

Convolutional Embedding Makes Hierarchical Vision Transformer Stronger

Cong Wang^{1,2}, Hongmin Xu¹, Xiong Zhang⁴, Li Wang², Zhitong Zheng¹, and Haifeng Liu^{1,3}

¹ Data & AI Engineering System, OPPO, Beijing, China

{wangcong3575, xhmjimmy}@gmail.com, {liam,blade}@oppo.com

² Beijing Key Lab of Urban Intelligent Traffic Control Technology, North China University of Technology, Beijing, China

li.wang@ncut.edu.cn

³ University of Science and Technology of China, Hefei, China

⁴ Neolix Autonomous Vehicle, Beijing, China

zhangxiong@neolix.cn

Abstract. Vision Transformers (ViTs) have recently dominated a range of computer vision tasks, yet it suffers from low training data efficiency and inferior local semantic representation capability without appropriate inductive bias. Convolutional neural networks (CNNs) inherently capture regional-aware semantics, inspiring researchers to introduce CNNs back into the architecture of the ViTs to provide desirable inductive bias for ViTs. However, *is the locality achieved by the micro-level CNNs embedded in ViTs good enough?* In this paper, we investigate the problem by profoundly exploring how the macro architecture of the hybrid CNNs/ViTs enhances the performances of hierarchical ViTs. Particularly, we study the role of token embedding layers, alias *convolutional embedding* (CE), and systemically reveal how CE injects desirable inductive bias in ViTs. Besides, we apply the optimal CE configuration to 4 recently released state-of-the-art ViTs, effectively boosting the corresponding performances. Finally, a family of efficient hybrid CNNs/ViTs, dubbed **CETNets**, are released, which may serve as generic vision backbones. Specifically, CETNets achieve 84.9% Top-1 accuracy on ImageNet-1K (training from scratch), 48.6% box mAP on the COCO benchmark, and 51.6% mIoU on the ADE20K, substantially improving the performances of the corresponding state-of-the-art baselines.

Keywords: Vision Transformers, Convolutional Neural Networks, Convolutional Embedding, Micro and Macro Design

1 Introduction

Over the last decades, convolutional neural networks (CNNs) significantly succeeded in the computer vision community due to their inherent properties, including the translation invariance, the locality attention, and the sharing weight design. Those characteristics prove critical for many tasks, such as image recognition [16, 25], semantic image segmentation[7, 70], and object detection[11, 42].

At the same time, researchers take a very different way in the natural language processing (NLP) field. Since the seminal work [54] demonstrated the extraordinary capability of the transformer by employing a unified yet simple architecture to tackle the machine translation task, transformers have become the de-facto architectures to resolve NLP tasks [40, 17].

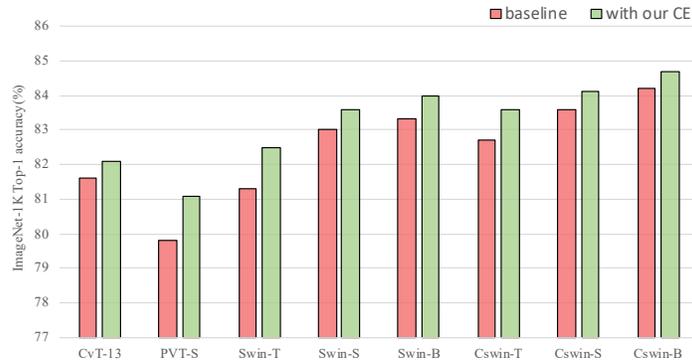


Fig. 1. Performance Improvements of the Convolutional Embedding (CE). The figure presents the performance gains of the CE among 4 SOTA ViTs, *i.e.*, *CvT* [59], *PVT* [56], *SWin* [35], and *CSWin* [18] on the ImageNet-1K dataset, indicating that the CE indeed significantly improves the baseline methods

In the computer vision domain, certain works [57, 3, 46, 29] successfully brought the key idea of transformer, *i.e.*, attention paradigm, into CNNs and achieved remarkable improvements. Naively transferring the transformer to image recognition, Dosovitskiy *et al.* [19] demonstrated that the vanilla Vision Transformers (ViTs) structure could achieve comparable performance compared with the state-of-the-art (SOTA) approaches on the ImageNet-1K dataset [19]. Further, pre-trained on the huge dataset JFT-300M [48], the vanilla ViTs outperformed all the SOTA CNNs by a large margin and substantially advanced the SOTA, suggesting that ViTs may have a higher performance ceiling.

ViTs rely on highly flexible multi-head self-attention layers to favor dynamic attention, capture global semantics, and achieve good generalization ability. Yet, recent works find that lacking proper inductive bias, locality, ViTs are substandard optimizability [60], low training sample-efficient [15], and poor to model the complex visual features and the local relation in an image [39, 59]. Most existing works attempt to introduce local mechanisms to ViTs in two paths. One line of work relieves the inductive bias problem via non-convolutional ways. Liu *et al.* [35, 18, 53] limits the attention computation in a local window to enable local receptive field for attention layers, making the overall network still maintain nearly pure attention-based architecture. Concurrently, as the CNNs are inherently efficient for their sliding-window manner, local receptive field, and inductive bias [2], another line of work directly integrates CNNs into ViTs design

to bring convolutional inductive bias into ViTs in a hard [59, 33, 65] or soft [15, 52] way. However, most of those works focus on modifying the micro design of ViTs to enable locality, which raises the question:

Is the inductive bias obtained via the micro design of ViTs good enough to empower the locality to ViTs? or Can the macro architecture design of the network further introduce the desirable inductive bias to ViTs?

We ask the questions mentioned above based on the following findings. The previous work, EarlyConv [60], notices that simply replacing the original patchify stem with a 5-layer convolutional stem can yield 1-2% top-1 accuracy on ImageNet-1K and improve the training stability of ViTs. Subsequently, CoAtNet [14] further explore the hybrid design of CNNs and ViTs based on the observation: the depth-wise convolution can be naturally integrated into the attention block. Meanwhile, those works suggest that the convolutional layers can efficiently introduce inductive bias in a shallow layer of the whole network. However, when we retrospect the roadmap of modern CNNs, at the late of 2012, after AlexNet [32] exhibited the potential of CNNs, subsequent studies, for instance, VGG [45], ResNet [25], DenseNet [30], EfficientNet [50, 51] ConvNet2020 [36], and *etc.* reveal that: convolutions can represent the complex visual features even in a deep layer of network effectively and efficiently. Our research explores the macro network design with hybrid CNNs/ViTs. We want to bridge the gap between the pure CNNs network and the pure ViT network and extend the limitation of the hybrid CNNs/ViTs network.

To test this hypothesis, we start from CNNs’ effective receptive field (ERF). As previous work of Luo *et al.* [38] points out, the output of CNNs considerably depends on their ERF. With larger ERF, CNNs can leave out no vital information, which leads to better prediction results or visual features. Under this perspective, our exploration is imposing strong and effective inductive bias for attention layers via the macro design of network architecture. We specifically focus on *patch embedding*, alias *convolutional embedding (CE)*, of the hierarchical ViTs architecture. CE locate at the beginning of each stage, as shown in Fig. 2. The CE aims to adjust the dimension and the number of tokens. Most following works also apply one or two convolutions embedding layers [59, 18, 68, 56, 33]. However, these embedding layers cannot offer enough ERF to capture more complex visual representations with desirable inductive bias. Since stacking more convolutional layers can increase the ERF [38], we construct a simple baseline with only 1-layer CE and gradually increase the number of convolutions layers in CE to get more variants. Meanwhile, keep changing FLOPs and parameter numbers as small as possible. We observe that the small change of CE in each stage results in a significant performance increase in the final model.

Based on extensive experiments, we further understand how CE affects the hybrid network design of CNNs/ViTs by injecting desirable inductive bias. We make several observations. 1) CNNs bring strong inductive bias even in deep layers of a network, making the whole network easier to train and capturing much complex visual features. At the same time, ViTs allow the entire network has a higher generalization ceiling. 2) CE can impose effective inductive bias, yet

different convolution layers show variable effectiveness. Besides, the large ERF is essential to designing CE or injecting the desirable inductive bias to ViTs, even though it is a traditional design in a pure CNNs network [51, 38]. 3) CNNs can help ViTs see better even in deep networks, providing valuable insights to guide how to design the hybrid CNNs/ViTs network. 4) It is beneficial to combine the macro and micro introduce inductive bias to obtain a higher generalization ceiling of ViTs-based network.

Our results advocate the importance of CE and deep hybrid CNNs/ViTs design for vision tasks. ViT is a general version of CNN [10], and tons of works have proven the high generalization of ViTs-based networks, which spurs researchers to line up to chase the performance ceiling with pure attention networks. After inductive bias is found to be crucial to significantly improve the training speed and sample-efficiency of ViTs, people’s efforts are mainly devoted to creating the micro design of ViTs to enhance it [35, 15]. Concurrently, EarlyConv [60] and CoAtNet [14] verify the efficiency of convolutions in shallow layers of the ViT-based network. Our study further pushes the boundary of the macro design of the hybrid CNNs/ViTs network. Our results also suggest that even in deep layers of the ViTs network, correctly choosing the combination of CNNs/ViTs design, one can further improve the upper-performance limitation of the whole network. Finally, we propose a family of models of hybrid CNNs/ViTs as a generic vision backbone.

To sum up, we hope our findings and discussions presented in this paper will deliver possible insights to the community and encourage people to rethink the value of CE in the hybrid CNNs/ViTs network design.

2 Related Work

Convolutional neural networks. Since the breakthrough performance of AlexNet[32], the computer vision field has been dominated by CNNs for many years. In the past decade, we have witnessed a steady stream of new ideas being proposed to make CNNs more effective and efficient[45, 49, 25, 30, 29, 22, 27, 43, 69, 50, 51, 41]. One line of work focuses on improving the individual convolutional layer except for the architectural advances. For example, the depthwise convolution [62] is widely used due to its lower computational cost and smaller parameter numbers, and the deformable convolution [12] can adapt to shape, size, and other geometric deformations of different objects by adding displacement variables. The dilated convolution introduces a new parameter called ”dilation rate” into the convolution layer, which can arbitrarily expand the receptive field without additional parameters cost. These CNNs architectures are still the primary backbones for computer vision tasks such as object detection[42], instance segmentation[24], semantic segmentation[37], and so on.

Vision Transformers. Transformer [54] has become a prevalent model architecture in natural language processing (NLP) [17, 4] for years. Inspired by the success of NLP, increasing effort on adapting Transformer to computer vision

tasks. Dosovitskiy *et al.* [19] is the pioneering work that proves pure Transformer-based architectures can attain very competitive results on image classification, shows strong potential of the Transformer architecture for handling computer vision tasks. The success of [19] further inspired the applications of Transformer to various vision tasks, such as image classification[52, 66, 63, 23, 59], object detection[5, 73, 71, 13] and semantic segmentation[55, 47]. Furthermore, some recent works [56, 67, 35] focusing on technique a general vision Transformer backbone for general-purpose vision tasks. They all follow a hierarchical architecture and develop different self-attention mechanisms. The hierarchical design can produce multi-scale features that are beneficial for dense prediction tasks. An efficient self-attention mechanism reduced the computation complexity and enhanced modeling ability.

Integrated CNNs and Transformer. CNNs are good at capturing local features and have the advantages of shift, scale, and distortion invariance, while Transformer have the properties of dynamic attention, global receptive field, and large model capacity. Combining convolutional and Transformer layers can achieve better model generalization and efficiency. Many researchers are trying to integrate CNNs and Transformer. Some methods[57, 3, 28, 58, 46, 44] attempt to augment CNNs backbones with self-attention modules or replace part of convolution blocks with Transformer layers. In comparison, inspired by the success of ViT[19], recent trials attempt to leverage some appropriate convolution properties to enhance Transformer backbones. ConViT[15] introduce a parallel convolution branch to impose convolutional inductive biases into ViT[19]. Localvit [33] adding a depth-wise convolution in Feed-Forward Networks(FFN) component to extract the locality, and CvT[59] employs convolutional projection to calculate self-attention matrices to achieve additional modeling of local spatial context. Besides the “internal” fusion, some method[14, 60] focus on structural combinations of Transformer and CNNs.

3 Hybrid CNNs/ViTs Network Design

A general hierarchical ViTs model architecture is illustrated in Figure 2. The convolutional stem is applied for an input image to extract the low-level feature, followed typically four stages to extract diverse-scale deep representations gradually. Each stage consists of a *convolutional embedding* (CE) block and a set of ViTs blocks. More specifically, the CE block is located at the beginning of each stage and aims to adjust the dimension, change the number, and reduce the resolution of input tokens. For two adjacent stages, the reduction factor is set to 2. After that, those tokens are fed to the following ViTs blocks to generate a global context. To sum up, the proposed network contains five stages (S_0, S_1, S_2, S_3, S_4), where S_0 is the convolutional stem.

The original hierarchical ViTs architecture [35] is following the traditional design of the CNNs network, VGG [45], targeting giving the tokens the ability to represent increasingly complex visual patterns over increasingly larger spatial

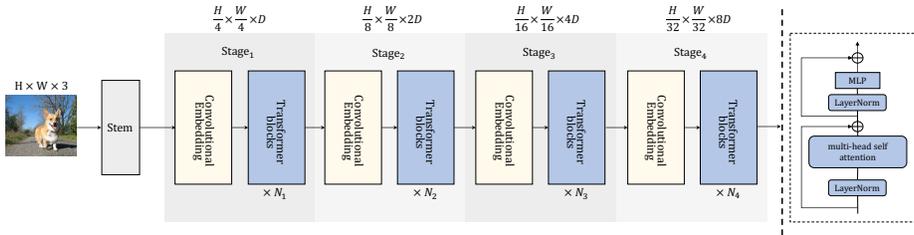


Fig. 2. Left: the overall architecture of a general hierarchical ViTs network, which consists of two main parts: *Stem* and *Stages*. Each stage contains a *convolutional embedding* and a set of *ViTs blocks*. Right: an example of a standard *ViTs block*[19]

footprints. Our goal is to impose the desirable inductive bias for ViTs blocks of the hierarchical architecture. Thus, our explorations are directed to two paths with pure CNNs structure: *Macro Design* and *Micro Design*.

3.1 Macro Design

Convolutional Stem. In ViTs-based network design, the *Stem* is concerned to extract the right inductive bias to following global attention modules. It is necessary that maintain enough effective receptive field (ERF) of CE to extract rich visual features at the early stage [60]. In consideration of parameter numbers and FLOPs budget, as well as we notice that S_0 can be merged with the CE of S_1 , our final *Stem* consists of 4 Fused-MBConvs [21] layers and 1 regular convolutional layer. Specifically, S_0 contains 2 Fused-MBConv with stride 2 and stride 1, respectively. The CE of S_1 is consists of the same as S_0 and followed by one 1×1 convolution at the end to match the channels and normalized by a layer normalization [1]. Another reason that we choose Fused-MBConv to compose *Stem* is that EfficientNetV2 [51] shows that Fused-MBConv is surprisingly effective in achieving better generalization and capacity as well as makes training convergence rate faster [60]. Besides, like EfficientNetV2, the hidden state’s expand ratios of Fused-MBConv and MBConv [43] are arranged to 1 or 2. Such a setting allows the convolutional *Stem* to have fewer parameters and FLOPs without the loss of accuracy. Please refer to the supplementary materials for more details about Fused-MBConv and MBConv.

Convolutional Embedding. In the following stages, S_2 , S_3 , and S_4 , each stage contains a CE block and a set of ViTs blocks. The CE block captures diverse deep representation with a convolutional inductive bias for subsequent attention modules. It is worth noting that EarlyConv [60] and CoAtNet [14] point out stacking convolutional structure may enhance the ViTs in early stage. However, as we show in Table 8, we argue that CNNs is also able to represent the same, if not better, deep feature as ViTs in the deep layer of a network. Meanwhile, maintaining the CNNs design of the embedding layer naturally introduces proper inductive bias to following ViTs and retains the sample-efficient learning

property. For the same computation resource constraint consideration, the CE adopts only effective and efficient convolutions. The CE of *S2*, *S3*, and *S4* adopts MBConv as the basic unit.

3.2 Micro Design

Locally Enhanced Window Self-Attention. Previous work [35] restricts the attention calculation with a local window to reduce computational complexity from quadratic to linear. Meanwhile, to some extent, the local window injects some inductive bias into ViTs. Concurrently, a set of works [59, 56] attempt to directly integrate CNNs into ViTs to bring in convolutional inductive bias, yet their computational complexity is still quadratic. We present a simple Locally Enhanced Window Self-Attention (LEWin) to take advantage of fast attention calculating and locality inductive bias. Inspired by CvT, we performed a *convolutional projection* on input tokens in a local shift window. The *convolutional projection* is implemented by a depth-wise separable convolution with kernel size 3×3 , stride 1, and padding 1. The LEWin can be formulated as:

$$\mathbf{x}^{i-1} = \text{Flatten}(\text{Conv2D}(\text{Reshape}(\mathbf{z}^{i-1}))), \quad (1)$$

$$\hat{\mathbf{z}}^i = \text{WHMSA}(\text{LN}(\mathbf{x}^{i-1})) + \mathbf{x}^{i-1}, \quad (2)$$

$$\mathbf{z}^i = \text{MLP}(\text{LN}(\hat{\mathbf{z}}^i)) + \hat{\mathbf{z}}^i, \quad (3)$$

$$\mathbf{x}^i = \text{Flatten}(\text{Conv2D}(\text{Reshape}(\mathbf{z}^i))), \quad (4)$$

$$\hat{\mathbf{z}}^{i+1} = \text{SWMSA}(\text{LN}(\mathbf{x}^i)) + \mathbf{x}^i, \quad (5)$$

$$\mathbf{z}^{i+1} = \text{MLP}(\text{LN}(\hat{\mathbf{z}}^{i+1})) + \hat{\mathbf{z}}^{i+1} \quad (6)$$

where \mathbf{x}^i is the unperturbed token in local window prior to the *convolutional projection*, $\hat{\mathbf{z}}^i$ and \mathbf{z}^i denote the output features of the WMSA or S-WMSA module and the MLP module for the i -th block, respectively. W-MSA and SW-MSA define window-based multi-head self-attention based on the regular and shifted windows from SWin [35], respectively.

3.3 CETNet Variants

We consider three different network configurations for CETNet to compare with other ViTs backbones under similar model size and computation complexity conditions. By changing the base channel dimension and the number of ViTs blocks of each stage, we build three variants, tiny, small, and base models, namely CETNet-T, CETNet-S, and CETNet-B. For more detailed configurations, please refer to our supplementary materials.

4 Experiments

To verify the ability of our CETNet as a general vision backbone. We conduct experiments on ImageNet-1K classification[16], COCO object detection[34], and ADE20K semantic segmentation[72]. In addition, comprehensive ablation studies are performed to validate the design of the proposed architecture.

Table 1. Comparison of image classification on ImageNet-1K for different models. The models are grouped based on the similar model size and computation complexity

ImageNet-1K 224 ² trained model											
Model	Params	FLOPs	Top-1(%)	Model	Params	FLOPs	Top-1(%)	Model	Params	FLOPs	Top-1(%)
ResNet-50[25]	25M	4.1G	76.2	ResNet-101[25]	45M	7.9G	77.4	ResNet-152[25]	60M	11.0G	78.3
RegNetY-4G[41]	21M	4.0G	80.0	RegNetY-8G[41]	39M	8.0G	81.7	RegNetY-16G[41]	84M	16.0G	82.9
DeiT-S[52]	22M	4.6G	79.8	PVT-M[56]	44M	6.7G	81.2	DeiT-B[52]	87M	17.5G	81.8
PVT-S[56]	25M	3.8G	79.8	T2T-19[66]	39M	8.9G	81.5	PiT-B[26]	74M	12.5G	82.0
T2T-14[66]	22M	5.2G	81.5	T2T _r -19[66]	39M	9.8G	82.2	T2T-24[66]	64M	14.1G	82.3
ViL-S[67]	25M	4.9G	82.0	ViL-M[67]	40M	8.7G	83.3	T2T _r -24[66]	64M	15.0G	82.6
TNT-S[23]	24M	5.2G	81.3	MViT-B[20]	37M	7.8G	81.0	CPVT-B[9]	88M	17.6G	82.3
CViT-15[6]	27M	5.6G	81.0	CViT-8[6]	43M	9.0G	82.5	TNT-B[23]	66M	14.1G	82.8
LViT-S[33]	22M	4.6G	80.8	CViT _c -18[6]	44M	9.5G	82.8	ViL-B[67]	56M	13.4G	83.2
CPVT-S[9]	23M	4.6G	81.9	Twins-B[8]	56M	8.3G	83.2	Twins-L[8]	99M	14.8G	83.7
Swin-T[35]	29M	4.5G	81.3	Swin-S[35]	50M	8.7G	83.0	Swin-B[35]	88M	15.4G	83.5
CvT-13[59]	20M	4.5G	81.6	CvT-21[59]	32M	7.1G	82.5	CETNet-B	75M	15.1G	83.8
CETNet-T	23M	4.3G	82.7	CETNet-S	34M	6.8G	83.4				

ImageNet-1K 384 ² finetuned model											
Model	Params	FLOPs	Top-1(%)	Model	Params	FLOPs	Top-1(%)	Model	Params	FLOPs	Top-1(%)
CvT-13[59]	25M	16.3G	83.0	CvT-21[59]	32M	24.9G	83.3	ViT-B/16[19]	86M	49.3G	77.9
T2T-14[66]	22M	17.1G	83.3	CViT _c -18[6]	45M	32.4G	83.9	DeiT-B[52]	86M	55.43G	83.1
CViT _c -15[6]	28M	21.4G	83.5	CETNet-S	34M	19.9G	84.6	Swin-B[35]	88M	47.0G	84.5
CETNet-T	24M	12.5G	84.2					CETNet-B	75M	44.50G	84.9

4.1 Image Classification

Settings For Image classification, we compare different methods on ImageNet-1K[16], with about 1.3M images and 1K classes, the training set, and the validation set containing 1.28M images and 50K images respectively. The top 1 accuracy on the validation set is reported to show the capacity of our CETNet. For a fair comparison, the experiment setting follows the training strategy in [35]. All model variants are trained for 300 epochs with a batch size of 1024 or 2048 and using a cosine decay learning rate scheduler with 20 epochs of linear warm-up. We adopt an AdamW[31] optimizer with an initial learning rate of 0.001, and a weight decay of 0.05 is used. We use the same data augmentation methods, and regularization strategies used in [35] for training. All models are trained with 224×224 input size, while a center crop is used during evaluation on the validation set. When fine-tuning on 384×384 input, we train the models for 30 epochs with a learning rate of $2e-5$, batch size of 512, and no center crop for evaluation.

Results Table 1 presents comparisons to other models on the image classification task, and the models are split into three groups based on the similar model size (Params) and computation complexity (FLOPs). It can be seen from the table, our models consistently exceed other methods by large margins. More specifically, CETNet-T achieves 82.7% Top-1 accuracy with only 23M parameters and 4.3G FLOPs, surpassing CvT-13, Swin-T, and DeiT-S by 1.1%, 1.4% and 2.9% separately. For small and base models, our CETNet-S and CETNet-B still achieve better performance with comparable Params and FLOPs. Compared with the state-of-the-art CNNs RegNet[41] which also trained with input size 224×224 with out extra data, our CETNet achieves better accuracy. On the 384×384 input, Our CETNet’s performance is still better than other

Backbone, achieve 84.2%, 84.6% and 84.9% respectively, well demonstrates the powerful learning capacity of our CETNet. As a result, we can conclude that an effective combination of CNNs and ViTs can optimize well in small(or middle) scale datasets.

Table 2. Object detection and instance segmentation performance on COCO 2017 validation set with the Mask R-CNN framework. The FLOPs are measured at a resolution of 800×1280 , and the backbones are pre-trained on the ImageNet-1K

Backbone	Mask R-CNN 1× schedule								Mask R-CNN 3× schedule					
	Params	FLOPs	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_{70}^m
Res50[25]	44M	260G	38	58.6	41.4	34.4	55.1	36.7	41	61.7	44.9	37.1	58.4	40.1
PVT-S[56]	44M	245G	40.4	62.9	43.8	37.8	60.1	40.3	43	65.3	46.9	39.9	62.5	42.8
ViL-S[67]	45M	218G	44.9	67.1	49.3	41.0	64.2	44.1	47.1	68.7	51.5	42.7	65.9	46.2
TwinsP-S[8]	44M	245G	42.9	65.8	47.1	40.0	62.7	42.9	46.8	69.3	51.8	42.6	66.3	46
Twins-S[8]	44M	228G	43.4	66.0	47.3	40.3	63.2	43.4	46.8	69.2	51.2	42.6	66.3	45.8
Swin-T[35]	48M	264G	43.7	64.6	46.2	39.1	61.6	42.0	46	68.2	50.2	41.6	65.1	44.8
CETNet-T	43M	261G	45.5	67.7	50.0	40.7	64.4	43.7	46.9	67.9	51.5	41.6	65	44.7
Res101[25]	63M	336G	40.4	61.1	44.2	36.4	57.7	38.8	42.8	63.2	47.1	38.5	60.1	41.3
X101-32[62]	63M	340G	41.9	62.5	45.9	37.5	59.4	40.2	42.8	63.2	47.1	38.5	60.1	41.3
PVT-M[56]	64M	302G	42.0	64.4	45.6	39.0	61.6	42.1	42.8	63.2	47.1	38.5	60.1	41.3
ViL-M[67]	60M	261G	43.4	-	-	39.7	-	-	44.6	66.3	48.5	40.7	63.8	43.7
TwinsP-B[8]	64M	302G	44.6	66.7	48.9	40.9	63.8	44.2	47.9	70.1	52.5	43.2	67.2	46.3
Twins-B[8]	76M	340G	45.2	67.6	49.3	41.5	64.5	44.8	48	69.5	52.7	43	66.8	46.6
Swin-S[35]	69M	354G	44.8	66.6	48.9	40.9	63.4	44.2	48.5	70.2	53.5	43.3	67.3	46.6
CETNet-S	53M	315G	46.6	68.7	51.4	41.6	65.4	44.8	48.6	69.8	53.5	43	66.9	46
X101-64[62]	101M	493G	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[56]	81M	364G	42.9	65.0	46.6	39.5	61.9	42.5	44.5	66.0	48.3	40.7	63.4	43.7
ViL-B[67]	76M	365G	45.1	-	-	41.0	-	-	45.7	67.2	49.9	41.3	64.4	44.5
TwinsP-L[8]	81M	364G	45.4	-	-	41.6	-	-	-	-	-	-	-	-
Twins-L[8]	111M	474G	45.9	-	-	41.6	-	-	-	-	-	-	-	-
Swin-B[35]	107M	496G	46.9	-	-	42.3	-	-	48.5	69.8	53.2	43.4	66.8	46.9
CETNet-B	94M	495G	47.9	70.3	53.0	42.5	67.2	45.6	48.6	69.5	53.7	43.1	66.9	46.4

4.2 Object Detection and Instance Segmentation

Settings For object detection and instance segmentation experiments, we evaluate our CETNet with the Mask R-CNN[24] framework on COCO 2017, which contains over 200K images with 80 classes. The models pre-trained on the ImageNet-1K dataset are used as visual backbones. We follow the standard to use two training schedules, 1× schedule(12 epochs with the learning rate decayed by 10× at epochs 8 and 11) and 3× schedule(36 epochs with the learning rate decayed by 10× at epochs 27 and 33). We utilize the multi-scale training strategy[5], resizing the input that the shorter side is between 480 and 800 while the longer side is no more than 1333. AdamW[31] optimizer with initial learning rate of 0.0001, weight decay of 0.05. All models are trained on the 118K training

images with a batch size of 16, and the results are reported on 5K images in COCO 2017 validation set.

Results As shown in Table 2, the results of the Mask R-CNN framework show that our CETNet variants clearly outperform all counterparts with $1\times$ schedule. In detail, our CETNet-T outperforms Swin-T by +1.8 box AP, +1.6 mask AP. On the small and base configurations, the performance gain can also be achieved with +1.8 box AP, +0.7 mask AP and +1.0 box AP, +0.2 mask AP respectively. For $3\times$ schedule, our model, CETNet-S and CETNet-B, can achieve competitive results with lower Params and FLOPs than the current state-of-the-art ViTs methods in small and base scenarios. However, like SWin, CETNet-T can't beat the SOTA performance. Also, in the small and base configurations, we notice the variants of our CETNet and SWin do not improve the performance. We conjecture such inferior performance may be because of insufficient data.

Table 3. Comparison of semantic segmentation on ADE20K with the Upernet framework, both single and multi-scale evaluations are reported in the last two columns. FLOPs are calculated with a resolution of 512×2048 , and the backbones are pre-trained on the ImageNet-1K

Backbone	Upernet 160k trained models			
	Params	FLOPs	mIoU(%)	MS mIoU(%)
TwinsP-S[8]	54.6M	919G	46.2	47.5
Twins-S[8]	54.4M	901G	46.2	47.1
Swin-T[35]	59.9M	945G	44.5	45.8
CETNet-T	53.2M	935G	46.5	47.9
Res101[25]	86.0M	1029G	-	44.9
TwinsP-B[8]	74.3M	977G	47.1	48.4
Twins-B[8]	88.5M	1020G	47.7	48.9
Swin-S[35]	81.3M	1038G	47.6	49.5
CETNet-S	63.4M	990G	48.9	50.6
TwinsP-L[8]	91.5M	1041G	48.6	49.8
Twins-L[8]	133.0M	1164G	48.8	50.2
Swin-B[35]	121.0M	1188G	48.1	49.7
CETNet-B	106.3M	1176G	50.2	51.6

4.3 Semantic Segmentation

Settings We further use the pre-trained models as the backbone to investigate the capability of our models for Semantic Segmentation on the ADE20K[72] dataset. ADE20K is a widely-used semantic segmentation dataset that contains 150 fine-grained semantic categories, with 20,210, 2,000, and 3,352 images for training, validation, and testing, respectively. We follow previous works[35] to employ UperNet[61] as the basic framework and follow the same setting for a fair comparison. In the training stage, we employ the AdamW[31] optimizer and set the initial learning rate to $6e-5$ and use a polynomial learning rate decay, and the weight decay is set to 0.01 and train Upernet 160k iterations with batch size of 16. The data augmentations adopt the default setting in mmsegmentation of random horizontal flipping, random re-scaling within ratio range [0.5, 2.0], and random photometric distortion. Both single and multi-scale inference are used for evaluation.

Results In Table 3, we list the results of Upernet 160k trained model on ADE20K. It can be seen that our CETNet models significantly outperform previous state-of-the-arts under different configurations. In details, CETNet-T achieves 46.5 mIoU and 47.9 multi-scale tested mIoU, +2.0 and + 2.1 higher

than Swin-T with similar computation cost. On the small and base configurations, CETNet-S and CETNet-B still achieve +1.3 and +2.1 higher mIOU and +1.1 and +1.9 multi-scale tested mIoU than the Swin counterparts. The performance gain is promising and demonstrates the effectiveness of our CETNet design.

4.4 Ablation Study

In this section, we ablate the critical elements of the proposed CETNet backbone using ImageNet-1K image classification. The experimental settings are the same as the settings in Sec. 4.1. For the attention mechanism, when replacing the micro design LEWin self-attention with the origin shifted window self-attention, the performance dropped 0.2%, demonstrating the effectiveness of our micro design. Furthermore, we find the “deep-narrow” architecture is better than the “shallow-wide” counterpart. Specifically, the deep-narrow model with [2,2,18,2] Transformer blocks for four stages with the base channel dimensions $D = 64$ and the shallow-wide model with [2,2,6,2] blocks for four stages with $D = 96$. As we can see from the 1st and 3rd rows, even with larger Params and FLOPs, the shallow-wide model performs worse than the deep-narrow design. The CE module is the key element in our models. To verify the effectiveness CE module, we compare it with the existing method used in hierarchical ViTs backbones, including the patch embedding and patch merging modules in SWin [35] and convolutional token embedding modules described in CvT [59]. For a fair comparison, we use the shallow-wide design mentioned above and apply these three methods in all these models with all other factors kept the same. As shown in the last three rows of Table 4, our CE module performs better than other existing methods.

Table 4. Ablation study of CETNet’s key elements. ‘Without LEAtten’ denotes replacing the micro design LEWin self-attention with the shifted window self-attention in SWin [35]. The ‘D’ and ‘S’ represents the deep-narrow and the shallow-wide model, respectively. The 1st and 3rd rows are the baseline ‘D’ and ‘S’ models.

Models	Param	FLOPs	Top-1(%)
CETNet-T (D with CE)	23.4M	4.3G	82.7
Without LEAtten (D with CE)	23.3M	4.2G	82.5
Shallow-wide (S with CE)	29M	4.6G	82.5
Patch embedding+patch merging [35] (S)	27M	4.6G	81.5
Convolutional token embedding [59] (S)	27M	4.6G	82.0

5 Investigating the role of Convolutional Embedding

This section systematically understands how the hybrid CNNs/ViT network benefits from the CE. As we need carefully make the trade-off among the computation budget, the model generalization, and the model capacity, the CE, two

design choices are mainly explored in this work: the number of CNNs layers of CE and the basic unit of CE. Besides, we show the superiority of CE via integrate CE into 4 popular hierarchical ViTs models. Finally, we further investigate the role of CNNs in hybrid CNNs/ViTs network design. Except for specific declaration, all configuration is follow the Image Classification setting as section 4.

5.1 Effect of the Stacking Number

We first explore how large effective receptive field (ERF) affect the computation budget and the performance of the network. As we mentioned in section 3.1, the $S0$ and the CE of $S1$ are combined as one whole CNNs stem. We refer it as the first CE (CE^{1st}) in rest of this section, CE^{2nd} , CE^{3rd} , and CE^{4th} represent the CE of $S2$, $S3$, $S4$, respectively. To explore the effect of the stacking number of the basic unit,

we slightly modify the SWin-Tiny [35] model by replacing its patch embedding and patch merging modules with pure CNNs. We choose MBConv as the basic unit of CE and gradually increase the layer number of CE from 1 to 3, 5, and 7. As shown in Table 5, as the CE contains more MBConvs, the Param and FLOPs grow, and the performance increases from 81.8% to 82.6%. The 3-layer CE model has slightly higher FLOPs than the 1-layer CE (4.48G vs. 4.55G) with near the same Param (28.13M vs. 28.11M). Besides, it is worth noticing that the performance grows negligibly from 5-layer CE to 7-layer CE. Thus, we employ 5-layer setting, offering large ERF [64, 51], as CETNet’s final configuration.

5.2 Effect of Different CNNs Blocks

We next explore how well convolutional inductive bias inject by different CNNs architectures. The CE layers aim to offer rich features for later attention modules. CNNs’ effective receptive field (ERF) determines the information it can cover and process. As Luo *et al.* mentioned, stacking more layers, subsampling, and CNNs with large ERF, such as dilated convolution [64] can enlarge the ERF of CNNs networks. For find an efficient basic unit for CE, we use the CETNet-Tiny (CETNet-T) model as a baseline and

Table 5. Performance of model with different layer numbers of CE. The baseline is slightly modified from SWin-Tiny [35]

Models	Param	FLOPs	Top-1(%)
Swin-T[35]	29M	4.5G	81.3
1-layer CE	28.1M	4.5G	81.8
3-layer CE	28.1M	4.6G	82.2
5-layer CE	29.1M	4.9G	82.4
7-layer CE	30.1M	5.4G	82.6

Table 6. Transfer some popular CNNs blocks to CETNet-T, including the DenseNet block, ShuffleNet block, ResNet block, SersNet block, and GhostNet block. PureMBConv represents to use MBConv block to replace the Fused-MBConv block in the early stage of CETNet-T

CNNs type	Params	FLOPs	Top-1(%)
CETNet-T	23.4M	4.3G	82.7
PureMBConvs[43]	23.4M	4.2G	82.6
GhostNetConvs[22]	24.6M	4.3G	82.6
DenseNetConvs[30]	22.3M	4.4G	82.5
ShuffleNetConvs[69]	23.4M	4.3G	82.4
ResNetConvs[25]	23.8M	4.3G	82.0
SersNetConvs[29]	24.0M	4.3G	82.0

replace the CE with some recent CNNs building blocks, such as MBCConv [43], DenseNet [30], ShuffleNet [69], ResNetConvs[25], SeresNetConvs [29], and GhostNetConvs [22]. All candidate convolutions layers stack to 5 layers based on previous finding. For a fair comparison, all model variants are constructed to have similar parameter numbers and FLOPs. As we can see from Table 6, when replaced the Fused-MBConvs in the early stage of CETNet-T by MBConvs, the top-1 accuracy increased 0.1%, and we observed a 12 percent decrease in training speed. Also, CETNet-T and PureMB model achieve higher performance than other candidate convolutions. We argue that may be the internal relation between depthwise convolution and ViTs as pointed out by CoAtNet [14], which is further verified by the GhostNet and shuffleNet model, which archive 82.6% and 82.4% top-1 accuracy. Besides, we noticed that under that CNNs help ViTs see better perspective, in CE case, dense connections in convolutions may not necessarily hurt performance. Since the result shows that our DenseNet model can also archive 82.5% top-1 accuracy, which is comparable with the performance of CETNet-T, PureMB, GhostNet, and shuffleNet model. However, ResNet and SeresNet show inferior performance. We conjecture that the basic units have different ERF with the same stacking number.

5.3 Generalization of CE

Then, we attempt to generalize the CE design to more ViTs backbones of CV. Here, we apply our CE design, 5-layer Fused-MBConv of CE^{1st}, and 5-layer MBCConv of CE^{2nd}, CE^{3rd}, and CE^{4th} respectively, to 4 prevalent backbones, CvT [59], PVT[56], SWin [35], and CSWin [18]. For a fair comparison, we slightly change the structure, removing some ViTs blocks, of 4 models to keeps their parameter numbers and FLOPs maintaining the similar level as their original version. Also, we modify the small-scale model variant CvT-13 and PVT-S of CvT and PVT. As shown in Table 7, those modified models outperform the original model 0.5% and 1.3% separately. Furthermore, when introducing our design into SWin and CSWin, the top-1 accuracy of all counterparts is improved even under lower parameter numbers and FLOPs scenarios. For details, the modified models of Swin counterparts gain 1.2%, 0.6% and 0.7%, and the CSwin counterparts gain 0.9%, 0.5% and 0.5% respectively. Those results demonstrated that CE could be easily integrated with other ViT models and significantly improve the performance of those ViT models.

Table 7. Generalize the CE module to 4 ViT backbones. All models are trained on ImageNet-1K dataset and compared with the original model under the same training scheme. Depths indicate the number of Transformer layers of each stage. FLOPs are calculated with a resolution of 224×224

Framework	Models	Channels	Depths	Param	FLOPs	Top-1
CvT[59]	CvT-13	64	[1, 2, 10]	20.0M	4.5G	81.6
	CE-CvT-13	64	[1, 2, 10]	20.0M	4.4G	82.1(0.5↑)
PVT[56]	PVT-S	64	[3, 4, 6, 3]	24.5M	3.8G	79.8
	CE-PVT-S	64	[3, 4, 4, 3]	22.6M	3.7G	81.1(1.3↑)
Swin[35]	Swin-T	96	[2, 2, 6, 2]	28.3M	4.5G	81.3
	CE-Swin-T	96	[2, 2, 4, 2]	23.4M	4.3G	82.5(1.2↑)
	Swin-S	96	[2, 2, 18, 2]	50.0M	8.7G	83.0
	CE-Swin-S	96	[2, 2, 16, 2]	48.2M	8.8G	83.6(0.6↑)
	Swin-B	128	[2, 2, 18, 2]	88.0M	15.4G	83.3
	CE-Swin-B	128	[2, 2, 16, 2]	85.2M	15.5G	84.0(0.7↑)
CSWin[18]	CSWin-T	64	[1, 2, 21, 1]	23.0M	4.3G	82.7
	CE-CSWin-T	64	[1, 2, 20, 1]	21.6M	4.2G	83.6(0.9↑)
	CSWin-S	64	[2, 4, 32, 2]	35.0M	6.9G	83.6
	CE-CSWin-S	64	[2, 4, 31, 2]	33.9M	6.6G	84.1(0.5↑)
	CSWin-B	96	[2, 4, 32, 2]	78.0M	15.0G	84.2
	CE-CSWin-B	96	[2, 4, 31, 2]	75.8M	14.7G	84.7(0.5↑)

5.4 Understanding the Role of CNNs in Hybrid CNNs/ViT’s Design

Finally, we explore how well CNNs in the deep layer of the hybrid CNNs/ViT’s network improves ViTs. Previous works [60, 14, 39] show the *shallow* CNNs structure is enough to bring the convolutional inductive bias to all following ViTs blocks. However, one may notice that the CE^{2nd} , CE^{3rd} , and CE^{4th} are **not** locate the *shallow* layer of network. To fully understand: 1) whether CNNs in the deep layer enhances the inductive bias for subsequent ViTs blocks; 2) how hybrid CNNs/ViT’s design affects the final performance of the network. We conduct the following experiments. From macro view, CETNet can be view as ‘C-T-C-T-C-T-C-T’, where C and T denote CE and ViTs blocks respectively, where CE^{1st} is Fused-MBConv, CE^{2nd} , CE^{3rd} , and CE^{4th} are MBConv. We conduct three main experiments: **CNNs to ViTs**, **ViTs to CNNs**, and **Others**. In **CNNs to ViTs** group, we gradually replace the convolutions with transformers. In **ViTs to CNNs** group, we do the reverse. As we can see, only adopting CNNs in *early* stage is not optimal. In addition, all hybrid models outperform the pure ViTs model in **CNNs to ViTs**. Besides, in comparison with **ViTs to CNNs**, one may notice that in deep layer architecture with more ViTs is superior to more CNNs. In addition, we have: hybrid CNNs/ViT’s \geq pure ViT \geq pure CNNs, in deep layer of network. In **Others** group, we further list some variants’ experiment results to the audience and hope that any possible insights may raise a rethinking of the hybrid CNNs/ViT’s network design.

Table 8. Comparison of different hybrid CNNs/ViT’s designs. ‘Arch’ represents architecture for short. C represents MBConvs(Fused-MBConvs in the early stage), and T represents the ViTs block mentioned in section 3.2

CNNs to ViTs			ViTs to CNNs			Others		
Arch	Param	FLOPs Top-1	Arch	Param	FLOPs Top-1	Arch	Param	FLOPs Top-1
C-T-C-T-C-T-C-T	23.4M	4.3G 82.7	C-T-C-T-C-T-C-T	23.4M	4.3G 82.7	C-T-C-T-C-T-C-T	23.4M	4.3G 82.7
C-T-C-T-C-T-T-T	24.0M	4.2G 82.8	C-T-C-T-C-T-C-C	23.7M	4.2G 82.0	C-C-C-T-C-T-C-T	23.5M	4.3G 82.7
C-T-C-T-T-T-T-T	24.1M	4.2G 82.5	C-T-C-T-C-C-C-C	24.4M	4.2G 79.6	C-C-C-T-T-T-T-T	25.5M	4.4G 81.8
C-T-T-T-T-T-T-T	24.1M	4.4G 82.3	C-T-C-C-C-C-C-C	24.3M	4.2G 79.2	T-T-T-T-C-C-C-C	23.4M	4.8G 76.3
T-T-T-T-T-T-T-T	24.3M	4.2G 80.1	C-C-C-C-C-C-C-C	24.6M	5.1G 79.0	T-C-T-C-T-C-T-C	24.5M	4.2G 79.8

6 Conclusions

This paper proposes a principled way to produce a hybrid CNNs/ViT’s architecture. With the idea of injecting desirable inductive bias in ViTs, we present 1) a conceptual understanding of combining CNNs/ViT’s into a single architecture, based on using a *convolutional embedding* and its effect on the inductive bias of the architecture. 2) a conceptual framework of micro and macro detail of an hybrid architecture, where different design decisions are made at the small and large levels of detail to impose an inductive bias into the architecture. Besides, we deliver a family of models, dubbed CETNets, which serve as a generic vision backbone and achieve the SOTA performance on various vision tasks under constrained data size. We hope that what we found could raise a rethinking of the network design and extend the limitation of the hybrid CNNs/ViT’s network.

References

1. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
2. Battaglia, P.W., Hamrick, J.B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al.: Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261 (2018)
3. Bello, I., Zoph, B., Vaswani, A., Shlens, J., Le, Q.V.: Attention augmented convolutional networks. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3286–3295 (2019)
4. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. arXiv preprint arXiv:2005.14165 (2020)
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European Conference on Computer Vision. pp. 213–229. Springer (2020)
6. Chen, C.F., Fan, Q., Panda, R.: Crossvit: Cross-attention multi-scale vision transformer for image classification. arXiv preprint arXiv:2103.14899 (2021)
7. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* **40**(4), 834–848 (2017)
8. Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C.: Twins: Revisiting spatial attention design in vision transformers. arXiv preprint arXiv:2104.13840 (2021)
9. Chu, X., Tian, Z., Zhang, B., Wang, X., Wei, X., Xia, H., Shen, C.: Conditional positional encodings for vision transformers. arXiv preprint arXiv:2102.10882 (2021)
10. Cordonnier, J.B., Loukas, A., Jaggi, M.: On the relationship between self-attention and convolutional layers. arXiv preprint arXiv:1911.03584 (2019)
11. Dai, J., Li, Y., He, K., Sun, J.: R-fcn: Object detection via region-based fully convolutional networks. *Advances in neural information processing systems* **29** (2016)
12. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 764–773 (2017)
13. Dai, Z., Cai, B., Lin, Y., Chen, J.: Up-detr: Unsupervised pre-training for object detection with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1601–1610 (2021)
14. Dai, Z., Liu, H., Le, Q.V., Tan, M.: Coatnet: Marrying convolution and attention for all data sizes. arXiv preprint arXiv:2106.04803 (2021)
15. d’Ascoli, S., Touvron, H., Leavitt, M., Morcos, A., Biroli, G., Sagun, L.: Convit: Improving vision transformers with soft convolutional inductive biases. arXiv preprint arXiv:2103.10697 (2021)
16. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition (CVPR). pp. 248–255. Ieee (2009)
17. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)

18. Dong, X., Bao, J., Chen, D., Zhang, W., Yu, N., Yuan, L., Chen, D., Guo, B.: Cswin transformer: A general vision transformer backbone with cross-shaped windows. arXiv preprint arXiv:2107.00652 (2021)
19. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
20. Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., Feichtenhofer, C.: Multiscale vision transformers. arXiv preprint arXiv:2104.11227 (2021)
21. Gupta, S., Tan, M.: Efficientnet-edgetpu: Creating accelerator-optimized neural networks with automl. Google AI Blog **2**, 1 (2019)
22. Han, K., Wang, Y., Tian, Q., Guo, J., Xu, C., Xu, C.: Ghostnet: More features from cheap operations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1580–1589 (2020)
23. Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. arXiv preprint arXiv:2103.00112 (2021)
24. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
25. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
26. Heo, B., Yun, S., Han, D., Chun, S., Choe, J., Oh, S.J.: Rethinking spatial dimensions of vision transformers. arXiv preprint arXiv:2103.16302 (2021)
27. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
28. Hu, H., Zhang, Z., Xie, Z., Lin, S.: Local relation networks for image recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3464–3473 (2019)
29. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7132–7141 (2018)
30. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)
31. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
32. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* **25**, 1097–1105 (2012)
33. Li, Y., Zhang, K., Cao, J., Timofte, R., Van Gool, L.: Localvit: Bringing locality to vision transformers. arXiv preprint arXiv:2104.05707 (2021)
34. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
35. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
36. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. arXiv preprint arXiv:2201.03545 (2022)

37. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3431–3440 (2015)
38. Luo, W., Li, Y., Urtasun, R., Zemel, R.: Understanding the effective receptive field in deep convolutional neural networks. *Advances in neural information processing systems* **29** (2016)
39. Marquardt, T.P., Jacks, A., Davis, B.L.: Token-to-token variability in developmental apraxia of speech: Three longitudinal case studies. *Clinical Linguistics & Phonetics* **18**(2), 127–144 (2004)
40. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding by generative pre-training (2018)
41. Radosavovic, I., Kosaraju, R.P., Girshick, R., He, K., Dollár, P.: Designing network design spaces. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10428–10436 (2020)
42. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28**, 91–99 (2015)
43. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4510–4520 (2018)
44. Shen, Z., Zhang, M., Zhao, H., Yi, S., Li, H.: Efficient attention: Attention with linear complexities. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3531–3539 (2021)
45. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
46. Srinivas, A., Lin, T.Y., Parmar, N., Shlens, J., Abbeel, P., Vaswani, A.: Bottleneck transformers for visual recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16519–16529 (2021)
47. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. arXiv preprint arXiv:2105.05633 (2021)
48. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: Proceedings of the IEEE international conference on computer vision. pp. 843–852 (2017)
49. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1–9 (2015)
50. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International Conference on Machine Learning. pp. 6105–6114. PMLR (2019)
51. Tan, M., Le, Q.V.: Efficientnetv2: Smaller models and faster training. arXiv preprint arXiv:2104.00298 (2021)
52. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning. pp. 10347–10357. PMLR (2021)
53. Vaswani, A., Ramachandran, P., Srinivas, A., Parmar, N., Hechtman, B., Shlens, J.: Scaling local self-attention for parameter efficient visual backbones. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12894–12904 (2021)
54. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Advances in neural information processing systems*. pp. 5998–6008 (2017)

55. Wang, H., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: Max-deeplab: End-to-end panoptic segmentation with mask transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5463–5474 (2021)
56. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. arXiv preprint arXiv:2102.12122 (2021)
57. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7794–7803 (2018)
58. Wu, B., Xu, C., Dai, X., Wan, A., Zhang, P., Yan, Z., Tomizuka, M., Gonzalez, J., Keutzer, K., Vajda, P.: Visual transformers: Token-based image representation and processing for computer vision. arXiv preprint arXiv:2006.03677 (2020)
59. Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L.: Cvt: Introducing convolutions to vision transformers. arXiv preprint arXiv:2103.15808 (2021)
60. Xiao, T., Dollar, P., Singh, M., Mintun, E., Darrell, T., Girshick, R.: Early convolutions help transformers see better. *Advances in Neural Information Processing Systems* **34** (2021)
61. Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 418–434 (2018)
62. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1492–1500 (2017)
63. Xu, W., Xu, Y., Chang, T., Tu, Z.: Co-scale conv-attentional image transformers. arXiv preprint arXiv:2104.06399 (2021)
64. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. arXiv preprint arXiv:1511.07122 (2015)
65. Yuan, K., Guo, S., Liu, Z., Zhou, A., Yu, F., Wu, W.: Incorporating convolution designs into visual transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 579–588 (2021)
66. Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token vit: Training vision transformers from scratch on imagenet. arXiv preprint arXiv:2101.11986 (2021)
67. Zhang, P., Dai, X., Yang, J., Xiao, B., Yuan, L., Zhang, L., Gao, J.: Multi-scale vision longformer: A new vision transformer for high-resolution image encoding. arXiv preprint arXiv:2103.15358 (2021)
68. Zhang, Q., Yang, Y.: Rest: An efficient transformer for visual recognition (2021)
69. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6848–6856 (2018)
70. Zhang, X., Xu, H., Mo, H., Tan, J., Yang, C., Wang, L., Ren, W.: Dcnas: Densely connected neural architecture search for semantic image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13956–13967 (2021)
71. Zheng, M., Gao, P., Zhang, R., Li, K., Wang, X., Li, H., Dong, H.: End-to-end object detection with adaptive clustering transformer. arXiv preprint arXiv:2011.09315 (2020)
72. Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torrallba, A.: Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision* **127**(3), 302–321 (2019)

73. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)