

## Mobile-Former: Bridging MobileNet and Transformer

Yinpeng Chen<sup>1</sup>   Xiyang Dai<sup>1</sup>   Dongdong Chen<sup>1</sup>   Mengchen Liu<sup>1</sup>   Xiaoyi Dong<sup>2</sup>  
Lu Yuan<sup>1</sup>   Zicheng Liu<sup>1</sup>

<sup>1</sup> Microsoft

<sup>2</sup> University of Science and Technology of China

{yiche, xidai, dochen, mengcliu, luyuan, zliu}@microsoft.com,

dlight@mail.ustc.edu.cn

### Abstract

We present *Mobile-Former*, a parallel design of *MobileNet* and transformer with a two-way bridge in between. This structure leverages the advantages of *MobileNet* at local processing and transformer at global interaction. And the bridge enables bidirectional fusion of local and global features. Different from recent works on vision transformer, the transformer in *Mobile-Former* contains very few tokens (e.g. 6 or fewer tokens) that are randomly initialized to learn global priors, resulting in low computational cost. Combining with the proposed light-weight cross attention to model the bridge, *Mobile-Former* is not only computationally efficient, but also has more representation power. It outperforms *MobileNetV3* at low FLOP regime from 25M to 500M FLOPs on ImageNet classification. For instance, *Mobile-Former* achieves 77.9% top-1 accuracy at 294M FLOPs, gaining 1.3% over *MobileNetV3* but saving 17% of computations. When transferring to object detection, *Mobile-Former* outperforms *MobileNetV3* by 8.6 AP in *RetinaNet* framework. Furthermore, we build an efficient end-to-end detector by replacing backbone, encoder and decoder in *DETR* with *Mobile-Former*, which outperforms *DETR* by 1.3 AP but saves 52% of computational cost and 36% of parameters. Code will be released at <https://github.com/aaboys/mobileformer>.

### 1. Introduction

Recently, vision transformer (ViT) [10, 34] demonstrates the advantage of global processing and achieves significant performance boost over CNNs. However, when constraining the computational budget within 1G FLOPs, the gain of ViT diminishes. If we further challenge the computational cost, *MobileNet* [16, 17, 28] and its extensions [12, 19] still dominate their backyard (e.g. fewer than 300M FLOPs for ImageNet classification) due to their efficiency in local processing filters via decomposition of depthwise and pointwise convolution. This in turn naturally raises a question:

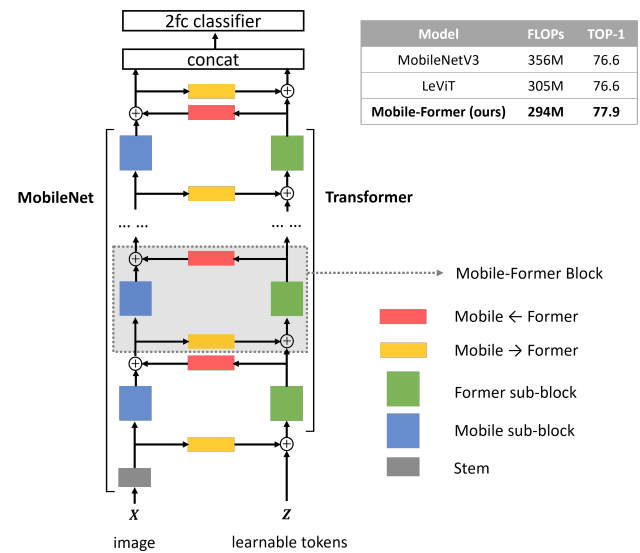


Figure 1. **Overview of Mobile-Former**, which parallelizes *MobileNet* [28] on the left side and *Transformer* [38] on the right side. Different from vision transformer [10] that uses image patches to form tokens, the transformer in *Mobile-Former* takes *very few learnable tokens* as input that are randomly initialized. *Mobile* (refers to *MobileNet*) and *Former* (refers to transformer) communicate through a bidirectional bridge, which is modeled by the proposed light-weight cross attention. Best viewed in color.

*How to design efficient networks to effectively encode both local processing and global interaction?*

A straightforward idea is to combine convolution and vision transformer. Recent works [11, 40, 41] show the benefit of combining convolution and vision transformer in *series*, either using convolution at the beginning or intertwining convolution into each transformer block.

In this paper, we shift the design paradigm from *series* to *parallel*, and propose a new network that parallelizes *MobileNet* and transformer with a two-way bridge in between (see Figure 1). We name it *Mobile-Former*, where *Mobile* refers to *MobileNet* and *Former* stands for transformer. *Mo-*




Figure 2. **Comparison among Mobile-Former, efficient CNNs and vision transformers**, in terms of accuracy over FLOPs. The comparison is performed on ImageNet classification. Mobile-Former consistently outperforms both efficient CNNs and vision transformers in low FLOP regime (from 25M to 500M MAAdd). Note that we implement Swin [23] and DeiT [34] at low computational budget from 100M to 2G FLOPs. Best viewed in color.

*mobile* takes an image as input and stacks mobile (or inverted bottleneck) blocks [28]. It leverages the efficient depthwise and pointwise convolution to extract local features. *Former* takes a few learnable tokens as input and stacks multi-head attention and feed-forward networks (FFN). These tokens are used to encode global features of the image.

*Mobile* and *Former* communicate through a two-way bridge to fuse local and global features. This is crucial since it feeds local features to *Former's* tokens as well as introduces global views to every pixel of feature map in *Mobile*. We propose a light-weight cross attention to model this bidirectional bridge by (a) performing the cross attention at the bottleneck of *Mobile* where the number of channels is low, and (b) removing projections on query, key and value ( $W^Q$ ,  $W^K$ ,  $W^V$ ) from *Mobile* side.

This parallel structure with a bidirectional bridge leverages the advantages of both MobileNet and transformer. Decoupling of local and global features in parallel leverages MobileNet's efficiency in extracting local features as well as transformer's power in modeling global interaction. More importantly, this is achieved in an efficient way via a thin transformer with very few tokens and a light-weight bridge to exchange local and global features between *Mobile* and *Former*. The bridge and *Former* consume less than 20% of the total computational cost, but significantly improve the representation capability. This showcases an efficient and effective implementation of part-whole hierarchy [15].

Mobile-Former achieves solid performance on both image classification and object detection. For example, it achieves 77.9% top-1 accuracy on ImageNet classification at 294M FLOPs, outperforming MobileNetV3 [16] and LeViT [11] by a clear margin (see Figure 1). More importantly, Mobile-Former consistently outperforms both ef-

ficient CNNs and vision transformers from 25M to 500M FLOPs (see Figure 2), showcasing the usage of transformer at the low FLOP regime where efficient CNNs dominate.

When transferring from image classification to object detection, Mobile-Former significantly outperforms MobileNetV3 as backbone in RetinaNet [21], gaining 8.6 AP (35.8 vs. 27.2 AP) with even less computational cost. In addition, we build an efficient end-to-end detector by using Mobile-Former to replace backbone and encoder/decoder in DETR [1]. Using the same number of object queries (100), it gains 1.3 AP over DETR (43.3 vs. 42.0 AP) but has significantly fewer FLOPs (41G vs. 86G) and smaller model size (26.3M vs. 41.3M).

## 2. Related Work

**Light-weight CNNs:** MobileNets [16, 17, 28] efficiently encode local features by stacking depthwise and pointwise convolutions. ShuffleNet [26, 47] uses group convolution and channel shuffle to simplify pointwise convolution. MicroNet [19] presents micro-factorized convolution to handle extremely low FLOPs by lowering node connectivity to enlarge network width. Dynamic operators [4, 5, 18, 44] boost performance for MobileNet with negligible computational cost. Other efficient operators include butterfly transform [36], cheap linear transformations in GhostNet [12], and using additions to trade multiplications in AdderNet [3]. In addition, MixConv [31] explores mixing up multiple kernel sizes, and Sandglass [49] flips the structure of inverted residual block. EfficientNet [30, 32] and TinyNet [13] study the compound scaling of depth, width and resolution.

**Vision transformers (ViTs):** Recently, ViT [10] and its follow-ups [9, 23, 34, 35, 37, 45, 50, 51] achieve impressive performance on multiple vision tasks. The original ViT requires training on large dataset such as JFT-300M. DeiT [34] introduces training strategies on the smaller ImageNet-1K dataset. Later, hierarchical transformers are proposed to handle high resolution images. Swin [23] computes self-attention within shifted local windows and CSWin [9] further improves it by introducing cross-shaped window. T2T-ViT [45] progressively converts an image to tokens by recursively aggregating neighboring tokens. HaloNet [37] improves speed, memory usage and accuracy by two extensions (blocked local attention and attention downsampling).

**Combination of CNN and ViT:** Recent works [7, 11, 29, 40, 41] show advantages of combining convolution and transformer. BoTNet [29] improves both instance segmentation and object detection by just using self-attention in the last three blocks of ResNet [14]. ConViT [7] presents a gated positional self-attention for soft convolutional inductive biases. CvT [40] introduces depthwise/pointwise convolution before multi-head attention. LeViT [11] and ViT<sub>C</sub> [41] use convolutional stem to replace the patchify stem and achieve

clear improvement. In contrast to these *series* design, our method *parallelizes* MobileNet and transformer with bidirectional cross attention in between, providing an efficient and effective solution at low FLOP regime.

**Parallel architectures:** compared to parallel architectures like Conformer [27] and CrossViT [2], our Mobile-Former uses significantly fewer tokens (6 vs. 196) that are initialized from learnable embeddings rather than linear projection of image patches. This significantly reduces the computational cost and showcases the usage of transformer in very low FLOPs (26M–508M).

### 3. Our Method: Mobile-Former

In this section, we first overview the design of Mobile-Former (see Figure 1), and then discuss details within a Mobile-Former block (see Figure 3). Finally, we show the network specification and variants with different FLOPs.

#### 3.1. Overview

**Parallel structure:** Mobile-Former parallelizes MobileNet and transformer, and connects them by bidirectional cross attention (see Figure 1). *Mobile* (refers to MobileNet) takes an image as input ( $\mathbf{X} \in \mathbb{R}^{HW \times 3}$ ) and applies inverted bottleneck blocks [28] to extract local features. *Former* (refers to transformer) takes learnable parameters (or tokens) as input, denoted as  $\mathbf{Z} \in \mathbb{R}^{M \times d}$  where  $M$  and  $d$  are the number and dimension of tokens, respectively. These tokens are randomly initialized. Different from vision transformer (ViT) [10], where tokens project the local image patch linearly, *Former* has significantly fewer tokens ( $M \leq 6$  in this paper), each represents a global prior of the image. This results in much less computational cost.

**Low cost two-way bridge:** *Mobile* and *Former* communicate through a two-way bridge where local and global features are fused bidirectionally. The two directions are denoted as *Mobile*→*Former* and *Mobile*←*Former*, respectively. We propose a light-weight cross attention to model them, in which the projections ( $\mathbf{W}^Q, \mathbf{W}^K, \mathbf{W}^V$ ) are removed from *Mobile* side to save computations, but kept at *Former* side. The cross attention is computed at the bottleneck of *Mobile* where the number of channels is low. Specifically, the light-weight cross attention from local feature map  $\mathbf{X}$  to global tokens  $\mathbf{Z}$  is computed as:

$$\mathcal{A}_{\mathbf{X} \rightarrow \mathbf{Z}} = \left[ \text{Attn}(\tilde{\mathbf{z}}_i \mathbf{W}_i^Q, \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_i) \right]_{i=1:h} \mathbf{W}^O, \quad (1)$$

where the local feature  $\mathbf{X}$  and global tokens  $\mathbf{Z}$  are split into  $h$  heads as  $\mathbf{X} = [\tilde{\mathbf{x}}_1 \cdots \tilde{\mathbf{x}}_h]$ ,  $\mathbf{Z} = [\tilde{\mathbf{z}}_1 \cdots \tilde{\mathbf{z}}_h]$  for multi-head attention. The split for the  $i^{\text{th}}$  head  $\tilde{\mathbf{z}}_i \in \mathbb{R}^{M \times \frac{d}{h}}$  is different to the  $i^{\text{th}}$  token  $\mathbf{z}_i \in \mathbb{R}^d$ .  $\mathbf{W}_i^Q$  is the query projection matrix for the  $i^{\text{th}}$  head.  $\mathbf{W}^O$  is used to combine multiple heads together.  $\text{Attn}(Q, K, V)$  is the standard attention function [38] over query  $Q$ , key  $K$ , and value  $V$

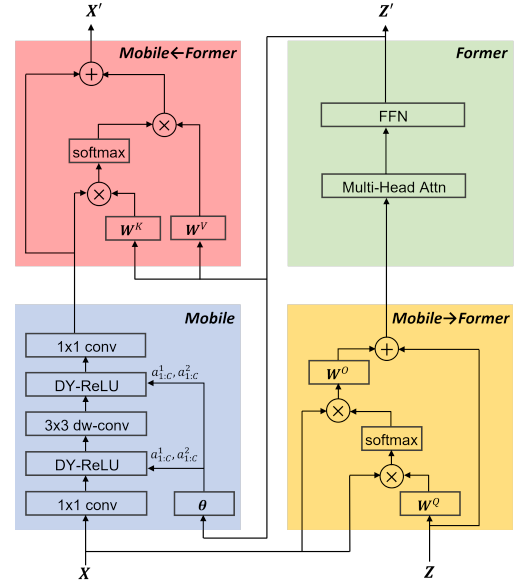


Figure 3. **Mobile-Former block** that includes four modules: *Mobile* sub-block modifies inverted bottleneck block in [28] by replacing ReLU with dynamic ReLU [5]. *Mobile*→*Former* uses light-weight cross attention to fuse local features into global features. *Former* sub-block is a standard transformer block including multi-head attention and FFN. Note that the output of *Former* is used to generate parameters for dynamic ReLU in *Mobile* sub-block. *Mobile*←*Former* bridges from global to local features.

as  $\text{softmax}(\frac{QK^T}{\sqrt{d_k}})V$ .  $[\cdot]_{1:h}$  denotes the concatenation of  $h$  elements. Note that the projection matrices for the key and value are removed from *Mobile* side, while the project matrix  $\mathbf{W}_i^Q$  for the query is kept at *Former* side. Similarly, the cross attention from global to local is computed as:

$$\mathcal{A}_{\mathbf{Z} \rightarrow \mathbf{X}} = \left[ \text{Attn}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{z}}_i \mathbf{W}_i^K, \tilde{\mathbf{z}}_i \mathbf{W}_i^V) \right]_{i=1:h}, \quad (2)$$

where  $\mathbf{W}_i^K$  and  $\mathbf{W}_i^V$  are the projection matrices for the key and value at *Former* side. The projection matrix of the query is removed from *Mobile* side.

#### 3.2. Mobile-Former Block

Mobile-Former consists of stacked *Mobile-Former blocks* (see Figure 1). Each block has four pillars: a *Mobile* sub-block, a *Former* sub-block, and two-way cross attention *Mobile*←*Former* and *Mobile*→*Former* (shown in Figure 3).

**Input and output:** Mobile-Former block has two inputs: (a) local feature map  $\mathbf{X} \in \mathbb{R}^{HW \times C}$ , which has  $C$  channels over height  $H$  and width  $W$ , and (b) global tokens  $\mathbf{Z} \in \mathbb{R}^{M \times d}$ , where  $M$  and  $d$  are the number and dimension of tokens, respectively. Note that  $M$  and  $d$  are identical across all blocks. Mobile-Former block outputs the updated local feature map  $\mathbf{X}'$  and global tokens  $\mathbf{Z}'$ , which are used as input for the next block.

**Mobile sub-block:** As shown in Figure 3, *Mobile* sub-block takes the feature map  $\mathbf{X}$  as input and its output is taken as the input for *Mobile* $\leftarrow$ *Former*. It is slightly different to the inverted bottleneck block in [28] by replacing ReLU with dynamic ReLU [5] as the activation function. Different from the original dynamic ReLU, in which the parameters are generated by applying two MLP layers on the average pooled feature, we save the average pooling by applying the two MLP layers ( $\theta$  in Figure 3) on the first global token output  $z'_1$  from *Former*. Note that the kernel size of depthwise convolution is  $3\times 3$  for all blocks.

**Former sub-block:** *Former* sub-block is a standard transformer block including a multi-head attention (MHA) and a feed-forward network (FFN). Expansion ratio 2 (instead of 4) is used in FFN. We follow [38] to use post layer normalization. *Former* is processed between *Mobile* $\rightarrow$ *Former* and *Mobile* $\leftarrow$ *Former* (see Figure 3).

**Mobile** $\rightarrow$ **Former:** The proposed light-weight cross attention (Equation 1) is used to fuse local features  $\mathbf{X}$  to global tokens  $\mathbf{Z}$ . Compared to the standard attention, the projection matrices for key  $\mathbf{W}^K$  and value  $\mathbf{W}^V$  (on the local features  $\mathbf{X}$ ) are removed to save computations (see Figure 3).

**Mobile** $\leftarrow$ **Former:** Here, the cross attention (Equation 2) is opposite to *Mobile* $\rightarrow$ *Former*, fusing global tokens (as key and value) to local features (as query). We keep the projection matrices for the key  $\mathbf{W}^K$  and value  $\mathbf{W}^V$ , but remove the projection matrix for the query  $\mathbf{W}^Q$  to save computations, as shown in Figure 3.

**Computational complexity:** The four pillars of a *Mobile-Former* block have different computational costs. Given an input feature map with size  $HW \times C$ , and  $M$  global tokens with  $d$  dimensions, *Mobile* consumes the most computations  $O(HWC^2)$ . *Former* and the two-way bridge have light weight, consuming less than 20% of the total computational cost. Specifically, *Former* has complexity  $O(M^2d + Md^2)$  for self attention and FFN. *Mobile* $\rightarrow$ *Former* and *Mobile* $\leftarrow$ *Former* share the complexity  $O(MHWC + MdC)$  for cross attention.

### 3.3. Network Specification

**Architecture:** Table 1 shows a *Mobile-Former* architecture with 294M FLOPs for image size  $224\times 224$ , which stacks 11 *Mobile-Former* blocks at different input resolutions. All blocks have six global tokens with dimension 192. It starts with a  $3\times 3$  convolution as stem and a lite bottleneck block [19] at stage 1, which expands and then squeezes the number of channels by stacking a  $3\times 3$  depthwise and a pointwise convolution. Stage 2–5 consists of *Mobile-Former* blocks. Each stage handles downsampling by a downsample variant denoted as *Mobile-Former* $^\downarrow$  (see supplementary material for details). The classification head applies average pooling on the local features, concatenates

Stage	Input	Operator	exp size	#out	Stride
tokens	$6\times 192$	–	–	–	–
stem	$224^2\times 3$	conv2d, $3\times 3$	–	16	2
1	$112^2\times 16$	bneck-lite	32	16	1
2	$112^2\times 16$	<i>Mobile-Former</i> $^\downarrow$	96	24	2
	$56^2\times 24$	<i>Mobile-Former</i>	96	24	1
3	$56^2\times 24$	<i>Mobile-Former</i> $^\downarrow$	144	48	2
	$28^2\times 48$	<i>Mobile-Former</i>	192	48	1
4	$28^2\times 48$	<i>Mobile-Former</i> $^\downarrow$	288	96	2
	$14^2\times 96$	<i>Mobile-Former</i>	384	96	1
	$14^2\times 96$	<i>Mobile-Former</i>	576	128	1
	$14^2\times 128$	<i>Mobile-Former</i>	768	128	1
5	$14^2\times 128$	<i>Mobile-Former</i> $^\downarrow$	768	192	2
	$7^2\times 192$	<i>Mobile-Former</i>	1152	192	1
	$7^2\times 192$	<i>Mobile-Former</i>	1152	192	1
	$7^2\times 192$	conv2d, $1\times 1$	–	1152	1
head	$7^2\times 1152$	pool, $7\times 7$	–	1152	–
	$1^2\times 1152$	concat w/ cls token	–	1344	–
	$1^2\times 1344$	FC	–	1920	–
	$1^2\times 1920$	FC	–	1000	–

Table 1. **Specification for *Mobile-Former-294M*.** “bneck-lite” denotes the lite bottleneck block [19]. “*Mobile-Former* $^\downarrow$ ” denotes the variant of downsample block.

with the first global token, and then passes through two fully connected layers with h-swish [16] in between.

**Mobile-Former variants:** *Mobile-Former* has seven models with different computational costs from 26M to 508M FLOPs. They share the similar architecture, but have different width and height. We follow [41] to refer our models by their FLOPs, e.g. *Mobile-Former-294M*, *Mobile-Former-96M*. The details of network architecture for these models are listed in the supplementary material.

## 4. Efficient End-to-End Object Detection

*Mobile-Former* can be easily used for object detection in both backbone and head, providing an *efficient end-to-end* detector. Using the same number of object queries (100), it outperforms DETR [1], but uses much lower FLOPs.

**Backbone–Head architecture:** We use *Mobile-Former* blocks in both backbone and head (see Figure 4), which have separate tokens. The backbone has six global tokens while the head has 100 object queries generated similarly to DETR [1]. Different from DETR that has a single scale ( $\frac{1}{32}$  or  $\frac{1}{16}$ ) in the head, *Mobile-Former* head employs multi-scales ( $\frac{1}{32}$ ,  $\frac{1}{16}$ ,  $\frac{1}{8}$ ) in low FLOPs due to its computational efficiency. The upsampling is achieved via bilinear interpolation followed by adding the feature output from the backbone (with the same resolution). All object queries progressively refine their representation across scales from coarse to fine, saving the manual process in FPN [20] to allocate objects across scales by size. We follow DETR to use prediction FFNs and auxiliary losses in the head during the training. The head is trained from scratch and the backbone is pretrained on ImageNet. Our end-to-end *Mobile-Former*

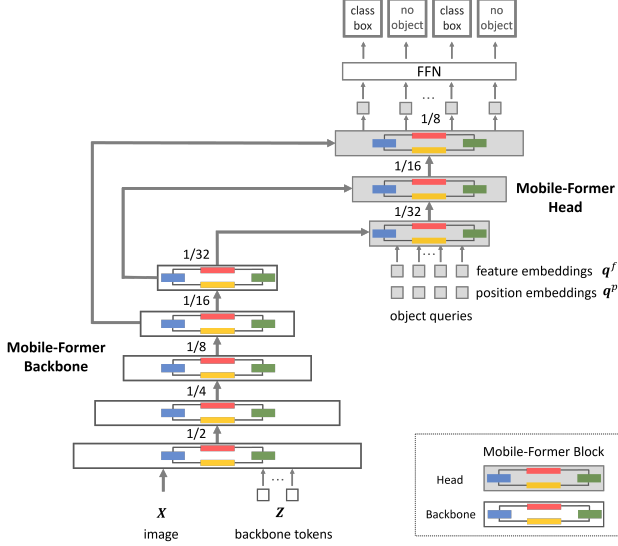


Figure 4. **Mobile-Former for object detection.** Both backbone and head use Mobile-Former blocks (see Figure 1, 3). The backbone has 6 global tokens while the head has 100 object queries. All object queries pass through multiple resolutions ( $\frac{1}{32}, \frac{1}{16}, \frac{1}{8}$ ) in the head. Similar to DETR [1], feed forward network (FFN) is used to predict class label and bounding box. Best viewed in color.

detector is computationally efficient. The total cost of E2E-MF-508M that uses Mobile-Former-508M as backbone and nine Mobile-Former blocks in the head is 41.4G FLOPs, significantly less than DETR (86G FLOPs). But it outperforms DETR by 1.3 AP (43.3 vs. 42.0 AP). The details of head structure is listed in supplementary material.

**Spatial-aware dynamic ReLU in backbone:** We extend dynamic ReLU in the backbone from *spatial-shared* to *spatial-aware* by involving all global tokens to generate parameters, rather than just using the first one, as these tokens have different spatial focuses. Let us denote the parameter generation of the spatial-shared dynamic ReLU as  $\theta = f(z_1)$ , where  $z_1$  is the first global token and  $f(\cdot)$  is modeled by two MLP layers with ReLU in between. By contrast, the spatial-aware dynamic ReLU generates parameters  $\theta_i$  per spatial position  $i$  in a feature map, by using all global tokens  $\{z_j\}$  as follows:

$$\theta_i = \sum_j \alpha_{i,j} f(z_j), s.t. \sum_j \alpha_{i,j} = 1, \quad (3)$$

where  $\alpha_{i,j}$  is the attention between the feature at position  $i$  and token  $z_j$ . Its calculation is cheap by just normalizing the cross attention obtained in *Mobile*→*Former* along tokens  $\{z_j\}$ . Spatial-aware dynamic ReLU is on par with its spatial-shared counterpart on image classification, but gains 1.1 AP on COCO object detection (see Table 10).

**Adapting position embedding in head:** Different from DETR [1] that shared the position embedding of object

queries across all decoder layers, we refine the position embedding after each block in the head as the feature map changes per block. Let us denote the feature and position embedding of a query at the  $k^{th}$  block as  $q_k^f$  and  $q_k^p$ , respectively. The sum of them ( $q_k^f + q_k^p$ ) is used to compute cross attention between object queries and feature map as well as self attention among object queries, after which the feature embedding is updated as the input for the next block  $q_{k+1}^f$ . Here, we adapt the position embedding based upon the feature embedding as:

$$q_{k+1}^p = q_k^p + g(q_{k+1}^f), \quad (4)$$

where the adaptation function  $g(\cdot)$  is implemented by two MLP layers with ReLU in between. Thus, object queries can adapt their positions across blocks based on the content.

## 5. Experimental Results

We evaluate the proposed Mobile-Former on both ImageNet classification [8], and COCO object detection [22].

### 5.1. ImageNet Classification

Image classification experiments are conducted on ImageNet [8] that has 1000 classes, including 1,281,167 images for training and 50,000 images for validation.

**Training setup:** The image resolution is  $224 \times 224$ . All models are trained from scratch using AdamW [24] optimizer for 450 epochs with cosine learning rate decay. A batch size of 1024 is used. Data augmentation includes Mixup [46], auto-augmentation [6], and random erasing [48]. Different combinations of initial learning rate, weight decay and dropout are used for models with different complexities, which are listed in the supplementary material.

**Comparison with efficient CNNs:** Table 2 shows the comparison between Mobile-Former and classic efficient CNNs: (a) MobileNetV3 [16], (b) EfficientNet [30], and (c) ShuffleNetV2 [26] and its extension WeightNet [25]. The comparison covers the FLOP range from 26M to 508M, organized in seven groups based on similar FLOPs. Mobile-Former consistently outperforms efficient CNNs with even less computational cost except the group around 150M FLOPs, where Mobile-Former costs slightly more FLOPs than ShuffleNet/WeightNet (151M vs. 146M/141M), but achieves significantly higher top-1 accuracy (75.2% vs. 69.4%/72.4%). This demonstrates that our parallel design improves the representation capability efficiently.

**Comparison with ViTs:** In Table 3, we compare Mobile-Former with multiple variants (DeiT [34], T2T-ViT [45], PVT [39], ConViT [7], CoaT [42], ViT<sub>C</sub> [41], Swin [23]) of vision transformer. All variants use image resolution  $224 \times 224$  and are trained *without* distillation from a teacher network. Mobile-Former achieves higher accuracy but uses

Model	Input	#Params	MAdds	Top-1
MobileNetV3 Small 1.0× [16]	160 <sup>2</sup>	2.5M	30M	62.8
<b>Mobile-Former-26M</b>	224 <sup>2</sup>	3.2M	<b>26M</b>	<b>64.0</b>
MobileNetV3 Small 1.0× [16]	224 <sup>2</sup>	2.5M	57M	67.5
<b>Mobile-Former-52M</b>	224 <sup>2</sup>	3.5M	<b>52M</b>	<b>68.7</b>
MobileNetV3 1.0× [16]	160 <sup>2</sup>	5.4M	112M	71.7
<b>Mobile-Former-96M</b>	224 <sup>2</sup>	4.6M	<b>96M</b>	<b>72.8</b>
ShuffleNetV2 1.0× [26]	224 <sup>2</sup>	2.3M	146M	69.4
ShuffleNetV2 1.0×+WeightNet 4× [25]	224 <sup>2</sup>	5.1M	<b>141M</b>	72.4
MobileNetV3 0.75× [16]	224 <sup>2</sup>	4.0M	155M	73.3
<b>Mobile-Former-151M</b>	224 <sup>2</sup>	7.6M	151M	<b>75.2</b>
MobileNetV3 1.0× [16]	224 <sup>2</sup>	5.4M	217M	75.2
<b>Mobile-Former-214M</b>	224 <sup>2</sup>	9.4M	<b>214M</b>	<b>76.7</b>
ShuffleNetV2 1.5× [26]	224 <sup>2</sup>	3.5M	299M	72.6
ShuffleNetV2 1.5×+WeightNet 4× [25]	224 <sup>2</sup>	9.6M	307M	75.0
MobileNetV3 1.25× [16]	224 <sup>2</sup>	7.5M	356M	76.6
EfficientNet-B0 [30]	224 <sup>2</sup>	5.3M	390M	77.1
<b>Mobile-Former-294M</b>	224 <sup>2</sup>	11.4M	<b>294M</b>	<b>77.9</b>
ShuffleNetV2 2× [26]	224 <sup>2</sup>	7.4M	591M	74.9
ShuffleNetV2 2×+WeightNet 4× [25]	224 <sup>2</sup>	18.1M	573M	76.5
<b>Mobile-Former-508M</b>	224 <sup>2</sup>	14.0M	<b>508M</b>	<b>79.3</b>

Table 2. **Comparing Mobile-Former with efficient CNNs** evaluated on ImageNet [8] classification.

Model	Input	#Params	MAdds	Top-1
T2T-ViT-7 [45]	224 <sup>2</sup>	4.3M	1.2G	71.7
DeiT-Tiny [34]	224 <sup>2</sup>	5.7M	1.2G	72.2
ConViT-Tiny [7]	224 <sup>2</sup>	6.0M	1.0G	73.1
ConT-Ti [43]	224 <sup>2</sup>	5.8M	0.8G	74.9
ViT <sub>C</sub> [41]	224 <sup>2</sup>	4.6M	1.1G	75.3
ConT-S [43]	224 <sup>2</sup>	10.1M	1.5G	76.5
Swin-1G [23] ‡	224 <sup>2</sup>	7.3M	1.0G	77.3
<b>Mobile-Former-294M</b>	224 <sup>2</sup>	11.4M	<b>294M</b>	<b>77.9</b>
PVT-Tiny [39]	224 <sup>2</sup>	13.2M	1.9G	75.1
T2T-ViT-12 [45]	224 <sup>2</sup>	6.9M	2.2G	76.5
CoaT-Lite Tiny [42]	224 <sup>2</sup>	5.7M	1.6G	76.6
ConViT-Tiny+ [7]	224 <sup>2</sup>	10.0M	2G	76.7
CrossViT-9† [2]	224 <sup>2</sup>	8.8M	2G	77.1
DeiT-2G [34] ‡	224 <sup>2</sup>	9.5M	2.0G	77.6
CoaT-Lite Mini [42]	224 <sup>2</sup>	11.0M	2.0G	78.9
BoT-S1-50 [29]	224 <sup>2</sup>	20.8M	4.3G	79.1
Swin-2G [23] ‡	224 <sup>2</sup>	12.8M	2.0G	79.2
<b>Mobile-Former-508M</b>	224 <sup>2</sup>	14.0M	<b>508M</b>	<b>79.3</b>

Table 3. **Comparing Mobile-Former with vision transformer variants** evaluated on ImageNet [8] classification. Here, we choose ViT variants that use image resolution 224×224 and are trained *without* distillation from a teacher network. We group ViT models based on FLOPs (using 1.5G as threshold) and rank them based on top-1 accuracy. ‡ indicates our implementation.

3~4 times less computational cost. This is because that Mobile-Former uses significantly fewer tokens to model global interaction and leverages MobileNet to extract local features efficiently. Note that our Mobile-Former (trained in 450 epochs without distillation) even outperforms LeViT [11] which leverages the distillation of a teacher network and much longer training (1000 epochs). Our method achieves higher top-1 accuracy (77.9% vs. 76.6%) but uses fewer FLOPs (294M vs. 305M) than LeViT.

Model	#Params	MAdds	Top-1	Top-5
<i>Mobile</i> (using ReLU)	6.1M	259M	74.2	91.8
+ <i>Former</i> and Bridge	10.1M	290M	76.8 <sub>(+2.6)</sub>	93.2 <sub>(+1.4)</sub>
+ DY-ReLU in <i>Mobile</i>	11.4M	294M	77.8 <sub>(+1.0)</sub>	93.7 <sub>(+0.5)</sub>

Table 4. **Ablation of *Former*+bridge and dynamic ReLU** evaluated on ImageNet classification. Mobile-Former-294M is used.

**Accuracy–FLOP plot:** Figure 2 compares Mobile-Former with more CNN models (e.g. GhostNet [12]) and vision transformer variants (e.g. Swin [23] and DeiT [34]) in one plot. We implement Swin and DeiT from 100M to 2G FLOPs, by carefully reducing network width and height. Mobile-Former clearly outperforms both CNNs and ViT variants, showcasing the parallel design to integrate MobileNet and transformer. This demonstrates that the transformer can also contribute to the low FLOP regime with proper architecture design.

## 5.2. Ablations

In this subsection, we show Mobile-Former is effective and efficient via several ablations performed on ImageNet classification. Here, Mobile-Former-294M is used and all models are trained for 300 epochs. Moreover, we summarize interesting observations on the visualization of the two-way bridge between *Mobile* and *Former*.

**Mobile-Former is effective:** Mobile-Former is more effective than MobileNet as it encodes global interaction via *Former*, resulting in more accurate prediction. As shown in Table 4, adding *Former* and bridge (*Mobile*→*Former* and *Mobile*←*Former*) only costs 12% of the computational cost, but gains 2.6% top-1 accuracy over the baseline that uses *Mobile* alone. In addition, using dynamic ReLU [5] in *Mobile* sub-block (see Figure 3) gains additional 1.0% top-1 accuracy. Note that the parameters in dynamic ReLU is generated by using the first global token. This validates our parallel design in Mobile-Former. We also find that increasing kernel size ( $3\times 3 \rightarrow 5\times 5$ ) of the depthwise convolution in *Mobile* only introduces negligible gain (see supplementary material), as the reception field of *Mobile* is enlarged by fusing global features from *Former*.

**Mobile-Former is efficient:** Mobile-Former is not only effective in encoding both local processing and global interaction, but achieves this *efficiently*. Ablations below show that *Former* only requires a few global tokens with low dimension. In addition, the efficient parallel design of Mobile-Former is stable when removing FFN in *Former* or replacing multi-head attention with position mixing MLP [33].

*Number of tokens in Former:* Table 5 shows the ImageNet classification results for using different number of global tokens in *Former*. The token dimension is 192. Interestingly, even a single global token achieves a good perfor-

#Tokens	#Params	MAdds	Top-1	Top-5
1	11.4M	269M	77.1	93.2
3	11.4M	279M	77.6	93.6
6	11.4M	294M	77.8	93.7
9	11.4M	309M	77.7	93.8

Table 5. **Ablation of the number of tokens** on ImageNet classification. Mobile-Former-294M is used.

Token Dimension	#Params	MAdds	Top-1	Top-5
64	7.3M	277M	76.8	93.1
128	9.1M	284M	77.3	93.5
192	11.4M	294M	77.8	93.7
256	14.3M	308M	77.8	93.7
320	17.9M	325M	77.6	93.6

Table 6. **Ablation of token dimension** on ImageNet classification. Mobile-Former-294M is used.

Attention	FFN	#Params	MAdds	Top-1	Top-5
MHA	✓	11.4M	294M	77.8	93.7
MHA	✗	9.8M	284M	77.5	93.6
Pos-Mix-MLP	✓	10.5M	284M	77.3	93.5

Table 7. **Ablation of multi-head attention (MHA) and FFN** on ImageNet classification. Mobile-Former-294M is used.

mance (77.1% top-1 accuracy). Additional improvement (0.5% and 0.7% top-1 accuracy) is achieved when using 3 and 6 tokens. The improvement stops when more than 6 tokens are used. Such compactness of global tokens is a key contributor to the efficiency of Mobile-Former.

*Token dimension:* Table 6 shows the results for different token lengths (or dimensions). Here, six global tokens are used in *Former*. The accuracy improves from 76.8% to 77.8% when token dimension increases from 64 to 192, but converges when higher dimension is used. This further supports the efficiency of *Former*. With six tokens of dimension 192, the total computational cost of *Former* and the bridge only consumes 12% of the overall budget (35M/294M).

*FFN in Former:* As shown in Table 7, removing FFN introduces a small drop in top-1 accuracy (−0.3%). Compared to the important role of FFN in the original vision transformer, FFN has limited contribution in Mobile-Former. We believe this is because FFN is not the only module for channel fusion in Mobile-Former. The  $1\times 1$  convolution in *Mobile* helps the channel fusion of local features, while the projection matrix  $W^O$  in *Mobile*→*Former* (see Equation 1) contributes to the fusion between local and global features.

*Multi-head attention (MHA) vs. position-mixing MLP:* Table 7 shows the results of replacing multi-head attention (MHA) with token/position mixing MLP [33] in both *Former* and bridge (*Mobile*→*Former* and *Mobile*←*Former*). The top-1 accuracy drops from 77.8% to 77.3%. MLP has more efficient implementation by a single matrix multipli-

Model	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	MAdds (G)	#Params (M)
Shuffle-V2 [26]	25.9	41.9	26.9	12.4	28.0	36.4	2.6 (161)	0.8 (10.4)
<b>MF-151M</b>	<b>34.2</b>	53.4	36.0	19.9	36.8	45.3	2.6 (161)	4.9 (14.4)
Mobile-V3 [16]	27.2	43.9	28.3	13.5	30.2	37.2	4.7 (162)	2.8 (12.3)
<b>MF-214M</b>	<b>35.8</b>	55.4	38.0	21.8	38.5	46.8	3.9 (162)	5.7 (15.2)
ResNet18 [14]	31.8	49.6	33.6	16.3	34.3	43.2	29 (181)	11.2 (21.3)
<b>MF-294M</b>	<b>36.6</b>	56.6	38.6	21.9	39.5	47.9	5.5 (164)	6.5 (16.1)
ResNet50 [14]	36.5	55.4	39.1	20.4	40.3	48.1	84 (239)	23.3 (37.7)
PVT-Tiny [39]	36.7	56.9	38.9	22.6	38.8	50.0	70 (221)	12.3 (23.0)
ConT-M [43]	37.9	58.1	40.2	23.0	40.6	50.4	65 (217)	16.8 (27.0)
<b>MF-508M</b>	<b>38.0</b>	58.3	40.3	22.9	41.2	49.7	9.8 (168)	8.4 (17.9)

Table 8. **COCO object detection results in RetinaNet framework.** All models are trained on `train2017` for 12 epochs ( $1\times$ ) from ImageNet pretrained weights, and tested on `val2017`. We use initial **MF** (e.g. MF-508M) to refer Mobile-Former. MAdds and #Params are in the format of “backbone (total)”. MAdds is based on the image size  $800\times 1333$ .

cation, but is not adaptive to different input images.

**Mobile-Former is explainable:** We observe three interesting patterns in the two-way bridge (*Mobile*→*Former* and *Mobile*←*Former*). First, from low to high levels, the global tokens change their focus from edges/corners, to foreground/background, and finally on the most discriminative region. Second, the cross attention has more diversity across tokens at lower levels than high levels. Thirdly, the separation between foreground and background is found at middle layers of *Mobile*←*Former*. The detailed visualization is shown in the supplementary material.

### 5.3. Object Detection

Object detection experiments are conducted on COCO 2017 [22] that contains 118K training and 5K validation images. We evaluate Mobile-Former in two detection frameworks: (a) comparing with other backbone networks in RetinaNet [21] that has dense proposals, and (b) end-to-end comparison with DETR [1] where proposals are sparse.

**RetinaNet training setup:** We follow standard settings of RetinaNet and replace the backbone with Mobile-Former to generate multi-scale feature maps. All models are trained for 12 epochs ( $1\times$ ) from ImageNet pretrained weights.

**DETR training setup:** All Mobile-Former models are trained for 300 epochs on 8 GPUs with 2 images per GPU. AdamW optimizer is used with initial learning rate  $1e-5$  for the backbone and  $1e-4$  for the head. The learning rate drops by a factor of 10 after 200 epochs. The weight decay is  $1e-4$  and dropout rate is 0.1. BatchNorm layers in the ImageNet pretrained Mobile-Former backbone are frozen. The head includes 100 object queries with 256 channels.

**Efficient and effective backbone in RetinaNet:** In Table 8, we compare Mobile-Former with both CNNs (ResNet [14], MobileNetV3 [16], ShuffleNetV2 [26]) and vision transformers (PVT [39] and ConT [43]). Mobile-Former outperforms MobileNetV3 and ShuffleNetV2 by 8.3+ AP

Model	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	MAdds (G)	#Params (M)
DETR [1]	42.0	62.4	44.2	20.5	45.8	61.1	86	41.3
DETR-DC5 [1]	<b>43.3</b>	<b>63.1</b>	45.9	22.5	<b>47.3</b>	<b>61.1</b>	187	41.3
<b>E2E-MF-508M</b>	<b>43.3</b>	61.8	<b>46.8</b>	<b>24.6</b>	47.0	60.4	<b>41.4</b>	<b>26.3</b>
<b>E2E-MF-294M</b>	40.5	58.8	43.5	20.6	44.0	56.9	<b>24.1</b>	<b>24.9</b>
<b>E2E-MF-214M</b>	39.3	57.3	42.1	19.9	42.4	56.6	<b>17.8</b>	<b>20.1</b>
<b>E2E-MF-151M</b>	37.2	54.5	39.9	17.4	39.8	54.9	<b>12.7</b>	<b>14.8</b>

Table 9. **End-to-end object detection results on COCO.** All models are trained on `train2017` and tested on `val2017`. DETR baselines are trained for 500 epochs, while our MobileFormer models are trained for 300 epochs. We use initial **E2E-MF** (e.g. E2E-MF-508M) to refer end-to-end MobileFormer detectors. MAdds is based on image size  $800 \times 1333$ .

under similar computational cost. Compared to ResNet and transformer variants, MobileFormer achieves higher AP with significantly less FLOPs. Specifically, MobileFormer-508M takes 9.8G FLOPs in backbone but achieves 38.0 AP, outperforming ResNet-50, PVT-Tiny, and ConT-M which consume 7 times more computation (65G to 84G FLOPs) in the backbone. This showcases MobileFormer as an effective and efficient backbone in object detection.

**Efficient and effective end-to-end detector:** Table 9 compares end-to-end MobileFormer detectors (denoted by prefix E2E-MF) with DETR [1]. All models use 100 object queries. Our E2E-MF-508M gains 1.3 AP over DETR but consumes fewer FLOPs (48%), fewer model parameters (64%) and fewer training epochs (300 vs. 500). It is even on par with DETR-DC5 that has four times more FLOPs. The other three MobileFormer variants achieve 40.5, 39.3 and 37.2 AP with 24.1G, 17.8G and 12.7G FLOPs respectively, providing more compact end-to-end detectors.

**Ablations of key components:** Table 10 shows the effects of three proposed components in the end-to-end MobileFormer detector: (a) spatial-aware dynamic ReLU in backbone, (b) multi-scale MobileFormer head, and (c) adapting position embedding. E2E-MF-508M is used and all models are trained for 300 epochs. Compared to DETR trained in 300 epochs, replacing ResNet-50 backbone with MF-508M saves computations but has 1.2 AP drop. Adding spatial-aware dynamic ReLU gains 1.1 AP ( $39.4 \rightarrow 40.5$ ). Then using multi-scale MobileFormer head to replace DETR’s encoder/decoder gains another 0.9 AP ( $40.5 \rightarrow 41.4$ ). Detection of small/medium objects is improved, while a slight drop is found in large objects. Finally, adapting position embedding gains additional 1.9 AP. The three components are complementary, gaining 3.9 AP ( $39.4 \rightarrow 43.3$ ) in total.

## 6. Limitations and Discussions

Although MobileFormer has faster inference than MobileNetV3 [16] and ShuffleNetV2 [26] for large images, it is slower as the image becomes smaller. Please see supplementary material for comparison among MobileFormer-

Backbone	SP-DY ReLU	Head	Adapt PE	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
ResNet-50		DETR		40.6	61.6	42.7	19.9	44.3	60.2
MF-508M		DETR		39.4	59.4	41.3	18.0	42.7	58.8
MF-508M	✓	DETR		40.5	60.7	42.4	19.0	44.0	60.0
MF-508M	✓	MF-Head		41.4	60.4	43.9	21.1	45.0	59.3
MF-508M	✓	MF-Head	✓	<b>43.3</b>	<b>61.8</b>	<b>46.8</b>	<b>24.6</b>	<b>47.0</b>	<b>60.4</b>

Table 10. **Ablation of end-to-end MobileFormer on COCO object detection.** All models are trained on `train2017` for 300 epochs and tested on `val2017`. The first line is the baseline of DETR [1] trained in 300 epochs. MF-508M refers to MobileFormer-508M backbone. SP-DY-ReLU refers to spatial-aware dynamic ReLU in backbone. DETR head includes 6 encoder and 6 decoder layers with resolution  $\frac{1}{32}$ . MF head includes 9 MobileFormer blocks (5, 2, 2 at resolution  $\frac{1}{32}$ ,  $\frac{1}{16}$ , and  $\frac{1}{8}$  respectively).

214M, MobileNetV3 Large, and ShuffleNetV2  $2 \times$  on CPU inference latency. MobileFormer is more accurate than the two baselines (76.7% vs. 75.2% vs. 74.9% top-1 on ImageNet). The comparison is performed on multiple image sizes due to the resolution variation across tasks (e.g. classification, detection). As image resolution decreases, MobileFormer loses its leading position to MobileNetV3. This is because *Former* and embedding projections in *Mobile* $\rightarrow$ *Former* and *Mobile* $\leftarrow$ *Former* are resolution independent, and their PyTorch implementations are not as efficient as convolution. Thus, the overhead is relative large when image is small, but becomes negligible as image size grows. The runtime of MobileFormer can be further improved by optimizing the implementation of these components. We will investigate these in the future work.

Another limitation is that MobileFormer is not efficient in parameters especially when performing image classification, due to the parameter-heavy classification head. For instance, the head of MobileFormer-294M consumes 4.6M of total 11.4M parameters (40%). This problem is mitigated when switching to object detection, due to the removal of image classification head. In addition, *Former* and two-way bridge are computationally but not parametrically efficient.

## 7. Conclusion

This paper presents MobileFormer, a new parallel design of MobileNet and Transformer with two-way bridge in between to communicate. It leverages the efficiency of MobileNet in local processing and the advantage of Transformer in encoding global interaction. This design is not only effective to boost accuracy, but also efficient to save computational cost. It outperforms both efficient CNNs and vision transformer variants with a clear margin on image classification and object detection in the low FLOP regime. Furthermore, we build an end-to-end MobileFormer detector that outperforms DETR but consumes significantly less computations and parameters. We hope MobileFormer encourage new design of efficient CNNs and transformers.



## References

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. [2](#), [4](#), [5](#), [7](#), [8](#)
- [2] Chun-Fu (Richard) Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 357–366, October 2021. [3](#), [6](#)
- [3] Hanting Chen, Yunhe Wang, Chunjing Xu, Boxin Shi, Chao Xu, Qi Tian, and Chang Xu. Addernet: Do we really need multiplications in deep learning? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [2](#)
- [4] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic convolution: Attention over convolution kernels. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [5] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic relu. In *ECCV*, 2020. [2](#), [3](#), [4](#), [6](#)
- [6] Ekin D. Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [5](#)
- [7] Stéphane d’Ascoli, Hugo Touvron, Matthew Leavitt, Ari Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. *arXiv preprint arXiv:2103.10697*, 2021. [2](#), [5](#), [6](#)
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [5](#), [6](#)
- [9] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. *arXiv preprint arXiv:2107.00652*, 2021. [2](#)
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [3](#)
- [11] Benjamin Graham, Alaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. *arXiv preprint arXiv:22104.01136*, 2021. [1](#), [2](#), [6](#)
- [12] Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, and Chang Xu. Ghostnet: More features from cheap operations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [1](#), [2](#), [6](#)
- [13] Kai Han, Yunhe Wang, Qiulin Zhang, Wei Zhang, Chunjing XU, and Tong Zhang. Model rubiks cube: Twisting resolution, depth and width for tinynets. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 19353–19364. Curran Associates, Inc., 2020. [2](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#), [7](#)
- [15] Geoffrey E. Hinton. How to represent part-whole hierarchies in a neural network. *CoRR*, abs/2102.12627, 2021. [2](#)
- [16] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [17] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. [1](#), [2](#)
- [18] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [2](#)
- [19] Yunsheng Li, Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Lu Yuan, Zicheng Liu, Lei Zhang, and Nuno Vasconcelos. Micronet: Improving image recognition with extremely low flops. In *International Conference on Computer Vision*, 2021. [1](#), [2](#), [4](#)
- [20] T. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 936–944, July 2017. [4](#)
- [21] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [2](#), [7](#)
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. [5](#), [7](#)
- [23] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021. [2](#), [5](#), [6](#)
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. [5](#)
- [25] Ningning Ma, X. Zhang, J. Huang, and J. Sun. Weightnet: Revisiting the design space of weight networks. volume abs/2007.11823, 2020. [5](#), [6](#)
- [26] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *The European Conference on Computer Vision (ECCV)*, September 2018. [2](#), [5](#), [6](#), [7](#), [8](#)

- [27] Zhiliang Peng, Wei Huang, Shanzhi Gu, Lingxi Xie, Yaowei Wang, Jianbin Jiao, and Qixiang Ye. Conformer: Local features coupling global representations for visual recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 367–376, October 2021. [3](#)
- [28] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. [1](#), [2](#), [3](#), [4](#)
- [29] Aravind Srinivas, Tsung-Yi Lin, Niki Parmar, Jonathon Shlens, Pieter Abbeel, and Ashish Vaswani. Bottleneck transformers for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16519–16529, June 2021. [2](#), [6](#)
- [30] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pages 6105–6114, Long Beach, California, USA, 09–15 Jun 2019. [2](#), [5](#), [6](#)
- [31] Mingxing Tan and Quoc V. Le. Mixconv: Mixed depthwise convolutional kernels. In *30th British Machine Vision Conference 2019*, 2019. [2](#)
- [32] Mingxing Tan, Ruoming Pang, and Quoc V. Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [2](#)
- [33] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision, 2021. [6](#), [7](#)
- [34] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers and distillation through attention. *arXiv preprint arXiv:2012.12877*, 2020. [1](#), [2](#), [5](#), [6](#)
- [35] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers, 2021. [2](#)
- [36] Keivan Alizadeh vahid, Anish Prabhu, Ali Farhadi, and Mohammad Rastegari. Butterfly transform: An efficient fft based neural architecture design. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [2](#)
- [37] Ashish Vaswani, Prajit Ramachandran, Aravind Srinivas, Niki Parmar, Blake Hechtman, and Jonathon Shlens. Scaling local self-attention for parameter efficient visual backbones. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12894–12904, June 2021. [2](#)
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. [1](#), [3](#), [4](#)
- [39] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions, 2021. [5](#), [6](#), [7](#)
- [40] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers, 2021. [1](#), [2](#)
- [41] Tete Xiao, Mannat Singh, Eric Mintun, Trevor Darrell, Piotr Dollár, and Ross B. Girshick. Early convolutions help transformers see better. *CoRR*, abs/2106.14881, 2021. [1](#), [2](#), [4](#), [5](#), [6](#)
- [42] Weijian Xu, Yifan Xu, Tyler Chang, and Zhuowen Tu. Co-scale conv-attentional image transformers, 2021. [5](#), [6](#)
- [43] Haotian Yan, Zhe Li, Weijian Li, Changhu Wang, Ming Wu, and Chuang Zhang. Contnet: Why not use convolution and transformer at the same time? *CoRR*, abs/2104.13497, 2021. [6](#), [7](#)
- [44] Brandon Yang, Gabriel Bender, Quoc V. Le, and Jiquan Ngiam. Condconv: Conditionally parameterized convolutions for efficient inference. In *NeurIPS*, 2019. [2](#)
- [45] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. *arXiv preprint arXiv:2101.11986*, 2021. [2](#), [5](#), [6](#)
- [46] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018. [5](#)
- [47] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [2](#)
- [48] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020. [5](#)
- [49] Daquan Zhou, Qi-Bin Hou, Y. Chen, Jiashi Feng, and S. Yan. Rethinking bottleneck structure for efficient mobile network design. In *ECCV*, August 2020. [2](#)
- [50] Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer, 2021. [2](#)
- [51] Daquan Zhou, Yujun Shi, Bingyi Kang, Weihao Yu, Zihang Jiang, Yuan Li, Xiaojie Jin, Qibin Hou, and Jiashi Feng. Refiner: Refining self-attention for vision transformers, 2021. [2](#)