

# SG-Former: Self-guided Transformer with Evolving Token Reallocation

Sucheng Ren, Xingyi Yang, Songhua Liu, Xinchao Wang\*  
National University of Singapore

## Abstract

Vision Transformer has demonstrated impressive success across various vision tasks. However, its heavy computation cost, which grows quadratically with respect to the token sequence length, largely limits its power in handling large feature maps. To alleviate the computation cost, previous works rely on either fine-grained self-attentions restricted to local small regions, or global self-attentions but to shorten the sequence length resulting in coarse granularity. In this paper, we propose a novel model, termed as Self-guided Transformer (SG-Former), towards effective global self-attention with adaptive fine granularity. At the heart of our approach is to utilize a significance map, which is estimated through hybrid-scale self-attention and evolves itself during training, to reallocate tokens based on the significance of each region. Intuitively, we assign more tokens to the salient regions for achieving fine-grained attention, while allocating fewer tokens to the minor regions in exchange for efficiency and global receptive fields. The proposed SG-Former achieves performance superior to state of the art: our base size model achieves **84.7%** Top-1 accuracy on ImageNet-1K, **51.2mAP** bbAP on CoCo, **52.7mIoU** on ADE20K surpassing the Swin Transformer by **+1.3%** / **+2.7 mAP** / **+3 mIoU**, with lower computation costs and fewer parameters. The code is available at <https://github.com/OliverRensu/SG-Former>

## 1. Introduction

Transformers [49], originated from natural language processing (NLP), have recently demonstrated state-of-the-art performance in visual learning. The pioneering work of Vision Transformer (ViT) [10] introduces the self-attention module and explicitly models the long-range dependency between image patches, which overcomes the inherent limitation of the local receptive field in convolution, thereby enhancing the performances of various tasks [9, 27, 56, 70, 74, 1].

Despite tremendous successes, the computation cost of self-attention grows quadratically with respect to the se-

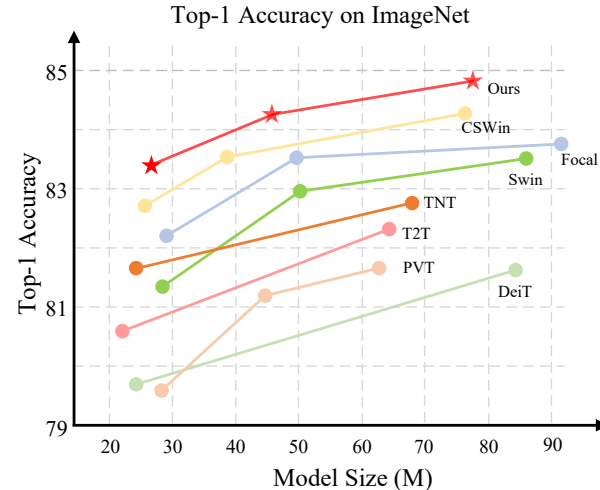


Figure 1. Top-1 accuracy on ImageNet of recent Transformer models. Our proposed SG-Former yields significant improvements over all the baselines, including Swin Transformer and the state-of-the-art CSWin Transformer.

quence length, which, in turn, greatly limits its applicability to large-scale inputs. To reduce the computation cost, ViT takes a large stride patch embedding to reduce the sequence lengths. However, such operations inevitably result in self-attention applicable only to small-size feature maps with coarse granularity. To compute self-attention at high-resolution features, some approaches of [27, 9, 48] are proposed to restrict the self-attention region into a local window instead of the whole feature map (i.e., *fine-grained local self-attention*). For example, Swin Transformer designs window attention, and CSWin designs cross-shape attention. As such, these methods [27, 9, 48] sacrifice the power of modeling global information in each self-attention layer. Another stream of methods [53, 59, 52] aims to aggregate tokens across the whole key-value feature maps to reduce the global sequence length (i.e., *coarse-grained global attention*). For instance, Pyramid vision Transformer (PVT) [53] uniformly aggregates tokens across the whole feature map with a large kernel of the large stride, which results in uniform coarse information over the whole feature map.

In this paper, we introduce a novel Transformer model,

\*Corresponding Author: xinchao@nus.edu.sg




Figure 2. Visualization of our core idea. (Right) A significance map from Transformer itself is used as guidance. Given an input sequence, SG-Former reallocates more tokens at salient regions (like the face of the dog) for fine-grained information and fewer tokens at backgrounds (like the wall) for the global receptive field with computation efficiency. (Left) PVT takes a predefined uniform map to aggregate tokens, regardless of their semantics.

termed as self-guided Transformer (SG-Former), towards *global attention with adaptive fine granularity* through an evolving self-attention design. The core idea of SG-Former lies in that, we preserve the long-range dependency across the whole feature map, while reallocating tokens based on the significance of image regions. Intuitively, we tend to assign more tokens to the salient regions so that each token may interact with salient regions at fine granularity, and meanwhile allocate fewer tokens over the minor regions in exchange for efficiency. SG-Former keeps estimating self-attention efficiently with global receptive fields and attending to fine-grained information at salient regions adaptively. As shown in Figure 2, our SG-Former reallocates more tokens at salient regions like the dog and fewer tokens at minor regions like the wall according to the significance maps obtained from itself. PVT, on the other hand, adopts predefined strategies to aggregate tokens uniformly.

Specifically, we keep the query tokens but reallocate the key and value tokens for efficient global self-attention. The significance of image regions, in the form of a score map, is *per se* estimated through a hybrid-scale self-attention and further utilized for guiding the token reallocation. In other words, given an input image, the token reallocation is accomplished through *self guidance*, indicating that each image undergoes a unique token reallocation customized only for itself. The reallocated tokens are, therefore, less influenced by humans prior. In addition, such self-guidance *evolves* due to progressively accurate significance map prediction during training. The significance maps greatly affect the efficacy of the reallocation, and as such, we propose a hybrid-scale self-attention that accounts for various granularity information within one layer with the same cost as Swin. The various granularity information in hybrid-scale self-attention is achieved by grouping heads and diversifying each group for different attention granularity. The hybrid-scale self-attention also provides hybrid-scale information to the whole Transformer.

Our contributions are therefore summarized as follows.

- We introduce a novel Transformer model, SG-Former, through unifying hybrid-scale information extraction, including fine-grained local and global coarse-grained information within one self-attention layer. With unified local global hybrid-scale information, we predict the significance map to identify regional significance.
- With the significance map, we model self-guided attention to automatically locate salient regions and keep salient regions fine-grained for accurate extract information and minor regions coarse-grained for low computation cost.
- We evaluate our proposed Transformer backbone on various downstream tasks, including classification, object detection, and segmentation. Experimental results demonstrate that SG-Former consistently outperforms previous Vision Transformers under similar model sizes.

## 2. Related Work

**Vision Transformers.** Vision Transformer (ViT) [11] originated from Transformer[49] for sequences in natural language processing (NLP) first introduces a pure Transformer architecture into computer vision. Such architecture achieves impressive improvements in image classification over previous convolution nets. The requirements of large-scale datasets like JFT-300M in training ViT limits its application on various tasks. Some works [45, 46, 36, 39, 38] propose complex data augmentation, regularization, knowledge distillation, and training strategies. Recently, some work design better Transformer architectures based on Vision Transformer for detection [55, 43], segmentation [42, 71, 37], restoration [35], retrieval [12], re-id [19], with improved scale-ability [41, 4] and robustness [22]. Different from these work, recent work [9, 27, 59, 68, 66, 47, 65, 23] aim at designing general Transformer backbone for general vision tasks instead of some specific tasks. The most significant difference between these Transformers and vanilla

ViT is that these Transformers take hierarchical architecture following the same idea of convolution instead of the single-scale features in ViT and more efficient self-attentions in dealing with large-size feature maps. Notably, it has been demonstrated that the overall architecture [63, 64], rather than the specialized attention operation, plays a crucial role in achieving superior performance. As a result, this insight has gives rise to more powerful and efficient architectures [3, 29].

**Efficient Deep Learning.** Designing efficient deep learning models with minimal training time [61, 32, 62], low inference cost [20, 67, 60, 13, 30, 14], and reduced data requirements [51, 25, 26] has been a long-standing pursuit. In the case of ViT, efficient self-attentions become crucial when dealing with a large number of tokens, such as processing long documents in NLP or high-resolution images in computer vision, where self-attention can incur quadratic computational and memory costs. On the one hand, local fine-grained self-attention limits the query token attending neighborhood tokens within a predefined window. Swin Transformer [27] split feature maps into window shape regions to limit the self-attention locally. Axial attention [21] propose calculating stripe window attention. Focal [59] Transformer predefined the coarse-fine granularity regions and attend different regions with different granularities. CSWin [9] designs windows with horizontal and vertical stripes. CSWin splits heads into parallel groups and applies different attention to different groups. On the other hand, global coarse-grained self-attention keeps the ability to model global dependency in self-attention but faces coarse-grained information processing. PVT [53] tries to keep the global receptive field with token reduction but leads to coarse-grained self-attention. Focal Transformer [59] predefines more attention to close surroundings and less attention to far away positions. However, such inductive bias limits the model’s ability to attend to salient objects far away from the current token position. Different from previous work, we propose to give the self-attention module more flexibility and let it find the best way to reduce the computation cost only with some predefined guidance instead of a specific method.

### 3. Method

#### 3.1. Overview

The overall pipeline of SG-Former is shown in Figure 3. SG-Former shares the same patch embedding layer and four-stage pyramid architecture as previous CNN and Transformer models [17, 53, 27, 9]. An image  $X \in \mathbb{R}^{H \times W \times 3}$  is first downsampled  $4\times$  by the patch embedding layer at the input level. A downsampling layer with a rate  $2\times$  exists between two stages. Therefore, the feature map is of size  $\frac{H}{2^{i+1}} \times \frac{W}{2^{i+1}}$  at the  $i$ -th stage. Except for the last stage, each

stage has  $N_i$  Transformer blocks, consisting of two types of blocks repeatedly: (i) a hybrid-scale Transformer block and (ii) a self-guided Transformer block. The hybrid-scale self-attention extracts hybrid-scale objects and multi-granularity information, guiding regional significance; its detail will be covered in Section 3.3. The self-guided self-attention models global information while keeping the fine-granularity at salient regions according to the significance information from the hybrid-scale Transformer block; its detail will be covered in Section 3.2.

#### 3.2. Self-Guided Attention

Despite the capability of modeling extended range information, the high computation cost and memory consumption of vanilla self-attention are quadratic versus the sequence length, limiting its applicability on large-size feature maps in various computer vision tasks, such as segmentation and detection. Recent works [53, 52] suggest reducing sequence length with token aggregation by merging a few tokens into one. However, such aggregation treats each token equally and ignores the inherent significance difference of different tokens. Such aggregation faces two problems: (i) information might lose or mix with irrelevant information at salient regions, and (ii) at a minor region or background region, a lot of tokens (a higher proportion of sequences) are redundant for simple semantics while demanding a lot of computation. Motivated by this observation, we propose self-guided attention, which uses significance as guidance to aggregate tokens. Namely, in salient regions, more tokens are preserved for fine-grained information, while in minor regions, fewer tokens are maintained to keep the global view of self-attention and reduce computation cost at the same time. As shown in Figure 4, "self-guided" indicates that the Transformer itself determines our computation cost reduction strategy during its training, instead of human-introduced prior knowledge, such as window attentions in Swin [27], cross-shape attention in CSWin [9], static spatial reduction in PVT [53].

The input feature map  $X \in \mathbb{R}^{h \times w \times c}$  are firstly projected to Query( $Q$ ), Key( $K$ ) and Value( $V$ ). Next,  $H$  independent self-attention heads compute self-attention in parallel. To reduce the computation cost while keeping the size of feature maps unchanged after attention, we fix the length of  $Q$  but aggregate tokens of  $K$  and  $V$  with importance guided aggregation module (IAM).

$$\begin{aligned} Q &= XW^Q, \\ K &= \text{IAM}(X, S, r)W^K, \\ V &= \text{IAM}(X, S, r)W^V. \end{aligned} \tag{1}$$

The goal of IAM is aggregate fewer tokens to one (i.e., preserve more) at salient regions and more tokens to one (i.e., preserve less) at background regions. In Eq. (1), the significance map  $S \in \mathbb{R}^{h \times w}$  contains the information of regional

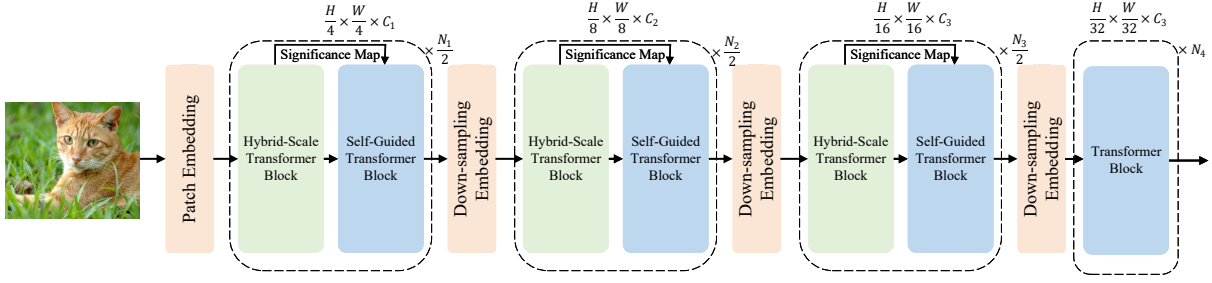


Figure 3. The overall architecture of SG-Former. Hybrid-Scale Transformer blocks extract hybrid-scale information and provide significance for self-guided attention.




Figure 4. Comparison of fusion. (A) PVT treats all tokens/regions to be equal, indicating that their significance is uniform across the whole map. All input images share the same aggregation strategies. (B) Focal Transformer attempts to preserve more tokens at central regions and fewer tokens at background regions, indicating that the significance is larger in the central regions and smaller at the background regions. All input images share the same aggregation strategies. (C) Ours takes the model-predicted significance maps as guidance for aggregation, and preserves more tokens (the blue ones) in salient regions and fewer tokens in minor regions (the green ones). The significance depends on the input images and the model.

significance with multiple granularities (see Section 3.3). We sort the value of the significance map in ascending order and evenly divide  $S$  into  $n$  sub-regions  $S^1, \dots, S^n$ . Consequently,  $S^n$  and  $S^1$  are the most important and minor regions, respectively. Meanwhile, all tokens in  $X$  are grouped into  $X^1, \dots, X^n$  according to  $S^1, \dots, S^n$ . In Eq. (1),  $r$  represents the aggregation rate that every  $r$  tokens are aggregated into one token. We set different aggregation rate  $r^1, \dots, r^n$  at areas of different significance, so that each sub-region has one aggregation rate, and the more important the sub region is, the smaller aggregation rate is. The specific values of  $r$  for different stages are listed in Table 1. Therefore, the IAM aggregate grouped input features  $X^1, \dots, X^n$  group by group each group with different aggregation rate reallocated

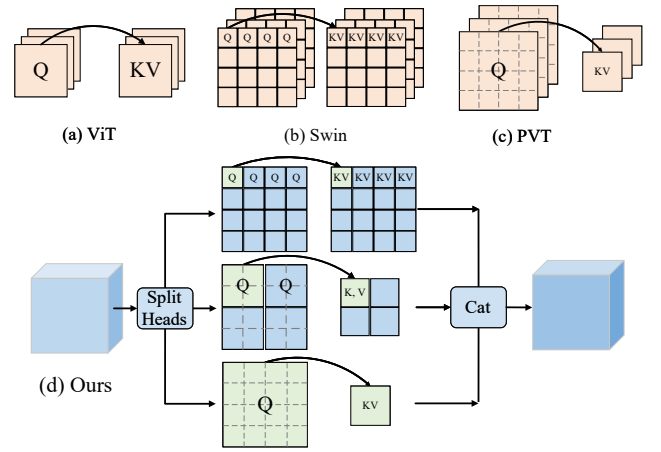


Figure 5. Comparison of Local Global information and hybrid-scale. (A) Self-attention in ViT is single scale and has huge computation cost, thereby applicable to only small feature maps. (B) Self-attention in Swin Transformer is single scale and is limited within small windows, which reduces computation cost and preserves fine-grained information; however, the attention is local. (C) Self-attention in PVT is single scale. It reduces tokens alongside computation cost, and preserves global receptive field; however, it captures coarse information only. (D) Our Self-attention is hybrid-scale and performs attention from local to global areas. It reduces computation cost, preserves global receptive field, and captures fine-grained hybrid-scale information.

tokens by concatenating them of each group.

$$\begin{aligned} \hat{X}^i &= F(X^i, r_i), \\ \hat{X} &= \text{Cat}(\hat{X}^1, \dots, \hat{X}^n) \end{aligned} \quad (2)$$

where  $F(X, r)$  is the aggregation function, we implement it by fully connected layers with input dimension of  $r$  and output dimension of 1. The number of tokens in  $\hat{X}_i$  equals the number of tokens in  $X^i$  divided by  $r_i$ .

### 3.3. Hybrid-scale Attention

The hybrid-scale attention serves for two purposes: (i) to extract hybrid-scale global and fine-grained information

	Size	SG-Former-S	SG-Former-M	SG-Former-B
Stage1	56x56	$s=1, 8$ $r=196, 56, 56, 28$ $C_1=64, head=2, N_1=2$	$s=1, 8$ $r=196, 56, 56, 28$ $C_1=64, head=2, N_1=2$	$s=1, 2, 4, 8$ $r=196, 56, 56, 28$ $C_1=96, head=4, N_1=4$
Stage2	28x28	$s=1, 4$ $r=49, 14, 14, 7$ $C_2=128, head=4, N_2=4$	$s=1, 4$ $r=49, 14, 14, 7$ $C_2=128, head=4, N_2=6$	$s=1, 2, 4$ $r=49, 14, 14, 7$ $C_2=192, head=6, N_2=6$
Stage3	14x14	$s=1, 2$ $r=2, 1$ $C_3=256, head=8, N_4=16$	$s=1, 2$ $r=2, 1$ $C_3=256, head=8, N_4=28$	$s=1, 2$ $r=2, 1$ $C_3=384, head=12, N_4=24$
Stage4	7x7	$C_4=512, head=16, N_4=1$	$C_4=512, head=16, N_4=2$	$C_4=768, head=24, N_4=2$

Table 1. Several model variants of our self-guided Transformer.  $C$  and  $N$  represent the dimension and number of blocks, respectively.  $head$  indicates the number of heads.

with no more computation cost than the window attentions in Swin Transformer, and (ii) to provide the significance for self-guided attention.

As shown in Figure 5, the input features  $X$  are projected into Query(Q), Key(K), and Values(V). Then the multi-head self-attention adopts  $H$  independent heads. Usually, these  $H$  independent heads perform within the same local regions and thus lack head diversity. By contrast, we evenly divide  $H$  heads into  $h$  groups and inject hybrid-scale and multi-receptive field attentions into these  $h$  groups, where there are  $\frac{H}{h}$  heads in each group. In the  $i$ -th head which belong to the  $j$ -th group, with scale  $s_j$  (where  $j = 1, \dots, h$ ), every  $s_j \times s_j$  tokens in  $\{K, V\}$  are merged to one token. Next, we split  $\{Q, K, V\}$  into windows. The window size of  $\{K, V\}$  is set to  $M$  and remains unchanged across all groups. To align the window size of  $\{Q\}$  and  $\{K, V\}$  with the token merge in  $\{K, V\}$ , window size of  $\{Q\}$  is chosen to be  $s_j M \times s_j M$ ,  $s_j$  times larger than that of  $\{K, V\}$ . The receptive field of attention is significantly enlarged by  $s_j$  times:

$$\begin{aligned} Q_i &= XW_i^Q, \\ K_i &= \text{Merge}(X, s_j)W_i^K, \\ V_i &= \text{Merge}(X, s_j)W_i^V, \end{aligned} \quad (3)$$

where  $\text{Merge}(X, s_j)$  indicates merging every  $s_j \times s_j$  tokens in  $X$  to one token, which is implemented with convolution of stride  $s_j$ . A special case is when  $s_j$  equals 1, where no token merges and  $\{Q, K, V\}$  have windows of the same sizes.

$$\begin{aligned} \text{Atten}_i &= \text{Softmax} \left( \frac{P(Q_i, s_j M)P(K_i^T, M)}{\sqrt{d_h}} \right), \\ h_i &= \text{Atten}_i P(V_i, M) \end{aligned} \quad (4)$$

where  $P(X, s_j M)$  indicates window partition with window size of  $s_j M \times s_j M$ .  $\text{Atten}_j$  is the attention map. There is a special case:  $s_j M \times s_j M$  equals to  $h \times w$ , where no window

partition is required and all tokens in  $\{K, V\}$  are attended by  $\{Q\}$ , leading to global information extraction.

The significance of the token is taken to be the sum of the similarities between all tokens and itself:

$$\begin{aligned} S_i &= \frac{1}{hw} \sum_{m=1}^{hw} \text{Atten}_i^{m,n}, \\ S &= \sum_{i=1}^h S_i, \end{aligned} \quad (5)$$

where  $S$  is the final significance map by summing over  $S_i$  for hybrid-scale guidance with both global and fine-grained information.

### 3.4. Transformer Block

With two self-attention mechanisms, we correspondingly design two types of Transformer blocks. These two Transformer blocks only differ at the attention layer, while all others are kept identical:

$$\begin{aligned} \hat{X}^i &= \text{MS-Attention}(\text{LN}(X^{i-1})) + X^{i-1}, \\ \hat{X}^i &= \text{MLP}(\text{LN}(X^{i-1})) + X^{i-1}, \end{aligned} \quad (6)$$

$$\begin{aligned} \hat{X}^i &= \text{SG-Attention}(\text{LN}(X^{i-1})) + X^{i-1}, \\ \hat{X}^i &= \text{MLP}(\text{LN}(X^{i-1})) + X^{i-1}, \end{aligned} \quad (7)$$

As shown in Figure 3, the first three stages are customized using our proposed hybrid-scale or self-guided Transformer blocks, while for the last stage, we use a vanilla Transformer block following previous Transformers [45, 11, 27, 9]. Note that the numbers of Transformer blocks at the first three stages (i.e.,  $N_1, N_2$ , and  $N_3$ ) are even, while that at the last stage (i.e.,  $N_4$ ) can be even or odd.



Model	Params (M)	FLOPs (G)	Top1 (%)
ResNet-50 [17]	25.0	4.1	76.2
RegNetY-4G [34]	20.6	4.0	79.4
EfficientNet-B4* [44]	19	4.2	82.9
DeiT-S [45]	22.1	4.6	79.9
DeepViT-S [73]	27.0	6.2	82.3
TNT-S [15]	23.8	5.2	81.3
CViT-15 [2]	27.4	5.6	81.5
PVT-S [53]	24.5	3.8	79.8
Swin-T [27]	28.3	4.5	81.2
Twin-S [6]	24	2.9	81.7
LITv2-S [33]	28	3.7	82.0
Focal-T [59]	29.1	4.9	82.2
ResTv2-T [69]	30	4.1	82.3
PerViT-S [31]	21	4.4	82.1
CSWin-T [9]	23	4.3	82.7
CETNet-T [50]	23	4.3	82.7
SG-Former-S	22.5	4.8	<b>83.2</b>
ResNet-101 [17]	45.0	7.9	77.4
EfficientNet-B5* [44]	30	9.9	83.6
Swin-S [27]	49.6	8.7	83.1
PVT-M [53]	44.2	8.7	81.2
Twin-B [6]	56	8.6	83.2
Focal-S [59]	51.1	9.1	83.5
CSWin-S [9]	35	6.9	83.6
SG-Former-M	38.7	7.5	<b>84.1</b>
EfficientNet-B6* [44]	43	19.0	84.0
ViT-B [11]	86.6	17.6	77.9
DeiT-B [45]	86.6	17.5	81.8
Swin-B [27]	87.8	15.4	83.4
Twin-L [6]	99.2	15.1	83.7
PVT-L [53]	61.4	9.8	81.7
LITv2-B [33]	87	13.2	83.6
CETNet-B [50]	75	15.1	83.8
Focal-B [59]	89.8	16.0	83.8
CSWin-B [9]	78	15.0	84.2
ResTv2-L [69]	87	13.8	84.2
SG-Former-B	77.9	15.6	<b>84.7</b>

Table 2. Comparison of different backbones on ImageNet-1K classification. Except EfficientNet (EffNet-B4\*), all models are trained and evaluated on the input size of  $224 \times 224$ . Our Transformer consistently outperforms previous state-of-the-art ones under all small, medium, and base model sizes.

### 3.5. Transformer Architecture Variants

We build three models of small (SG-Former-S), medium (SG-Former-M), and base (SG-Former-B) sizes to fairly compare with other Transformers. Their differences including channels, heads, and number of blocks in each stage are detailed in Table 1. Both small and medium models have two scales in each stage, while the base model has four scales at the first stage and three scales in the second stage. Following [9, 40], we also adopt convolution

Model	Params (M)	FLOPs (G)	Top1 (%)
CvT-13	20	16.3	83.0
T2T-14	22	17.1	83.3
CViT-15	28	21.4	83.5
CSWin-T	23	14.0	84.3
SG-Former-S	23	15.1	84.7
CvT-21	32	24.9	83.3
CViT-18	45	32.4	83.9
CSWin-T	35	22.0	85.0
SG-Former-M	39	23.4	85.3
ViT-B/15	86	48.3	77.9
DeiT-B	86	55.4	83.1
Swin-B	88	47.0	84.2
CSWin-B	78	47.0	85.4
SG-Former-B	78	47.9	85.8

Table 3. Classification results on ImageNet-1K with input resolution of  $384 \times 384$ .

patch embedding, locally-enhanced detail-specific feedforward layer, and for simplicity, these details are not shown in equations.

## 4. Experiments

To demonstrate the architectural superiority of SG-Former backbone, we conduct experiments on ImageNet-1K classification [8], COCO object detection [24], and ADE20K semantic segmentation [72]. Besides, we ablate the influence of single scale and hybrid-scale self-attention, as well as the influence of different guidance in self-evolving attention.

### 4.1. Classification on ImageNet-1K

To fairly compare with previous methods, we follow the same training setting [45, 9]. Specifically, all models are trained for 300 epochs on ImageNet-1K with the input resolution of  $224 \times 224$  and batch size of 1024. For optimization, we take the AdamW optimizer with the weight decay of 0.05 and cosine learning rate decay with linear warm-up with 20 epochs, and the peak learning rate is set to  $1e-3$ . For data augmentation and regularization, we keep the same as CSWin [9], including RandAugment, Exponential Moving Average, data mixing, label smoothing, stochastic depth, and Random Erasing. During inference, we resize and center crop to  $224 \times 224$ .

In Table 2, we compare SG-Former with state-of-the-art CNN and Transformer architectures. SG-Former consistently outperforms competitors under similar parameters and computation costs. Specifically, SG-Former achieves 83.2%, 84.1%, and 84.7% Top-1 Accuracy under small, medium, and base models, respectively. Comparing the

Backbone	Params (M)	Mask R-CNN 1× schedule						Mask R-CNN 3× schedule + MS					
		$AP^b$	$AP_{50}^b$	$AP_{75}^b$	$AP^m$	$AP_{50}^m$	$AP_{75}^m$	$AP^b$	$AP_{50}^b$	$AP_{75}^b$	$AP^m$	$AP_{50}^m$	$AP_{75}^m$
Res50 [17]	44	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-S [53]	44	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 [27]	48	42.2	64.6	46.2	39.1	61.6	42.0	46.0	68.2	50.2	41.6	65.1	44.8
Twin-S [6]	44	43.4	66.0	47.3	40.3	63.2	43.4	46.8	69.2	51.2	42.6	66.3	45.8
Focal-T [59]	49	44.8	67.7	49.2	41.0	64.7	44.2	47.2	69.4	51.9	42.7	66.5	45.9
CSWin-T [9]	42	46.7	68.6	51.3	42.2	65.6	45.4	49.0	70.7	53.7	43.6	67.9	46.6
SG-Former-S	41	<b>47.4</b>	<b>69.0</b>	<b>52.0</b>	<b>42.6</b>	<b>65.9</b>	<b>46.0</b>	<b>49.6</b>	<b>71.1</b>	<b>54.5</b>	<b>44.0</b>	<b>68.3</b>	<b>46.9</b>
Res101 [17]	63	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-M [53]	64	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 [27]	69	44.8	66.6	48.9	40.9	63.4	44.2	48.5	70.2	53.5	43.3	67.3	46.6
Twin-B [6]	76	45.2	67.6	49.3	41.5	64.5	44.8	48.0	69.5	52.7	43.0	66.8	46.6
Focal-S [59]	71	47.4	69.8	51.9	42.8	66.6	46.1	48.8	70.5	53.6	43.8	67.7	47.2
CSWin-S [9]	54	47.9	70.1	52.6	43.2	<b>67.1</b>	46.2	50.0	71.3	54.7	44.5	68.4	47.7
SG-Former-M	51	<b>48.2</b>	<b>70.3</b>	<b>53.1</b>	<b>43.6</b>	66.9	<b>47.0</b>	<b>50.5</b>	<b>71.5</b>	<b>54.9</b>	<b>45.4</b>	<b>68.8</b>	<b>48.2</b>
X101-64 [58]	101	42.8	63.8	47.3	38.4	60.6	41.3	44.4	64.9	48.8	39.7	61.9	42.6
PVT-L [54]	81	42.9	65.0	46.6	39.5	61.9	42.5	44.5	66.0	48.3	40.7	63.4	43.7
Twins-L [5]	111	45.9	—	—	41.6	—	—	—	—	—	—	—	—
Swin-B [28]	107	46.9	—	—	42.3	—	—	48.5	69.8	53.2	43.4	66.8	46.9
CSWin-B [9]	97	48.7	70.4	53.9	43.9	67.8	47.3	50.8	72.1	55.8	44.9	69.1	48.3
SG-Former-B	95	<b>49.2</b>	<b>70.6</b>	<b>54.3</b>	<b>68.2</b>	<b>68.1</b>	<b>47.7</b>	<b>51.3</b>	<b>72.4</b>	<b>56.0</b>	<b>45.2</b>	<b>69.6</b>	<b>48.8</b>

Table 4. Object detection and instance segmentation with Mask R-CNN on COCO. Only 3× schedule has the multi-scale training. All backbones are pretrained on ImageNet-1K.

SG-Former backbone with the CNN backbone, the SG-Former backbone is the first Transformer backbone that outperforms the previous state-of-the-art CNN-based backbone EfficientNet, which has much larger input resolutions. Although EfficientNet comes from neural network architecture search, our manually designed SG-Former-S/M/B still makes +0.3/+0.5/+0.6 Top-1 accuracy improvements.

Compared with baseline Transformers DieT, our small-size model even significantly outperforms the base size model by +1.4 with only a quarter of parameters (86.6M→24.2M) and computation cost(17.6G→4.8G). SG-Former-S/M/B achieve +2.0/+1.0/+1.2 improvements over Swin-T/S/B. Compared with state-of-the-art Transformer CSWin, SG-Former-S/M/B makes +0.5/+0.5/+0.4 improvements under similar parameters.

**Throughput and Performance.** As shown in the right figure, we evaluate the throughputs of different models on a single RTX2080Ti. Our SG-Former is the *fastest* under similar performance or has the *best* performance under similar inference speed.

**Classification with Higher Input Resolution.** Following the same settings of [9, 27], we finetune our model with input resolution of 384×384. With pretrained model on ImageNet-1K with 224×224 input resolution, we finetune it for 30 epochs. We take AdamW as optimizer with weight decay of 1e-8. The learning rate is set to 5e-6 and the

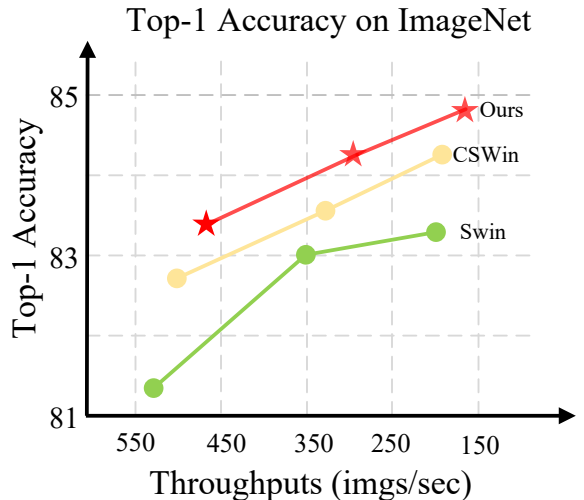


Figure 6. Top-1 accuracy on ImageNet and throughputs of recent Transformer models.

stochastic drop path is set to 0.4.

As shown in Table 3, SG-Former significantly outperforms its counterparts under similar parameters. Specifically, our base size model outperforms Swin-B by 1.6. SG-Former-S/M/B outperform previous state-of-the-art CSWin by +0.4, +0.3, +0.4 respectively.

## 4.2. Object Detection and Instance Segmentation

We evaluate the SG-Former backbone on object detection and instance segmentation with ImageNet pretraining as initialization and Mask R-CNN [16] as detection pipelines. All models are trained and evaluated on COCO2017 [24] with 118k images for training and 5K images for validation. We follow the same settings [27, 53, 9] to use  $1\times$  schedule with 12 epoch training and  $3\times$  with 36 epochs training. For  $1\times$  schedule, we take single scale training with the input images of the shorter side to 800 while keeping the longer side no more than 1333. For  $3\times$  schedule, we use multi-scale training by randomly resizing its shorter side to the range of [480, 800]. We compare SG-Former backbone with CNN backbones: ResNet [17], ResNeXt(X) [57] and Transformer backbone Swin[27], PVT[53], Focal[59], CSWin[9], Twin[6].

In table 4, we report the results with the MaskRCNN [16] pipeline and take bounding box mAP and mask mAP as an evaluation metric for object detection and instance segmentation, respectively. Compared with the CNN backbone, our method makes incredible improvements over the ResNeXt by +6.4 with  $1\times$  schedule and by +6.9 with  $3\times$  schedule. Compared with the Transformer backbones, SG-Former also shows superiority. SG-Former improves over Swin Transformer (the most well-known Transformer) and CSWin Transformer (the state-of-the-art one). Such supremacy comes from the ability to preserve more fine-granularity information and the global receptive field in our model.

## 4.3. Semantic Segmentation on ADE20K

We further evaluate SG-Former backbone on semantic segmentation with ADE20K dataset. We take semantic FPN, UperNet as basic framework. ADE20K is a widely-used semantic segmentation dataset with 150 semantic categories and contains 20K images for training and 2K images for validation. Following the same settings [9], we take Semantic FPN and Upernet as the framework and implement them by mmsegmentation [7]. For SemanticFPN, we take AdamW with weight decay of 0.0001 as the optimizer and the learning rate to be also 0.0001 for 80K iterations. We take mIoU as evaluation metric. For UperNet, we take AdamW with weight decay of 0.01 as the optimizer for 160K iterations with the batch size of 16. The learning rate is  $6 \times 10e-5$  with 1500 iteration warmup at the beginning of training and linear learning rate decay. The augmentations include random flipping, random scaling and random photo-metric distortion. The input size is  $512 \times 512$  in training. For SemanticFPN, we take AdamW with weight decay of 0.0001 as the optimizer and the learning rate to be also 0.0001 for 80K iterations. We take mIoU as the evaluation metric, and we report single-scale test results of Semantic FPN, both single-scale and multi-scale (MS) testing results

of UperNet.

The results are reported in Table 5, SG-Former outperforms all competitors under similar parameters. Specifically, when using Semantic FPN, SG-Former achieves 50.6 mIoU and outperforms Swin Transformer by 4.6 mIoU with 10% model size down. When the framework is UperNet, our SG-Former achieves 52.7 mIoU and outperforms Swin by 3 mIoU.

## 4.4. Ablation Study

**Single Scale and Hybrid Scale Self-Attention.** Our hybrid-scale Transformer blocks have two functions: 1) capturing and identifying objects of different sizes with hybrid-granularity information. 2) Proving hybrid-scale significance to self-guided attention for aggregating background tokens and preserving more fine-granularity information at salient regions.

We conduct ablation studies on the function of capturing and identifying objects of different sizes with hybrid-granularity information in Table 6. Single Scale (Local) indicates there is no token aggregation but just window attention and there is only one scale, while Single Scale (Global) indicates the aggregation rate is large enough  $sM \times sM = H \times W$ , the self-attention performs globally, and there is also only one scale. Our hybrid-scale self-attention outperforms local and global attention by **1.0** and **0.6** respectively

**Significance Guidance** To evaluate the influence of the significance map, in each module, we keep the first hybrid-scale Transformer block the same but change the source of important score maps in the second self-guided Transformer block. We manually define the significance map to be equal at all positions across the whole significance map as a comparison. The hybrid-scale self-attention layer provides significance maps with hybrid-scale information, yet we use significance from the global scale or the local scale only.

As shown in Table 7, the hybrid-scale significance outperforms the manually defined baseline which regards all region importance equally by **+0.9**. The local-scale significance which contains local fine-grained region importance information has better performance than that of the global scale which contains global coarse-grained region importance information, while hybrid-scale significance maps show much better performance over all competitors.

## 5. Conclusion

In this paper, we propose a novel Transformer backbone, termed as Self-guided Transformer (SG-Former). The key motivation is to aggregate tokens according to the region or token significance with hybrid-scale guidance, which explicitly accounts for long-range dependency crosses the whole feature maps and attends more foreground tokens for fine-granularity information while merging more background tokens for computation efficiency.



Backbone	Semantic FPN 80k			Upernet 160k			
	#Param.	FLOPs	mIoU	#Param.	FLOPs	mIoU	MS mIoU
Res50 [18]	28.5	183	36.7	—	—	—	—
PVT-S [54]	28.2	161	39.8	—	—	—	—
Twins-S [5]	28.3	144	43.2	54.4	901	46.2	47.1
Swin-T [28]	31.9	182	41.5	59.9	945	44.5	45.8
Focal-T [59]	-	-	-	62.0	998	45.8	47.0
CSWin-T	26.1	202	48.2	59.9	959	49.3	50.7
SG-Former-S	25.4	205	<b>49.0</b>	52.5	989	<b>49.9</b>	<b>51.5</b>
Res101 [18]	47.5	260	38.8	86.0	1029	—	44.9
PVT-M [54]	48.0	219	41.6	—	—	—	—
Twins-B [5]	60.4	261	45.3	88.5	1020	47.7	48.9
Swin-S [28]	53.2	274	45.2	81.3	1038	47.6	49.5
Focal-S [59]	-	-	-	85.0	1130	48.0	50.0
CSWin-S	38.5	271	49.2	64.6	1027	50.4	51.5
SG-Former-M	38.2	273	<b>50.1</b>	68.3	1114	<b>51.2</b>	<b>52.1</b>
X101-64 [58]	86.4	—	40.2	—	—	—	—
PVT-L [54]	65.1	283	42.1	—	—	—	—
Twins-L [5]	103.7	404	46.7	133.0	1164	48.8	50.2
Swin-B [28]	91.2	422	46.0	121.0	1188	48.1	49.7
Focal-B [59]	-	-	-	126.3	1354	49.0	50.5
CSWin-B	81.2	464	49.9	109.2	1222	51.1	52.2
SG-Former-B	81.0	475	<b>50.6</b>	109.3	1304	<b>52.0</b>	<b>52.7</b>

Table 5. Performance comparison of different backbones on the ADE20K segmentation task. Two different frameworks, semantic FPN and Upernet, are used.

Model	Param. (M)	Flops (G)	Top-1 (%)
Single Scale (Local)	25.6	4.9	81.6
Single Scale (Global)	28.4	4.6	82.0
Hybrid-Scale (Local-Global)	26.5	4.8	82.6

Table 6. Comparison of different scales in hybrid-scale self-attention.

Guidance	Param. (M)	Flops (G)	Top-1 (%)
Manual defined	21.6	4.6	82.3
Local	20.9	4.9	82.9
Global	24.6	4.6	82.4
Hybrid-Scale	22.5	4.8	83.2

Table 7. Comparison of different significance maps. Manual defined indicate the significance equals at all position cross the whole significance map.

Experiments on classification, detection, and segmentation demonstrate that, the proposed SG-Former yields gratifying results: our base size model achieves a 84.7% Top-1 accuracy on ImageNet-1K, 51.2mAP bbAP on CoCo, 52.7mIoU on ADE20K, surpassing the Swin Transformer by +1.3% /

+2.7 mAP/ +3 mIoU, yet with a lower computation cost.

## Acknowledgement

This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2021-023).

## References

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020. **1**
- [2] Chun-Fu Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. *arXiv preprint arXiv:2103.14899*, 2021. **6**
- [3] Hanqing Chen, Yunhe Wang, Jianyuan Guo, and Dacheng Tao. Vanillanet: the power of minimalism in deep learning. *arXiv preprint arXiv: 2305.12972*, 2023. **3**
- [4] Tianlong Chen, Zhenyu Zhang, Ajay Kumar Jaiswal, Shiwei Liu, and Zhangyang Wang. Sparse moe as the new dropout: Scaling dense and self-slimmable transformers. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. **2**

- [5] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting spatial attention design in vision transformers. *arXiv preprint arXiv:2104.13840*, 2021. 7, 9
- [6] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. In *NeurIPS*, 2021. 6, 7, 8
- [7] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 8
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [9] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *CVPR*, pages 12124–12134, 2022. 1, 2, 3, 5, 6, 7, 8
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 5, 6
- [12] Alaaeldin El-Nouby, Natalia Neverova, Ivan Laptev, and Hervé Jégou. Training vision transformers for image retrieval. *arXiv preprint arXiv:2102.05644*, 2021. 2
- [13] Gongfan Fang, Xinyin Ma, Mingli Song, Michael Bi Mi, and Xinchao Wang. Depgraph: Towards any structural pruning. In *CVPR*, pages 16091–16101, 2023. 3
- [14] Gongfan Fang, Xinyin Ma, and Xinchao Wang. Structural pruning for diffusion models. *arXiv preprint arXiv:2305.10924*, 2023. 3
- [15] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. *arXiv preprint arXiv:2103.00112*, 2021. 6
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 8
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 3, 6, 7, 8
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 9
- [19] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. *arXiv preprint arXiv:2102.04378*, 2021. 2
- [20] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531, 2015. 3
- [21] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019. 3
- [22] Yanxi Li and Chang Xu. Trade-off between robustness and accuracy of vision transformers. In *CVPR*, pages 7558–7568, 2023. 2
- [23] Yawei Li, Kai Zhang, Jiezhong Cao, Radu Timofte, and Luc Van Gool. Localvit: Bringing locality to vision transformers. *arXiv:2104.05707*, 2021. 2
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6, 8
- [25] Songhua Liu, Kai Wang, Xingyi Yang, Jingwen Ye, and Xinchao Wang. Dataset distillation via factorization. In *Advances in neural information processing systems*, 2022. 3
- [26] Songhua Liu, Jingwen Ye, Rungeng Yu, and Xinchao Wang. Slimmable dataset condensation. In *CVPR*, 2023. 3
- [27] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *IEEE International Conference on Computer Vision*, 2021. 1, 2, 3, 5, 6, 7, 8
- [28] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021. 7, 9
- [29] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 11966–11976. IEEE, 2022. 3
- [30] Xinyin Ma, Gongfan Fang, and Xinchao Wang. LLM-Pruner: On the Structural Pruning of Large Language Models. *arXiv preprint arXiv:2305.11627*, 2023. 3
- [31] Juhong Min, Yucheng Zhao, Chong Luo, and Minsu Cho. Peripheral vision transformer. In *Neurips*, 2022. 6
- [32] Arild Nøkland and Lars Hiller Eidnes. Training neural networks with local error signals. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 4839–4850. PMLR, 2019. 3
- [33] Zizheng Pan, Jianfei Cai, and Bohan Zhuang. Fast vision transformers with hilo attention. In *Neurips*, 2022. 6
- [34] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, pages 10428–10436, 2020. 6
- [35] Jingjing Ren, Qingqing Zheng, Yuanyuan Zhao, Xuemiao Xu, and Chen Li. Diformer: Discrete latent transformer for video inpainting. In *CVPR*, pages 3511–3520, June 2022. 2
- [36] Sucheng Ren, Zhengqi Gao, Tianyu Hua, Zihui Xue, Yonglong Tian, Shengfeng He, and Hang Zhao. Co-advise: Cross inductive bias distillation. In *CVPR*, pages 16773–16782, June 2022. 2

- [37] Sucheng Ren, Wenxi Liu, Yongtuo Liu, Haoxin Chen, Guoqiang Han, and Shengfeng He. Reciprocal transformations for unsupervised video object segmentation. In *CVPR*, pages 15455–15464, June 2021. [2](#)
- [38] Sucheng Ren, Huiyu Wang, Zhengqi Gao, Shengfeng He, Alan Yuille, Yuyin Zhou, and Cihang Xie. A simple data mixing prior for improving self-supervised learning. In *CVPR*, 2022. [2](#)
- [39] Sucheng Ren, Fangyun Wei, Zheng Zhang, and Han Hu. Tinyvim: An empirical study of distilling mim pre-trained models. In *CVPR*, pages 3687–3697, June 2023. [2](#)
- [40] Sucheng Ren, Daquan Zhou, Shengfeng He, Jiashi Feng, and Xinchao Wang. Shunted self-attention via multi-scale token aggregation. In *CVPR*, pages 10853–10862, June 2022. [6](#)
- [41] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 8583–8595, 2021. [2](#)
- [42] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. *arXiv preprint arXiv:2105.05633*, 2021. [2](#)
- [43] Zhiqing Sun, Shengcao Cao, Yiming Yang, and Kris M Kitani. Rethinking transformer-based set prediction for object detection. In *ICCV*, pages 3611–3620, 2021. [2](#)
- [44] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pages 6105–6114. PMLR, 2019. [6](#)
- [45] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357. PMLR, 2021. [2](#), [5](#), [6](#)
- [46] Hugo Touvron, Matthieu Cord, and Herve Jegou. Deit iii: Revenge of the vit. *arXiv preprint arXiv:2204.07118*, 2022. [2](#)
- [47] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. *arXiv preprint arXiv:2103.17239*, 2021. [2](#)
- [48] Ashish Vaswani, Prajit Ramachandran, Aravind Srinivas, Niki Parmar, Blake Hechtman, and Jonathon Shlens. Scaling local self-attention for parameter efficient visual backbones. *arXiv preprint arXiv:2103.12731*, 2021. [1](#)
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [1](#), [2](#)
- [50] Cong Wang, Hongmin Xu, Xiong Zhang, Li Wang, Zhitong Zheng, and Haifeng Liu. Convolutional embedding makes hierarchical vision transformer stronger. In *ECCV*, 2022. [6](#)
- [51] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A. Efros. Dataset distillation, 2020. [3](#)
- [52] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvtv2: Improved baselines with pyramid vision transformer. *arXiv preprint arXiv:2106.13797*, 2021. [1](#), [3](#)
- [53] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *IEEE International Conference on Computer Vision*, 2021. [1](#), [3](#), [6](#), [7](#), [8](#)
- [54] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. *arXiv preprint arXiv:2102.12122*, 2021. [7](#), [9](#)
- [55] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. *arXiv preprint arXiv:2011.14503*, 2020. [2](#)
- [56] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *arXiv preprint arXiv:2105.15203*, 2021. [1](#)
- [57] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. [8](#)
- [58] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks, 2017. [7](#), [9](#)
- [59] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers, 2021. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#), [9](#)
- [60] Xingyi Yang, Jingwen Ye, and Xinchao Wang. Factorizing knowledge in neural networks. In *European Conference on Computer Vision*, 2022. [3](#)
- [61] Xingyi Yang, Daquan Zhou, Songhua Liu, Jingwen Ye, and Xinchao Wang. Deep model reassembly. In *Advances in neural information processing systems*, 2022. [3](#)
- [62] Tao Yu, Zhihe Lu, Xin Jin, Zhibo Chen, and Xinchao Wang. Task residual for tuning vision-language models. In *CVPR*, 2023. [3](#)
- [63] Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, and Shuicheng Yan. Metaformer is actually what you need for vision. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10809–10819, 2022. [3](#)
- [64] Weihao Yu, Chenyang Si, Pan Zhou, Mi Luo, Yichen Zhou, Jiashi Feng, Shuicheng Yan, and Xinchao Wang. Metaformer baselines for vision. *arXiv preprint arXiv:2210.13452*, 2022. [3](#)
- [65] Kun Yuan, Shaopeng Guo, Ziwei Liu, Aojun Zhou, Fengwei Yu, and Wei Wu. Incorporating convolution designs into visual transformers. *arXiv preprint arXiv:2103.11816*, 2021. [2](#)
- [66] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zihang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng

- Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. *arXiv preprint arXiv:2101.11986*, 2021. [2](#)
- [67] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. [3](#)
- [68] Pengchuan Zhang, Xiyang Dai, Jianwei Yang, Bin Xiao, Lu Yuan, Lei Zhang, and Jianfeng Gao. Multi-scale vision longformer: A new vision transformer for high-resolution image encoding. *arXiv preprint arXiv:2103.15358*, 2021. [2](#)
- [69] Qinglong Zhang and Yu-Bin Yang. Rest v2: simpler, faster and stronger. *Neurips*, 2022. [6](#)
- [70] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. *arXiv preprint arXiv:2012.15840*, 2020. [1](#)
- [71] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *CVPR*, pages 6881–6890, 2021. [2](#)
- [72] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017. [6](#)
- [73] Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*, 2021. [6](#)
- [74] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [1](#)