Focal Attention for Long-Range Interactions in Vision Transformers

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recently, Vision Transformers and its variants have shown promise on various vision tasks. The ability to capture short- and long-range visual dependencies through self-attention is arguably the secret sauce for its success but by nature brings difficulty due to quadratic computational overhead, especially for the high-resolution tasks (*e.g.*, object detection). A variety of works have attempted to address this by applying either coarse-grained global attention or fine-grained local attention. However, both paradigms discount the original power of transformer layers. In this paper, we present a focal self-attention mechanism, which simultaneously models the local and global interactions in a transformer layer and helps to capture both short and long-range visual dependencies efficiently *and* effectively. With focal self-attention, our built vision transformer models, called Focal Transformers achieve superior performance over the state-of-the-art methods in various standard benchmark settings on image classification and high-resolution object detection tasks. Specifically, a focal transformer model with a moderate size of 51.1M achieves 83.5 Top-1 accuracy on ImageNet classification at 224×224 resolution. Using our Focal Transformers as the backbones, we demonstrate consistent and substantial improvements over the current state-of-the-art Swin Transformers [23] on 6 different detection methods trained at both 1x and 3x schedule.

1 Introduction

Nowadays, Transformer [35] has become a prevalent model architecture in natural language processing (NLP) [12, 4]. In the light of its success in NLP, there is an increasing effort on adapting it to computer vision [26, 29]. Since its promise was first demonstrated in Vision Transformer (ViT) [13], we have witnessed a flourish of full-Transformer models for image classification [33, 38, 40, 23, 46, 34], object detection [6, 51, 49, 10] and semantic segmentation [36, 39]. Beyond these static image tasks, it has also been applied on various temporal tasks, such as action recognition [20, 48, 7], object tracking [8, 37], scene flow estimation [19], etc.

In Transformers, self-attention is the key component making it unique from the widely used convolutional neural networks (CNNs) [18]. It enables the global content-dependent interactions among different image regions for modeling both short- and long-range dependencies at each Transformer layer. Through the visualization of self-attention in DeiT-Tiny model[1], we indeed observe on the left side of Fig. 1 that it learns to attend local surroundings (like CNN) and the global contexts at the same time. Nevertheless, when it comes to high-resolution images for dense predictions such as object detection or segmentation, a global and fine-grained self-attention becomes non-trivial due to the quadratic computational cost with respect to the dimension of feature map. Recent works alternatively exploited either a coarse-grained global self-attention [38, 40] or a fine-grained local

[1] Pre-trained checkpoint downloaded from `https://github.com/facebookresearch/deit`.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

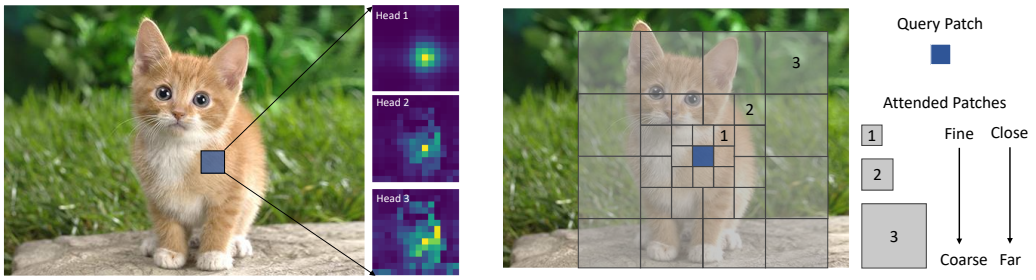

Figure 1: Left: Visualization of the attention maps of the three heads at the given query patch (blue) in the first layer of the DeiT-Tiny model. Right: An illustrative depiction of focal self-attention mechanism. Three granularity levels are used to compose the attention region for the blue query.

self-attention [23, 46, 34] to reduce the computational burden. However, both of them hurt the power of original self-attention after sacrificing one of its two merits, *i.e.*, the ability to simultaneously model short- and long-range visual dependencies.

In this paper, we present a new self-attention mechanism to capture both local and global interactions in Transformer layers. As discussed above, it is difficult to model the fine-grained global interactions for high-resolution inputs. We instead propose to perform self-attention at fine-grain locally whereas coarse-grain globally. As depicted in the right side of Fig. 1, for a query token in the feature map, it attends its closest surrounding at the finest grain as itself. When it goes to farther surroundings, we use summarized tokens to represent the coarser regions. Depending on how many levels we want to extend, we can further cover the whole feature map using more coarser grain at even farther distance. We call this mechanism focal self-attention considering each token attends other tokens in a focal manner. This focal self-attention has the ability to capture both short- and long-range dependencies that approximates the original self-attention while introducing much less computational overhead after coarsening the tokens in the feature map.

Based on the proposed focal self-attention, we develop focal vision Transformers for image classification and object detection. It shares some common properties as previous works for dense prediction tasks: 1) We exploit a multi-stage architecture to maintain a reasonable computational cost for high-resolution images, following [38, 40, 23, 46]; 2) Instead of performing focal self-attention for each token, we split the feature map into multiple windows in which tokens share the same surroundings, similar to the strategies used in [34, 46, 23]. As a result, we focus our study on the effectiveness of our proposed focal self-attention. Extensive experiments show that our Focal Transformers with similar sizes and complexities outperform the current state-of-the-art Transformer models consistently across various settings, particularly on object detection. These results demonstrate the focal self-attention stand-alone is an effective strategy for modeling the local-global interactions in vision Transformers.

2 Related work

Vision Transformers. The Vision Transformer (ViT) [13] applies a standard Transformer, originally developed for natural language processing (NLP), for image encoding by treating an image as a word sequence, i.e., splitting an image into patches (words) and using the linear embeddings of these patches as an input sequence. ViT has shown to outperform convolution neural network (CNN) models such as the ResNet [16], achieving state-of-the-art performance on multiple image classification benchmarks, where training data is sufficient. ViT has been improved from different perspectives, such as data-efficient training [33], improved patch embedding/encoding [9, 44, 14], introducing convolutional projects into transformers [40, 43], multi-stage ViT architecture and efficient attention mechanisms for high-resolution vision tasks [38, 40, 23, 46]. Our Focal ViT follows the multi-stage ViT architecture and introduces a new efficient attention mechanism – focal self-attention.

Local fine-grain and global coarse-grain attention. When there are large number of tokens in Transformer models, such as long document processing in NLP and high-resolution vision tasks in CV, efficient attention mechanisms are required to overcome the quadratic computational and memory increase in the vanilla self-attention mechanisms. Local fine-grain attention, i.e., attending

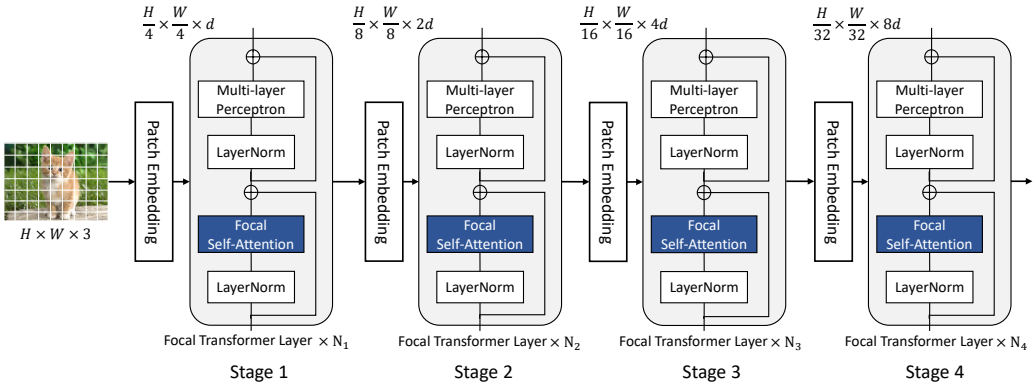

Figure 2: Model architecture for our focal vision Transformers. As highlighted by blue box, the main innovation part is the proposed focal self-attention module in each Transformer layer.

to neighboring tokens within a constant window size, is the most popular efficient attention in both NLP (see, e.g., [2, 45, 1]) and CV (see, e.g., [40, 23, 46]) scenarios. Global coarse-grain attention, i.e., attending to a small number of downsampled/summarized tokens, is another effective efficient attention in preserving the global communication among tokens in the vanilla self-attention; see, e.g., [28, 25] in NLP and [38, 40, 14] in CV. We argue that both types of attentions are important and the vanilla ViT model indeed had learned both of them, as shown in Fig. 1 left. Our proposed focal self-attention captures both types of attentions and approximates the vanilla full attention effectively. Finally, we refer to [32, 31, 46] for a comprehensive survey and benchmarks of various efficient attention mechanims in NLP and CV applications.

3 Method

3.1 Model Architecture

To accommodate the high-resolution vision tasks, our model architecture takes similar multi-scale design used in [38, 46, 23] which allows us to obtain high-resolution feature maps at earlier stages. As shown in Fig. 2, an image $I \in \mathcal{R}^{H \times W \times 3}$ is first partitioned into patches of size 4×4, resulting in $\frac{H}{4} \times \frac{W}{4}$ visual tokens with dimension $4 \times 4 \times 3$. Then, we use a patch embedding layer which consists of a convolutional layer with filter size and stride both equal to 4, to project these patches into hidden features with dimension d. Given this spatial feature map, we then pass it to four stages of focal Transformer blocks. At each stage $i \in \{1, 2, 3, 4\}$, the focal Transformer block consists of N_i focal Transformer layers. After each stage, we use another patch embedding layer to reduce the spatial size of feature map by factor 2, while the feature dimension is increased by 2. For image classification task, we take the average of the output from last stage and send it to a classification layer. For object detection, the feature maps from last 3 or all 4 stages are fed to the detector head depending on the particular detection method we use. The model capacity can be customized by varying the input feature dimension d and the number of focal Transformer layers at each stage $\{N_1, N_2, N_3, N_4\}$.

Standard self-attention can model both short- and long-range interactions at fine-grain, but it suffers from high computational cost when it performs the attention on high-resolution feature maps. Take stage 1 in Fig. 2 as the example. For the feature map of size $\frac{H}{4} \times \frac{W}{4} \times d$, the complexity of self-attention is $\mathcal{O}((\frac{H}{4} \times \frac{W}{4})^2 d)$, resulting in an explosion of time and memory cost considering $\min(H, W) = 800$ typically for object detection. In the next, we will explain how we address this by using focal self-attention.

3.2 Focal Self-Attention

In this paper, we propose focal self-attention to make Transformer layers scalable to high-resolution inputs. Instead of attending all tokens at fine-grain, we propose to attend the fine-grain tokens only locally but the summarized ones globally. As such, it can cover as much region as standard

self-attention but with much less cost. In Fig. 3, we show the area of receptive field for standard self-attention and our focal self-attention. As we can see, for a query position, when we use gradually coarser-grain for its far surroundings, it can have significantly larger receptive field at the cost of attending the same number of visual tokens.

Theoretically, our focal mechanism enables global self-attention with much less time and memory cost because it attends much less number of surrounding (summarized) tokens. In practice, however, extracting the surrounding tokens for each query position suffers from high time and memory cost since we need to duplicate each token for all queries that can get access to it. This practical issue has been exhaustively noted by a number of previous works [34, 46, 23] and the common solution is to partition the input feature map into windows. Inspired by them, we resort to perform focal attention at the window level. Given a feature map of $x \in \mathcal{R}^{M \times N \times d}$ with spatial size $M \times N$, we first partition it into a grid of windows with size $s_p \times s_p$. Then, we find the surroundings for each window rather than individual tokens. In the following, we elaborate the window-wise focal self-attention.

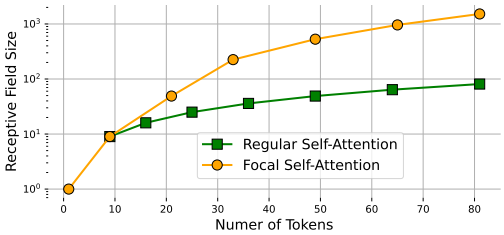

Figure 3: The receptive field size with the increase of tokens for standard self-attention and our focal self-attention. For focal self-attention, we increase the window granularity by factor 2 gradually but no more than 8. Note that the y-axis is logarithmic.

3.2.1 Window-wise Attention

An illustration of the proposed window-wise focal self-attention is shown in Fig. 4. Before we explain the details, we first define three terms for clarity:

- **Focal levels** L – the number of granularity levels we extract the tokens for our focal attention. In, Fig. 1, we show 3 focal levels in total for example.
- **Focal window size** s_w^l – the size of sub-window on which we get the summarized tokens at level $l \in \{1, ..., L\}$, which are 1, 2 and 4 for the three levels in Fig. 1.
- **Focal region size** s_r^l – the number of sub-windows horizontally and vertically in attended region at level l, and they are 3, 4 and 4 from level 1 to 3 in Fig. 1.

With the above three terms $\{L, s_w, s_r\}$, we can specify our focal self-attention module. Below we explain its two main steps:

Sub-window pooling. Given the input feature map $x \in \mathcal{R}^{M \times N \times d}$, we perform sub-window pooling for all L levels. For the focal level l, we first split the input feature map x into a grid of sub-windows with size $s_w^l \times s_w^l$. Then we use a simple linear layer f_p^l to pool the sub-windows spatially by:

$$x^l = f_p^l(\hat{x}) \in \mathcal{R}^{\frac{M}{s_w^l} \times \frac{N}{s_w^l} \times d}, \quad \hat{x} = Reshape(x) \in \mathcal{R}^{(\frac{M}{s_w^l} \times \frac{N}{s_w^l} \times d) \times (s_w^l \times s_w^l)}, \tag{1}$$

The pooled feature maps $\{x^l\}_1^L$ at different levels l provide rich information at both fine-grain and coarse-grain. Since we set $s_w^l = 1$ for the first focal level which has the same granularity as the input feature map, there is no need to perform any sub-window pooling. Considering the focal window size is usually very small (7 maximally in our settings), the number of extra parameters introduced by these sub-window pooling are fairly neglectable.

Attention computation. Once we obtain the pooled feature maps $\{x^l\}_1^L$ at all L levels, we compute the query at the first level and key and value for all levels using three linear projection layers f_q, f_k and f_v:

$$Q = f_q(x^1), \quad K = \{K^l\}_1^L = f_k(\{x^1, ..., x^L\}), \quad V = \{V^l\}_1^L = f_v(\{x^1, ..., x^L\}) \tag{2}$$

To perform focal self-attention, we need to first extract the surrounding tokens for each query token in the feature map. As we mentioned earlier, tokens inside a window partition $s_p \times s_p$ share the same set of surroundings. For the queries inside the i-th window $Q_i \in \mathcal{R}^{s_p \times s_p \times d}$, we extract the $s_r^l \times s_r^l$ keys and values from K^l and V^l around the window where the query lies in, and then gather the keys and values from all L to obtain $K_i = \{K_i^1, ..., K_i^L\} \in \mathcal{R}^{s \times d}$ and $V_i = \{V_i^1, ..., V_i^L\} \in \mathcal{R}^{s \times d}$, where s is the sum of focal region from all levels, *i.e.*,, $s = \sum_{l=1}^{L} (s_r^l)^2$. Note that a strict version

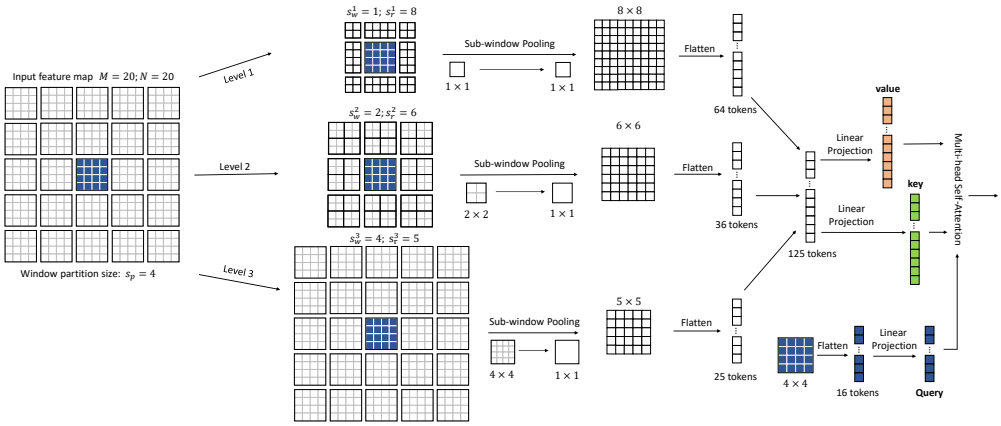

Figure 4: An illustration of our focal self-attention at window level. We suppose an input feature map of 20×20 and 4×4 window size. Each of the finest square cell represents a visual token either from the original feature map or the squeezed ones. Suppose we have an input feature map of size 20×20. We first partition it into 5×5 windows of size 4×4. Take the 4×4 blue window in the middle as the query, we extract its surroundings tokens at multiple granularity levels as its keys and values. For the first level, we extract the 8×8 tokens which are closest to the blue window at the finest grain. Then at the second level, we expand the attention region and pool the surrounding 2×2 sub-windows, which results in 6×6 pooled tokens. At the final level, we attend even larger region covering the whole feature map and pool 4×4 sub-windows. Finally, these three levels of tokens are concatenated to compute the keys and values for the $4 \times 4 = 16$ tokens (queries) in the blue window.

of focal self-attention following Fig. 1 requires to exclude the overlapped regions across different levels. In our model, we intentionally keep them in order to capture the pyramid information for the overlapped regions. Finally, we follow [23] to include a relative position bias and compute the focal self-attention for Q_i by:

$$Attention(Q_i, K_i, V_i) = Softmax(\frac{Q_i K_i^T}{\sqrt{d}} + B)V_i, \tag{3}$$

where $B = \{B^l\}_1^L$ is the learnable relative position bias. It consists of L subsets for L focal levels. Similar to [23], for the first level, we parameterize it to $B^1 \in \mathcal{R}^{(2s_p-1)\times(2s_p-1)}$, considering the horizontal and vertical position range are both in $[-s_p + 1, s_p - 1]$. For the other focal levels, considering they have different granularity to the queries, we treat all the queries inside a window equally and use $B^l \in \mathcal{R}^{s_r^l \times s_r^l}$ to represent the relative position bias between the query window and each of $s_r^l \times s_r^l$ pooled tokens. Since the focal self-attention for each window is independent of others, we can compute Eq. (3) in parallel. Once we complete it for the whole input feature map, we send it to the MLP block for proceeding computation as usual.

3.2.2 Complexity Analysis

We analyze the computational complexity for the two main steps discussed above. For the input feature map $x \in \mathcal{R}^{M\times N\times d}$, we have $\frac{M}{s_w^l} \times \frac{N}{s_w^l}$ sub-windows at focal level l. For each sub-window, the pooling operation in Eq.1 has the complexity of $\mathcal{O}((s_w^l)^2 d)$. Aggregating all sub-windows brings us $\mathcal{O}((MN)d)$. Then for all focal levels, we have the complexity of $\mathcal{O}(L(MN)d)$ in total, which is independent of the sub-window size at each focal level. Regarding the attention computation in Eq. 3, the computational cost for a query window $s_p \times s_p$ is $\mathcal{O}((s_p)^2 \sum_l (s_r^l)^2 d)$, and $\mathcal{O}(\sum_l (s_r^l)^2 (MN)d)$ for the whole input feature map. To sum up, the overall computational cost for our focal self-attention becomes $\mathcal{O}((L + \sum_l (s_r^l)^2)(MN)d)$. In an extreme case, one can set $s_r^L = 2\max(M, N)/s_w^L$ to ensure global receptive field for all queries (including both corner and middle queries) in this layer.

3.3 Model Configuration

We customize three different model capacities for our focal Transformers. Here, we simply follow the design strategy suggested by previous works [38, 40, 23], though we believe there should be a better

	Output Size	Layer Name	Focal-Tiny	Focal-Small	Focal-Base
stage 1	56×56	Patch Embedding	$p_1 = 4; c_1 = 96$	$p_1 = 4; c_1 = 96$	$p_1 = 4; c_1 = 128$
	56×56	Transformer Block	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 7\} \end{bmatrix} \times 2$	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 7\} \end{bmatrix} \times 2$	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 7\} \end{bmatrix} \times 2$
stage 2	28×28	Patch Embedding	$p_2 = 2; c_2 = 192$	$p_2 = 2; c_2 = 192$	$p_2 = 2; c_2 = 256$
	28×28	Transformer Block	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 5\} \end{bmatrix} \times 2$	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 5\} \end{bmatrix} \times 2$	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 5\} \end{bmatrix} \times 2$
stage 3	14×14	Patch Embedding	$p_3 = 2; c_3 = 384$	$p_3 = 2; c_3 = 384$	$p_3 = 2; c_3 = 512$
	14×14	Transformer Block	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 3\} \end{bmatrix} \times 6$	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 3\} \end{bmatrix} \times 18$	$\begin{bmatrix} s^0_{w,r} = \{1, 13\} \\ s^1_{w,r} = \{7, 3\} \end{bmatrix} \times 18$
stage 4	7×7	Patch Embedding	$p_4 = 2; c_4 = 768$	$p_4 = 2; c_4 = 768$	$p_4 = 2; c_4 = 1024$
	7×7	Transformer Block	$\begin{bmatrix} s^0_{w,r} = \{1, 7\} \\ s^1_{w,r} = \{7, 1\} \end{bmatrix} \times 2$	$\begin{bmatrix} s^0_{w,r} = \{1, 7\} \\ s^1_{w,r} = \{7, 1\} \end{bmatrix} \times 2$	$\begin{bmatrix} s^0_{w,r} = \{1, 7\} \\ s^1_{w,r} = \{7, 1\} \end{bmatrix} \times 2$

Table 1: Model configurations for our focal Transformers. We introduce three configurations Focal-Tiny, Focal-Small and Focal-Base with different model capacities.

configuration specifically for our focal Transformers. Specifically, we use similar design to the tiny, small and base models in Swin Transformer [23], as shown in Table 1. Our models take 224×224 images as inputs and the window partition size is also set to 7 to make our models comparable to the Swin Transformers. For the focal self-attention layer, we introduce two levels, one for fine-grain local attention and one for coarse-grain global attention. Expect for the last stage, the focal region size is consistently set to 13 for the window partition size of 7, which means that we expand 3 tokens for each window partition. For the last stage, since the whole feature map is 7×7, the focal region size at level 0 is set to 7, which is sufficient to cover the entire feature map. For the coarse-grain global attention, we set its focal window size same to the window partition size 7, but gradually decrease the focal region size to get $\{7, 5, 3, 1\}$ for the four stages. For the patch embedding layer, the spatial reduction ratio p_i for four stages are all $\{4, 2, 2, 2\}$, while Focal-Base has a higher hidden dimension compared with Focal-Tiny and Focal-Small.

4 Experiments

4.1 Image Classification

We compare different methods on ImageNet-1K [11]. For fair comparison, we follow the training recipes in [33, 38]. All models are trained for 300 epochs with a batch size 1024. The initial learning rate is set to 10^{-3} with 20 epochs of linear warm-up starting from 10^{-5}. For optimization, we use AdamW [24] as the optimizer with a cosine learning rate scheduler. The weight decay is set to 0.05 and the maximal gradient norm is clipped to 5.0. We use the same set of data augmentation and regularization strategies used in [33] after excluding random erasing [50], repeated augmentation [3, 17] and exponential moving average (EMA) [27]. The stochastic depth drop rates are set to 0.2, 0.2 and 0.3 for our tiny, small and base models, respectively. During training, we crop images randomly to 224×224, while a center crop is used during evaluation on the validation set.

In Table 2, we summarize the results for baseline models and the current state-of-the-art models on image classification task. We can find our Focal Transformers consistently outperforms other methods with similar model size (#Params.) and computational complexity (GFLOPs). Specifically, Focal-Tiny improves over the Transformer baseline DeiT-Small/16 by 2.0%. Meanwhile, using the same model configure (2-2-6-2) and a few extra parameters and computations, our Focal-Tiny improves over Swin-Tiny by 0.7 points (81.2 $\rightarrow$ 81.9). When we increase the window size from 7 to 14 to match the settings in ViL-Small [46], the performance can be further improved to 82.3. For small and base models, our Focal Transformers still achieves slightly better performance than the others. Notably, our Focal-Small can even reach 83.5 which is better than all counterpart small and base models using much less parameters (51.1M). We refer the readers to our appendix for more detailed comparisons.

Model	#Params.	FLOPs	Top-1 (%)
ResNet-50 [16]	25.0	4.1	76.2
DeiT-Small/16 [33]	22.1	4.6	79.9
PVT-Small [38]	24.5	3.8	79.8
ViL-Small [46]	24.6	5.1	82.0
CvT-13 [40]	20.0	4.5	81.6
Swin-Tiny [23]	28.3	4.5	81.2
Focal-Tiny (Ours)	29.1	4.9	82.2
ResNet-101 [16]	45.0	7.9	77.4
PVT-Medium [38]	44.2	6.7	81.2
CvT-21 [40]	32.0	7.1	82.5
ViL-Medium [46]	39.7	9.1	83.3
Swin-Small [23]	49.6	8.7	83.1
Focal-Small (Ours)	51.1	9.1	**83.5**
ResNet-152 [16]	60.0	11.0	78.3
ViT-Base/16 [13]	86.6	17.6	77.9
DeiT-Base/16 [33]	86.6	17.5	81.8
PVT-Large [38]	61.4	9.8	81.7
ViL-Base [46]	55.7	13.4	83.2
Swin-Base [23]	87.8	15.4	83.4
Focal-Base (Ours)	89.8	16.0	**83.5**

Table 2: Comparison of image classification on ImageNet-1K for different models. Except for ViT-Base/16, all other models are trained and evaluated on 224×224 resolution.

Backbone	RetinaNet	Mask R-CNN	
	AP^b	AP^b	AP^m
ResNet-50 [16]	36.3	38.0	34.4
PVT-Small	40.4	40.4	37.8
ViL-Small [46]	41.6	41.8	38.5
Swin-Tiny [23]	42.0	43.7	39.8
Focal-Tiny (Ours)	**43.7** (+1.7)	**44.8** (+1.1)	**41.0** (+1.3)
ResNet-101 [16]	38.5	40.4	36.4
ResNeXt101-32x4d [41]	39.9	41.9	37.5
PVT-Medium [38]	41.9	42.0	39.0
ViL-Medium [46]	42.9	43.4	39.7
Swin-Small [23]	45.0	46.5	42.1
Focal-Small (Ours)	**45.6** (+0.6)	**47.4** (+0.9)	**42.8** (+0.7)
ResNeXt101-64x4d [41]	41.0	42.8	38.4
PVT-Large [38]	42.6	42.9	39.5
ViL-Base [46]	44.3	45.1	41.0
Swin-Base [23]	45.0	46.9	42.3
Focal-Base (Ours)	**46.3** (+1.3)	**47.8** (+0.9)	**43.2** (+0.9)

Table 3: Comparisons with CNN and Transformer baselines and state-of-the-art 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 1x schedule. More detailed comparisons with 3x schedule are in Table 4.

4.2 Object Detection and Segmentation

We benchmark our models on object detection with COCO 2017 [22]. The pretrained models are used as visual backbones and then plug into two representative pipelines, RetinaNet [21] and Mask R-CNN [15]. All models are trained on the 118k training images and results reported on 5K validation set. We follow the standard to use two training schedules, 1x schedule with 12 epochs and 3x schedule with 36 epochs. For 1x schedule, we resize image's shorter side to 800 while keeping its longer side no more than 1333. For 3x schedule, we use multi-scale training strategy by randomly resizing its shorter side to the range of $[480, 800]$. Considering this higher input resolution, we adaptively increase the focal sizes at four stages to $(15, 13, 9, 7)$, to ensures the focal attention covers more than half of the image region (first two stages) to the whole image (last two stages). With the focal size increased, the relative position biases are accordingly up-sampled to corresponding sizes using bilinear interpolation. During training, we use AdamW [24] for optimization with initial learning rate 10^{-4} and weight decay 0.05. Similarly, we use 0.2,0.2 and 0.3 stochastic depth drop rates to regularize the training for our tiny, small and base models. Since Swin Transformer does not report the numbers on RetinaNet, we train it by ourselves using their official code with the same hyperparameters to our Focal Transformers.

In Table 3, we show the performance for both CNN-based models and the current Transformer-based state-of-the-arts methods. The bbox mAP (AP^b) and mask mAP (AP^m) are reported. As we can see, our Focal Transformers outperform the CNN-based models consistently with the gap of 4.8-7.1 points. Compared with the other methods which also use multi-scale Transformer architectures, we still observe substantial gain across all settings and metrics. Particularly, our Focal Transformers brings 0.7-1.7 points of mAP against the current best approach Swin Transformer [23] at comparable settings. Different from the other multi-scale Transformer models, our method can simultaneously enable both short-range fine-grain and long-range coarse-grain interactions for each visual token, and thus capture richer visual contexts at each layer for better dense predictions. To get better understanding on the models, we further train them with 3x schedule and show the detailed numbers for RetinaNet and Mask R-CNN in Table 4. For comprehension, we also list the number of parameters and the associated computational cost for each model. As we can see, even for 3x schedule, our models can still achieve 0.3-1.1 gain over the best Swin Transformer models at comparable settings.

To further verify the effectiveness of our proposed Focal Transformers, we follow [23] to train four different object detectors including Cascade R-CNN [5], ATSS [47], RepPoints [42] and Sparse R-CNN [30]. We use Focal-Tiny as the backbone and training all four models using 3x schedule. The box mAPs on COCO validation set are reported in Table 5. As we can see, our Focal-Tiny

Backbone	#Params (M)	FLOPs (G)	RetinaNet 3x schedule + MS						Mask R-CNN 3x schedule + MS					
			AP^b	AP^b_{50}	AP^b_{75}	AP_S	AP_M	AP_L	AP^b	AP^b_{50}	AP^b_{75}	AP^m	AP^m_{50}	AP^m_{75}
ResNet50 [16]	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[38]	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 [46]	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 [23]	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 (Ours)	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**
ResNet101 [16]	56.7/63.2	315/336	40.9	60.1	44.0	23.7	45.0	53.8	42.8	63.2	47.1	38.5	60.1	41.3
ResNeXt101-32x4d [41]	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 [38]	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 [46]	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 [23]	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 (Ours)	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**
ResNeXt101-64x4d [41]	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[38]	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 [46]	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 [23]	98.4/107	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 (Ours)	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**

Table 4: COCO object detection and segmentation results with RetinaNet [21] and Mask R-CNN [16]. All models are trained with 3x schedule and multi-scale inputs (MS). The numbers before and after "/" at column 2 and 3 are the model size and complexity for RetinaNet and Mask R-CNN, respectively.

Method	Backbone	#Param	FLOPs	AP^b	AP^b_{50}	AP^b_{75}
Cascade Mask R-CNN [5]	R-50	82.0	739	46.3	64.3	50.5
	Swin-T	85.6	742	50.5	69.3	54.9
	Focal-T	86.7	770	**51.5** (+1.0)	**70.6**	**55.9**
ATSS [47]	R-50	32.1	205	43.5	61.9	47.0
	Swin-T	35.7	212	47.2	66.5	51.3
	Focal-T	36.8	239	**49.5** (+2.3)	**68.8**	**53.9**
RepPointsV2 [42]	R-50	43.4	431	46.5	64.6	50.3
	Swin-T	44.1	437	50.0	68.5	54.2
	Focal-T	45.4	491	**51.2** (+1.2)	**70.4**	**54.9**
Sparse R-CNN [30]	R-50	106.1	166	44.5	63.4	48.2
	Swin-T	109.7	172	47.9	67.3	52.3
	Focal-T	110.8	196	**49.0** (+1.1)	**69.1**	**53.2**

Table 5: Comparison with ResNet-50 and Swin-Tiny across different object detection methods. We use Focal-Tiny as the backbone and train all models using 3x schedule.

Model	W-Size	FLOPs	Top-1 (%)	AP^b	AP^m
Swin-Tiny	7	4.5	81.2	43.7	39.8
	14	4.9	82.1	44.0	40.5
Focal-Tiny	7	4.9	82.2	44.9	41.1
	14	5.2	82.3	45.5	41.5

Table 6: Model performance with different window sizes.

Model	W-Shift	Top-1 (%)	AP^b	AP^m
Swin-Tiny	-	80.2	38.8	36.4
	✓	81.2	43.7	39.8
Focal-Tiny	-	82.2	44.8	41.0
	✓	81.9	44.9	41.1

Table 7: Model performance without and with window shift.

exceeds Swin-Tiny by 1.0-2.3 points on all methods. These significant and consistent improvements over different detection methods in addition to RetinaNet and Mask R-CNN suggest that our Focal Transformer can be used as a generic backbone for a variety of object detection methods.

4.3 Ablation Studies

Above we have shown the superior performance of our Focal Transformer. Here we conduct ablation studies to inspect the model's capacity from different aspects. We use Focal-Tiny as the target and experiment it on both image classification and object detection.

Contributions of short- and long-range interaction. We attempt to factorize the effect of short-range fine-grain and long-range coarse-grain interactions in our Focal Transformers. We ablate the original Focal-Tiny model to: a) Focal-Tiny-Window merely performing attention inside each window; b) Focal-Tiny-Local attending the additional fine-grain surrounding tokens and c) Focal-Tiny-Global attending the extra coarse-grain squeezed tokens. We train them using the same setting as Focal-Tiny and report their performance on image classification and object detection using Mask R-CNN 1x schedule. As we can see from Fig. 5, Focal-Tiny-Window suffers from significant drop on both image classification (82.2→80.1) and object detection (44.8→38.3). This is expected since the communication across windows are totally cut off at each Transformer layer. After we enable either the local fine-grain or global coarse-grain interactions (middle two columns), we observe significant jumps. Though they prompt richer interactions from different paths, finally both of them enable the

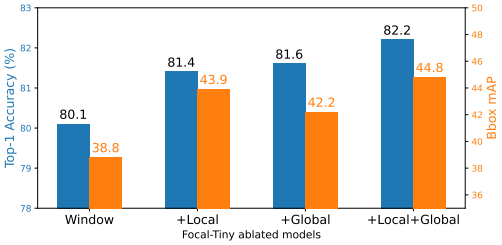

Figure 5: Ablating Focal-Tiny model by removing local, global and both interactions.

Depths	Model	#Params.	FLOPs	Top-1 (%)	AP^b	AP^m
2-2-2-2	Swin	21.2	3.1	78.7	38.2	35.7
	Focal	21.7	3.4	79.9	40.5	37.6
2-2-4-2	Swin	24.7	3.8	80.2	41.2	38.1
	Focal	25.4	4.1	81.4	43.3	39.8
2-2-6-2	Swin	28.3	4.5	81.2	43.7	39.8
	Focal	29.1	4.9	82.2	44.8	41.0

Table 8: Model performance with the change of model depth.

model to capture more contextual information. When we combine them together, we observe further improvements on both tasks. This implies that these two type of interactions are complementary to each other and both of them should be enabled in our model. Another observation is that adding long-range tokens can bring more relative improvement for image classification than object detection and vice versa for local tokens. We suspect that dense predictions like object detection more rely on fine-grained local context while image classification favors more the global information.

Effect of varying window size. Above we have demonstrated that both short- and long-range interactions are necessary. Based on this, a natural question is that whether increasing the window size can further help the model learning giving an enlarged receptive field. In Table 6, we show the performance of Swin-Tiny and Focal-Tiny with window size 7 and 14. Clearly, a larger window size brings gain for both methods on all three metrics and our Focal-Tiny model consistently outperforms Swin-Tiny using both window sizes. Comparing the second and third row, we find our model beats Swin even using much smaller window size (7 *v.s.* 14). We suspect the long-range interactions in our model is the source of this gain.

Model capacity against model depth. Considering our focal attention prompts local and global interactions at each Transformer layer, one question is that whether it needs less number of layers to obtain similar modeling capacity as those without global interactions. To answer this, we conduct the experiments by reducing the number of Transformer layers at stage 3 in Swin-Tiny and Focal-Tiny from the original 6 to 4 and 2. In Table 8, we show the performance and model complexity for each variant. First, we can find our model outperforms Swin model consistently with the same depth. More importantly, using two less layers, our model achieves comparable performance to Swin Transformer. Particularly, Focal-Tiny with 4 layers achieves 81.4 on image classification which is even better than original Swin-Tiny model with 6 layers (highlighted by blue color). Though we do not explore different architectures for our Focal Transformer, these results suggest that we can potential find even more efficient *and* effective architectures.

The necessary of window shift. In [23], the authors proposed window shift operation to enable the cross-window interactions using two successive layers. In contrast, the visual tokens in our Focal Transformer can always communicate with those in other windows at both fine- and coarse-grain. Then a natural question is whether adding the window shift to our Focal Transformers can further bring us improvements. To investigate, we remove the window shift from Swin Transformer while adding it to our Focal Transformer. As shown in Table 7, Swin Transformer has a severe degradation after removing the window shift. However, our Focal Transformer is even hurt on classification task. These results indicate that the window shift is not a necessary ingredient in our model. As such, our model can get rid of the constraint in Swin Transformer that there should be even number of layers at each stage for the alternative window shift operation.

5 Conclusion

In this paper, we presented a new attention mechanism called focal attention to enable efficient long-range interactions in vision transformers. Our design performs the local attention at fine-grain but global attention at coarse-grain which results in an effective way to capture richer context at a reasonable cost. By plugging this to a multi-scale vision transformer, we designed a new focal transformer and demonstrated its superiority over the state-of-the-art methods on both image classification and object detection.

References

[1] Joshua Ainslie, Santiago Ontanon, Chris Alberti, Philip Pham, Anirudh Ravula, and Sumit Sanghai. Etc: Encoding long and structured data in transformers. *arXiv preprint arXiv:2004.08483*, 2020.

[2] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.

[3] Maxim Berman, Hervé Jégou, Andrea Vedaldi, Iasonas Kokkinos, and Matthijs Douze. Multigrain: a unified image embedding for classes and instances. *arXiv preprint arXiv:1902.05509*, 2019.

[4] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.

[5] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6154–6162, 2018.

[6] 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.

[7] Shuning Chang, Pichao Wang, Fan Wang, Hao Li, and Jiashi Feng. Augmented transformer with adaptive graph for temporal action proposal generation. *arXiv preprint arXiv:2103.16024*, 2021.

[8] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. *arXiv preprint arXiv:2103.15436*, 2021.

[9] Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Conditional positional encodings for vision transformers. *Arxiv preprint 2102.10882*, 2021.

[10] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. Up-detr: Unsupervised pre-training for object detection with transformers. *arXiv preprint arXiv:2011.09094*, 2020.

[11] 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.

[12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL*, 2019.

[13] 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.

[14] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer, 2021.

[15] 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.

[16] 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.

[17] Elad Hoffer, Tal Ben-Nun, Itay Hubara, Niv Giladi, Torsten Hoefler, and Daniel Soudry. Augment your batch: Improving generalization through instance repetition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8129–8138, 2020.

[18] Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.

[19] Bing Li, Cheng Zheng, Silvio Giancola, and Bernard Ghanem. Sctn: Sparse convolution-transformer network for scene flow estimation. *arXiv preprint arXiv:2105.04447*, 2021.

[20] Xiangyu Li, Yonghong Hou, Pichao Wang, Zhimin Gao, Mingliang Xu, and Wanqing Li. Trear: Transformer-based rgb-d egocentric action recognition. *arXiv preprint arXiv:2101.03904*, 2021.

[21] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.

[22] 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 *ECCV*, 2014.

[23] 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.

[24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[25] R. Pappagari, P. Zelasko, J. Villalba, Y. Carmiel, and N. Dehak. Hierarchical transformers for long document classification. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 838–844, 2019.

[26] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *International Conference on Machine Learning*, pages 4055–4064. PMLR, 2018.

[27] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.

[28] Jack W Rae, Anna Potapenko, Siddhant M Jayakumar, and Timothy P Lillicrap. Compressive transformers for long-range sequence modelling. *arXiv preprint arXiv:1911.05507*, 2019.

[29] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jonathon Shlens. Stand-alone self-attention in vision models. *arXiv preprint arXiv:1906.05909*, 2019.

[30] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. *arXiv preprint arXiv:2011.12450*, 2020.

[31] Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena: A benchmark for efficient transformers. *arXiv preprint arXiv:2011.04006*, 2020.

[32] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*, 2020.

[33] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877*, 2020.

[34] 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.

[35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

[36] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. *arXiv preprint arXiv:2012.00759*, 2020.

[37] Ning Wang, Wengang Zhou, Jie Wang, and Houqaing Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. *arXiv preprint arXiv:2103.11681*, 2021.

[38] 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.

[39] 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.

[40] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. *arXiv preprint arXiv:2103.15808*, 2021.

[41] 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.

[42] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9657–9666, 2019.

[43] 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.

[44] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, 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.

[45] Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. *arXiv preprint arXiv:2007.14062*, 2020.

[46] 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.

[47] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9759–9768, 2020.

[48] Jiaojiao Zhao, Xinyu Li, Chunhui Liu, Shuai Bing, Hao Chen, Cees GM Snoek, and Joseph Tighe. Tuber: Tube-transformer for action detection. *arXiv preprint arXiv:2104.00969*, 2021.

[49] Minghang Zheng, Peng Gao, Xiaogang Wang, Hongsheng Li, and Hao Dong. End-to-end object detection with adaptive clustering transformer. *arXiv preprint arXiv:2011.09315*, 2020.

[50] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13001–13008, 2020.

[51] 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.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]

(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Appendix

Optionally include extra information (complete proofs, additional experiments and plots) in the appendix. This section will often be part of the supplemental material.