there are still deficiencies in global visual perception. This global view is essential for downstream tasks, such as object detection and semantic segmentation. Recently, ViT [10] and its follow-ups [31,24,5,36] employed transformer encoders to address the image task and achieved comparable performance against their CNN counterparts because of the global receptive field. However, the global

<sup>&</sup>lt;sup>†</sup>Equal contribution. <sup>‡</sup>Corresponding author. <sup>\*</sup>This work was partly done while Rui Yang interned at ByteDance. Code: https://github.com/Yangr116/ScalableViT



Fig. 1: Visualization for feature maps in the Vision Transformer. We show the feature maps after the second Transformer blocks in Window-based Self-Attention (WSA) [5,24] and Scalable Self-Attention (SSA). The activation from WSA is discontinuous because of a limited mapping dimension. SSA ( $r_c \equiv 1$ ) reduces computational overhead while retaining a global perception, ensuring its feature map is nearly continuous.  $r_c \equiv 1$  denotes no scaling factors in the channel dimension. SSA introduces scale factors to spatial and channel dimensions, modeling a holistic representation and a context-oriented generalization.

perception of the Transformer entails an unaffordable computation since selfattention (the primary operation of the Transformer) is quadratically computed on the whole sequence. To alleviate this overhead, typical Swin transformer [24] employed Window-based Self-Attention (WSA), which partitioned a feature map into many non-overlapped sub-regions and enabled it to process large-scale images with linear complexity. They also proposed a novel Shifted Window-based Self-Attention (SWSA) to compensate for losses of potential long-range dependency. Twins [5] combined the WSA with Global Sub-sampled Attention (GSA) for better performance.

To gain an insight into the WSA [5,24], we visualize feature maps after the second block. As shown in Fig. 1, features captured by the WSA are dispersed, and their responses incline to partial rather than object-oriented. It may attribute to an invariably fixed dimension that results in limited learning ability, thereby the final performance of the model being highly determined by the difficulty of input data. To alleviate this problem, we develop a novel self-attention mechanism, termed Scalable Self-Attention (SSA), which simultaneously introduces two scaling factors  $(r_n \text{ and } r_c)$  to spatial and channel dimensions. Namely, SSA selectively applies these factors to query, key, and value matrices (Q, K, and V), ensuring the dimension is more elastic and no longer deeply bound by the input. On the one hand, SSA aggregates redundant tokens with similar semantic information to a more compact one via spatial scalability. Consequently,

unnecessary intermediate multiplication operations are eliminated, and the computational complexity is reduced significantly. In the third row of Fig. 1, we can easily observe that spatial scalability can bring nearly contiguous visual modeling for objects, but some contextual cues are still lost. Hence, on the other hand, we expand the channel dimension to learn a more graphic representation. As depicted in the last row of Fig. 1, SSA successfully obtains complete object activation while maintaining context-oriented generalization via channel scalability. For instance, the contextual cues of the cat in the last column are represented in detail. Such scaling factors also restore the output dimension to align with the input, which makes the residual connection feasible.

Moreover, we propose an Interactive Window-based Self-Attention (IWSA) that consists of a regular WSA and a local interactive module (LIM). The IWSA establishes information connections by re-merging independent *value* tokens and aggregating spatial information from adjacent windows. Therefore, it no longer limits the self-attention to local windows, particularly non-overlapping windows. Such characteristic enhances the desired global receptive field and takes good advantage of the most significant superiority of the Transformer in a single layer. The effectiveness of LIM for WSA is validated in Tab. 5b. To achieve a more efficient backbone for general vision tasks, we adopt a hierarchical design [14,29] and propose a new Vision Transformer architecture, termed **ScalableViT**, which alternately arranges IWSA and SSA blocks in each stage.

Main contributions of our ScalableViT lie in two aspects:

- For the global self-attention, we propose SSA to supply context-oriented generalization in the vanilla self-attention block, which significantly reduces computational overhead without sacrificing contextual expressiveness.
- For the local self-attention, we design LIM to enhance the global perception ability of WSA.

Both SSA and IWSA can model long-range dependency in a single layer instead of stacking more self-attention layers; hence, the ScalableViT is more suitable for visual tasks. We employ ScalableViT on several vision tasks, including image-level classification on ImageNet [8], pixel-level object detection and instance segmentation on COCO [23], and semantic segmentation on ADE20K [47]. Extensive experiments demonstrate that the ScalableViT outperforms other state-of-the-art Vision Transformers with similar or less computational cost. For example, ScalableViT-S achieves +1.4% gains against Twins-SVT-S and +1.8% gains against Swin-T on ImageNet-1K classification.

## 2 Related Work

The Transformer architecture [33] has become a common template for natural language processing (NLP) tasks due to its solid global modeling capabilities and convenient parallelization ability. Inspired by this, many researchers tried to equip CNNs with the self-attention to modulate and augment outputs of convolutions [2,43,39]. DETR [3] employed the self-attention mechanism to model relations between objects for end-to-end detection. Others [18,1,37] combined selfattention with convolutions for full-image contextual information. Recently, the emergence of ViT [10], DeiT [31], and a series of follow-ups [24,9,45,35,38,44,5] proved the bright prospect of the Vision Transformer.

#### 2.1 Vision Transformer

ViT [10] applied standard Transformer encoders to build a convolution-free image classifier by decomposing the image into a sequence of non-overlapping patches directly. Although it harvested promising results, a gap still existed between data-hungry Transformers and top-performing CNNs [30] when only training on the midsize ImageNet-1K [8] from scratch. In order to bridge this gap, DeiT [31] proposed a token-based distillation procedure and a data-efficient training strategy to optimize the Transformer effectively. Later, the follow-ups improved different aspects of the ViT, making them more suitable for vision tasks. T2T-ViT [45] optimized the tokenization by concatenating the neighboring tokens into one token. DynamicViT [28] pruned the tokens of less importance in a dynamic way for a better lightweight module. Cvt [38], CeiT [44] incorporated the convolution designs into the self-attention or the FFN to enhance the locality. CPVT [6] utilized the implicit position representation ability from convolutions (with zero padding) to encode the conditional position information for inputs with the arbitrary size. Then, hierarchical pyramid structures [35,24,9,5] were performed by progressively shrinking the number of tokens and replacing the class token with the average pooling. Thus, the Transformer, supported by multi-level features [21], can handle object detection and image segmentation tasks conveniently. In this paper, we develop a Vision Transformer, ScalableViT, which achieves a better accuracy and cost trade-off on visual tasks.

#### 2.2 Local Self-Attention

The computational complexity of the self-attention mechanism is a barrier that confines it in only downsampled feature maps or small images. Thus, several previous studies [26,15,34,18,16] proposed decomposing the global self-attention into much paralleled local self-attention to handle expensive computation burdens. However, this local self-attention limits the receptive field that is critical to dense predict tasks. [18,34] proposed generating the sparse attention map on a criss-cross path to realize global interaction. [16] captured the information from all the other positions via interlacing elements between different local windows. HaloNet [32] used the overlapped local windows to add the interactions between independent windows. After ViT [10] showed competitiveness, several followups [24,5,17,46] applied the self-attention within non-overlapped local windows for linear computational complexity. To compensate for lost information, Swin Transformer [24] introduced a novel shifted window strategy, and Twins Transformer [5] baked sparse global attention [35] after WSA. We design the IWSA. which can aggregate information from a collection of discrete *value* tokens and enable local self-attention to model long-range dependency in a single block.



Fig. 2: The architecture of the ScalableViT. IW-MSA and S-MSA, the multi-head format of IWSA and SSA, are organized alternately in each stage. A PEG [6] is placed between two blocks in the front of each stage to encode implicit position information dynamically.

## 3 Method

In this section, we elaborately introduce the architecture of ScalableViT and mainly focus on SSA and IWSA mechanisms. SSA simultaneously introduces different scale factors into spatial and channel dimensions to maintain contextoriented generalization while reducing computational overhead. IWSA enhances the receptive field of local self-attention by aggregating information from a set of discrete *value* tokens. Both have linear computational complexity and can learn long-range dependency in a single layer.

## 3.1 Overall Architecture

The architecture of ScalableViT is illustrated in Fig. 2. For an input image with size  $H \times W \times 3$ , a convolutional patch embedding layer  $(7 \times 7, \text{stride 4})$  is used to obtain a collection of tokens  $(\frac{H}{4} \times \frac{W}{4})$  and project the channel dimension to C. Then, these initial tokens will pass through four stages which contain a series of Transformer blocks. Between two adjacent stages, another convolutional patch embedding layer  $(3 \times 3, \text{stride 2})$  is utilized to merge tokens and double the channel dimension. For the  $i^{th}$  stage, there are  $\frac{H}{2i+1} \times \frac{W}{2i+1}$  input tokens with  $2^{i-1}C$  channels and  $L_i$  Transformer blocks. As a result, the quantity of tokens will eventually be reduced to  $\frac{H}{32} \times \frac{W}{32}$ . This architecture enables us to obtain a hierarchical representation similar to the typical backbones based on CNNs [14,29]. This merit allows ScalableViT to naturally migrate to various vision tasks, such as object detection and segmentation. In each stage, we devise an alternate arrangement of IW-MSA and S-MSA blocks to organize the topological structure. In the front of each stage, a position encoding generator (PEG) [6] is inserted between two Transformer blocks to generate position embedding dynamically.





Fig. 3: The diagram of SSA. Two scale factors,  $r_n$  and  $r_c$ , are introduced to the spatial and channel dimension for better computational efficiency and representational sufficiency.

Fig. 4: The diagram of LIM. It merges a series of discrete  $V_i$  by the function  $\mathcal{F}$ , and outputs an interacted Y.

## 3.2 Scalable Self-Attention

Self-attention is a critical mechanism in the Transformer, and the vanilla selfattention can be calculated as:

$$Z = A(X)V(X) = Softmax(\frac{Q(X)K(X)^T}{\sqrt{d_k}})V(X),$$
(1)

where A(X) is the attention matrix of the input X;  $Q(X), K(X), V(X) \in \mathbb{R}^{N \times C}$ are the query, key, and value matrices;  $d_k$  is the channel dimension of query or key matrix; N is the number of tokens in each matrix, and C is the channel dimension. The original self-attention mechanism obtains a global receptive field by establishing associations between all input tokens, which is a vital advantage of the Transformer over CNNs. However, it has quadratic computational overhead with N, leading to inefficiency in the intermediate multiplication operations.

Generally, there is much homologous information in natural images, but vanilla self-attention still calculates their similarity. Notably, not all information is necessary to calculate self-attention in the Vision Transformer. For example, similar background tokens should be aggregated as one representative token to attend to other foreground tokens. Namely, the dimension of Q(X), K(X), and V(X) should not be bounded with the input X. More importantly, the fixed dimension results in limited learning ability. Thus, we develop the Scalable Self-Attention (SSA), where two scaling factors  $(r_n \text{ and } r_c)$  are introduced to spatial and channel dimensions, respectively, resulting a more efficient intermediate calculation than the vanilla one. As illustrated in Fig. 3, the spatial dimension Nand channel dimension C are selectively scaled to  $N \times r_n$  and  $C \times r_c$ , respectively, by three transformation functions  $f_q(\cdot)$ ,  $f_k(\cdot)$ , and  $f_v(\cdot)$ . These scaling factors can also restore the output dimension to align with the input, making the subsequent FFN layers and residual connections feasible. As a result, the intermediate dimension is more elastic and no longer deeply bound with the input X. The model can reap context-oriented generalization while dwindling computational overhead significantly. SSA can be naturally written as:

$$Z' = A'(X)V'(X) = Softmax(\frac{Q'(X)K'(X)^{T}}{\sqrt{d'_{k}}})V'(X),$$
(2)

$$Q'(X) = f_q(X), \quad K'(X) = f_k(X), \quad V'(X) = f_v(V),$$
 (3)

where  $Q'(X) \in \mathbb{R}^{N \times Cr_c}, K'(X) \in \mathbb{R}^{Nr_n \times Cr_c}, V'(X) \in \mathbb{R}^{Nr_n \times C}$  are the scaled query, key and value matrices of the input  $X \in \mathbb{R}^{H \times W \times C}$ ;  $A'(X) \in \mathbb{R}^{N \times Nr_n}$  is the scaled attention matrix; Z' is the weighted sum of V'(X). The transformation  $f_q(\cdot)$  scales the channel dimension of query from C to  $C_{r_c}$ .  $f_k(\cdot)$  is the scaling function for key, which scales the spatial dimension from N to  $N_{r_n}$  while scaling the channel dimension from C to  $C_{r_c}$ .  $f_v(\cdot)$  is the scaling function for value, which scales the spatial dimension from N to  $N_{r_n}$ . Hence, some unnecessary intermediate multiplication is decreased significantly. The computation complexity of the proposed SSA is equal to  $\mathcal{O}(NN_{r_n}C + NN_{r_n}C_{r_c})$  that is linear with the input size  $(N = H \times W)$ . For utility and briefness, three transformations are operated by convolutions and cooperates with linear projections to get the scaled query, key, and value matrices. The efficient SSA does not change the size of Zand can be expanded to the Scalable Multi-Head Self-Attention (S-MSA) easily.

More importantly, the introduced spatial and channel scalability can bring context-oriented generalization (see Fig. 1). If only spatial scalability is introduced ( $r_c \equiv 1$ ), there would realize nearly contiguous visual modeling for objects but a lack of critical graphic representation. When further introducing channel scalability, SSA can successfully maintain contextual cues and obtain complete object activation, which is essential in visual tasks. The values of these two scaling factors vary with model configurations and different network stages. As the network gradually deepens, the quantity of tokens shrinks, and the degree of redundancy is also dropped. Thus,  $r_n$  is largen with the stage depth. Similarly, the channel dimension does not always mismatch with spatial dimension in the self-attention operation. Thus, we set  $r_c \geq 1$  in ScalableViT-S and ScalableViT-B. Because of a too-large channel dimension, we set  $r_c \leq 1$  in ScalableViT-L. Details about two scale factors are displayed in Table 1.

## 3.3 Interactive Window-based Self-Attention

Besides the efficient self-attention [35], earlier researches have developed the local self-attention [24,5] to avoid the quadratic computational complexity with the number of tokens. For example, WSA divides an image  $(H \times W \times C)$  into multiple partial windows which contains  $M \times M$  tokens. Then, the self-attention would be calculated in every isolated window and produce a set of discrete outputs  $\{Z_n\}_{n=1}^{\frac{H}{M} \times \frac{W}{M}}$ , where  $Z_n$  can be calculated as:

$$Z_n = A_n(X_n)V_n(X_n) = Softmax(\frac{Q_n(X_n)K_n(X_n)^T}{\sqrt{d_k}})V_n(X_n), \qquad (4)$$

in which  $X_n \in \Omega_n^{M \times M \times C}$  is the partial window field;  $Q_n(X_n), K_n(X_n), V_n(X_n) \in \mathbb{R}^{M^2 \times C}$  are the query, key, and value matrices of the discrete window  $X_n$ ,

8

respectively.  $d_k$  is equal to the channel dimension of discrete query/key tokens. Finally, a collection of discrete  $\{Z_n\}_{n=1}^{\frac{M}{M} \times \frac{W}{M}}$  is merged back to  $Z \in \mathbb{R}^{N \times C}$ . Thus, for an image, the computational complexity of attention would be reduced from  $\mathcal{O}(2H^2W^2C)$  to  $\mathcal{O}(2M^2HWC)$ . The WSA can be suitable for various vision tasks that require high-resolution input due to its linear complexity.

However, such computationally efficient WSA yields a feature map with an integrated shape but isolated activation (see Fig. 1), which ascribed to the missed global receptive field in a single layer. This is far from the initial aims of self-attention. To alleviate above problem, we propose Interactive Window-based Self-Attention (IWSA) that incorporates a local interactive module (LIM) into WSA, as illustrated in Fig. 4. After getting a collection of discrete values  $\{V_n(X_n)\}_{n=1}^{\frac{H}{N}\times \frac{W}{N}}$ , the LIM reshapes them into  $M \times M \times C$  and merges them into a shape-integrated value map  $V \in \mathbb{R}^{H \times W \times C}$ . Subsequently, a function  $\mathcal{F}(x)$  is employed to establish marriages and connections between adjacent  $V_n(X_n)$ s. As a result, the output  $Y = \mathcal{F}(V)$  is an integrated feature map with global information. Finally, this feature map is added on Z as the final output Z'. Without loss of generality, the IWSA is calculated as:

$$Z' = Z + \mathcal{F}(V), \tag{5}$$

where  $Z \in \mathbb{R}^{N \times C}$  is merged by a set of  $\{Z_n\}_{n=1}^{\frac{M}{M} \times \frac{W}{M}}$ . In order to be implemented friendly, a depth-wise convolution with zero padding is employed to take the place of function  $\mathcal{F}(x)$ . If the kernel size of this depth-wise convolution is  $k \times k$  (set to 3 by default), the computational cost from the LIM is negligible in practice. Additionally, [19] demonstrated that the convolution with zero padding could implicitly encode position information through experiments. Thus, IWSA allows self-attention to benefit from the translation invariance. Furthermore, IWSA can be easily expanded to Interactive Window-based Multi-head Self-Attention (IW-MSA) format easily if calculated in different heads.

CoaT [42] also introduced a depth-wise convolution into self-attention. However, they only considered the convolution as a positional encoding method and inserted it deeply into the calculation. If this convolution is expanded into the WSA, it would be limited in the discrete  $V_n(X_n)$ , which is denoted as local enhanced module (LEM). Differently, we regard our LIM as a matchmaker, which is applied on the spliced *value* map V and parallels with self-attention. By making the sufficient ablation study in Section 4.4, we demonstrate that LIM is capable of delivering stable improvements, especially for downstream tasks.

#### 3.4 Position Encoding

Besides the position information introduced by LIM, we utilize the positional encoding generator (PEG) [6], composed of a convolution layer with fixed weights, to acquire implicit positional information. As illustrated in Fig. 2, it is plugged between two consecutive Transformer blocks, with only one in the front of each stage. After the PEG, input tokens are sent to subsequent blocks where position bias could enable the Transformer to realize the input permutation.

Table 1: Detailed configurations of ScalableViT series.  $r_c$  and  $r_n$  are scale factors for the channel and the spatial dimensions, respectively. '#Blocks' and '#Heads' refer to the number of blocks ( $[L_1, L_2, L_3, L_4]$ ) and heads in four stages, respectively. '#Channels' refers to the channel dimension of the first stage.

Models	#Channels	#Blocks	# Heads	$ $ $r_c$	$r_n$
ScalableViT-S ScalableViT-B ScalableViT-L	64 96 128	$\begin{array}{ } [2,2,20,2] \\ [2,2,14,6] \\ [2,6,12,4] \end{array}$	$\begin{array}{c} [2,4,8,16] \\ [3,6,12,24] \\ [4,8,16,32] \end{array}$	$ \begin{bmatrix} 1.25, 1.25, 1.25, 1.0 \\ 2.0, 1.25, 1.25, 1.0 \end{bmatrix} \\ \begin{bmatrix} 0.25, 0.5, 1.0, 1.0 \end{bmatrix} $	$\frac{\left \left[\frac{1}{64}, \frac{1}{16}, \frac{1}{4}, 1\right]\right }{\left[\frac{1}{64}, \frac{1}{16}, \frac{1}{4}, 1\right]} \\ \left[\frac{1}{64}, \frac{1}{16}, \frac{1}{4}, 1\right]}$

### 3.5 Architecture Variants

In order to fairly compare with other models under similar computation complexity, we set three models: ScalableViT-S, ScalableViT-B, and ScalableViT-L. The detailed configurations are provided in Table 1, where  $r_c$  and  $r_n$  denote expansion or reduction factors for channel and spatial dimensions, respectively, as described in Section 3.2. Due to the varying representational capability, we set different  $r_c$  for three models. Additionally, the number of blocks, channels, and heads varies with the computational cost.

# 4 Experiments

In the following, we compare the proposed model with other state-of-the-art works on ImageNet-1K [8], COCO [23], and ADE20K [47]. Then, we conduct ablation studies on the upgraded parts to verify their effectiveness.

## 4.1 Image Classification on ImageNet-1K

Settings. Image classification experiments are conducted on the ImageNet-1K [8] dataset. All settings mainly follow DeiT [31]. During training, we apply data augmentation and regularization strategies in [31]. We employ the AdamW optimizer [25] to train models for 300 epochs from scratch. The learning rate is set to 0.001 initially and varies with the cosine scheduler. The global batchsize is set to 1024 on 8 V100 GPUs. During testing on the validation set, the shorter side of an input image is first resized to 256, and a center crop of  $224 \times 224$  is used to evaluate the classification accuracy.

**Result.** Classification results on ImageNet-1K are reported in Table 2, where all models are divided into small (around 4G), base (around 9G), and large (around 15G) levels according to computation complexity (FLOPs). ScalableViT-S with a two-layer head outperforms comparable models (1.4% better than Twins-SVT-S, and 1.8% better than Swin-T). Moreover, it can even approach or exceed other base models. For the base level, ScalableViT-B surpasses Twins-SVT-B by 0.9% and SWin-S by 1.1% with similar FLOPs. ScalableViT-L also achieves a prominent accuracy-cost trade-off. Additionally, our ScalableViT outperforms the EfficientNet by 0.2%, 0.5%, and 0.4% under three magnitude receptively.

#### 10 Rui Yang, Hailong Ma, Jie Wu, et al.

Table 2: Comparison with different state-of-the-art backbones on ImageNet-1K classification. Except for EfficientNet, other models are trained and evaluated on  $224 \times 224$  input size. Top-1 refers to top-1 accuracy (%).

		-		- * ( )			
Method	#Param	. FLOPs	Top-1	Method	#Param	. FLOPs	Top-1
Conv	vNet			Transf	ormer		
RegNetY-4G [27]	21M	4.0G	80.0	T2T-ViT-19 [45]	39M	8.9G	81.9
RegNetY-8G [27]	39M	8.0G	81.7	CoaT(S) [42]	22M	12.6G	82.1
RegNetY-16G [27]	84M	16.0G	82.9	CoaT-Lite(M) [42]	45M	9.8G	83.6
EfficientNet-B4 [30]	19M	4.2G	82.9	PVT-Medium [35]	44M	6.7G	81.2
EfficientNet-B5 [30]	30M	9.9G	83.6	Swin-S [24]	50M	8.7G	83.0
EfficientNet-B6 [30]	43M	19.0G	84.0	CvT-21 [38]	32M	7.1G	82.5
Transf	ormer			Twins-SVT-B [5]	56M	8.6G	83.2
				CrossFormer-B [36]	52M	9.2G	83.4
DeiT-Small/16 [31]	22M	4.6G	79.9	ScalableViT-B(ours)	81M	8.6G	84.1
T2T-ViT-14 [45]	22M	5.2G	81.5				
TNT-S [12]	24M	5.2G	81.3	Deit-Base/16 $[31]$	86M	17.6G	81.8
CoaT-Lite(S) [42]	20M	4.0G	81.9	T2T-ViT-24 [45]	64M	14.1G	82.3
PVT-Small [35]	25M	3.8G	79.8	TNT-B [12]	66M	14.1G	82.8
Swin-T [24]	29M	4.5G	81.3	PVT-Large [35]	61M	9.8G	81.7
CvT-13 [38]	20M	4.5G	81.6	Swin-B [24]	88M	15.4G	83.3
Twins-SVT-S [5]	24M	2.9G	81.7	Twins-SVT-L [5]	99M	15.1G	83.7
CrossFormer-S [36]	31M	4.9G	82.5	CrossFormer-L [36]	92M	16.1G	84.0
ScalableViT-S(ours)	32M	4.2G	83.1	${\bf ScalableViT-L(ours)}$	104M	14.7G	84.4

#### 4.2 Object Detection on COCO

Settings. Object detection experiments are conducted on COCO 2017 [23] dataset. We verify the model effectiveness on RetinaNet [22] and Mask R-CNN [13] detection frameworks using the MMDetection [4]. Before training, we initialize the backbone with the weight pre-trained on ImageNet-1K, FPN with Xavier [11] scheme, and other new layers with Normal scheme (std = 0.01). All models utilize the same settings as [5]: AdamW [25] optimizer,  $1 \times (12 \text{ epochs})$ , and  $3 \times (36 \text{ epochs})$  schedules with a global batchsize of 16 on 8 GPUs. For the  $1 \times$  schedule, the short side of images is resized to 800 pixels, and the long side is never more than 1333 pixels. The learning rate is declined at the 8th and 11th epoch with a decay rate of 0.1. For the  $3 \times$  schedule, we adopt the multi-scale training, which randomly resizes the short side of images within the range of [480, 800] while keeping the longer side at most 1333. The learning rate is declined at the 27th and 33rd with a decay rate of 0.1.

**Result.** We present results of RetinaNet and Mask R-CNN frameworks in Table 3, where  $AP^b$  and  $AP^m$  refer to box mAP and mask mAP, respectively. For object detection with RetinaNet, ScalableViT performs a notable advantage against its CNN and Transformer counterparts. With the 1× schedule, our ScalableViT brings 7.3-8.9  $AP^b$  against ResNet at comparable settings. Compared with the popular Swin and Twins Transformers, our ScalableViT performs 3.5-3.7  $AP^b$  and 0.5-2.2  $AP^b$  improvements, respectively. With the 3× schedule, our ScalableViT still achieves competitive performance. For Mask R-CNN, our ScalableViT-S outperforms ResNet-50 by 7.8  $AP^b$  and 7.3  $AP^m$  with the 1×

Table 3: Results on COCO object detection using the RetinaNet [22] and Mask R-CNN [13] framework.  $1 \times$  refers to 12 epochs, and  $3 \times$  refers to 36 epochs. MS means multi-scale training. AP<sup>b</sup> and AP<sup>m</sup> denotes box mAP and mask mAP, respectively. FLOPs are measured at resolution 800 × 1280.

Backbone	#Param	. FLOPs		I	Retina	Net 1	×			Reti	naNet	$3 \times +$	- MS	
Ducinouno	(M)	(G)	$ AP^b $	$AP_{50}^b$	$AP_{75}^b$	$AP_S^b$	$AP_M^b$	$\operatorname{AP}_L^b$	$ AP^b $	$AP_{50}^b$	$AP_{75}^b$	$\mathrm{AP}^b_S$	$\operatorname{AP}^b_M$	$\mathrm{AP}^b_L$
ResNet50 [14]	38	239	36.3	55.3	38.6	19.3	40.0	48.8	39.0	58.4	41.8	22.4	42.8	51.6
PVT-Small [35]	34	226	40.4	61.3	43.0	25.0	42.9	55.7	42.2	62.7	45.0	26.2	45.2	57.2
Swin-T [24]	39	245	41.5	62.1	44.2	25.1	44.9	55.5	43.9	64.8	47.1	28.4	47.2	57.8
Twins-SVT-S [5]	34	210	43.0	64.2	46.3	28.0	46.4	57.5	45.6	67.1	48.6	29.8	49.3	60.0
CrossFormer-S [36]	41	272	44.4	65.8	47.4	28.2	48.4	59.4	—					
ScalableViT-S(ours)	36	238	45.2	66.5	48.4	29.2	49.1	60.3	47.8	69.2	51.2	31.4	51.5	63.4
ResNet101 [14]	58	315	38.5	57.8	41.2	21.4	42.6	51.1	40.9	60.1	44.0	23.7	45.0	53.8
PVT-Medium [35]	54	283	41.9	63.1	44.3	25.0	44.9	57.6	43.2	63.8	46.1	27.3	46.3	58.9
Swin-S [24]	60	335	44.5	65.7	47.5	27.4	48.0	59.9	46.3	67.4	49.8	31.1	50.3	60.9
Twins-SVT-B [5]	67	326	45.3	66.7	48.1	28.5	48.9	60.6	46.9	68.0	50.2	31.7	50.3	61.8
CrossFormer-B [36]	62	389	46.2	67.8	49.5	30.1	49.9	61.8	—	—	—	_		
ScalableViT-B(ours)	85	330	45.8	67.3	49.2	29.9	49.5	61.0	48.0	69.3	51.4	32.8	51.6	62.4
Backbone	#Param.	FLOPs		Mε	sk R-	CNN	$1 \times$		:	Mask	R-CN	N $3\times$	+ MS	1
	(M)	(G)	$AP^b$	$AP_{50}^b$	$AP_{75}^{b}$	$AP^m$	$\operatorname{AP}_{50}^m$	$AP_{75}^m$	$ AP^b $	$AP_{50}^b$	$AP_{75}^b$	$\mathbf{AP}^m$	$AP_{50}^m$	$AP_{75}^m$
ResNet50 [14]	44	260	38.0	58.6	41.4	34.4	55.1	36.7	41.0	61.7	44.9	37.1	58.4	40.1
PVT-Small [35]	44	245	40.4	62.9	43.8	37.8	60.1	40.3	43.0	65.3	46.9	39.9	62.5	42.8
Swin-T [24]	48	264	42.2	64.4	46.2	39.1	64.6	42.0	46.0	68.2	50.2	41.6	65.1	44.8
Twins-SVT-S [5]	44	228	43.4	66.0	47.3	40.3	63.2	43.4	46.8	69.2	51.2	42.6	66.3	45.8
CoaT-Lite(S) [42]	40	—	45.2		_	40.7			45.7		—	41.1		
CrossFormer-S [36]	50	301	45.4	68.0	49.7	41.4	64.8	44.6	—	—	—	—	—	—
ScalableViT-S(ours)	46	256	45.8	67.6	50.0	41.7	64.7	44.8	48.7	70.1	53.6	43.6	67.2	47.2
ResNet101 [14]	63	336	40.4	61.1	44.2	36.4	57.7	38.8	42.8	63.2	47.1	38.5	60.1	41.3
PVT-Medium [35]	64	302	42.0	64.4	45.6	39.0	61.6	42.1	44.2	66.0	48.2	40.5	63.1	43.5
Swin-S [24]	69	354	44.8	66.6	48.9	40.9	63.4	44.2	48.5	70.2	53.5	43.3	67.3	46.6
Twins-SVT-B [5]	76	340	45.2	67.6	49.3	41.5	64.5	44.8	48.0	69.5	52.7	43.0	66.8	46.6
CoaT(S) [42]	42	—	46.5		—	41.8			49.0	—	—	43.7		
CrossFormer-B [36]	72	408	47.2	69.9	51.8	42.7	66.6	46.2	—	—	—	—	—	
${\bf ScalableViT}\text{-}{\bf B}({\rm ours})$	95	349	46.8	68.7	51.5	42.5	65.8	45.9	49.0	70.3	53.6	<b>43.8</b>	67.4	47.5

schedule. ScalableViT-S achieves 3.6  $AP^b$  and 2.6  $AP^m$  gains than Swin-T. With the 3× schedule, ScalableViT-S brings 7.7  $AP^b$  and 6.5  $AP^m$  against ResNet-50. Similarly, it also surpasses Swin-T and Twins-SVT-S Transformers. Under base level, there is also a similar improvement, demonstrating its stronger contextoriented generalization. Additionally, Fig. 5 depicts some qualitative object detection and instance segmentation results from ScalableViT-S-based RetinaNet and Mask R-CNN, which show that contextual representation from the backbone enables the model to detect objects better.

#### 4.3 Semantic Segmentation on ADE20K

**Settings.** Semantic segmentation experiments are conducted on the challenging ADE20K [47] dataset. We use the typical Semantic FPN [20] and the Uper-Net [40] as segmentation frameworks to evaluate our models. We use the MM-Segmentation [7] to implement all related experiments, and the settings follow



Fig. 5: Qualitative results based on ScalableViT-S. (a), (b) and (c) are yielded by RetinaNet [22], Mask R-CNN [13], and Semantic FPN [20], respectively.

[5,24,35]. For the Semantic FPN, we train 80K iterations with a batch size 16 on 4 GPUs. For the UperNet, we train 160K iterations with a batch size 16 on 8 GPUs. During training, we first resize the short side of input images to 512 pixels, and the long side is never more than 2048 pixels, then they are randomly cropped to  $512 \times 512$ . During testing, we resize input images as the training phase but without cropping. We also use the test time augmentation for UperNet, including multi-scale test ([0.5, 0.75, 1.0, 1.25, 1.5, 1.75] × resolution) and flip.

**Result.** Table 4 reports the segmentation results. For the Semantic FPN, our ScalableViT outperforms Swin Transformer by +3.4 mIoU, +3.2 mIoU, and +3.4 mIoU, respectively, under three FLOPs levels. Compared with CrossFormer-S [36], ScalableViT-S performs a modest mIoU but has a fewer computation. When equipped into the UperNet, the ScalableViT achieves +4 mIoU, +1.9 mIoU, and +1.6 mIoU gains than Swin Transformer under different model sizes. The same competitive results are achieved when test time augmentation is adopted. In addition, ScalableViT-S outperforms CrossFormer-S by +0.9 mIoU and achieves comparable performance on the base and large size. Fig. 5(c) shows some qualitative results from ScalableViT-S-based Semantic FPN on validation split. These results indicate that the ScalableViT can obtain high-quality semantic segmentation results under contextual-oriented generalization.

## 4.4 Ablation Study

Analysis for Self-Attention mechanisms. Our ScalableViT contains two important designs: SSA and IWSA. We ablate their benefits in Table 5a. Firstly, all attention modules in ScalableViT-S are replaced with the regular windowbased self-attention (WSA). Although WSA achieves 82.4% top-1 accuracy, the dispersed feature (see Fig. 1) hinders it from better performance on the downstream visual task. Then, we substituted all attention modules with our IWSA and SSA, respectively. Both of them outperform WSA 0.4% top-1 accuracy. More

Backbone	Sema #Param	antic FPI FLOPs	N 80k mIoU(%)	#Param	UperN . FLOPs	Net 160k mIoU/MS mIoU(%)
ResNet50 [14]	29M	183G	36.7	_		/
PVT-Small [35]	28M	161G	39.8	_	_	_/
Swin-T [24]	32M	182G	41.5	60M	945G	44.5/45.8
Twins-SVT-S [5]	28M	144G	43.2	54M	901G	46.2/47.1
CrossFormer-S [36]	34M	221G	<b>46.0</b>	62M	980G	47.6/48.4
ScalableViT-S(ours)	30M	174G	44.9	57M	931G	48.5/49.4
ResNet101 [14]	48M	260G	38.8	86M	1092G	-/44.9
PVT-Medium [35]	48M	219G	41.6	_	_	—/—
Swin-S [24]	53M	274G	45.2	81M	1038G	47.6/49.5
Twins-SVT-B [5]	60M	261G	45.3	89M	1020G	47.7/48.9
CrossFormer-B [36]	56M	331G	47.7	84M	1090G	49.7/50.6
ScalableViT-B(ours)	79M	270G	<b>48.4</b>	107M	1029G	49.5/50.4
ResNeXt101-64×4d [41]	86M		40.2			_/_
PVT-Large [35]	65M	283G	42.1	_		—/—
Swin-B [24]	91M	422G	46.0	121M	1188G	48.1/49.7
Twins-SVT-L [5]	104M	404G	46.7	133M	1164G	48.8/50.2
CrossFormer-L [36]	95M	497G	48.7	126M	1258G	50.4/51.4
ScalableViT-L(ours)	105M	402G	<b>49.4</b>	135M	1162G	49.8/50.7

Table 4: Results on ADE20K segmentation using the Semantic FPN [20] and UperNet [40] framework. FLOPs are measured at resolution  $512 \times 2048$ . MS refers to the test time augmentation, including flip and multi-scale test.

importantly, they bring +4.6 mIoU and +5.5 mIoU improvements on ADE20K because of the ability modeling long-range dependency. With spatial scalability  $(r_c \equiv 1)$ , SSA only achieve 82.6% top-1 accuracy and 43.7 mIoU. Thus, the context-oriented generalization from the cooperation between spatial and channel scalability plays a critical role in visual tasks. Additionally, we examine the topology by rearranging IWSA and SSA. Results demonstrate that prioritizing IWSA followed by SSA performs best. We also compare IWSA with SWSA [24] in ScalableViT, where our IWSA is more appropriate than SWSA.

**Speed analysis.** Following [24], we measure throughput of the ScalableViT-S on single 3090 GPU with a batch size of 64 in Table 6. ScalableViT-S achieves 859.0 img/s, which perform better speed-accuracy trade-offs than Swin-S.

Effectiveness of Local Interactive Module. We examine the effectiveness of LIM in Table 5b. The ScalableViT-S without position encoding generator (PEG), locally enhanced module (LEM), or LIM is regarded as a baseline model which achieves 82.7% top-1 accuracy on ImageNet. Then, three modules are inserted and yield +0.2%, +0.1%, and +0.3% gains than baseline, respectively. It demonstrates that the reasonable convolution can help the model perform better. Due to the window connection, LIM outperforms LEM by +0.2% top-1 accuracy, proving the significance of the information interaction. Additionally, we combine PEG with LEM or LIM, whose results are better than only using a single module. Note that the combination of PEG and LIM outperforms the PEG and LEM under the same overhead. LIM aims to bring global perception into the single Transformer block. Its effectiveness is greatly demonstrated on

Table 5: Ablation study for different self-attention mechanisms and LIM using ScalableViT-S. Top-1 refers to top-1 accuracy (%) on ImageNet-1K. Semantic segmentation results are yielded from Semantic FPN on ADE20K.

(a) Analysis for different self-attention (b) Analysis for local interactive module and mechanisms. positional encoding generator.

Method  #Param. FLOPs Top-1 mIoU(%)			PEG	LEM	LIM	#Param.	FLOPs	Top-1	mIoU(%)		
$\begin{tabular}{ c c c c } \hline WSA \\ IWSA \\ SSA \\ SSA \\ SSA \\ (r_c \equiv 1) \\ SWSA & SSA \\ \hline \end{tabular}$	30M 30M 34M 32M 32M	4.3G 4.3G 4.1G 3.9G 4.2G	82.4 82.9 82.9 82.6 82.9	38.9 43.5 44.4 43.7	<b>v</b>	√ (	✓	32M 32M 32M 32M 32M	4.2G 4.2G 4.2G 4.2G 4.2G	82.7 82.9 82.8 83.0	41.7 43.2  43.7
SSA & IWSA IWSA & SSA	32M 32M	4.2G 4.2G	83.0 83.1	44.9	✓ ✓	V	$\checkmark$	32M 32M	4.2G 4.2G	83.0 83.1	44.9

Table 6: Comparison speed with state-of-the-art models.

Method	#Param.	FLOPs	Top-1	throughput(img/s)
Swin-T [24]	29M	4.5G	81.3	975.0
Swin-S $[24]$	50M	8.7G	83.0	589.2
CrossFormer-S [36]	31M	4.9G	82.5	859.0
ScalableViT-S(ours)	32M	4.2G	83.1	832.9

downstream tasks. Using Semantic FPN with ScalableViT-S on ADE20K, LIM obtains +2.0 mIoU, and associating PEG with LIM brings +3.2 mIoU gains.

# 5 Conclusion

In this paper, we have presented a Vision Transformer backbone named ScalableViT, composed of two highly effective self-attention mechanisms (SSA and IWSA). SSA employs two cooperated scaling factors in spatial and channel dimensions for context-oriented generalization, which maintains more contextual cues and learns graphic representations. IWSA develops a local interactive module to establish information connections between independent windows. Both of them owns the capability to model long-range dependency in a single layer. The proposed ScalableViT alternately stakes these two self-attention modules. It pushes the whole framework into a more effective trade-off state and achieves state-of-the-art performance on various vision tasks.

## Acknowledgements

This work was supported by the National Key R&D Program of China 505 (Grant No.2020AAA0108303), the National Natural Science Foundation of China (Grant No.41876098) and the Shenzhen Science and Technology Project (Grant No.JCYJ20200109143041798).

# References

- Bello, I., Zoph, B., Le, Q., Vaswani, A., Shlens, J.: Attention augmented convolutional networks. In: ICCV. pp. 3285–3294. IEEE (2019)
- Cao, Y., Xu, J., Lin, S., Wei, F., Hu, H.: Gcnet: Non-local networks meet squeezeexcitation networks and beyond. In: ICCV. pp. 1971–1980 (2019)
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: Endto-end object detection with transformers. In: ECCV. Lecture Notes in Computer Science, vol. 12346, pp. 213–229 (2020)
- Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: MMDetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 (2019)
- 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. arXiv preprint arXiv:2104.13840 (2021)
- 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)
- 7. Contributors, M.: MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. https://github.com/open-mmlab/mmsegmentation (2020)
- Deng, J., Dong, W., Socher, R., Li, L., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255 (2009)
- 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)
- 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: ICLR (2021)
- Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Teh, Y.W., Titterington, D.M. (eds.) Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. vol. 9, pp. 249–256 (2010)
- Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. arXiv preprint arXiv:2103.00112 (2021)
- He, K., Gkioxari, G., Dollár, P., Girshick, R.B.: Mask R-CNN. IEEE TPAMI 42(2), 386–397 (2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778. IEEE Computer Society (2016)
- Hu, H., Zhang, Z., Xie, Z., Lin, S.: Local relation networks for image recognition. In: ICCV. pp. 3463–3472 (2019)
- Huang, L., Yuan, Y., Guo, J., Zhang, C., Chen, X., Wang, J.: Interlaced sparse self-attention for semantic segmentation. arXiv preprint arXiv:1907.12273 (2019)
- 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)
- Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: Ccnet: Criss-cross attention for semantic segmentation. In: ICCV. pp. 603–612 (2019)

- 16 Rui Yang, Hailong Ma, Jie Wu, et al.
- Islam, M.A., Jia, S., Bruce, N.D.B.: How much position information do convolutional neural networks encode? In: ICLR (2020)
- Kirillov, A., Girshick, R.B., He, K., Dollár, P.: Panoptic feature pyramid networks. In: CVPR. pp. 6399–6408 (2019)
- Lin, T., Dollár, P., Girshick, R.B., He, K., Hariharan, B., Belongie, S.J.: Feature pyramid networks for object detection. In: CVPR. pp. 936–944 (2017)
- Lin, T., Goyal, P., Girshick, R.B., He, K., Dollár, P.: Focal loss for dense object detection. IEEE TPAMI 42(2), 318–327 (2020)
- Lin, T., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: ECCV. Lecture Notes in Computer Science, vol. 8693, pp. 740–755 (2014)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 (2021)
- 25. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (2019)
- Parmar, N., Ramachandran, P., Vaswani, A., Bello, I., Levskaya, A., Shlens, J.: Stand-alone self-attention in vision models. In: NeurIPS. pp. 68–80 (2019)
- Radosavovic, I., Kosaraju, R.P., Girshick, R.B., He, K., Dollár, P.: Designing network design spaces. In: CVPR. pp. 10425–10433 (2020)
- Rao, Y., Zhao, W., Liu, B., Lu, J., Zhou, J., Hsieh, C.J.: Dynamicvit: Efficient vision transformers with dynamic token sparsification. arXiv preprint arXiv:2106.02034 (2021)
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: Bengio, Y., LeCun, Y. (eds.) ICLR (2015)
- 30. Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning, ICML. vol. 97, pp. 6105–6114 (2019)
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: Proceedings of the 38th International Conference on Machine Learning, ICML. vol. 139, pp. 10347–10357 (2021)
- Vaswani, A., Ramachandran, P., Srinivas, A., Parmar, N., Hechtman, B.A., Shlens, J.: Scaling local self-attention for parameter efficient visual backbones. In: CVPR. pp. 12894–12904 (2021)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS. pp. 5998–6008 (2017)
- Wang, H., Zhu, Y., Green, B., Adam, H., Yuille, A.L., Chen, L.: Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In: ECCV. Lecture Notes in Computer Science, vol. 12349, pp. 108–126 (2020)
- 35. 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. arXiv preprint arXiv:2102.12122 (2021)
- 36. Wang, W., Yao, L., Chen, L., Lin, B., Cai, D., He, X., Liu, W.: Crossformer: A versatile vision transformer hinging on cross-scale attention. arXiv preprint arXiv:2108.00154 (2021)
- Wang, X., Girshick, R.B., Gupta, A., He, K.: Non-local neural networks. In: CVPR. pp. 7794–7803 (2018)
- Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L.: Cvt: Introducing convolutions to vision transformers. arXiv preprint arXiv:2103.15808 (2021)
- Xia, X., Li, J., Wu, J., Wang, X., Wang, M., Xiao, X., Zheng, M., Wang, R.: Trt-vit: Tensorrt-oriented vision transformer. arXiv preprint arXiv:2205.09579 (2022)

- Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: ECCV. Lecture Notes in Computer Science, vol. 11209, pp. 432–448 (2018)
- Xie, S., Girshick, R.B., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: CVPR. pp. 5987–5995 (2017)
- Xu, W., Xu, Y., Chang, T., Tu, Z.: Co-scale conv-attentional image transformers. arXiv preprint arXiv:2104.06399 (2021)
- Yin, M., Yao, Z., Cao, Y., Li, X., Zhang, Z., Lin, S., Hu, H.: Disentangled non-local neural networks. In: ECCV. vol. 12360, pp. 191–207 (2020)
- 44. Yuan, K., Guo, S., Liu, Z., Zhou, A., Yu, F., Wu, W.: Incorporating convolution designs into visual transformers. arXiv preprint arXiv:2103.11816 (2021)
- Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token vit: Training vision transformers from scratch on imagenet. arXiv preprint arXiv:2101.11986 (2021)
- 46. Yuan, Y., Fu, R., Huang, L., Lin, W., Zhang, C., Chen, X., Wang, J.: Hrformer: High-resolution transformer for dense prediction. arXiv preprint arXiv:2110.09408 (2021)
- 47. Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ADE20K dataset. Int. J. Comput. Vis. 127(3), 302–321 (2019)