

# DaViT: Dual Attention Vision Transformers

Mingyu Ding<sup>1\*</sup>, Bin Xiao<sup>2†</sup>, Noel Codella<sup>2</sup>, Ping Luo<sup>1†</sup>, Jingdong Wang<sup>3</sup>, Lu Yuan<sup>2</sup>

<sup>1</sup>The University of Hong Kong   <sup>2</sup>Microsoft   <sup>3</sup>Baidu  
mingyuding@hku.hk   {bixi,ncodella,luyuan}@microsoft.com  
pluo@cs.hku.hk   wangjingdong@outlook.com

**Abstract.** In this work, we introduce Dual Attention Vision Transformers (DaViT), a simple yet effective vision transformer architecture that is able to capture global context while maintaining computational efficiency. We propose approaching the problem from an orthogonal angle: exploiting self-attention mechanisms with both “*spatial tokens*” and “*channel tokens*”. With spatial tokens, the spatial dimension defines the token scope, and the channel dimension defines the token feature dimension. With channel tokens, we have the inverse: the channel dimension defines the token scope, and the spatial dimension defines the token feature dimension. We further group tokens along the sequence direction for both spatial and channel tokens to maintain the linear complexity of the entire model. We show that these two self-attentions complement each other: (i) since each channel token contains an abstract representation of the entire image, the channel attention naturally captures global interactions and representations by taking all spatial positions into account when computing attention scores between channels; (ii) the spatial attention refines the local representations by performing fine-grained interactions across spatial locations, which in turn helps the global information modeling in channel attention. Extensive experiments show DaViT backbones achieve state-of-the-art performance on four different tasks. Specially, DaViT-Tiny, DaViT-Small, and DaViT-Base achieve 82.8%, 84.2%, and 84.6% top-1 accuracy on ImageNet-1K without extra training data, using 28.3M, 49.7M, and 87.9M parameters, respectively. When we further scale up DaViT with 1.5B weakly supervised image and text pairs, DaViT-Giant reaches 90.4% top-1 accuracy on ImageNet-1K. Code is available at <https://github.com/microsoft/DaViT>.


## 1 Introduction

The global context is essential for many computer vision approaches, such as image classification and semantic segmentation. Convolutional neural networks (CNNs) [38] gradually obtain a global receptive field by multi-layer architectures and down-sampling operators. Recently, vision transformers [7,18], which directly capture long-range visual dependencies with a single self-attention layer, have drawn much attention. While these methods present strong capabilities to

---

\*This work is done when Mingyu was an intern at Microsoft.

†Corresponding authors




**Fig. 1.** (a) Spatial window multihead self-attention splits the spatial dimension into local windows, where each window contains multiple spatial tokens. Each token is also divided into multiple heads. (b) Channel group single-head self-attention groups channel tokens into multi groups. Attention is performed in each channel group with an entire image-level channel as a token. A channel-wise token that captures global information is also highlighted in (a). In this work, we alternately use these two types of attention to obtain both local fine-grained and global features.

model the global context, their computational complexity grows quadratically with the token length, limiting their ability to scale up to high-resolution scenarios.

Designing an architecture that can capture global contexts while maintaining efficiency to learn from high-resolution inputs is still an open research problem. A substantial body of work has been dedicated to developing vision transformers toward this goal. iGPT [7] first utilized a standard transformer to solve vision tasks by treating the image as sequences of pixels and performing pixel-level interactions. After that, ViT [18] used non-overlapped image patches as tokens to model the relationship between small image patches instead of pixels, showing promising performance on middle-resolution tasks such as classification. To further reduce the computational cost, local attention [37,75,52] that limits attention in a spatially local window, and squeezed projection [58,55] that performs attention on downsampled tokens, were proposed. Though local attention methods benefit from linear complexity with the spatial size, operators like “Shift” [37], “Overlapping Patch” [58,55], “ConvFFN” [55,72,61] are indispensable to compensate for the loss of global contextual information.

The general pattern across all prior works is that they attain various tradeoffs between resolution, global context, and computational complexity: pixel-level [7] and patch-level [18,37,58,56] self-attentions suffer either the cost of quadratic computational overhead or loss of global contextual information. Beyond variations of pixel-level and patch-level self-attentions, can we design an image-level self-attention mechanism that captures global information but is still efficient concerning the spatial size?

In this work, we introduce such a self-attention mechanism that is able to capture global context while maintaining computational efficiency. In addition to “*spatial tokens*” defined by existing works in Fig 1(a) representing the feature of an image patch, we introduce “*channel tokens*” by applying self-attention to the *transpose* of the token matrix, as shown in Fig 1(b). With channel tokens, the channel dimension defines the token scope, and the spatial dimension defines the token feature dimension. In this way, each channel token is global on



**Fig. 2.** Comparisons of the efficiency (*i.e.*, FLOPs) and performance (*e.g.*, Acc, mIoU, mAP) between the proposed approach and existing SoTA methods [37,65,9] on four computer vision tasks. Each method is represented by a circle, whose size represents the number of parameters. Our approach achieves superior performance under similar FLOPs than its counterparts on all four benchmarks.

the spatial dimension, containing an abstract representation of the entire image. Correspondingly, performing self-attention on such channel tokens further captures the global interaction by taking all spatial positions into account when computing attention scores between channels. Compared to conventional self-attention that performs global interactions over local pixels or patch tokens at a quadratic computational cost, the information exchange of channel self-attention is naturally imposed from a global perspective rather than a pixel/patch-wise one. Based on the global receptive field of the channel token, it fuses the representations to produce new global tokens and passes the information to the following layers. Thus, one can take such channel self-attention as a dynamic feature fusion over a series of abstract representations of the entire image.

Although this approach presents many advantages, a few challenges must be overcome. First, the computational complexity suddenly increases quadratically with the channel dimension, limiting the representation power of the layers. Inspired by spatial local attention [37,75,52], we propose *channel group attention* by dividing the feature channels into several groups and performing image-level interactions within each group. By group attention, we reduce the complexity to linear with respect to both the spatial and the channel dimensions. Since each channel token contains an abstract representation of the entire image, self-attention in this setting naturally captures the global interaction even we apply it locally along the channel dimensions.

Second, though channel-wise self-attention can capture global information easily, image-level tokens hinder local interactions across spatial locations. To solve this problem, we introduce Dual Attention Vision Transformers (DaViT) that alternately applies spatial window attention and channel group attention to capture both short-range and long-range visual dependencies, as shown in Figure 1. Our results show that these two structures complement each other: the channel group attention provides a global receptive field on the spatial dimension and extracts high-level global-image representations by dynamic feature fusion across global channel tokens; the spatial window attention refines the local representations by performing fine-grained local interactions across spatial locations, which in turn helps the global information modeling in channel attention.


To summarize, we propose DaViT that contains two seamlessly integrated self-attentions: spatial window self-attention with “*spatial tokens*”, and channel group self-attention with “*channel tokens*”. The two forms of attention are complementary and alternatively arranged, providing both local fine-grained and global interactions in a computationally efficient manner. We evaluate the effectiveness of the proposed DaViT via comprehensive empirical studies on image classification, object detection, and segmentation. Results in Figure 2 show that DaViT consistently outperforms the SoTA vision transformers across three benchmarks and four tasks with even fewer computational costs.

## 2 Related Work

**Vision Transformers.** Transformers [53,14] have dominated a wide range of natural language processing tasks. In computer vision, pioneering works iGPT [7] and ViT [18] apply attention directly to a sequence of image pixels or patches. Similarly, follow-up works [25,49,56,47,59,20,73,43,44] model global relationships on the patch-level tokens. All these works apply attention to capture interactions over all the local tokens, which can be pixel-level or patch-level. In this work, we model the global relationship from an orthogonal perspective and apply attention mechanisms to both spatial tokens as well as their transpose, which we refer to as image-level (global) channel tokens. In this manner, we capture both fine-grained structural patterns and global interactions.

**Hierarchical Vision Transformers.** Hierarchical designs are widely adopted to transformers [37,56,58,16,75,52,40,69,33,1,31,77,48,52,8,30,67,28,64,32,78] in vision. PVT [56] and CvT [58] perform attention on the squeezed tokens to reduce the computational cost. Swin Transformer [37], ViL [75], and HaloNet [52] apply local windows attention to the patch tokens, which capture fine-grained features and reduce the quadratic complexity to linear, but lose the ability of global modeling. To compensate for the loss of global context, Swin Transformer [37] conducts attention on the shifted local windows alternatively between consecutive blocks, and ViL [75] and HaloNet [52] play on overlapped windows. In this work, the proposed approach shares merits of hierarchical architectures and fine-grained local attention, meanwhile our proposed group channel attention still efficiently models the global context.

**Channel-wise Attentions.** The author of [27] first proposes a Squeeze-and-Excitation (SE) block as channel-wise attention to re-calibrating the channel-wise features through the squeezed global feature. Other operators in CNNs related to our work are Dynamic Head [11] and DANet [19]. They apply attention along different feature dimensions on top of the CNN backbone for a specific task. Some transformer architectures involve channel-wise operations as well to reduce the computational costs. LambdaNetworks [2] first transforms the context into a linear function lambda that is applied to the corresponding query. XcIT [1] proposes cross-covariance attention (XCA) for efficient processing of high-resolution images. Similarly, CoaT [63] introduces a factorized attention mechanism that works efficiently in a multi-branch transformer backbone. We



**Fig. 3.** Model architecture for our dual attention block. It contains two transformer blocks: spatial window self-attention and channel group self-attention blocks. By alternately using the two types of attention, our model enjoys the benefit of capturing both local fine-grained and global image-level interactions.

propose channel group attention to capture global information in transformers and demonstrate its power when combined with spatial window attention, forming our dual attention mechanism. Furthermore, we analyze in detail how our dual attention obtains global interactions as well as fine-grained local features.

### 3 Methodology

We propose Dual Attention Vision Transformers (DaViT), a clean, efficient, yet effective transformer backbone containing both local fine-grained features and global representations. In this section, we first introduce the hierarchical layout of our model. We then detail our channel group attention and the combination with spatial window attention [37].

#### 3.1 Overview

We divide the model into four stages, where a patch embedding layer is inserted at the beginning of each stage. We stack our dual attention blocks in each stage with the resolution and feature dimension kept the same. Figure 3(a) illustrates the architecture of our dual attention block, consisting of a spatial window attention block and a channel group attention block.

**Preliminaries.** Let us assume a  $\mathbb{R}^{P \times C}$  dimensional visual feature, where  $P$  is the number of total patches and  $C$  is the number of total channels. Simply applying the standard global self-attention leads to a complexity of  $O(2P^2C +$

$4PC^2$ ). It is defined as:

$$\begin{aligned} \mathcal{A}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{Concat}(\text{head}_1, \dots, \text{head}_{N_h}) \\ \text{where head}_i &= \text{Attention}(\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i) \\ &= \text{softmax} \left[ \frac{\mathbf{Q}_i(\mathbf{K}_i)^T}{\sqrt{C_h}} \right] \mathbf{V}_i \end{aligned} \quad (1)$$

where  $\mathbf{Q}_i = \mathbf{X}_i \mathbf{W}_i^Q$ ,  $\mathbf{K}_i = \mathbf{X}_i \mathbf{W}_i^K$ , and  $\mathbf{V}_i = \mathbf{X}_i \mathbf{W}_i^V$  are  $\mathbb{R}^{P \times C_h}$  dimensional visual features with  $N_h$  heads,  $\mathbf{X}_i$  denotes the  $i_{th}$  head of the input feature and  $\mathbf{W}_i$  denotes the projection weights of the  $i_{th}$  head for  $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ , and  $C = C_h * N_h$ . Please note that the output projection  $\mathbf{W}^O$  is omitted here. Considering  $P$  can be very large, *e.g.*,  $128 \times 128$ , the computational cost is immoderate.

### 3.2 Spatial Window Attention

Window attention computes self-attention within local windows, as shown in Figure 1(a). The windows are arranged to partition the image in a non-overlapping manner evenly. Supposing there are  $N_w$  different windows with each window containing  $P_w$  patches, where  $P = P_w * N_w$ . Then window attention can be represented by:

$$\mathcal{A}_{window}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \{\mathcal{A}(\mathbf{Q}_j, \mathbf{K}_j, \mathbf{V}_j)\}_{j=0}^{N_w} \quad (2)$$

where  $\mathbf{Q}_j, \mathbf{K}_j, \mathbf{V}_j \in \mathbb{R}^{P_w \times C_h}$  are local window queries, keys, and values. The computational complexity of a window-based self-attention is  $O(2PP_wC + 4PC^2)$  with a linear complexity with the spatial size  $P$ . More details of window attention are shown in Figure 3(b).

Though the computation is reduced, window attention loses the ability to model the global information. We will show that our proposed channel attention naturally solves this problem and mutually benefits with window attention.

### 3.3 Channel Group Attention

We visit self-attention from another perspective and propose channel-wise attention, as shown in Figure 3(c). Previous self-attentions [53,37,58,52,74,7] in vision define tokens with pixels or patches, and gather the information along spatial dimensions. Instead of performing attention on pixel-level or patch-level, we apply attention mechanisms on the transpose of patch-level tokens. To obtain global information in the spatial dimension, we set the number of heads equal to 1. We argue that each transposed token abstracts the global information. In this way, channel tokens interact with global information on the channel dimension in linear spatial-wise complexity, as shown in Figure 1(b).

Simply transposing the feature can obtain a vanilla channel-level attention with a complexity of  $O(6PC^2)$ . To further reduce the computational complexity, we group channels into multiple groups and perform self-attention within each

group. Formally, let  $N_g$  denotes the number of groups and  $C_g$  denotes the number of channels in each group, we have  $C = N_g * C_g$ . In this way, our channel group attention is global, with image-level tokens interacting across a group of channels. It is defined as:

$$\begin{aligned} \mathcal{A}_{channel}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \{\mathcal{A}_{group}(\mathbf{Q}_j, \mathbf{K}_j, \mathbf{V}_j)^T\}_{j=0}^{N_g} \\ \mathcal{A}_{group}(\mathbf{Q}_j, \mathbf{K}_j, \mathbf{V}_j) &= \text{softmax} \left[ \frac{\mathbf{Q}_j^T \mathbf{K}_j}{\sqrt{C_g}} \right] \mathbf{V}_j^T \end{aligned} \quad (3)$$


where  $\mathbf{Q}_j, \mathbf{K}_j, \mathbf{V}_j \in \mathbb{R}^{P \times C_g}$  are grouped channel-wise image-level queries, keys, and values. Note that although we transpose the tokens in channel attention, the projection layers  $\mathbf{W}$  and the scaling factor  $\frac{1}{\sqrt{C_g}}$  remain performed and computed along the channel dimension, rather than the spatial one. Considering that the number of spatial patches varies with the image size, the above design ensures our model can generalize to any image size.

**Complexity analysis.** Our channel group attention is performed on the image-level tokens across the channel dimension. Compared to window attention that produces an attention map with size  $P_w \times P_w$ , the channel-wise attention map is of  $C_g \times C_g$ -dimensional. The overall computational complexity of our model includes  $O(2PC(P_w + C_g))$  for window and channel attentions,  $O(8PC^2)$  for linear projections, and  $O(16PC^2)$  for FFNs (expand ratio is 4). It can be seen that our dual attention is computationally efficient with linear complexity to both the spatial and channel dimensions. FFN dominates the number of FLOPs and model parameters. Considering our dual attention has both channel-wise and spatial-wise interactions, in this work, we conduct an initial exploration to show the potential of the pure-attention structure without FFNs. Details can be found in Appendix.

**Global interactions in channel attention.** Channel attention naturally captures global information and interactions for visual recognition tasks. (i) After transposing the feature, each channel token itself is global on the spatial dimension, providing a global view of the image. (ii) Given  $C_g$  tokens with dimension  $P$ , the  $C_g \times C_g$ -dimensional attention map is computed by involving all spatial locations, *i.e.*,  $(C_g \times P) \cdot (P \times C_g)$ . (iii) With such a global attentive map, channel attention fuses multiple global views of the image dynamically, producing new global tokens and passing the information to the following spatial-wise layers. Compared to spatial-wise global attentions [49,18] that perform interactions across spatial locations, the information exchange of our channel self-attention is imposed from a global perspective rather than a patch-wise one, complementing the spatial window attention. Detailed analysis can be found in Sec. 4.

### 3.4 Model Instantiation

In this work, we follow the design strategy suggested by previous works [37,65]. Take an image with  $H \times W$ , a  $C$ -dimensional feature with a resolution of  $\frac{H}{4} \times \frac{W}{4}$  is obtained after the first patch embedding layer. And its resolution is further



**Fig. 4.** Illustrating how channel attention gathering global information by visualizing attended feature maps. The first and second columns denote the original image and a feature channel after our channel attention; while the other columns are the channel tokens with top-7 highest attention scores before the attention. Channel attention is able to select globally important regions and suppress background regions for better recognition. The third network stage of the classification model is used for visualization.

reduced into  $\frac{H}{8} \times \frac{W}{8}$ ,  $\frac{H}{16} \times \frac{W}{16}$ , and  $\frac{H}{32} \times \frac{W}{32}$  with the feature dimension increasing to  $2C$ ,  $4C$ , and  $8C$  after the other three patch embedding layer, respectively. Here, our patch embedding layer is implemented by stride convolution. The convolutional kernels and stride values of our four patch embedding layers are  $\{7, 2, 2, 2\}$  and  $\{4, 2, 2, 2\}$ , respectively.

We consider three different network configurations for image classification, objection detection, and segmentation:

- DaViT-Tiny:  $C = 96, L = \{1, 1, 3, 1\}, N_g = N_h = \{3, 6, 12, 24\}$
- DaViT-Small:  $C = 96, L = \{1, 1, 9, 1\}, N_g = N_h = \{3, 6, 12, 24\}$
- DaViT-Base:  $C = 128, L = \{1, 1, 9, 1\}, N_g = N_h = \{4, 8, 16, 32\}$ ,

where  $L$  is the layer numbers,  $N_g$  is the number of groups in channel attention, and  $N_h$  is the number of heads in window attention for each stage.


When more training data involved, we further scale up DaViT to large, huge, and giant size to validate the scaling ability of the proposed architecture for image classification:

- DaViT-Large:  $C = 192, L = \{1, 1, 9, 1\}, N_g = N_h = \{6, 12, 24, 48\}$
- DaViT-Huge:  $C = 256, L = \{1, 1, 9, 1\}, N_g = N_h = \{8, 16, 32, 64\}$
- DaViT-Giant:  $C = 384, L = \{1, 1, 12, 3\}, N_g = N_h = \{12, 24, 48, 96\}$ .

See Appendix for more details of model configurations.

## 4 Analysis

**Interpretation of global interactions.** Global interactions in transformers can be summarized into different types. Vanilla ViT [18] and DeiT [49] perform information exchange between different patches among the whole image; Focal Transformer [65] proposes interactions between tokens with different scales to get a larger receptive field; PVT [55] leverages a spatial reduction mechanism



**Fig. 5.** Feature map visualization of our DaViT, Swin [37], and DeiT [49] at four different network stages (from top to bottom representing stages 1 to 4). DeiT [49] is divided into 4 stages by 2, 2, 6, 2 layers, respectively. For all network stages, a random feature channel is visualized.

to obtain a coarse approximation of global attention; Swin [37] stacks multiple hierarchical layers to get global information eventually. Unlike them, a single block with our channel attention is able to learn the global interactions from another perspective by taking all spatial positions into account when computing the attention scores, as in  $(C_g \times P) \cdot (P \times C_g)$ . It captures the information from multiple global tokens, which represent different abstract views of the entire image. For example, different channels may contain information from different parts of an object; such part information can be aggregated into a global view.

Figure 4 illustrates how our channel attention works by visualizing the featuremaps before and after the attention. For each image in the first column, we randomly choose an output channel (the second column) in the third stage of the network and its corresponding top-7 relevant input channels for visualization. We can see that the channel attention fuses information from multiple tokens, selects globally important regions, and suppresses unimportant regions.

**Relation to Swin and DeiT.** To show the effectiveness of our dual self-attention, we make detailed comparisons with two representative clean baselines: Swin [37] with fine-grained local features, and DeiT [49] with global coarse features. Note that though our work and Swin both use window attention as a network element, the key design of Swin, *i.e.*, the use of “shifted window partitions” between successive layers to increase the receptive field, is not utilized in our work. We also simplify its relative position encoding [10] with a depth-wise convolution before layernorm to get a cleaner structure for arbitrary input sizes. To further keep our architecture clean and efficient, we do not use additional operators like “Overlapping Patch” [58,55] and “ConvFFN” [72,61]. See Appendix for detailed throughput of our model compared with Swin.

Figure 5 shows the effectiveness of our channel group attention. We randomly visualize a feature channel in each stage of DaViT, Swin [37], and DeiT [49]. We observe that: (i) Swin captures fine-grained details but no focus in the first two

**Table 1.** Comparison of image classification on ImageNet-1K for different models. All models are trained and evaluated with  $224 \times 224$  resolution on ImageNet-1K by default, unless otherwise noted. For fair comparison, token labeling [29] and distillation [49] are not used for all models and their counterparts. † and ‡ denote the model is evaluated with resolution of  $384 \times 384$  and  $512 \times 512$ , respectively.

Model	#Params (M)	FLOPs (G)	Top-1 (%)	Model	#Params (M)	FLOPs (G)	Top-1 (%)
ResNet-50 [24]	25.0	4.1	76.2	ResNet-152 [24]	60.0	11.0	78.3
DeiT-Small/16 [49]	22.1	4.5	79.8	PVT-Large [56]	61.4	<b>9.8</b>	81.7
PVT-Small [56]	24.5	3.8	79.8	DeiT-Base/16 [49]	86.7	17.4	81.8
ConvMixer-768/32 [51]	21.1	–	80.2	CrossViT-Base [5]	104.7	21.2	82.2
CrossViT-Small [5]	26.7	5.6	81.0	T2T-ViT-24 [68]	64.1	14.1	82.3
Swin-Tiny [37]	28.3	4.5	81.2	CPVT-Base [10]	88.0	17.6	82.3
CvT-13 [58]	<b>20.0</b>	4.5	81.6	TNT-Base [22]	65.6	14.1	82.8
CoAtNet-0 [12]	25.0	4.2	81.6	ViL-Base [75]	<b>55.7</b>	13.4	83.2
CaiT-XS-24 [50]	26.6	5.4	81.8	UFO-ViT-B [46]	64.0	11.9	83.3
ViL-Small [75]	24.6	5.1	82.0	Swin-Base [37]	87.8	15.4	83.4
PVTv2-B2 [55]	25.4	4.0	82.0	CaiT-M24 [50]	185.9	36.0	83.4
UFO-ViT-S [46]	21.0	<b>3.7</b>	82.0	NFNet-F0 [4]	71.5	12.4	83.6
Focal-Tiny [65]	29.1	4.9	82.2	PVTv2-B5 [55]	82.0	11.8	83.8
DaViT-Tiny (Ours)	28.3	4.5	<b>82.8</b>	Focal-Base [65]	89.8	16.0	83.8
ResNet-101 [24]	45.0	7.9	77.4	CoAtNet-2 [12]	75.0	15.7	84.1
PVT-Medium [56]	44.2	<b>6.7</b>	81.2	CSwin-B [17]	78.0	15.0	84.2
CvT-21 [58]	<b>32.0</b>	7.1	82.5	DaViT-Base (Ours)	87.9	15.5	<b>84.6</b>
UFO-ViT-M [46]	37.0	7.0	82.8	Pre-trained on ImageNet-22k			
Swin-Small [37]	49.6	8.7	83.1	Swin-Large [37] †	197.0	103.9	86.4
ViL-Medium [75]	39.7	9.1	83.3	CSwin-B [17] †	78.0	47.0	87.0
CaiT-S36 [50]	68.0	13.9	83.3	CSwin-L [17] †	173.0	96.8	87.5
CoAtNet-1 [12]	42.0	8.4	83.3	CoAtNet-3 [12] †	168.0	107.4	87.6
Focal-Small [65]	51.1	9.1	83.5	DaViT-Base (Ours) †	87.9	46.4	86.9
CSwin-S [17]	35.0	6.9	83.6	DaViT-Large (Ours) †	196.8	103.0	87.5
VAN-Large [21]	44.8	9.0	83.9	Pre-trained on 1.5B image and text pairs			
UniFormer-B [31]	50.0	8.3	83.9	DaViT-Huge (Ours) ‡	362	334	90.2
DaViT-Small (Ours)	49.7	8.8	<b>84.2</b>	DaViT-Giant (Ours) ‡	1437	1038	<b>90.4</b>

stages as it lacks global information. It can not focus on the main object until the last stage. (ii) DeiT learns coarse-grained global features over the image but loses details hence difficult to focus on the main content. (iii) Our DaViT captures both short-range and long-range visual dependencies by combining two types of self-attention. It shows strong global modeling capabilities by finding out fine-grained details of the main content in stage 1, and further focusing on some keypoints in stage 2. It then gradually refines the regions of interest from both global and local perspectives for final recognition.

## 5 Experiments

We conduct experiments on ImageNet-1K image classification [13], COCO object detection [36], and ADE20K semantic segmentation [76]. Neither token labeling [29] nor distillation [49] is used in all experiments and comparisons.

**Table 2.** Comparisons with CNN and Transformer baselines and SoTA methods on COCO object detection. The box mAP ( $AP^b$ ) and mask mAP ( $AP^m$ ) are reported for RetinaNet and Mask R-CNN trained with  $1\times$  schedule. FLOPs are measured by  $800 \times 1280$ . More detailed comparisons with  $3\times$  schedule are in Table 3.

Backbone	FLOPs (G)	RetinaNet	Mask R-CNN	
		$AP^b$	$AP^b$	$AP^m$
ResNet-50 [24]	239/260	36.3	38.0	34.4
PVT-Small [56]	226/245	40.4	40.4	37.8
ViL-Small [75]	252/174	41.6	41.8	38.5
Swin-Tiny [37]	245/264	42.0	43.7	39.8
Focal-Tiny [65]	265/291	43.7	44.8	41.0
DaViT-Tiny (Ours)	244/263	<b>44.0</b>	<b>45.0</b>	<b>41.1</b>
ResNeXt101-32x4d [62]	319/340	39.9	41.9	37.5
PVT-Medium [56]	283/302	41.9	42.0	39.0
ViL-Medium [75]	339/261	42.9	43.4	39.7
Swin-Small [37]	335/354	45.0	46.5	42.1
Focal-Small [65]	367/401	45.6	47.4	42.8
DaViT-Small (Ours)	332/351	<b>46.0</b>	<b>47.7</b>	<b>42.9</b>
ResNeXt101-64x4d [62]	473/493	41.0	42.8	38.4
PVT-Large [56]	345/364	42.6	42.9	39.5
ViL-Base [75]	443/365	44.3	45.1	41.0
Swin-Base [37]	477/496	45.0	46.9	42.3
Focal-Base [65]	514/533	46.3	47.8	43.2
DaViT-Base (Ours)	471/491	<b>46.7</b>	<b>48.2</b>	<b>43.3</b>

## 5.1 Image Classification

We compare different methods on ImageNet-1K [13]. We implement our DaViT on the timm framework [57]. Following [37,35,58,75,15], we use the same set of data augmentation and regularization strategies used in [49] after excluding repeated augmentation [3,26] and exponential moving average (EMA) [41]. We train all the models for 300 epochs with a batch size 2048 and use AdamW [39] as the optimizer. The weight decay is set to 0.05 and the maximal gradient norm is clipped to 1.0. We use a simple triangular learning rate schedule [45] as in [51]. The stochastic depth drop rates are set to 0.1, 0.2, and 0.4 for our tiny, small, and base models, respectively. During training, we crop images randomly to  $224 \times 224$ , while a center crop is used during evaluation on the validation set.

In Table 1, we summarize the results for baseline models and current state-of-the-art models on the image classification task. We can find our DaViT achieves new state-of-the-art and consistently outperforms other methods with similar model size (#Params.) and computational complexity (GFLOPs). Specifically, DaViT-Tiny, Small, and Base improve over the Swin Transformer [37] by 1.5%, 1.1%, and 1.2%, respectively. Notably, our DaViT-Small with 49.7M parameters reaches 84.2%, which surpasses all counterpart -Base models using much fewer parameters. For example, our DaViT-Small achieves 0.4% and 0.8% higher accuracy than Focal-Base and Swin-Base, respectively, using near half computations.

Following [58,18,37], when 13M images from ImageNet-22k [13] involved for pre-training, DaViT-Base and DaViT-Large obtained 86.9% and 87.5% top-1 accuracy, respectively. Furthermore, when we further scale up DaViT with 1.5B privately collected weakly supervised image-text pairs data and pre-train DaViT

**Table 3.** COCO object detection and segmentation results with RetinaNet [35] and Mask R-CNN [24]. All models are trained with 3× schedule and multi-scale inputs. The numbers before and after “/” at column 2 and 3 are the model size and complexity for RetinaNet and Mask R-CNN, respectively. FLOPs are measured by 800 × 1280.

Backbone	#Params (M)	FLOPs (G)	RetinaNet 3x						Mask R-CNN 3x					
			$AP^b$	$AP_{50}^b$	$AP_{75}^b$	$AP_S$	$AP_M$	$AP_L$	$AP^b$	$AP_{50}^b$	$AP_{75}^b$	$AP^m$	$AP_{50}^m$	$AP_{75}^m$
ResNet50 [24]	37.7/44.2	239/260	39.0	58.4	41.8	22.4	42.8	51.6	41.0	61.7	44.9	37.1	58.4	40.1
PVT-Small[56]	34.2/44.1	226/245	42.2	62.7	45.0	26.2	45.2	57.2	43.0	65.3	46.9	39.9	62.5	42.8
ViL-Small [75]	35.7/45.0	252/174	42.9	63.8	45.6	27.8	46.4	56.3	43.4	64.9	47.0	39.6	62.1	42.4
Swin-Tiny [37]	38.5/47.8	245/264	45.0	65.9	48.4	29.7	48.9	58.1	46.0	68.1	50.3	41.6	65.1	44.9
Focal-Tiny [65]	39.4/48.8	265/291	45.5	66.3	48.8	31.2	49.2	58.7	47.2	69.4	51.9	42.7	66.5	45.9
DaViT-Tiny (Ours)	38.5/47.8	244/263	<b>46.5</b>	<b>68.1</b>	<b>49.6</b>	<b>32.3</b>	<b>50.6</b>	<b>59.9</b>	<b>47.4</b>	<b>69.5</b>	<b>52.0</b>	<b>42.9</b>	<b>66.8</b>	<b>46.4</b>
ResNeXt101-32x4d [62]	56.4/62.8	319/340	41.4	61.0	44.3	23.9	45.5	53.7	44.0	64.4	48.0	39.2	61.4	41.9
PVT-Medium [56]	53.9/63.9	283/302	43.2	63.8	46.1	27.3	46.3	58.9	44.2	66.0	48.2	40.5	63.1	43.5
ViL-Medium [75]	50.8/60.1	339/261	43.7	64.6	46.4	27.9	47.1	56.9	44.6	66.3	48.5	40.7	63.8	43.7
Swin-Small [37]	59.8/69.1	335/354	46.4	67.0	50.1	31.0	50.1	60.3	48.5	70.2	53.5	43.3	67.3	46.6
Focal-Small [65]	61.7/71.2	367/401	47.3	67.8	51.0	31.6	50.9	61.1	48.8	70.5	53.6	43.8	67.7	47.2
DaViT-Small (Ours)	59.9/69.2	332/351	<b>48.2</b>	<b>69.7</b>	<b>51.7</b>	<b>32.7</b>	<b>52.2</b>	<b>62.9</b>	<b>49.5</b>	<b>71.4</b>	<b>54.7</b>	<b>44.3</b>	<b>68.4</b>	<b>47.6</b>
ResNeXt101-64x4d [62]	95.5/102	473/493	41.8	61.5	44.4	25.2	45.4	54.6	44.4	64.9	48.8	39.7	61.9	42.6
PVT-Large[56]	71.1/81.0	345/364	43.4	63.6	46.1	26.1	46.0	59.5	44.5	66.0	48.3	40.7	63.4	43.7
ViL-Base [75]	66.7/76.1	443/365	44.7	65.5	47.6	29.9	48.0	58.1	45.7	67.2	49.9	41.3	64.4	44.5
Swin-Base [37]	98.4/107.0	477/496	45.8	66.4	49.1	29.9	49.4	60.3	48.5	69.8	53.2	43.4	66.8	46.9
Focal-Base [65]	100.8/110.0	514/533	46.9	67.8	50.3	31.9	50.3	61.5	49.0	70.1	53.6	43.7	67.6	47.0
DaViT-Base (Ours)	98.5/107.3	471/491	<b>48.7</b>	<b>70.0</b>	<b>52.3</b>	<b>33.7</b>	<b>52.8</b>	<b>62.9</b>	<b>49.9</b>	<b>71.5</b>	<b>54.6</b>	<b>44.6</b>	<b>68.8</b>	<b>47.8</b>

with unified contrastive learning [70,66] approach, DaViT-Huge and DaViT-Giant reach 90.2% and 90.4% top-1 accuracy on ImageNet with 362M and 1.4B parameters, respectively.

## 5.2 Object Detection and Instance Segmentation

We benchmark our models on object detection with COCO 2017 [36]. The pre-trained models are used as visual backbones and then plugged into two representative pipelines, RetinaNet [35] and Mask R-CNN [23]. All models are pre-trained on ImageNet-1K [13] data, and are then trained on the COCO 2017 [36] training set. Results are reported on COCO 2017 [36] validation set. We follow the standard to use two training schedules, 1× schedule with 12 epochs and 3× schedule with 36 epochs. The same multi-scale training strategy as in [37] by randomly resizing the shorter side of the image to the range of [480, 800] is used. During training, we use AdamW [39] for optimization with initial learning rate  $10^{-4}$  and weight decay 0.05. We use 0.1, 0.2, and 0.3 stochastic depth drop rates to regularize the training for our tiny, small, and base models, respectively. The numbers of counterparts [37,75] are borrowed from [65].

In Table 2 and Table 3, we show the performance of our models against several state-of-the-art counterparts. The bbox mAP ( $AP^b$ ) and mask mAP ( $AP^m$ ) are reported. Results of 1× schedule shown in Table 2 have demonstrated the effectiveness of our method. We observe substantial gains across all settings and metrics compared with several strong transformer baselines.

To have more comprehensive comparisons, we further train our models with 3× schedule and show the detailed numbers, #parameters, and associated computational costs for RetinaNet and Mask R-CNN in Table 3. As we can see, even

**Table 4.** Comparison with SoTA methods for semantic segmentation on ADE20K [76] val set. Single-scale evaluation is used. FLOPs are measured by  $512 \times 2048$ .

Backbone	Method	#Params (M)	FLOPs (G)	mIoU (%)
Swin-Tiny [37]	UperNet [60]	60	945	44.5
PVT-Large [56]	SemanticFPN [34]	65	318	44.8
HRNet-w48 [54]	OCRNet [71]	71	664	45.7
Focal-Tiny [65]	UperNet [60]	62	998	45.8
XCiT-S12/16 [1]	UperNet [60]	52	-	45.9
Twins-SVT-Small [9]	UperNet [60]	54	912	46.2
DaViT-Tiny (Ours)	UperNet [60]	60	940	<b>46.3</b>
ResNet-101 [24]	UperNet [60]	86	1029	44.9
XCiT-S24/16 [1]	UperNet [60]	73	-	46.9
Swin-Small [37]	UperNet [60]	81	1038	47.6
Twins-SVT-Base [9]	UperNet [60]	88	1044	47.7
Focal-Small [65]	UperNet [60]	85	1130	48.0
ResNeSt-200 [74]	DLab.v3+ [6]	88	1381	48.4
DaViT-Small (Ours)	UperNet [60]	81	1030	<b>48.8</b>
Swin-Base [37]	UperNet [60]	121	1188	48.1
XCiT-M24/8 [1]	UperNet [60]	109	-	48.4
Twins-SVT-Large [9]	UperNet [60]	133	1188	48.8
ViT-Hybrid [42]	DPT [42]	124	1231	49.0
Focal-Base [65]	UperNet [60]	126	1354	49.0
DaViT-Base (Ours)	UperNet [60]	121	1175	<b>49.4</b>

for the  $3\times$  schedule, our models can still achieve 1.5-2.9% gains on RetinaNet  $3\times$  and 1.0-1.4% gains on Mask R-CNN  $3\times$  over Swin Transformer models.

Moreover, from Table 3 we observe a saturated and even degraded mAP in Swin Transformer [37] and Focal Transformer [65] from small to base model, while the mAP of our model is continuously increased with larger model size, showing a better scale-up ability. Our base model outperforms the state-of-the-art [65] by 1.8% on RetinaNet  $3\times$  and 0.9% on Mask R-CNN  $3\times$ .

### 5.3 Semantic Segmentation on ADE20k

Besides the instance segmentation results above, we further evaluate our model on semantic segmentation, a task that usually requires high-resolution input and long-range interactions. We benchmark our method on ADE20K [76]. Specifically, we use UperNet [60] as the segmentation method and our DaViT as the backbone. We train three models with DaViT-Tiny, DaViT-Small, DaViT-Base, respectively. For all models, we use a standard recipe by setting the input size to  $512 \times 512$  and train the model for 160k iterations with batch size 16. All the models are pre-trained on ImageNet-1K [13] set. In Table 4, we show the comparisons to previous works. As we can see, our tiny, small, and base models consistently outperform recent SoTAs, such as 1.2-1.8% gains over Swin Transformers [37] with a similar number of parameters and FLOPs.

### 5.4 Ablation Study

**Evaluation of channel group attention at different stages.** We make comparisons by inserting a dual attention block at different stages of a window

**Table 5.** The effect of channel group attention at different network stages. Taking a transformer layout with all spatial window attention blocks as the baseline (n/a), we replace two spatial attention blocks at different stages with a dual attention block to show its effectiveness. The first two spatial attention blocks are selected in the third stage to compare with other stages fairly.

Stage	#Params (M)	FLOPs (G)	Top-1 (%)	Top-5 (%)
n/a	28.3	4.6	81.1	95.6
1	28.3	4.5	81.7	95.9
2	28.3	4.5	82.2	96.1
3	28.3	4.6	82.1	96.0
4	28.3	4.6	81.9	95.9
1-4	28.3	4.6	<b>82.8</b>	<b>96.2</b>

**Table 6.** Quantitative comparisons of different dual attention layouts on ImageNet.

Model	#Params (M)	FLOPs (G)	Top-1 (%)
Window → Channel	28.3	4.5	<b>82.8</b>
Channel → Window	28.3	4.5	82.6
Hybrid (parallel)	28.3	4.5	82.6

transformer. From the quantitative results in Table 5, we observe: (i) The dual attention module consistently boosts performance at each stage. (ii) Dual attention in the second stage improves the most, as the earlier stage requires more global information. We speculate that the relatively small improvement in the first stage is that local texture features dominate the shallow part of the network. (iii) We achieve the best results when adding dual attention in all four stages.

**Dual attention layout.** We conduct experiments on the layout of our dual attention. There are three options with similar computations: (i) window attention first; (ii) channel attention first; and (iii) two types of attention are paralleled arranged (*i.e.*, half of the token dims are used for channel attention, and the other half for spatial attention, then their outputs are concatenated and fed to an FFN). The comparison is shown in Table 6. We can see that the three strategies achieve similar performance, with ‘window attention first’ slightly better.

## 6 Conclusion

This work introduces the dual attention mechanism, containing spatial window attention and channel group attention, to capture global contexts while maintaining computational efficiency. We show that these two self-attentions complement each other. We further visualize how our channel group attention captures global interactions and demonstrate its effectiveness in various benchmarks.

**Acknowledgement.** Ping Luo is supported by the General Research Fund of HK No.27208720, No.17212120, and No.17200622.

## References

1. Ali, 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. *NeurIPS* **34** (2021) [4](#), [13](#)
2. Bello, I.: Lambdanetworks: Modeling long-range interactions without attention. *arXiv preprint arXiv:2102.08602* (2021) [4](#)
3. Berman, M., Jégou, H., Vedaldi, A., Kokkinos, I., Douze, M.: Multigrain: a unified image embedding for classes and instances. *arXiv preprint arXiv:1902.05509* (2019) [11](#)
4. Brock, A., De, S., Smith, S.L., Simonyan, K.: High-performance large-scale image recognition without normalization. *arXiv: Computer Vision and Pattern Recognition* (2021) [10](#)
5. Chen, C.F., Fan, Q., Panda, R.: Crossvit: Cross-attention multi-scale vision transformer for image classification. In: *ICCV* (2021) [10](#)
6. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: *ECCV*. pp. 801–818 (2018) [13](#)
7. Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: *ICML*. pp. 1691–1703 (2020) [1](#), [2](#), [4](#), [6](#)
8. Chen, Z., Xie, L., Niu, J., Liu, X., Wei, L., Tian, Q.: Visformer: The vision-friendly transformer. In: *ICCV*. pp. 589–598 (2021) [4](#)
9. Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C.: Twins: Revisiting spatial attention design in vision transformers. *arXiv preprint arXiv:2104.13840* (2021) [3](#), [13](#)
10. Chu, X., Tian, Z., Zhang, B., Wang, X., Wei, X., Xia, H., Shen, C.: Conditional positional encodings for vision transformers. *Arxiv preprint 2102.10882* (2021) [9](#), [10](#)
11. Dai, X., Chen, Y., Xiao, B., Chen, D., Liu, M., Yuan, L., Zhang, L.: Dynamic head: Unifying object detection heads with attentions. In: *CVPR*. pp. 7373–7382 (2021) [4](#)
12. Dai, Z., Liu, H., Le, Q., Tan, M.: Coatnet: Marrying convolution and attention for all data sizes. *Advances in Neural Information Processing Systems* **34** (2021) [10](#)
13. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *CVPR*. pp. 248–255. *Ieee* (2009) [10](#), [11](#), [12](#), [13](#)
14. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: *NAACL-HLT* (1) (2019) [4](#)
15. Ding, M., Huo, Y., Lu, H., Yang, L., Wang, Z., Lu, Z., Wang, J., Luo, P.: Learning versatile neural architectures by propagating network codes. In: *ICLR* (2022) [11](#)
16. Ding, M., Lian, X., Yang, L., Wang, P., Jin, X., Lu, Z., Luo, P.: Hr-nas: Searching efficient high-resolution neural architectures with lightweight transformers. In: *CVPR* (2021) [4](#)
17. 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) [10](#)
18. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR* (2021) [1](#), [2](#), [4](#), [7](#), [8](#), [11](#)

19. Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: CVPR. pp. 3146–3154 (2019) [4](#)
20. Graham, B., El-Nouby, A., Touvron, H., Stock, P., Joulin, A., Jégou, H., Douze, M.: Levit: a vision transformer in convnet’s clothing for faster inference. In: ICCV. pp. 12259–12269 (2021) [4](#)
21. Guo, M.H., Lu, C.Z., Liu, Z.N., Cheng, M.M., Hu, S.M.: Visual attention network. arXiv preprint arXiv:2202.09741 (2022) [10](#)
22. Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. NeurIPS **34** (2021) [10](#)
23. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV. pp. 2961–2969 (2017) [12](#)
24. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016) [10](#), [11](#), [12](#), [13](#)
25. Heo, B., Yun, S., Han, D., Chun, S., Choe, J., Oh, S.J.: Rethinking spatial dimensions of vision transformers. In: ICCV. pp. 11936–11945 (2021) [4](#)
26. Hoffer, E., Ben-Nun, T., Hubara, I., Giladi, N., Hoefler, T., Soudry, D.: Augment your batch: Improving generalization through instance repetition. In: CVPR. pp. 8129–8138 (2020) [11](#)
27. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: CVPR. pp. 7132–7141 (2018) [4](#)
28. 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) [4](#)
29. Jiang, Z.H., Hou, Q., Yuan, L., Zhou, D., Shi, Y., Jin, X., Wang, A., Feng, J.: All tokens matter: Token labeling for training better vision transformers. NeurIPS **34** (2021) [10](#)
30. Li, C., Tang, T., Wang, G., Peng, J., Wang, B., Liang, X., Chang, X.: Bossnas: Exploring hybrid cnn-transformers with block-wisely self-supervised neural architecture search. In: ICCV. pp. 12281–12291 (2021) [4](#)
31. Li, K., Wang, Y., Zhang, J., Gao, P., Song, G., Liu, Y., Li, H., Qiao, Y.: Uniformer: Unifying convolution and self-attention for visual recognition. arXiv preprint arXiv:2201.09450 (2022) [4](#), [10](#)
32. Li, Y., Wu, C.Y., Fan, H., Mangalam, K., Xiong, B., Malik, J., Feichtenhofer, C.: Improved multiscale vision transformers for classification and detection. arXiv: Computer Vision and Pattern Recognition (2021) [4](#)
33. Li, Y., Zhang, K., Cao, J., Timofte, R., Van Gool, L.: Localvit: Bringing locality to vision transformers. arXiv preprint arXiv:2104.05707 (2021) [4](#)
34. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR. pp. 2117–2125 (2017) [13](#)
35. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. pp. 2980–2988 (2017) [11](#), [12](#)
36. 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: ECCV. pp. 740–755. Springer (2014) [10](#), [12](#)
37. 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: ICCV (2021) [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [9](#), [10](#), [11](#), [12](#), [13](#)
38. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. arXiv preprint arXiv:2201.03545 (2022) [1](#)
39. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017) [11](#), [12](#)

40. Pan, Z., Zhuang, B., Liu, J., He, H., Cai, J.: Scalable visual transformers with hierarchical pooling. arXiv preprint arXiv:2103.10619 (2021) [4](#)
41. Polyak, B.T., Juditsky, A.B.: Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization* **30**(4), 838–855 (1992) [11](#)
42. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: ICCV. pp. 12179–12188 (2021) [13](#)
43. Riquelme, C., Puigcerver, J., Mustafa, B., Neumann, M., Jenatton, R., Sussano Pinto, A., Keysers, D., Houlsby, N.: Scaling vision with sparse mixture of experts. *NeurIPS* **34** (2021) [4](#)
44. Ryoo, M.S., Piergiovanni, A., Arnab, A., Dehghani, M., Angelova, A.: Token-learner: What can 8 learned tokens do for images and videos? arXiv: Computer Vision and Pattern Recognition (2021) [4](#)
45. Smith, L.N., Topin, N.: Super-convergence: Very fast training of neural networks using large learning rates. In: *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*. vol. 11006, p. 1100612. International Society for Optics and Photonics (2019) [11](#)
46. Song, J.g.: Ufo-vit: High performance linear vision transformer without softmax. arXiv preprint arXiv:2109.14382 (2021) [10](#)
47. Srinivas, A., Lin, T.Y., Parmar, N., Shlens, J., Abbeel, P., Vaswani, A.: Bottleneck transformers for visual recognition. In: CVPR. pp. 16519–16529 (2021) [4](#)
48. Tang, S., Zhang, J., Zhu, S., Tan, P.: Quadtree attention for vision transformers. arXiv preprint arXiv:2201.02767 (2022) [4](#)
49. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: ICML. pp. 10347–10357. PMLR (2021) [4](#), [7](#), [8](#), [9](#), [10](#), [11](#)
50. Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., Jégou, H.: Going deeper with image transformers. arXiv preprint arXiv:2103.17239 (2021) [10](#)
51. Trockman, A., Kolter, J.Z.: Patches are all you need? arXiv preprint arXiv:2201.09792 (2022) [10](#), [11](#)
52. Vaswani, A., Ramachandran, P., Srinivas, A., Parmar, N., Hechtman, B., Shlens, J.: Scaling local self-attention for parameter efficient visual backbones. In: CVPR. pp. 12894–12904 (2021) [2](#), [3](#), [4](#), [6](#)
53. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *NeurIPS* (2017) [4](#), [6](#)
54. Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., et al.: Deep high-resolution representation learning for visual recognition. *IEEE TPAMI* (2020) [13](#)
55. 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. arXiv preprint arXiv:2106.13797 (2021) [2](#), [8](#), [9](#), [10](#)
56. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: ICCV (2021) [2](#), [4](#), [10](#), [11](#), [12](#), [13](#)
57. Wightman, R.: Pytorch image models. URL <https://github.com/rwightman/pytorch-image-models>.(cited on p.) (2019) [11](#)
58. 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) [2](#), [4](#), [6](#), [9](#), [10](#), [11](#)
59. Wu, K., Peng, H., Chen, M., Fu, J., Chao, H.: Rethinking and improving relative position encoding for vision transformer. In: ICCV. pp. 10033–10041 (2021) [4](#)

60. Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: ECCV. pp. 418–434 (2018) [13](#)
61. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. In: NeurIPS (2021) [2](#), [9](#)
62. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: CVPR. pp. 1492–1500 (2017) [11](#), [12](#)
63. Xu, W., Xu, Y., Chang, T., Tu, Z.: Co-scale conv-attentional image transformers. In: ICCV. pp. 9981–9990 (2021) [4](#)
64. Xu, Y., Zhang, Q., Zhang, J., Tao, D.: Vitae: Vision transformer advanced by exploring intrinsic inductive bias. NeurIPS **34** (2021) [4](#)
65. Yang, J., Li, C., Zhang, P., Dai, X., Xiao, B., Yuan, L., Gao, J.: Focal self-attention for local-global interactions in vision transformers. arXiv preprint arXiv:2107.00641 (2021) [3](#), [7](#), [8](#), [10](#), [11](#), [12](#), [13](#)
66. Yang, J., Li, C., Zhang, P., Xiao, B., Liu, C., Yuan, L., Gao, J.: Unified contrastive learning in image-text-label space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19163–19173 (June 2022) [12](#)
67. Yu, Q., Xia, Y., Bai, Y., Lu, Y., Yuille, A.L., Shen, W.: Glance-and-gaze vision transformer. NeurIPS **34** (2021) [4](#)
68. Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token vit: Training vision transformers from scratch on imagenet. arXiv preprint arXiv:2101.11986 (2021) [10](#)
69. Yuan, L., Hou, Q., Jiang, Z., Feng, J., Yan, S.: Volo: Vision outlooker for visual recognition. arXiv preprint arXiv:2106.13112 (2021) [4](#)
70. Yuan, L., Chen, D., Chen, Y.L., Codella, N., Dai, X., Gao, J., Hu, H., Huang, X., Li, B., Li, C., et al.: Florence: A new foundation model for computer vision. arXiv preprint arXiv:2111.11432 (2021) [12](#)
71. Yuan, Y., Chen, X., Wang, J.: Object-contextual representations for semantic segmentation. arXiv preprint arXiv:1909.11065 (2019) [13](#)
72. Yuan, Y., Fu, R., Huang, L., Lin, W., Zhang, C., Chen, X., Wang, J.: Hrformer: High-resolution transformer for dense prediction. In: NeurIPS (2021) [2](#), [9](#)
73. Zhai, X., Kolesnikov, A., Houlsby, N., Beyer, L.: Scaling vision transformers. arXiv: Computer Vision and Pattern Recognition (2021) [4](#)
74. Zhang, H., Wu, C., Zhang, Z., Zhu, Y., Lin, H., Zhang, Z., Sun, Y., He, T., Mueller, J., Manmatha, R., et al.: Resnest: Split-attention networks. arXiv preprint arXiv:2004.08955 (2020) [6](#), [13](#)
75. 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: ICCV (2021) [2](#), [3](#), [4](#), [10](#), [11](#), [12](#)
76. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: CVPR. pp. 633–641 (2017) [10](#), [13](#)
77. Zhou, D., Kang, B., Jin, X., Yang, L., Lian, X., Jiang, Z., Hou, Q., Feng, J.: Deepvit: Towards deeper vision transformer. arXiv preprint arXiv:2103.11886 (2021) [4](#)
78. Zhou, J., Wang, P., Wang, F., Liu, Q., Li, H., Jin, R.: Elsa: Enhanced local self-attention for vision transformer. arXiv: Computer Vision and Pattern Recognition (2021) [4](#)