

ParC-Net: Position Aware Circular Convolution with Merits from ConvNets and Transformer

Haokui Zhang, Wenze Hu, and Xiaoyu Wang

Intellifusion, Shenzhen, China

hkzhang1991@mail.nwpu.edu.cn, windsor.hwu@gmail.com, fanghuaxue@gmail.com

Abstract. Recently, vision transformers started to show impressive results which outperform large convolution based models significantly. However, in the area of small models for mobile or resource constrained devices, ConvNet still has its own advantages in both performance and model complexity. We propose ParC-Net, a pure ConvNet based backbone model that further strengthens these advantages by fusing the merits of vision transformers into ConvNets. Specifically, we propose *position aware circular convolution* (ParC), a light-weight convolution op which boasts a global receptive field while producing location sensitive features as in local convolutions. We combine the ParCs and squeeze-excitation ops to form a meta-former like model block, which further has the attention mechanism like transformers. The aforementioned block can be used in plug-and-play manner to replace relevant blocks in ConvNets or transformers. Experiment results show that the proposed ParC-Net achieves better performance than popular light-weight ConvNets and vision transformer based models in common vision tasks and datasets, while having fewer parameters and faster inference speed. For classification on ImageNet-1k, ParC-Net achieves 78.6% top-1 accuracy with about 5.0 million parameters, saving 11% parameters and 13% computational cost but gaining 0.2% higher accuracy and 23% faster inference speed (on ARM based Rockchip RK3288) compared with MobileViT, and uses only $0.5\times$ parameters but gaining 2.7% accuracy compared with DeiT. On MS-COCO object detection and PASCAL VOC segmentation tasks, ParC-Net also shows better performance. Source code is available at <https://github.com/hkzhang91/ParC-Net>

Keywords: Light-weight; edge devices; pure ConvNet; vision transformer

1 Introduction

Recently, various vision transformers (ViTs) models have achieved remarkable results in many vision tasks, forming strong alternatives to convolutional neural networks (ConvNets) [5] [33] [21].

However, we believe both ViTs and ConvNets are indispensable for the following reasons: 1) From application perspective, both ViTs and ConvNets have their advantages and disadvantages. ViT models generally have better performance but usually suffer from high computational cost and are difficult to train [33].

Compared with ViTs, ConvNets may show inferior performance, but they still have some unique advantages. For instance, ConvNets have better hardware support and are easy to train. In addition, as is summarized in [9] and our experiments, ConvNets still dominate in the area of small models for mobile or edge devices. 2) From the information processing perspective, both ViTs and ConvNets have unique features. ViTs are good at extracting global information and use attention mechanism to extract information from different locations driven by input data [3] [25]. ConvNets focus on modeling local relationships and have strong prior by inductive bias [4]. The above analysis naturally raise a question: *can we learn from ViTs to improve ConvNets for mobile or edge computing applications?*.

In this paper, we aim to design new light-weight pure ConvNets that further enhance its strength in the area of mobile and edge computing friendly models.

Pure convolution is more mobile friendly because convolutions are highly optimized by existing tool chains that are widely used to deploy model into these resource constrained devices. Even more, because of the huge popularity of ConvNets in the past few years, some existing neural network accelerators are designed mainly around convolution style operations, and the complex non-linear operations such as softmax and data bus bandwidth demanding large matrix multiplications are not efficiently supported. These hardware and software constraints make a pure convolutional light-weight model more preferable even if a ViT based model is equally competitive in other aspects.

To design such a ConvNet, we compare ConvNets with ViTs and summarize three main differences between them: 1) ViTs are good at extracting global features [3] [25] [4]; 2) ViTs adopt Meta-former block [40]; 3) Information aggregations in ViTs are data driven (data dependent dynamic computation). Corresponding to these three points, we design our ParC block. 1) We propose the position aware circular convolution (ParC) to extract global features; 2) Based on the proposed ParC, we build a pure ConvNet Meta-former block as the basic outer structure; 3) We add channel wise attention module to the feature forward network (FFN) part of meta-former, which makes our proposed ParC block adapt kernel weights according inputs. Finally, inspired by CoatNet [4] and MobileViT [25], we use a bifurcate structure (section 3.2) as the outer frame to build a complete network ParC-Net.

Experiment results show that the proposed ParC-Net achieves solid performance on three popular vision tasks, including image classification, object detection and semantic segmentation. Taking experiment results of image classification as an example, ParC-Net achieves 78.6% top-1 accuracy with about 5.0 million parameters, saving 11% parameters and 13% computational cost but gaining 0.2% higher accuracy and 23% faster inference speed (on Rockchip RK3288) compared with MobileViT [25]. For experiments of object detection and semantic segmentation, compared with other light-weight models, the proposed ParC-Net achieves higher mAP and mIOU, while having fewer parameters.

Our main contributions are summarized as follows:

- To overcome the restriction that traditional convolutions have limited perception fields, we propose position aware circular convolution (ParC), where base-instance kernel and position embedding strategies are used to handle input size variations and inject location information to output feature maps respectively. We jointly use the proposed ParC and conventional convolution operations to extract local-global features, which brings higher accuracy.
- We propose ParC-Net, a pure ConvNet for mobile and edge computing applications. The proposed ParC-Net inherits advantages of ConvNets and ViTs. To our knowledge, this is the first attempt that combines strengths of ConvNets and ViTs to design a light-weight ConvNet.
- We apply the proposed ParC-Net on three vision tasks. Compared with the baseline model, the proposed ParC-Net achieves better performance on all three tasks, while having fewer parameters, lower computational cost and higher inference speed.

2 Related work

2.1 Vision transformers

Vaswani et al. firstly proposed transformer [34] for natural language processing (NLP) tasks. Compared with recurrent neural network (RNN) models, transformer has much higher computational efficiency and it is good at capturing relationship from any pair of elements in the input sequence. As a result, transformers replaced RNNs and dominate the NLP field.

In 2020, Dosovitskiy et al. introduced transformer into vision tasks and proposed vision transformer (ViT) [5], where each image is cropped into a sequence of patches to meet the input requirement of transformer and PE is adopted to ensure the model is sensitive to position information of the input patches. With pre-training on huge datasets such as JFT-300M [29], ViT achieves impressive performance on various vision tasks. However, the original ViT model has some restrictions, for instance, it is heavy-weight, having low computational efficiency and hard to train. Subsequent variants of ViTs are proposed to overcome these problems. From the point of improving training strategy, Touvron et al. [33] proposed to use knowledge distillation to train ViT models, and achieved competitive accuracy with less pre-training data. To further improve the model architecture, some researchers attempted to optimize ViTs by learning from ConvNets. Among them, PVT [35] and CVT [37] insert convolutional operations into each stage of ViT model to reduce the number of tokens, and build hierarchical multi-stage structures. Swin transformer [21] computes self attention within shifted local windows. PiT [11] jointly use pooling layer and depth wise convolution layer to achieve channel multiplication and spatial reduction. CCNet [15] propose a simplified version of self attention mechanism criss-cross attention and inserted it into ConvNet to build ConvNet which has global receptive field. These papers clearly show that some techniques of ConvNets can be applied on vision transformers to design better vision transformer models.

2.2 Hybrid structures combining ConvNet and vision transformers

Another popular line of research is combining elements of ViTs and ConvNets to design new backbones. Graham et al. mixed ConvNet and transformer in their LeViT model, which significantly outperforms previous ConvNet and ViT models with respect to the speed/accuracy tradeoff [8]. BoTNet [28] replaces the standard convolution with multi-head attention in the last several blocks of ResNet. ViT-C [38] adds early convolutional stem to vanilla ViT. ConViT [6] incorporates soft convolutional inductive biases via a gated positional self-attention. The CMT [9] block consists of depth wise convolution based local perception unit and a light-weight transformer module. CoatNet [4] merges convolution and self-attention to design a new transformer module, which focuses on both local and global information. After comprehensive comparison, we find that these hybrid models simultaneously employed similar structure, that is using convolutional stem to extract local features in the beginning stages and transformer style models later to extract global or local-global features. We choose a similar structure when designing our pure convolutional model.

2.3 Light-weight ConvNets and ViTs

Since 2017, light-weight ConvNets attract much attentions as more and more applications needs to run ConvNet models on mobile devices. Now, there are a lot of light-weight ConvNets, such as ShuffleNets [24] [24], MobileNets [13] [27] [12], MicroNet [18], GhostNet [10], EfficientNet [32], TinyNet [2] and MnasNet [31]. Compared with standard ConvNets, light-weight ConvNets have fewer parameters, lower computational cost and faster inference speed. In addition, light-weight ConvNets can be applied on a wide range of devices. Despite these benefits, these light-weight models have inferior performance compared with heavy-weight models. Very recently, following the research line of combining strengths of ConvNet and ViT, some researcher attempted to build light-weight hybrid models for mobile vision tasks. MobileFormer presents a parallel design of MobileNet and transformer, which leverages the advantages of MobileNet at extracting local features and transformer at capturing global information [3]. Mehta and Rastegari proposed MobileViT, where the upper stages of MobileNetv2 [27] are replaced with MobileViT block [25]. In MobileViT block, local representations extracted by convolution and global representations are concatenated to generate local-global representations.

In terms of purpose, our proposed ParC-Net is related to MobileFormer and MobileViT. Different from these two models which still keep transformer blocks, our proposed ParC-Net is pure ConvNet, which makes our proposed ParC-Net more mobile friendly. Our experiments of deploying models on low power platform confirm this point. In terms of designing a pure ConvNet via learning from ViTs, our work is most closely related to a parallel work ConvNext [22]. The two major differences are: 1) Ideas and architectures are different. The ConvNext modernizes a standard ResNet toward the design of a vision transformer by introducing a series (more than ten) of incremented but effective designs.

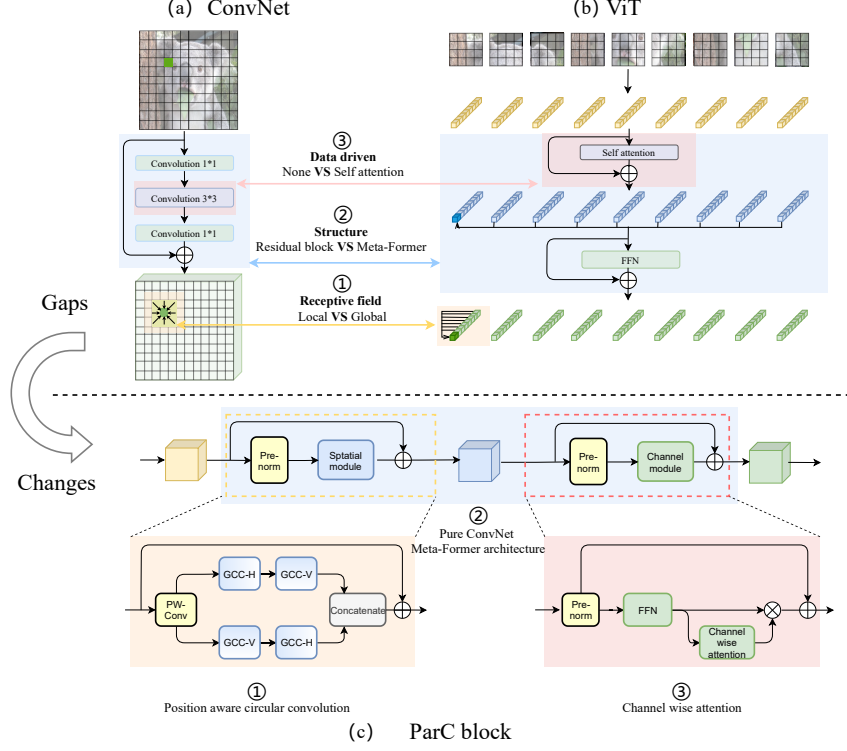


Fig. 1. ParC block. (a) A residual block that is widely used in ConvNets; (b) A ViT block; (c) An ParC block

Our proposed ParC-Net starts from three main differences between ConvNets and ViTs and fills the gaps from macro level. As the ideas are different, the corresponding structures are also different; 2) They are proposed for different purposes. Our ParC-Net is proposed for mobile devices. Compared with ConvNext, the proposed ParC-Net shows advantages when constraining models as light-weight models.

3 The proposed method

In this section, we will introduce our ParC-Net in two parts, the details of the building block (ParC block) and the overall model structure (ParC-Net).

3.1 ParC block

The upper part of Figure 1 shows three major differences between common ConvNets and ViTs. The bottom half of Figure 1 illustrates the architecture of

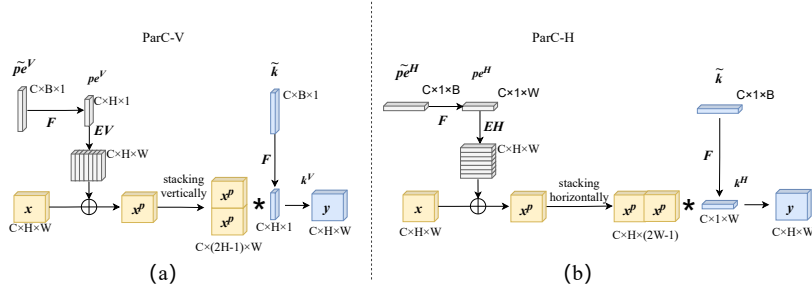


Fig. 2. Illustration of the position aware circular convolution. (a) ParC-V; (b) ParC-H. F , EV and EH are explained in equations 1 and 2

our proposed ParC block. In the following, we will explain the motivation and the specific structure of each component of the proposed ParC block.

Extracting global features with ParC. In ConvNets, feature is calculated as $y_i = \sum_{j \in \mathcal{L}(i)} w_{i-j} x_j$, where x_i , y_i are the input and output at position i respectively, and $\mathcal{L}(i)$ denotes a local neighborhood of i . In ViTs, self-attention modules extracts features based on formula $y_i = \sum_{j \in \mathcal{G}} \frac{e^{(x_i^T x_j)}}{\sum_{k \in \mathcal{G}} e^{(x_i^T x_k)}} x_j$, where \mathcal{G} means the global spatial space. Comparing these two formulas, we can see that self attention learns global features from the entire spatial locations but convolution gathers information from a local receptive field.

To overcome this issue, we propose the position aware circular convolution (ParC). As shown in Figure 2, our proposed ParC has two types, one is ParC of vertical direction (ParC-V) and the other one is ParC of horizontal direction (ParC-H). The receptive field of the ParC-V and ParC-H covers all pixel in the same column and the same row, respectively. Jointly using ParC-V and ParC-H can extract global features from all input pixels. For notational simplicity, we assume the input x has only one channel and the corresponding shape is $1 \times h \times w$. The output of ParC-V at location (i, j) is computed with:

$$\begin{aligned}
 pe^V &= F(\tilde{pe}^V) = [pe_0^V, pe_1^V, \dots, pe_{h-1}^V]^T \\
 pe_e^V &= EV(pe^V, w) \\
 k^V &= F(\tilde{k}) = [k_0^V, k_1^V, \dots, k_{h-1}^V] \\
 x^p &= x + pe_e^V \\
 y_{i,j} &= \sum_{t \in (0, h-1)} k_t^V x_{((i+t) \bmod h, j)}^p
 \end{aligned} \tag{1}$$

where, pe^V is instance position embedding (PE) and it is generated from a base embedding \tilde{pe}^V via bilinear interpolation function $F()$. Here $F()$ is used to adapt the size of position embedding to the size of input features. pe_e^V is expanded PE. k^V is instance kernel. $EV()$ is an expand function of vertical direction. After copying the input vector w times, $EV()$ concatenates these copied vectors along




Fig. 3. Illustration of global circular convolution on horizontal direction.

horizontal direction to generate a $h \times w$ -sized PE matrix. Similarly, the output of ParC-H at location (i, j) can be expressed as:

$$\begin{aligned} x^p &= x + pe_e^H \\ y_{i,j} &= \sum_{t \in (0, w-1)} k_t^H x_{(i, (j+t) \bmod w)}^p \end{aligned} \quad (2)$$

where $pe_e^H = EH(pe^H, h)$ and $EH()$ is an expand function. $EH()$ expands input vector along the vertical direction. Implementing the ParC in modern deep learning libraries is straightforward. Taking the most complicated part $y_{i,j} = \sum_{t \in (0, w-1)} k_t^H x_{(i, (j+t) \bmod w)}^p$ as an example, it can be implemented with one line of code: $y = F.conv2D(torch.cat(x^p, x^p, dim = 3), k^H)$. Figure 3 illustrates the computational process in the case that the input is an one dimensional vector. From Figure 3, we can see that ParC-H perform convolutions along a circle generated by connecting the start and the end of the input. So, we name the proposed convolution as the circular convolution. The proposed ParC introduces three modifications:

- The receptive field is increased to global spatial space. Note that, increasing the kernel size of tradition local convolution to full input size does not extract global features. In local convolution, zero padding is usually used to keep the size of convolutional feature the same with that of the input. Even if we increase the kernel size to global size, the global kernel only covers part pixels coming from input. Especially for extracting feature in edge portion, only about half of pixels that covered by global kernel are from input actual input, while others are simply zeros.
- The PE is used to keep the output feature sensitive to spatial location. Circular convolution can extract global features, but it disturbed the spatial structure of the original input. For classification, keeping spatial structure may not be a big issue. But, as is shown in ablation study, for location sensitive tasks such as segmentation and detection, keeping spatial structure does matter. Here, following the design in ViTs, we introduce PE to keep spatial structure. Experiment results in ablation study show that PE is useful in segmentation and detection tasks
- The kernel and PE are dynamically generated according to the input size. In ParC, the sizes of kernels and PE codes must be consistent with that of instance inputs. To handle the case that inputs have different spatial resolution, we generate instance kernels and PE codes via interpolation functions.




Fig. 4. Three main hybrid structures. (a) serial structure; (b) parallel structure; (c) bifurcate structure

Designing ParC block with ParC. From ConvNets to ViTs, a considerable modification is meta-former block replaced residual block (the blue two-way arrow). A Meta-former block generally consists of a sequence of two components: a token mixer and a channel mixer. The token mixer is for exchanging information among tokens in different spatial locations. The channel mixer is for mixing information among different channels. Both two components use residual learning structure.

Inspired by this, we insert ParC into Meta-former like block to build our ParC block. Specifically, we replace self-attention module with the proposed ParC to build a new spatial module to replace token mixer part. Here, we do this for two main reasons: 1) ParC can extract global features and interacts information among pixels from global space, which meets the requirement of token mixer module; 2) the computation complexity of self attention module is quadratic. Replacing this part with ParC can reduce computational cost significantly, which helps achieving our goal of designing a light-weight ConvNet. Based on the proposed ParC, we build a pure ConvNet meta-former like block.

Adding channel wise attention in channel mixer part. In ViTs, self attention module can adapt weights according input, which makes ViTs data driven models. By adopting attention mechanism, data driven models can focus on important features and suppress unnecessary ones, which brings better performance. Previous literature [14][36][16] already explained the importance of keep model data driven.

By replacing the self-attention with the proposed global circular convolution, we get a pure ConvNet which can extract global features. But the replaced model is no longer a data driven model. To compensate, we insert channel wise attention module into channel mixer part, as shown in Figure 1(c). Following SENet [14], we first aggregate spatial information of input features $x \in \mathbb{R}^{c \times h \times w}$ via global average pooling and get aggregated feature $x_a \in \mathbb{R}^{c \times 1 \times 1}$; Then we feed x_a into a multi-layer perceptron to generate channel wise weight $a \in \mathbb{R}^{c \times 1 \times 1}$. The a is multiplied with x channel wise to generate the final output.

3.2 ParC-Net

In section 3.1, we have presented the ParC block, which is a basic block and can be inserted into most of the current existing models. In this section, we select an outer frame for it and build the complete network ParC-Net.

Currently, as shown in Figure 4, existing hybrid structures can be basically divided into three main structures, including serial structure (Figure 4(a)) [8][38], parallel structure (Figure 4(b)) [3] and bifurcate structure (Figure 4(c)) [25][4]. Among all three structures, the third one achieves best performance for now. At present, bifurcate model CoatNet [4] achieves the highest classification accuracy on Imagenet-1k. Mobile device aimed model MobileViT [25] also adopts the third structure.

Inspired by this, we adopt bifurcate structure as our outer frame and build our final outer frame based on MobileViT. Specifically, taking the outer frame adopted in MobileViT as baseline, we further make some improvements:

- MobileViT consists of two major types of modules. Shallow stages consist of MobileNetV2 blocks, which have local receptive field. Deep stages are made up of ViT blocks, which enjoy global receptive field. We keep all MobileNetV2 blocks and replacing ViT blocks with corresponding ParC blocks. This replacement converts the model from hybrid structure to pure ConvNet.
- We appropriately increase the widths of ParC blocks. Even so, the replaced model still has fewer parameters and less computational cost.
- As show in Figure 4(c), the bifurcate structure contains some interaction modules, which are in charge of interacting information between local and global feature modules. In the original MobileViT, ViT blocks are the most heavy modules. After replacing ViT blocks with ParC blocks, the cost of these interaction modules becomes prominent. So, we introduce group convolution and point wise convolution into these modules, which decreases number of parameters without hurting performance.

4 Experiment results

In experiments, we show the overall advantages of the proposed ParC-Net on three typical vision tasks, and then conduct detailed study to show the value of our design choices, the model scaling characteristics, and its speed advantage on low power devices.

4.1 Image classification

We conduct image classification experiments on ImageNet-1k, the most widely used benchmark dataset for this task. We train the proposed ParC-Net models on the training set of ImageNet-1K, and report top-1 accuracy on the validation set.

Training setting. As we adopt MobileViT like structure as our outer framework, we train our model using a very similar training strategy as well. To be specific, we train each model for 300 epochs on 8 V100 or A100 GPUs with AdamW optimizer [23], where the maximum learning rate, minimum learning rate, weight decay and batchsize are set to 0.004, 0.0004, 0.025 and 1024 respectively. Optimizer momentum β_1 and β_2 of the AdamW optimizer are set to

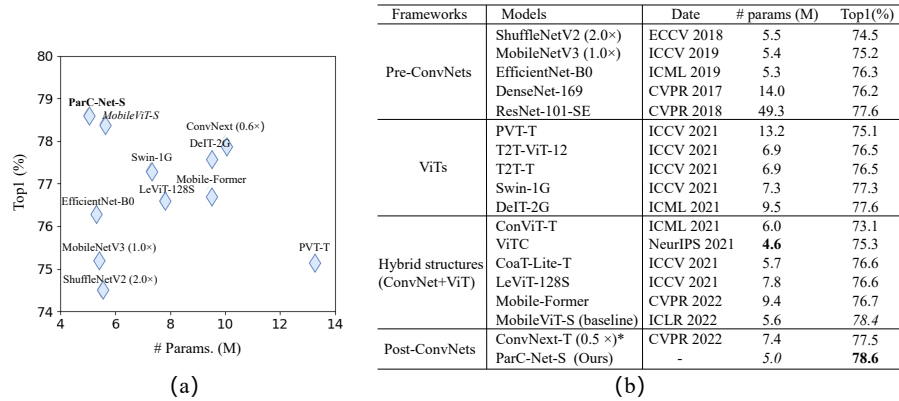


Fig. 5. Classification experiment results on ImageNet-1K. (a) Accuracy vs model size. Here we only keep part of comparison models for clarity. (b) Comparison of results on image classification. * indicates our implementation. Pre-ConvNets indicate classical ConvNets appeared before ViTs. Post-ConvNets denote ConvNets which integrate merits of ViTs but still keep pure ConvNet structures.

0.9 and 0.999 respectively. We use the first 3000 iterations as warm up stage. We adjust learning rate following the cosine schedule. For data augmentation, we use random cropping, horizontal flipping and multi-scale sampler. We use label smoothing [30] to regularize the networks and set smoothing factor to 0.1. We use Exponential Moving Average (EMA) [26]. More details of the training settings and *link to source code will be provided in supplementary materials*.

Comparison results. The experiment results of image classification are listed in Figure 5. Figure 5 (a) shows that ParC-Net-S and MobileViT-S beat other model by a clear margin. Figure 5 (b) shows comparison with more models. The proposed ParC-Net-S achieves highest classification accuracy, and have fewer parameters than most models. Compared with the second best model MobileViT-S, our ParC-Net-S decreases the number of parameters by 11% and increases the top 1 accuracy by 0.2 percentage points.

Light-weight models. Table 1 shows comparison results among light-weight models, which confirms our ideas and answers the question proposed in introduction.


Firstly, comparing results of light-weight ConvNets with that of ViTs, light-weight ConvNets show much better performance.

Secondly, comparing the popular ConvNets before ViT appears (pre-ConvNets), ViTs and hybrid structures, hybrid structures achieve the best performance. Therefore improving ConvNets by learning from the merits of ViT is feasible.

Finally, the proposed ParC-Net achieves the best performance among all comparison models. So indeed by learning from ViT design, performance of pure light-weight ConvNets can be improved significantly.

Table 1. Comparisons of light-weight models on ImageNet-1K classification

Frameworks	Models	Date	# params (M)	Top1(%)
Pre-CNNs	ShuffleNetV2(2.0×)	ECCV 2018	5.5	74.5
	MobileNetV3(1.0×)	ICCV 2019	5.4	75.2
	EfficientNet-B0	ICML 2019	5.3	76.3
ViTs	T2T-ViT-7	ICCV 2021	4.3	71.7
	DeiT-T	ICML 2021	5.7	72.2
Hybrid structures	ViT-C	NeurIPS 2021	4.6	75.3
	CoaT-Lite-T	ICCV 2021	5.7	76.6
	MobileViT-S	ICLR 2022	5.6	78.4
Post-CNN	ParC-Net-S	-	5.0	78.6



(a)

Models	# params (M)	mAP
MobileNetV3	4.9	22.0
MobileNetV2	4.3	22.1
MobileNetV1	5.1	22.2
MixNet	4.5	22.3
MNASNet	4.9	23.0
MobileViT-XS	2.7	24.8
VGG	35.6	25.1
ResNet50	22.9	25.2
MobileViT-S (baseline)	5.7	27.7
ParC-Net-S (Ours)	5.2	28.8

(b)

Fig. 6. Object detection results on MS-COCO. (a) mAP vs model size. (b) Comparison results

4.2 Object detection

We use MS-COCO [19] datasets and its evaluation protocol for object detection experiments. Following [25][27], we take single shot object detection (SSD) [20] as the detection framework and use separable convolution to replace the standard convolutions in the detection head.

Experiment setting. Taking models pretrained on ImageNet-1K as backbone, we finetune detection models on training set of MS-COCO with AdamW optimizer for 200 epochs. Batchsize and weight decay are set to 128 and 0.01. We use the first 500 iterations as warm up stage, where the learning rate is increased from 0.000001 to 0.0009. Both label smoothing and EMA are used during training.

Comparison results. Figure 6 lists the corresponding results. Similar to results in image classification, MobileViT-S and ParC-Net-S achieve the the second best and the best in terms of mAP. Compared with the second best model, ParC-Net-S shows advantages in both model size and detection accuracy.

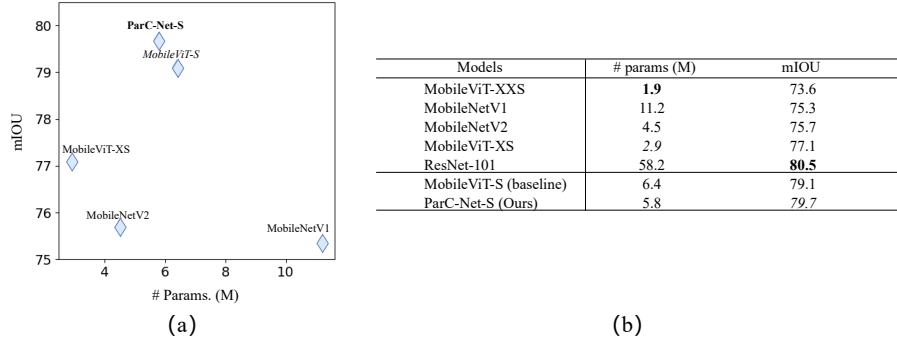


Fig. 7. Semantic segmentation experiments on PASCAL VOC. (a) mIOU vs model size. (b) Comparison results with more models.

Table 2. Ablation study. BK, MF, CA and PE denote big kernel, meta-former architecture, channel wise attention and position embedding. BK 1/4 and BK 1/2 means the kernel size is set to 1/4 and 1/2 of the input features, respectively.

Row	Task	Kernel	MF	CA	PE	# params (M)	Top1/mAP/mIOU
1	classification	Baseline	-	-	-	5.6	78.35
2	classification	BK 1/4	Y	Y	N	5.0	78.46
3	classification	BK 1/2	Y	Y	N	5.0	78.45
4	classification	ParC	N	Y	Y	5.3	76.00
5	classification	ParC	Y	N	Y	5.0	78.50
6	classification	ParC	Y	Y	N	5.0	78.63
7	classification	ParC	Y	Y	Y	5.0	78.63
8	detection	Baseline	-	-	-	5.7	27.7
9	detection	ParC	Y	Y	N	5.2	27.5
10	detection	ParC	Y	Y	Y	5.2	28.5
11	segmentation	Baseline	-	-	-	6.4	79.1
12	segmentation	ParC	Y	Y	N	5.8	79.2
13	segmentation	ParC	Y	Y	Y	5.8	79.7

4.3 Semantic segmentation

Experiment settings. DeepLabV3 is adopted as the semantic segmentation framework. We fine tune segmentation models on training set of PASCAL VOC [7] and COCO dataset, then evaluate trained models on validation set of PASCAL VOC using mean intersection over union (mIOU) and report the final results for comparison. We fine tune each model for 50 epochs with AdamW. Readers may refer to more details about training settings in supplementary materials.

Comparison results. Results are summarized in Figure 7. We can see that MobileViT-S and ParC-Net-S have the best trade-off between model scale and mIOU. Compared with ResNet-101, MobileViT-S and ParC-Net-S achieve competitive mIOU, while having much fewer parameters.

4.4 Ablation study

Using the MobileViT as a baseline model, we further conduct ablation analysis on three components proposed in our ParC-Net.

- **Position aware circular convolution.** The proposed ParC has two major characteristics: 1) Circular convolution brings global receptive field; 2) PE keeps spatial structure information. Experiment results confirm that both characteristics are important. 1) Results in rows 1-3 show that, using big kernel can also improve accuracy, but the benefit of it reaches a saturation point when kernel size reaches a certain level. This results are consistent with the statement claimed in [22]. Using ParC can further improve accuracy, as shown in rows 2-3 and 6-7. 2) Introducing PE to ParC is necessary. As we explained in section 3.1, using circular convolution alone can indeed capture global features but it disturbs the original spatial structures. For classification task, PE has no impact (rows 6 and 7). However, for detection and segmentation tasks which are sensitive to spatial location, abandoning PE hurts performances (rows 9-10 and 12-13).
- **Meta-former architecture.** In experiments of abandoning Meta-former architecture, we integrate ParC with the ResNeXt block [39] to replace Meta-former architecture. By comparing row 4 and 7, we can see that using the proposed pure ConvNet meta-former architecture is useful.
- **Channel wise attention.** Results in rows 5 and 7 show that using channel wise attention can improves performance. Compared with ParC, channel wise attention brings less benefit.

In summary, all three components are useful. Connecting them as a whole achieves the best performance.

4.5 Inference speed on low power devices.

In this section, we conduct experiments to verify two points: 1) as we mentioned in introduction, the ParC-Net is proposed for edge computing devices. To verify whether the proposed ParC-Net meets our requirements, we deploy the proposed ParC-Net on a widely used low power chip Rockchip RK3288 and an in house low power neural network processor DP2000, compare it with baseline. We use ONNX [1] and MNN[17] to port these models to chips and time each model for 100 iterations to measure the average inference speed; 2) The proposed ParC block is an plug-and-play block, it can be inserted into other models. We replaced convolutions in the last few blocks of typical CNNs with our proposed ParC (with PE and kernel generation etc.) Comparison results are listed in Table 3.

As shown in rows 1-4 of Table 3, compared with baseline, ParC-Net is 23% faster on Rockchip RK3288 and $3.77 \times$ faster On DP2000. Besides less FLOPs operations, we believe this speed improvement is also brought by two factors: 1) Convolutions are highly optimized by existing tool chains that are widely used to deploy models into these resource constrained devices; 2) Compared with

Table 3. Applying ParC-Net designs on different backbones. CPU used here is Xeon E5-2680 v4. DP2000 is the code name of a in house unpublished low power neural network processor that highly optimizes the convolutions. *denotes the models are trained under convnext hyperparameters settings, which may not be the optimal. W means network width. Latency is measured with batch size 1.

Row	Models	# param	FLOPs	Devices	Speed (ms)	Top1 (%)
1	MobileViT-S	5.6M	4.0G	RK3288	457	78.4
2	ParC-Net-S	5.0M	3.5G	RK3288	353	78.6
3	MobileViT-S	5.6M	4.0G	DP2000	368	78.4
4	ParC-Net-S	5.0M	3.5G	DP2000	98	78.6
5	ResNet50 *	26 M	4.1G	CPU	98	78.8
6	ParC-ResNet50 *	24 M	4.0G	CPU	98	79.6
7	MobileNetV2*	3.5M	0.6G	CPU	24	70.2
8	ParC-MobileNetV2*	3.5M	0.6G	CPU	27	71.1
9	ConvNext-T($0.5 \times W$)*	7.4M	1.1G	CPU	47	77.5
10	ParC-ConvNext-T($0.5 \times W$)*	7.4M	1.1G	CPU	48	78.3

convolutions, transformers are more data bandwidth demanding as computing the attention map involves two large matrices K and Q , whereas in convolutions the kernel is a rather small matrix compared with the input feature map. In case the bandwidth requirement exceeds that of the chip design, the CPU will be left idle waiting for data, resulting in lower CPU utilization and overall slower inference speed;

Results in rows 3-10 show that our ParC-Net universally improves performances of typical light weight models. MobileViT-S has much higher FLOPs but achieves good trade-off between model size and accuracy, which excels in its own application purpose. By applying our ParC-Net designs on MobileViT-S, ParC-Net-S achieves better balance between model size, FLOPs and accuracy. Results on ResNet50, MobileNetV2 and ConvNext-T shows that models which focus on optimizing FLOPs-accuracy trade-offs can also benefit from ParC-Net designs.

5 Conclusion

In this paper, for edge computing devices, we present ParC-Net, a pure ConvNet, which inherits advantages of ConvNet and integrated structure characteristics of ViT. To evaluate the performances, we apply the proposed model on three popular vision tasks, image classification, object detection and semantic segmentation. The proposed model achieves better performance on all three tasks, while having fewer parameters compared with other ConvNet, ViT and hybrid models. Experimental results on low power devices Rockchip RK3288 and our in house processor DP2000 show that the proposed ParC-Net does inherit ConvNets and it is well supported by edge computing devices.

References

1. Bai, J., Lu, F., Zhang, K., et al.: Onnx: Open neural network exchange. <https://github.com/onnx/onnx> (2019)
2. Chen, G., Wang, Y., Li, H., Dong, W.: Tinynet: a lightweight, modular, and unified network architecture for the internet of things. In: Proceedings of the ACM SIGCOMM 2019 conference posters and demos. pp. 9–11 (2019)
3. Chen, Y., Dai, X., Chen, D., Liu, M., Dong, X., Yuan, L., Liu, Z.: Mobile-former: Bridging mobilenet and transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5270–5279 (2022)
4. Dai, Z., Liu, H., Le, Q., Tan, M.: Coatnet: Marrying convolution and attention for all data sizes. *Advances in Neural Information Processing Systems* **34** (2021)
5. 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)
6. d’Ascoli, S., Touvron, H., Leavitt, M.L., Morcos, A.S., Biroli, G., Sagun, L.: Convit: Improving vision transformers with soft convolutional inductive biases. In: *International Conference on Machine Learning*. pp. 2286–2296. PMLR (2021)
7. Everingham, M., Eslami, S., Gool, L.V., Williams, C., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision* **111**(1), 98–136 (2015)
8. Graham, B., El-Nouby, A., Touvron, H., Stock, P., Joulin, A., Jégou, H., Douze, M.: Levit: a vision transformer in convnet’s clothing for faster inference. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12259–12269 (2021)
9. Guo, J., Han, K., Wu, H., Xu, C., Tang, Y., Xu, C., Wang, Y.: Cmt: Convolutional neural networks meet vision transformers. *arXiv preprint arXiv:2107.06263* (2021)
10. 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)
11. Heo, B., Yun, S., Han, D., Chun, S., Choe, J., Oh, S.J.: Rethinking spatial dimensions of vision transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 11936–11945 (2021)
12. Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al.: Searching for mobilenetv3. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1314–1324 (2019)
13. 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)
14. 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)
15. Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: Ccnet: Criss-cross attention for semantic segmentation. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 603–612 (2019)
16. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. *Advances in neural information processing systems* **28** (2015)
17. Jiang, X., Wang, H., Chen, Y., Wu, Z., Wang, L., Zou, B., Yang, Y., Cui, Z., Cai, Y., Yu, T., et al.: Mnn: A universal and efficient inference engine. *Proceedings of Machine Learning and Systems* **2**, 1–13 (2020)

18. Li, Y., Chen, Y., Dai, X., Chen, D., Liu, M., Yuan, L., Liu, Z., Zhang, L., Vasconcelos, N.: Micronet: Improving image recognition with extremely low flops. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 468–477 (2021)
19. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Zitnick, C.L.: Microsoft coco: Common objects in context. Springer International Publishing (2014)
20. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. Springer, Cham (2016)
21. 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*. pp. 10012–10022 (2021)
22. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11976–11986 (2022)
23. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *ICLR* (2019)
24. Ma, N., Zhang, X., Zheng, H.T., Sun, J.: Shufflenet v2: Practical guidelines for efficient cnn architecture design. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 116–131 (2018)
25. Mehta, S., Rastegari, M.: Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer. *ICLR* (2022)
26. Polyak, B.T., Juditsky, A.B.: Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization* **30**(4), 838–855 (1992)
27. 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)
28. 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)
29. 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)
30. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2818–2826 (2016)
31. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: Mnasnet: Platform-aware neural architecture search for mobile. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2820–2828 (2019)
32. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: *International conference on machine learning*. pp. 6105–6114. PMLR (2019)
33. 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)
34. 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)
35. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction with-

- out convolutions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 568–578 (2021)
36. Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: Proceedings of the European conference on computer vision (ECCV). pp. 3–19 (2018)
 37. Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L.: Cvt: Introducing convolutions to vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22–31 (2021)
 38. 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)
 39. 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)
 40. Yu, W., Luo, M., Zhou, P., Si, C., Zhou, Y., Wang, X., Feng, J., Yan, S.: Metaformer is actually what you need for vision. *arXiv preprint arXiv:2111.11418* (2021)