

# EdgeViTs: Competing Light-weight CNNs on Mobile Devices with Vision Transformers

Junting Pan<sup>1\*</sup>, Adrian Bulat<sup>2</sup>, Fuwen Tan<sup>2</sup>, Xiatian Zhu<sup>2</sup>, Lukasz Dudziak<sup>2</sup>,  
Hongsheng Li<sup>1</sup>, Georgios Tzimiropoulos<sup>2,3</sup> and Brais Martinez<sup>2</sup>

<sup>1</sup> The Chinese University of Hong Kong

<sup>2</sup> Samsung AI Cambridge

<sup>3</sup> Queen Mary University of London


**Abstract.** Self-attention based models such as vision transformers (ViTs) have emerged as a very competitive architecture alternative to convolutional neural networks (CNNs) in computer vision. Despite increasingly stronger variants with ever higher recognition accuracies, due to the quadratic complexity of self-attention, existing ViTs are typically demanding in computation and model size. Although several successful design choices (e.g., the convolutions and hierarchical multi-stage structure) of prior CNNs have been reintroduced into recent ViTs, they are still not sufficient to meet the limited resource requirements of mobile devices. This motivates a very recent attempt to develop light ViTs based on the state-of-the-art MobileNet-v2, but still leaves a performance gap behind. In this work, pushing further along this under-studied direction we introduce **EdgeViTs**, a new family of light-weight ViTs that, for the first time, enable attention based vision models to compete with the best light-weight CNNs in the tradeoff between accuracy and on-device efficiency. This is realized by introducing a highly cost-effective *local-global-local* (LGL) information exchange bottleneck based on optimal integration of self-attention and convolutions. For device-dedicated evaluation, rather than relying on inaccurate proxies like the number of FLOPs or parameters, we adopt a practical approach of focusing directly on on-device latency and, for the first time, energy efficiency. Extensive experiments on image classification, object detection and semantic segmentation validate high efficiency of our EdgeViTs when compared to the state-of-the-art efficient CNNs and ViTs in terms of accuracy-efficiency tradeoff on mobile hardware. Specifically, we show that our models are Pareto-optimal when both accuracy-latency and accuracy-energy tradeoffs are considered, achieving strict dominance over other ViTs in almost all cases and competing with the most efficient CNNs. Code is available at <https://github.com/saic-fi/edgevit>.

## 1 Introduction

Vision transformers (ViTs) have rapidly superseded convolutional neural networks (CNNs) on a variety of visual recognition tasks [9,51], particularly when

---

\* Work done during an internship at Samsung AI Cambridge.



**Fig. 1.** Our EdgeViTs yield comparable or superior tradeoff between accuracy and efficiency (e.g., run speed, latency) on ImageNet-1K, in comparison to state-of-the-art efficient CNNs [42,47], representative generic ViTs [55,51] and latest MobileViTs [36]. Note that, EdgeViTs outperform MobileViTs, the only specially designed model for mobile device, across all the metrics by a large margin. Whilst our lightest EdgeViT consumes more GFLOPs than EfficientNet B0 [47], it runs faster (33.3 vs. 52.1 ms per image), with a reduced gap in *on-device energy consumption*. *Testing device:* Samsung Galaxy S21 (latency), Snapdragon 888 Hardware Development Kit (energy).

the priors and successful designs of previous CNNs are reintroduced for leveraging the induction bias of visual data such as local grid structures [54,30,7,15]. Due to the quadratic complexity of ViTs and the high-dimension property of visual data, it is indispensable that the computational cost needs to be taken into account in design [54]. Three representative designs to make computationally viable ViTs are (1) the use of a hierarchical architecture with the spatial resolution (i.e., the token sequence length) progressively down-sampled across the stages [11,54,7], (2) the locally-grouped self-attention mechanisms for controlling the length of input token sequences and parameter sharing [30,15], (3) the pooling attention schemes to subsample the **key** and **value** by a factor [53,34,59,54]. The general trend has been on designing more complicated and stronger ViTs to challenge the dominance of top-performing CNNs [19,39,47] in computer vision by achieving ever higher accuracies [52,65]. These advances however are still insufficient to satisfy the design requirements and constraints for mobile and edge platforms (e.g., smart phones, robotics, self-driving cars, AR/VR devices), where the vision tasks need to carry out in a timely manner under certain computational budgets. Prior efficient CNNs (e.g., MobileNets [22,42,21], ShuffleNets [35,63], EfficientNets [47,50,48], and etc.) remain the state-of-the-art network architectures for such platforms in the tradeoff between running latency and recognition accuracy (Fig. 1).

In this work, we focus on the development of largely under-studied *efficient ViTs* with the aim to surpass the CNN counterparts on mobile devices. We consider a collection of very practical design requirements for running a ViT model on a target real-world platform as follows: **(1)** *Inference efficiency* needs to be high (e.g., *low latency and energy consumption*) so that the running cost becomes generically affordable and more on-device applications can be supported. This is a direct metric that we really care about in practice. In contrast, the often-used efficiency metric, FLOPs (i.e., the number of multiply-adds), cannot

directly translate into the latency and energy consumption on a specific device, with several conditional factors including memory access cost, degree of parallelism, and the platform’s characteristics [35]. This is, *not* all operations of a model can be carried out at the same speed and energy cost on a device. Hence, FLOPs is merely an approximate and indirect metric of efficiency. **(2) Model size** (i.e., parameter number) is affordable for modern average devices. Given the availability of ever cheaper and larger storage spaces, this constraint has been relaxed significantly. For example, an average smart phone often comes with 32GB or more storage. As a consequence, using it as a threshold metric is no longer valid in most cases. **(3) Implementational friendliness** is also critical in real-world applications. For a wider range of deployment, it is necessary that a model can be implemented efficiently using the standard computing operations supported and optimized in the generic deep learning frameworks (e.g., ONNX, TensorRT, and TorchScript), without costly per-framework specialization. Otherwise, the on-device speed of a model might be unsatisfactory even with low FLOPs. For instance, the cyclic shift and its reverse operations introduced in Swin Transformers [30] are rarely supported by the mainstream frameworks, i.e., deployment unfriendly. In the literature, very recent MobileViTs [36] are the only series of ViTs designed for mobile devices. In architecture design, they are a straightforward combination of MobileNetv2 [42] and ViTs [9]. As a very initial attempt in this direction, MobileViTs still lag behind CNN counterparts. Further, its evaluation protocol takes the *model size* (i.e. the parameter number) as the competitor selection criteria (i.e., comparing the accuracy of models only with similar parameter numbers), which however is no longer a hard constraint with modern hardware as discussed above and is hence out of date.

We present a family of light-weight attention based vision models, dubbed as **EdgeViTs**, for the first time, enabling ViTs to compete with the best light-weight CNNs (e.g., MobileNetv2 [42] and EfficientNets [47]) in terms of accuracy-efficiency tradeoff on mobile devices. This sets a milestone in the landscape of light-weight ViTs *vs.* CNNs in the low resource regime. Our EdgeViTs are based on a novel factorization of the standard self-attention for more cost-effective information exchange within every individual layer. This is made possible by introducing a highly light-weight and easy-to-implement *local-global-local* (LGL) information exchange bottleneck characterized with three operations: **(i)** Local information aggregation from neighbor tokens (each corresponding to a specific patch) using efficient depth-wise convolutions; **(ii)** Forming a sparse set of evenly distributed *delegate tokens* for long-range information exchange by self-attention; **(iii)** Diffusing updated information from *delegate tokens* to the *non-delegate tokens* in local neighborhoods via transposed convolutions. As we show in experiments, this design presents a favorable hybrid of self-attention, convolutions, and transposed convolutions, achieving the best accuracy-efficiency tradeoff. It is efficient in that the self-attention is applied to a sparse set of delegate tokens. To support a variety of computational budgets, with our primitive module we establish a family of EdgeViT variants with three computational complexities: small (S), extra-small (XS), extra-extra-small (XXS).

We make the following **contributions**: **(1)** We investigate the design of light-weight ViTs from the practical on-device deployment and execution perspective. **(2)** For best scalability and deployment, we present a novel family of efficient ViTs, termed as EdgeViTs, designed based on an optimal decomposition of self-attention using standard primitive operations. **(3)** Regarding on-device performance, towards relevance for real-world deployment, we directly consider latency and energy consumption of different models rather than relying on high-level proxies like number of FLOPs or parameters. Our results experimentally verify efficiency of our models in a practical setting and refute some of the claims made in the existing literature. More specifically, extensive experiments on three visual tasks show that our EdgeViTs can match or surpass state-of-the-art light-weight CNNs, whilst consistently outperform the recent MobileViTs in accuracy-efficiency tradeoff, including largely ignored on-device energy evaluation. Importantly, EdgeViTs are consistently Pareto-optimal in terms of both latency and energy efficiency, achieving strict dominance over other ViTs in almost all cases and competing with the most efficient CNNs. On ImageNet classification our EdgeViT-XXS outperforms MobileNetv2 by 2.2% subject to the similar energy-aware efficiency.

## 2 Related Work

**Efficient CNNs.** Since the advent of modern CNN architectures [19,45], there has been a steady stream of works focusing on efficient architecture design for on-device deployment. The first widely adopted families bring depthwise separable convolutions in a ResNet-like structure, e.g., MobileNets [22,42], ShuffleNets [63,35]. These works define a space of well-performing efficient architectures, resulting in widespread usage. Successive works further exploit this design space by automating the architectural design choices [47,21,46,49]. As a parallel line of research, net pruning creates efficient architectures by removing spurious parts of a larger network with close-to-zero weights [16,56], or via first training a super-network that is further slimmed to meet a pre-specified computational budget [31,3]. Dynamic computing has also been explored, consisting of the mechanisms that condition the network parameters on the input data [61,6]. Finally, using low bit-width is a very critical technique that can offer different tradeoffs between the accuracy and efficiency [16,23,5].

**Vision transformers.** ViTs [9] quickly popularize transformer-based architectures for computer vision. A series of works followed instantly, offering large improvements to the original ViTs in terms of data efficiency [51,29] and architecture design [30,62,11,4]. Among these works, one of the main modifications is to introduce hierarchical designs in multiple stages from convolutional architectures [30,7,54,55]. Several works also focus on improving the positional encoding by using a relative positional embedding [44,43], making it learnable [13], or even replacing it by an attention bias element [14]. All these approaches mostly aim to improve the model performance.

Recently, more efforts have been made towards finding efficient alternatives to the multi-head self-attention (MHSA) module, which is typically the computational bottleneck in the ViT architectures. A particularly effective solution is to reduce the internal spatial dimensions within the MHSA. The MHSA involves projecting the input tensor into key, query and value tensors. Several recent works, e.g. [7,54,55], find that the key and value tensors could be downsampled with a limited loss in accuracy, leading to a better efficiency-accuracy tradeoff. Our work extends this idea by also downsampling the query tensors, which further improves the efficiency, as shown in Fig. 2.

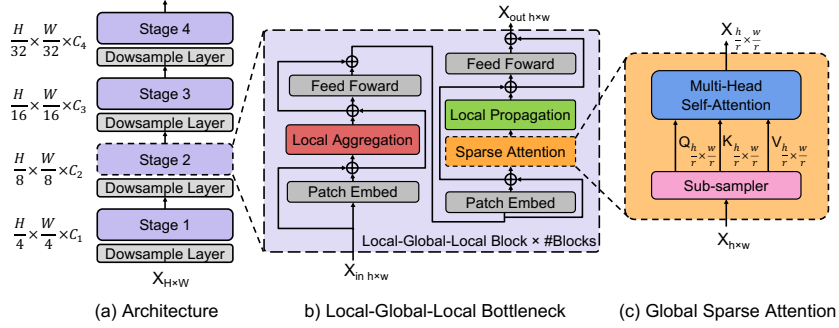
There are also alternative approaches reducing the number of tokens dynamically [40,60,37,12]. That is, in the forward pass, tokens deemed to not contain the important information for the target task are pruned or pooled together, reducing the overall complexity thereafter. Finally, encouraged by their potential complementarity, many works have attempted to combine convolutional designs with self-attentions. This ranges from using convolutions at the stem [58], integrating convolutional operations into the MHSA block [25,57], or incorporating the MHSA block into ResNet-like architectures [44]. It is interesting to note that even the original ViTs explored similar tradeoffs. [9].

**Vision transformers for mobile devices.** Whilst the efficiency issue has been taken into account in designing the ViT variants discussed above, they are still not dedicated and satisfactory architectures for on-device applications. There is only one exception, MobileViTs [36], which are introduced very recently. However, compared to the current best light-weight CNNs such as MobileNets [42,21] and EfficientNets [47], these ViTs are still clearly inferior in terms of the on-device accuracy-efficiency tradeoff. In this work, we present the first family of efficient ViTs that can deliver comparable or even superior tradeoffs in comparison to the best CNNs and ViTs. We also extensively carry out the critical yet largely lacking on-device evaluations with energy consumption analysis.

### 3 EdgeViTs

#### 3.1 Overview

For designing light-weight ViTs suitable for mobile/edge devices, we adopt a hierarchical pyramid network structure (Fig. 2(a)) used in recent ViT variants [54,8,7,55,11]. A pyramid transformer model typically reduces the spatial resolution but expands the channel dimension across different stages. Each stage consists of multiple transformer-based blocks processing tensors of the same shape, mimicking the ResNet-like networks. The transformer-based blocks heavily rely on the self-attention operations at a quadratic complexity w.r.t the spatial resolution of the visual features. By progressively aggregating the spatial tokens, pyramid vision transformers are potentially more efficient than isotropic models [9]. In this work, we dive deeper into the transformer-based block and introduce a cost-effective bottleneck, *Local-Global-Local* (LGL) (Fig. 2(b)). LGL further reduces the overhead of self-attention with a sparse attention module (Fig. 2(c)), achieving better accuracy-latency balancing.




**Fig. 2.** (a) Schematic overview of our four stages EdgeViT architecture, with each stage consisting of a stack of (b) *Local-Global-Local* (LGL) blocks constructed with local aggregation module, sparse-self-attention and local propagation module, patch embedding (PE) and Feed Forward Network (FFN). In this example,  $h$  and  $w$  refer to input height and width of stage-2:  $h = \frac{H}{8}$  and  $w = \frac{W}{8}$ .  $C_i$  refers to the number of channels for stage- $i$  and  $r$  denotes the sub-sampling rate.

### 3.2 *Local-Global-Local* bottleneck

Self-attention has been shown to be very effective for learning the global context or long-range spatial dependency of an image, which is critical for visual recognition. On the other hand, as images have high spatial redundancy (e.g., nearby patches are semantically similar) [17], applying attention to all the tokens, even in a down-sampled feature map, is inefficient. There is hence an opportunity to reduce the scope of tokens whilst still preserving the underlying information flows that model the global and local contexts. In contrast to previous transformer blocks that perform self-attention at each spatial location, our LGL bottleneck only computes self-attention for a subset of the tokens but enables full spatial interactions, as in the standard multi-head self-attention (MHSA) [9].

To achieve this, we decompose the self-attention into consecutive modules that process the spatial tokens within different ranges (Fig. 2(b)). We introduce three efficient operations: i) *Local aggregation* that integrates signals only from locally proximate tokens; ii) *Global sparse attention* that model long-range relations among a set of delegate tokens where each of them is treated as a representative for a local window; iii) *Local propagation* that diffuses the global contextual information learned by the delegates to the non-delegate tokens with the same window. Combining these, our LGL bottleneck enables information exchanges between any pair of tokens in the same feature map at a low-compute cost. Each of these components is described in detail below:

- *Local aggregation*: for each token, we leverage depth-wise and point-wise convolutions to aggregate information in local windows with a size of  $k \times k$  (Fig. 3(a)).
- *Global sparse attention*: we sample a sparse set of *delegate tokens* distributed evenly across the space, one *token* for each  $r \times r$  window. Here,  $r$  denotes the



**Fig. 3.** Illustration of three key operations involved in the proposed *Local-Global-Local* (LGL) transformer block. In this example, we showcase how *the target token* (the orange square) at the center conducts information exchange with all the others in three sequential steps: (a) Local information from neighbor tokens within the yellow area is first aggregated to the target token. (b) Global sparse attention is then computed among the target token and other selected delegates in orange color. (c) Global context information encoded in the target token is last propagated to its neighbor *non-delegate* tokens within the pink area.

sub-sample rate. We then apply self-attention on these selected tokens only (Fig. 3(b)). This is distinct from all the existing ViTs [7,54,55] where all the spatial tokens are involved as queries in the self-attention computation.

- *Local propagation*: We propagate the global contextual information encoded in the delegate tokens to their neighbor tokens by transposed convolutions (Fig. 3(c)).

Formally, our LGL bottleneck can be formulated as:

$$\begin{aligned}
 X &= \text{LocalAgg}(\text{Norm}(X_{in})) + X_{in}, \\
 Y &= \text{FFN}(\text{Norm}(X)) + X, \\
 Z &= \text{LocalProp}(\text{GlobalSparseAttn}(\text{Norm}(Y))) + Y, \\
 X_{out} &= \text{FFN}(\text{Norm}(Z)) + Z.
 \end{aligned} \tag{1}$$

Here  $X_{in} \in \mathcal{R}^{H \times W \times C}$  indicates the input tensors. **Norm** is the layer normalization operation [2]. **LocalAgg** represents the local aggregation operator, **FFN** is a two-layer perceptron, similar to the position-wise feed-forward network introduced in [9]. **GlobalSparseAttn** is the global sparse self-attention. **LocalProp** is the local propagation operator. For simplicity, positional encoding is omitted. Note that, all these operators can be implemented by commonly used and highly optimized operations in the standard deep learning platforms. Hence, our LGL bottleneck is implementation friendly.

**Comparisons to existing designs.** Our LGL bottleneck shares a similar goal with the recent PVTs [54,55] and Twins-SVTs [7] models that attempt to reduce the self-attention overhead. However, they differ in the core design. PVTs [54,55] perform self-attention where the number of **keys** and **values** are reduced by strided-convolutions, whilst the number of **queries** remains the same.

Model	#Channels	#Blocks	#Heads	FLOPs	#Param
EdgeViT-XXS	[36, 72, 144, 288]	[1,1,3,2]	[1,2,4,8]	0.56G	4.1M
EdgeViT-XS	[48, 96, 240, 384]	[1,1,2,2]	[1,2,4,8]	1.1G	6.7M
EdgeViT-S	[48, 96, 240, 384]	[1,2,3,2]	[1,2,4,8]	1.9G	11.1M

**Table 1. Configuration of three EdgeViT variants.** “#Channels”: number of channels per stage. “#Blocks”: number of LGL blocks per stage. “#Heads”: number of attention heads in MHSA. “#Param”: number of parameters.

In other words, PVTs still perform self-attention at each grid location. In this work, we question the necessity of positional-wise self-attention and explore to what extent the information exchange enabled by our LGL bottleneck could approximate the standard MHSA (see Section 4 for more details). Twins-SVTs [7] combine local-window self-attention [30] with global pooled attention from PVTs [54]. This is different from our hybrid design using both self-attention and convolution operations distributed in a series of *local-global-local* operations. As demonstrated in the experiments (Table 2 and 3), our design achieve a better tradeoff between the model performance and the computation overhead (e.g. latency, energy consumption, etc).

### 3.3 Architectures

We build a family of EdgeViTs with the proposed LGL bottleneck at different computational complexities (i.e. 0.5G, 1G, and 2G). The configurations are summarized in Table 1. Following the hierarchical ViTs [7,54,30,25], EdgeViTs consist of four stages with the spatial resolution (*i.e.*, the token sequence length) gradually reduced throughout, and their self-attention module replaced with our LGL bottleneck. For the stage-wise down-sampling, we use a conv-layer with a kernel size of  $2 \times 2$  and stride 2, except for the first stage where we down-sample the input feature by  $\times 4$ , and use a  $4 \times 4$  kernel and a stride of 4. We adopt the conditional positional encoding [8] that has been shown to be superior to the absolute positional encoding. This can be implemented using 2D depth-wise convolutions with a residual connection. In our model, we use  $3 \times 3$  depth-wise convolutions with zero paddings. It is placed before the local aggregation and global sparse self-attention. The FFN consists of two linear layers with GeLU non-linearity [20] placed in-between. Our local aggregation operator is implemented as a stack of pointwise and depthwise convolutions. The global sparse attention is composed of a spatial uniform sampler with sample rates of (4, 2, 2, 1) for the four stages, and a standard MHSA. The local propagation is implemented with a depthwise separable transposed convolution with the kernel size and stride equal to the sample rate used in the global sparse attention. The exact architecture for the LGL bottleneck is described in the **supplementary material**.



## 4 Experiments

We benchmark EdgeViTs on visual recognition tasks. We pre-train EdgeViTs on the Imagenet1K recognition task [41], comparing the performances and computation overheads against alternative approaches. We also evaluate the generalization capacity of EdgeViTs on downstream dense prediction tasks: object detection and instance segmentation on the COCO benchmark [28], and semantic segmentation on the ADE20K Scene Parsing benchmark [64]. For on-device execution, we report execution time (latency) and energy consumption of all relevant models on ImageNet. We do not report on-device measurements on downstream tasks as they reuse ImageNet models.

### 4.1 Image Classification on ImageNet-1K

**Training Settings.** ImageNet-1K [41] provides 1.28 million training images and 50,000 validation images from 1000 categories. We follow the training recipe introduced in DeiT [51]. We optimize the models using AdamW [32] with a batch size of 1024, weight decay of  $5 \times 10^{-2}$ , and momentum of 0.9. The models are trained from *scratch* for 300 epochs with a linear warm-up during the first 5 epochs. Our base learning rate is set as  $1 \times 10^{-3}$ , and decay after the warm-up using a cosine schedule [33]. We apply the same data augmentations as in [51,7,55,25] which include random cropping, random horizontal flipping, mixup, random erasing and label-smoothing. During training, the images are randomly cropped to  $224 \times 224$ . During testing, we use a single center crop of  $224 \times 224$ . We report the top-1 accuracy on the validation set.

**Benchmarking Settings.** For latency measurements, we use a Samsung Galaxy S21 mobile phone equipped with a Snapdragon 888 chipset. All relevant models are benchmarked by running a forward pass 50 times using TorchScript lite interpreter via the Android benchmarking app provided by PyTorch [38]. We use CPU implementation, full precision and batch 1 to execute all operations. This choice comes from the fact that this is the only combination that was able to robustly execute all of the models from our paper. In general, more efficient implementations exist, such as those utilizing specialized hardware like Neural Processing Units (NPUs). However, these put more restrictions on what can be executed and many models failed to run in our experiments when trying to use different hardware targets.

For energy measurements, we use a Monsoon High Voltage Power Monitor connected to a Snapdragon 888 Hardware Development Kit (HDK8350) to obtain accurate power readings over the course of running a forward pass of each test model 50 times. The same TorchScript runtime is used as in latency measurements. From the power signal reported by the monitor, we derive the average per-inference power and energy consumption by first subtracting background power consumption (i.e., power readings when not running any model) and then identifying 50 continuous regions of significantly higher power draw. Each region like that is considered a single inference and we calculate its total energy

Model	#Params	FLOPs	CPU(ms)	Acc Top-1 (%)
MobileNet-v2[42]	3.4M	0.3G	33.3 $\pm$ 5.3	72.0
<b>MobileNet-v3 0.75</b> [21]	<b>4.0M</b>	<b>0.16G</b>	<b>23.0<math>\pm</math>3.7</b>	<b>73.3</b>
<b>EfficientNet-B0</b> [47]	<b>5.3M</b>	<b>0.4G</b>	<b>52.1<math>\pm</math>7.4</b>	<b>77.1</b>
MobileViT-XXS [36]	1.3M	0.4G	69.5 $\pm$ 5.1	69.0
PVT-v2-B0[55]	3.4M	0.6G	26.0 $\pm$ 6.9	70.5
Uniformer-Tiny*[25]	3.9M	0.6G	40.5 $\pm$ 3.1	74.1
Twins-SVT-Tiny*[7]	4.1M	0.6G	36.9 $\pm$ 2.3	71.2
<b>EdgeViT-XXS</b>	<b>4.1M</b>	<b>0.6G</b>	<b>32.8<math>\pm</math>2.7</b>	<b>74.4</b>
T2T-ViT-7 [62]	4.3M	1.1G	48.8 $\pm$ 6.5	71.7
MobileViT-XS [36]	2.4M	1.1G	150.1 $\pm$ 6.1	74.7
DeiT-Tiny [51]	5.7M	1.3G	46.2 $\pm$ 13.6	72.2
TNT-Tiny [15]	6.1M	1.4G	86.4 $\pm$ 6.0	73.9
<b>EdgeViT-XS</b>	<b>6.7M</b>	<b>1.1G</b>	<b>54.1<math>\pm</math>2.2</b>	<b>77.5</b>
T2T-ViT-12 [62]	6.9M	1.9G	69.9 $\pm$ 5.6	76.5
PVT-v2-B1[55]	14M	2.1G	75.4 $\pm$ 2.3	78.7
MobileViT-S [36]	5.6M	2.0G	221.3 $\pm$ 9.3	78.3
<b>LeViT-384<math>\dagger</math></b> [14]	<b>39.1M</b>	<b>2.4G</b>	<b>71.3<math>\pm</math>2.3</b>	<b>79.5</b>
<b>EdgeViT-S</b>	<b>11.1M</b>	<b>1.9G</b>	<b>85.3<math>\pm</math>3.9</b>	<b>81.0</b>

**Table 2. Results on ImageNet-1K validation set.** All models are tested on input scale of  $224 \times 224$ , except for MobileViTs [36] that are tested with  $256 \times 256$  according to their original implementation. \* indicates down-scaled architectures beyond original definitions by authors to fit the mobile compute budget. LeViT-384 $\dagger$  denotes the LeViT model retrained under the same setting as our EdgeViT.

as the integral over the individual power samples. Analogously we also calculate average power consumption by averaging over the same set of samples. After energy and power are calculated for each inference, the final statistics of a model are obtained by again averaging over the 50 identified runs. Our methodology follows what can be found in the literature [1].

**Results.** We compare EdgeViTs to a variety of baseline models, including the classic efficient CNNs, e.g. MobileNetV2 [42], MobileNetV3 [21], EfficientNet [47], and the state-of-the-art ViTs, e.g. MobileViT [36], PVT-v2 [55], DeiT [51], LeViT [14]. As the original LeViT [14] was optimized in a large-scale setting (i.e. 1000 epochs) with knowledge distillation, we perform a comparison by re-training LeViT under the same setting (300 epochs) as EdgeViT without knowledge distillation. We denote the retrained LeViTs as LeViT-384 $\dagger$ . We select the baselines with a complexity of less than 2 GFLOPs as i) in real-world applications, the computational cost remains the top concern; ii) whilst FLOPs is an indirect metric for the latency, it is the most used cost metric in prior works. This selection criterion is different from [36] that instead uses the model size (i.e. the parameter number) which however has become a less restricted facet in mobile devices.

From Table 2, we can learn: i) EdgeViTs significantly outperform other lightweight *ViTs* at a similar level of GFLOPs complexity. Compared to the PVT-

Model	Top-1 (%)	CPU (ms)	Energy (mJ)	Power (W)	Efficiency (%/msW)
MobileNet-v2[42]	72.0	33.3	85.7 $\pm$ 7.4	3.31 $\pm$ 0.26	0.841
MobileNet-v3 0.75[21]	<b>73.3</b>	<b>23.0</b>	<b>63.0<math>\pm</math>9.6</b>	<b>3.46<math>\pm</math>0.4</b>	<b>1.164</b>
EfficientNet-B0[47]	<b>77.1</b>	<b>52.1</b>	<b>159.0<math>\pm</math>26.2</b>	<b>3.62<math>\pm</math>0.45</b>	<b>0.485</b>
PVT-v2-B0[55]	70.5	26.0	91.7 $\pm$ 19.7	3.94 $\pm$ 0.68	0.769
PVT-v2-B1[55]	<b>78.7</b>	75.4	309.0 $\pm$ 65.8	4.63 $\pm$ 0.71	0.255
Twins-SVT-Tiny*[7]	71.2	36.9	114.5 $\pm$ 17.3	3.71 $\pm$ 0.24	0.622
DeiT-Tiny [51]	72.2	46.2	187.2 $\pm$ 7.6	4.77 $\pm$ 0.21	0.386
Uniformer-Tiny*[25]	74.1	40.5	134.7 $\pm$ 27.3	4.1 $\pm$ 0.71	0.55
T2T-ViT-12 [62]	76.5	69.9	266.2 $\pm$ 42.6	4.37 $\pm$ 0.36	0.287
TNT-Tiny [15]	73.9	86.4	308.7 $\pm$ 70.5	3.94 $\pm$ 0.63	0.239
LeViT-384† [14]	<b>79.5</b>	<b>71.3<math>\pm</math>2.2</b>	<b>455.2<math>\pm</math>125.8</b>	<b>6.18<math>\pm</math>0.74</b>	<b>0.173</b>
MobileViT-XXS [36]	69.0	69.5	175.3 $\pm$ 28.7	2.77 $\pm$ 0.24	0.394
MobileViT-XS [36]	74.7	150.1	251.5 $\pm$ 81.1	2.63 $\pm$ 0.61	0.297
MobileViT-S [36]	78.3	221.3	503.6 $\pm$ 117.0	2.76 $\pm$ 0.21	0.155
<b>EdgeViT-XXS</b>	<b>74.4</b>	<b>32.8</b>	<b>127.4<math>\pm</math>27.3</b>	<b>4.27<math>\pm</math>0.67</b>	<b>0.584</b>
<b>EdgeViT-XS</b>	<b>77.5</b>	<b>54.1</b>	<b>234.6<math>\pm</math>44.0</b>	<b>4.77<math>\pm</math>0.84</b>	<b>0.33</b>
<b>EdgeViT-S</b>	<b>81.0</b>	<b>85.3</b>	<b>386.7<math>\pm</math>43.5</b>	<b>4.8<math>\pm</math>0.26</b>	<b>0.209</b>

**Table 3. On-device energy evaluation on ImageNet-1K.** All relevant metrics are reported as mean values per forward pass across 50 executions. For facilitating comparison, we define an energy-aware *efficiency* metric as the average gain in top-1 accuracy from each 1W run for 1ms (equivalent to consuming 1mJ of energy). (*Pareto-optimal models* are highlighted in bold in the last column).

v2 family [55], our EdgeViT-XXS/EdgeViT-S achieve 3.9%/2.3% improvements over PVT-v2-B0/PVT-v2-B1. Compared to MobileViTs, EdgeViTs achieve 5.4%, 2.8% and 2.7% gains in the three complexity settings. ii) **ViTs vs. CNNs:** Our EdgeViTs lift the performance of efficient ViTs to approach the level of well-established efficient CNNs. For example, the EdgeViT-XXS performs superior to MobileNet-v2 and MobileNet-v3-0.75 at a similar level of model size, but requires more GFLOPs. However, we observe that the efficient CNNs still surpass efficient ViTs in the accuracy-FLOPs tradeoff by a small margin.

On the other hand, as discussed early, numbers of FLOPs or parameters are merely indicative but do not fully reflect the on-device efficiency [35,10]. We further consider on-device latency and energy consumption directly. Other than the representative ViTs and CNNs, we also compare two recent ViT variants [25,7] with the number of channels and layers re-scaled to fit the complexity need. As presented in Table 2, EdgeViTs demonstrate strong performance with latencies comparable to MobileNets: EdgeViT-XXS achieves a gain of 2.4% over MobileNet-V2 while running slightly faster. EdgeViT-XXS also surpasses MobileNet-V3 by 1.1% but at the cost of being 9.8ms slower. EdgeViT-XS performs on par with the auto-searched EfficientNet-B0 model. We believe our models could also benefit from the automatic architecture search techniques as use

	CPU	Top1	CPU	Top1	CPU	Top1
w/o LA	33.9ms	72.7 max	34.8ms	74.3 w/o LP	<b>32.4ms</b>	73.9
LA(LSA)	36.1ms	74.0 avg	34.5ms	<b>74.5</b> LP(Bilinear)	34.1ms	74.1
LA(Ours)	<b>32.8ms</b>	<b>74.4</b> center	<b>32.8ms</b>	74.4 LP(Ours)	32.8ms	<b>74.4</b>

(a) **Local Aggregation** (b) **Global attention** (c) **Local Propagation****Table 4. Ablation on ImageNet-1K.** LA: the local aggregation operator. LP: the local propagation operator. LSA: the Locally-grouped Self-Attention used in [7].

in MobileNet-V3 and EfficientNets. Our models yield clear advantages over alternative ViT models. Compared to MobileViTs in the three GFLOPs settings, EdgeViTs excel by 5.4%, 2.8%, and 2.7% while being  $\times 2$ ,  $\times 2.7$ ,  $\times 2.6$  faster.

Energy results are presented in Table 3. In addition to the raw energy and power numbers, for comparison simplicity, we define an energy-aware *efficiency* metric as the average gain in top-1 accuracy (in percentages) from each consumed 1mJ of energy. We observe that less accurate models tends to be more efficient. This is not a surprise in that improvements in accuracy scale sublinearly with model complexity. However, this also means that comparing efficiency of models with very different top-1 scores might be severely biased by the sole difficulty of achieving certain accuracy levels, which is independent from a model. Therefore, we limit our comparison to identifying *pareto-optimal models*, those upon which no other models can improve in either accuracy or energy efficiency without degrading other metrics. We can see that our EdgeViT family is able to dominate almost all other ViTs, with the only exception being LeViT-384 $\dagger$  whose accuracy and efficiency fall between our EdgeViT-S and EdgeViT-XS. When compared to CNNs, our EdgeViTs compete with MobileNet-v3 and EfficientNet-B0 that are more efficient but also less accurate. MobileNet-v2 achieves decent results but is dominated by its newer version, MobileNet-v3. PVT-v2-B0, although high on the efficiency side, is rather inaccurate and hence is favored by highly efficient CNNs. Visibly at the end of the spectrum are the latest MobileViT models which turn out to be neither efficient nor accurate, when compared to the rest. Unlike them, our EdgeViT models, although not as efficient as best CNNs in the absolute sense, exhibit favourable trade-off between efficiency and accuracy by being rather highly accurate while not sacrificing efficiency.

**Ablation study.** We conduct detailed ablations to validate our design choices in the LGL bottleneck. We use EdgeViT-XXS as the base model and re-scale the alternative designs to  $\sim 0.5$ GFLOPs for fair comparison.

*Local aggregation.* We compare our local aggregation (LA) operation to the Locally-grouped Self-Attention (LSA) used in [7,30]. It is shown in Table 4a that applying LA consistently improve the performance. Our convolutional LA mod-

Backbone	RetinaNet 1×							Mask R-CNN 1×						
	#Par.	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	#Par.	AP <sup>b</sup>	AP <sub>50</sub> <sup>b</sup>	AP <sub>75</sub> <sup>b</sup>	AP <sup>m</sup>	AP <sub>50</sub> <sup>m</sup>	AP <sub>75</sub> <sup>m</sup>
PVTv2-B0 [55]	<b>13.0</b>	37.2	57.2	39.5	<b>23.1</b>	40.4	49.7	<b>23.5</b>	38.2	60.5	40.7	36.2	57.8	38.6
<b>EdgeViT-XXS</b>	13.1	<b>38.7</b>	<b>59.0</b>	<b>41.0</b>	22.4	<b>42.0</b>	<b>51.6</b>	23.8	<b>39.9</b>	<b>62.0</b>	<b>43.1</b>	<b>36.9</b>	<b>59.0</b>	<b>39.4</b>
<b>EdgeViT-XS</b>	16.3	40.6	61.3	43.3	25.2	43.9	54.6	26.5	41.4	63.7	45.0	38.3	60.9	41.3
ResNet18 [19]	<b>21.3</b>	31.8	49.6	33.6	16.3	34.3	43.2	<b>31.2</b>	34.0	54.0	36.7	31.2	51.0	32.7
PVTv1-Tiny [54]	23.0	36.7	56.9	38.9	22.6	38.8	50.0	32.9	36.7	59.2	39.3	35.1	56.7	37.3
PVTv2-B1 [55]	23.8	41.2	61.9	43.9	25.4	44.5	54.3	33.7	41.8	64.3	45.9	38.8	61.2	41.6
<b>EdgeViT-S</b>	22.6	<b>43.4</b>	<b>64.9</b>	<b>46.5</b>	<b>26.9</b>	<b>47.5</b>	<b>58.1</b>	32.8	<b>44.8</b>	<b>67.4</b>	<b>48.9</b>	<b>41.0</b>	<b>64.2</b>	<b>43.8</b>

**Table 5.** Comparison to other visual backbones using RetinaNet and Mask-RCNN on COCO val2017 object detection and instance segmentation. “#Par.” refers to number of parameters in million. AP<sup>b</sup> and AP<sup>m</sup> indicate bounding box AP and mask AP.

ule performs better than the self-attention based operator (LSA). This validates our choice of using **depth-wise convolutions** in LA for local context learning.

*Global sparse attention.* We explore three options for delegate token sampling: **max**, **avg**, and **center**. All choices perform similarly in terms of accuracy, with our default design **center** being slightly faster.

*Local propagation.* We investigate two alternatives to the local propagation operator: i) **w/o LP**: We simply remove the local propagation. Note that EdgeViTsw/o LP has similar complexity to standard EdgeViTs. ii) **Bilinear**: we use the bilinear interpolation, instead of the transposed convolution, to up-sample the delegate tokens. Table 4c shows that adding LP improves the top-1 accuracy by 0.5%, with only 0.4ms overhead.

## 4.2 Dense Prediction

Following [54,55], we also evaluate the proposed EdgeViTs on COCO Objection Detection/Instance Segmentation [28] and ADE20K Scene Parsing [64]. Here, we use the EdgeViTs as the feature extractor for the main model and initialize it with the ImageNet1K-pretrained weights obtained in our previous experiments.

**COCO Object Detection/Instance Segmentation.** We demonstrate the performance of our model in main-stream object detection and instance segmentation frameworks: RetinaNet [27] for object detection, Mask R-CNN[18] with the FPN [26] for instance segmentation. Following the training protocol in [55,7,25], we resize the training images to have a shorter side of 800 pixels while keeping the longer side to be smaller than 1333 pixels. During testing, the images are re-scaled to have a shorter size of 800 pixels. The models are fine-tuned with 1× schedule (i.e. 12 epochs) by AdamW[32] using an initial learning rate of  $1 \times 10^{-4}$  and a batch size of 16. We train the models on the COCO 2017 training set and report the mAP@100 score on the COCO 2017 validation set.

**Results.** In Table 5, our EdgeViTs perform consistently better than other visual backbones on RetinaNet [27] and Mask R-CNN[18]. Our smallest vari-

Backbone	Semantic FPN		
	#Param (M)	GFLOPs	mIoU (%)
PVTv2-B0[55]	<b>7.6</b>	<b>25.0</b>	37.2
<b>EdgeViT-XXS</b>	7.9	24.4	<b>39.7</b>
<b>EdgeViT-XS</b>	10.6	27.7	41.4
ResNet18 [19]	<b>15.5</b>	32.2	32.9
PVTv1-Tiny [54]	17.0	33.2	35.7
PVTv2-B1[55]	17.8	34.2	42.5
<b>EdgeViT-S</b>	16.9	<b>32.1</b>	<b>45.9</b>

**Table 6. Semantic segmentation results on the validation set of ADE20K.** Segmentation model: Semantic FPN [24]. GFLOPs: Calculated at  $512 \times 512$  input size.

ant EdgeViT-XXS, when used on RetinaNet [27], achieves 1.5 higher AP than PVTv2-B0. When used on Mask R-CNN [18], EdgeViT-XXS also surpasses PVTv2-B0 by 1.7 on the bounding box detection task ( $AP^b$ ), and by 0.7 on the mask segmentation task ( $AP^m$ ). For EdgeViT-S, we observe even larger gains when comparing to PVTv2-B1: +2.2 on RetinaNet [27], +3.0  $AP^b$  and +1.2  $AP^m$  on Mask R-CNN[18].

**ADE20K Scene Parsing.** We incorporate the pretrained EdgeViT in the Semantic FPN segmentation model [24]. As in [54,55], we create  $512 \times 512$  random crops of the images during training and resize the images to have a shorter side of 512 pixels during inference. The models are finetuned by AdamW [32] using an initial learning rate of  $1 \times 10^{-4}$  and a batch size of 16. We train the models for 80K iterations on the ADE20K training set, and report the mean Intersection over Union (mIoU) score on the validation set.

**Results** In Table 6, we compare EdgeViTs to both CNN (ResNet-18 [19]) and ViT backbones (PVTs[54,55]) for FPN based Semantic Segmentation [24]. EdgeViTs achieves better performance than all counterparts at similar compute costs. Particularly, EdgeViT-XXS outperforms PVTv2-B0 by 2.5% in mIoU, EdgeViT-S surpasses PVTv2-B1 by a margin of 3.4%.

## 5 Conclusion

In this work, we investigate the design of efficient ViTs from the on-device deployment perspective. By introducing a novel decomposition of self-attention, we present a family of EdgeViTs that, for the first time, achieve comparable or even superior accuracy-efficiency tradeoff on generic visual recognition tasks, in comparison to a variety of state-of-the-art efficient CNNs and ViTs. We conduct extensive on-device experiments using practically critical and previously underestimated metrics (e.g., energy-aware efficiency) and reveal new insights and observations in the comparison of light-weight CNN and ViT models.

## References

1. Almeida, M., Laskaridis, S., Mehrotra, A., Dudziak, L., Leontiadis, I., Lane, N.D.: Smart at what cost? characterising mobile deep neural networks in the wild. In: ACM Internet Measurement Conference (2021)
2. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
3. Berman, M., Pishchulin, L., Xu, N., Blaschko, M.B., Medioni, G.: AOWS: adaptive and optimal network width search with latency constraints. In: IEEE Conference on Computer Vision and Pattern Recognition (2020)
4. Bulat, A., Perez Rúa, J.M., Sudhakaran, S., Martinez, B., Tzimiropoulos, G.: Space-time mixing attention for video transformer. Advances on Neural Information Processing Systems (2021)
5. Bulat, A., Tzimiropoulos, G.: Bit-Mixer: Mixed-precision networks with runtime bit-width selection. In: IEEE International Conference on Computer Vision (2021)
6. Chen, Y., Dai, X., Liu, M., Chen, D., Yuan, L., Liu, Z.: Dynamic convolution: Attention over convolution kernels. In: IEEE Conference on Computer Vision and Pattern Recognition (2020)
7. Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C.: Twins: Revisiting the design of spatial attention in vision transformers. Advances on Neural Information Processing Systems (2021)
8. Chu, X., Tian, Z., Zhang, B., Wang, X., Wei, X., Xia, H., Shen, C.: Conditional positional encodings for vision transformers. arXiv preprint arXiv:2102.10882 (2021)
9. 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. In: International Conference on Learning Representations (2021)
10. Dudziak, L., Chau, T., Abdelfattah, M.S., Lee, R., Kim, H., Lane, N.D.: BRP-NAS: Prediction-based NAS using GCNs. In: Advances on Neural Information Processing Systems (2020)
11. Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., Feichtenhofer, C.: Multiscale vision transformers. In: IEEE International Conference on Computer Vision (2021)
12. Fayyaz, M., Koohpayegani, S.A., Jafari, F.R., Sommerlade, E., Joze, H.R.V., Pirsiavash, H., Gall, J.: ATS: adaptive token sampling for efficient vision transformers. arXiv preprint arXiv:2111.15667 (2021)
13. Gehring, J., Auli, M., Grangier, D., Yarats, D., Dauphin, Y.N.: Convolutional sequence to sequence learning. In: International Conference on Machine Learning (2017)
14. 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: IEEE International Conference on Computer Vision (2021)
15. Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. Advances on Neural Information Processing Systems (2021)
16. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. In: International Conference on Learning Representations (2016)
17. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: IEEE Conference on Computer Vision and Pattern Recognition (2022)

18. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: IEEE International Conference on Computer Vision (2017)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (2016)
20. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)
21. Howard, A., Pang, R., Adam, H., Le, Q.V., Sandler, M., Chen, B., Wang, W., Chen, L., Tan, M., Chu, G., Vasudevan, V., Zhu, Y.: Searching for MobileNetV3. In: IEEE International Conference on Computer Vision (2019)
22. 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)
23. Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., Kalenichenko, D.: Quantization and training of neural networks for efficient integer-arithmetic-only inference. In: IEEE Conference on Computer Vision and Pattern Recognition (2018)
24. Kirillov, A., Girshick, R., He, K., Dollár, P.: Panoptic feature pyramid networks. In: IEEE Conference on Computer Vision and Pattern Recognition (2019)
25. Li, K., Wang, Y., Peng, G., Song, G., Liu, Y., Li, H., Qiao, Y.: Uniformer: Unified transformer for efficient spatial-temporal representation learning. In: International Conference on Learning Representations (2022)
26. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: IEEE Conference on Computer Vision and Pattern Recognition (2017)
27. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: IEEE International Conference on Computer Vision (2017)
28. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: European Conference on Computer Vision (2014)
29. Liu, Y., Sangineto, E., Bi, W., Sebe, N., Lepri, B., De Nadai, M.: Efficient training of visual transformers with small datasets. *Advances on Neural Information Processing Systems* (2021)
30. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. *IEEE International Conference on Computer Vision* (2021)
31. Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., Zhang, C.: Learning efficient convolutional networks through network slimming. In: IEEE International Conference on Computer Vision (2017)
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
33. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with warm restarts. In: *International Conference on Learning Representations* (2017)
34. Lu, J., Yao, J., Zhang, J., Zhu, X., Xu, H., Gao, W., Xu, C., Xiang, T., Zhang, L.: Soft: Softmax-free transformer with linear complexity. *Advances on Neural Information Processing Systems* (2021)
35. Ma, N., Zhang, X., Zheng, H.T., Sun, J.: ShuffleNetV2: Practical guidelines for efficient CNN architecture design. In: *European Conference on Computer Vision* (2018)
36. Mehta, S., Rastegari, M.: MobileViT: light-weight, general-purpose, and mobile-friendly vision transformer. *International Conference on Learning Representations* (2022)



37. Pan, B., Jiang, Y., Panda, R., Wang, Z., Feris, R., Oliva, A.: IA-RED<sup>2</sup>: Interpretability-aware redundancy reduction for vision transformers. In: *Advances on Neural Information Processing Systems* (2021)
38. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An imperative style, high-performance deep learning library. In: *Advances on Neural Information Processing Systems* (2019)
39. Radosavovic, I., Kosaraju, R.P., Girshick, R., He, K., Dollár, P.: Designing network design spaces. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2020)
40. Rao, Y., Zhao, W., Liu, B., Lu, J., Zhou, J., Hsieh, C.J.: DynamicViT: Efficient vision transformers with dynamic token sparsification. *Advances on Neural Information Processing Systems* (2021)
41. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal on Computer Vision* (2015)
42. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: MobileNetV2: Inverted residuals and linear bottlenecks. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2018)
43. Shaw, P., Uszkoreit, J., Vaswani, A.: Self-attention with relative position representations. In: *North American Chapter of the Association for Computational Linguistics* (2018)
44. Srinivas, A., Lin, T.Y., Parmar, N., Shlens, J., Abbeel, P., Vaswani, A.: Bottleneck transformers for visual recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2021)
45. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2016)
46. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: MnasNet: Platform-aware neural architecture search for mobile. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2019)
47. Tan, M., Le, Q.: EfficientNet: Rethinking model scaling for convolutional neural networks. In: *International Conference on Machine Learning* (2019)
48. Tan, M., Le, Q.: EfficientNetV2: Smaller models and faster training. In: *International Conference on Machine Learning* (2021)
49. Tan, M., Le, Q.V.: Mixconv: Mixed depthwise convolutional kernels. In: *British Machine Vision Conference* (2019)
50. Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2020)
51. 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* (2021)
52. Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., Jégou, H.: Going deeper with image transformers. *arXiv preprint arXiv:2103.17239* (2021)
53. Wang, S., Li, B., Khabsa, M., Fang, H., Ma, H.: Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768* (2020)
54. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: *IEEE International Conference on Computer Vision* (2021)

55. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: PVTv2: Improved baselines with pyramid vision transformer. *Computational Visual Media* (2022)
56. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. In: *Advances on Neural Information Processing Systems* (2016)
57. Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L.: CvT: Introducing convolutions to vision transformers. *IEEE International Conference on Computer Vision* (2021)
58. Xiao, T., Singh, M., Mintun, E., Darrell, T., Dollár, P., Girshick, R.: Early convolutions help transformers see better. *Advances on Neural Information Processing Systems* (2021)
59. Xiong, Y., Zeng, Z., Chakraborty, R., Tan, M., Fung, G., Li, Y., Singh, V.: Nyströmformer: A nyström-based algorithm for approximating self-attention. In: *AAAI Conference on Artificial Intelligence* (2021)
60. Xu, Y., Zhang, Z., Zhang, M., Sheng, K., Li, K., Dong, W., Zhang, L., Xu, C., Sun, X.: Evo-vit: Slow-fast token evolution for dynamic vision transformer. In: *AAAI Conference on Artificial Intelligence* (2022)
61. Yang, B., Bender, G., Le, Q.V., Ngiam, J.: CondConv: Conditionally parameterized convolutions for efficient inference. In: *Advances on Neural Information Processing Systems* (2019)
62. Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token ViT: Training vision transformers from scratch on ImageNet. In: *IEEE International Conference on Computer Vision* (2021)
63. Zhang, X., Zhou, X., Lin, M., Sun, J.: ShuffleNet: An extremely efficient convolutional neural network for mobile devices. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2018)
64. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2017)
65. Zhou, D., Kang, B., Jin, X., Yang, L., Lian, X., Jiang, Z., Hou, Q., Feng, J.: Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886* (2021)