# EfficientViT: Multi-Scale Linear Attention for High-Resolution Dense Prediction

Han Cai<sup>1</sup>, Junyan Li<sup>2</sup>, Muyan Hu<sup>3</sup>, Chuang Gan<sup>4</sup>, Song Han<sup>1</sup> <sup>1</sup>MIT, <sup>2</sup>Zhejiang University, <sup>3</sup>Tsinghua University, <sup>4</sup>MIT-IBM Watson AI Lab https://github.com/mit-han-lab/efficientvit

## Abstract

High-resolution dense prediction enables many appealing real-world applications, such as computational photography, autonomous driving, etc. However, the vast computational cost makes deploying state-of-the-art highresolution dense prediction models on hardware devices difficult. This work presents EfficientViT, a new family of highresolution vision models with novel multi-scale linear attention. Unlike prior high-resolution dense prediction models that rely on heavy softmax attention, hardware-inefficient large-kernel convolution, or complicated topology structure to obtain good performances, our multi-scale linear attention achieves the global receptive field and multi-scale learning (two desirable features for high-resolution dense prediction) with only lightweight and hardware-efficient operations. As such, EfficientViT delivers remarkable performance gains over previous state-of-the-art models with significant speedup on diverse hardware platforms, including mobile CPU, edge GPU, and cloud GPU. Without performance loss on Cityscapes, our EfficientViT provides up to  $13.9 \times$  and  $6.2 \times$  GPU latency reduction over Seg-Former and SegNeXt, respectively. For super-resolution, EfficientViT delivers up to  $6.4 \times$  speedup over Restormer while providing 0.11dB gain in PSNR. For Segment Anything, EfficientViT delivers similar zero-shot image segmentation quality as ViT-Huge with  $84 \times$  higher throughput on GPU.

# 1. Introduction

High-resolution dense prediction is a fundamental task in computer vision and has broad applications in the real world, including autonomous driving, medical image processing, computational photography, etc. Therefore, deploying state-of-the-art (SOTA) high-resolution dense prediction models on hardware devices can benefit many use cases.

However, there is a large gap between the computational cost required by SOTA high-resolution dense prediction models and the limited resources of hardware devices. It makes using these models in real-world applications impractical. In particular, high-resolution dense prediction models require high-resolution images and strong context information extraction ability to work well [1, 2, 3, 4, 5, 6]. Therefore, directly porting efficient model architectures from image classification is unsuitable for high-resolution dense prediction.

This work introduces EfficientViT, a new family of vision transformer models for efficient high-resolution dense prediction. The core of EfficientViT is a new multi-scale linear attention module that enables the global receptive field and multi-scale learning with hardware-efficient operations. Our module is motivated by prior SOTA highresolution dense prediction models. They demonstrate that the multi-scale learning [3, 4] and global receptive field [7] are critical in improving models' performances. However, they do not consider hardware efficiency when designing their models, which is essential for real-world applications. For example, SegFormer [7] introduces softmax attention [8] into the backbone to have a global receptive field. However, its computational complexity is quadratic to the input resolution, making it unable to handle high-resolution images efficiently. SegNeXt [9] proposes a multi-branch module with large-kernel convolutions (kernel size up to 21) to enable a large receptive field and multiscale learning. However, large-kernel convolution requires exceptional support on hardware to achieve good efficiency [10, 11], which is usually unavailable on hardware devices.

Hence, the design principle of our module is to enable these two critical features while avoiding hardwareinefficient operations. Specifically, we propose substituting the inefficient softmax attention with lightweight ReLU linear attention [12] to have the global receptive field. By leveraging the associative property of matrix multiplication, ReLU linear attention can reduce the computational complexity from quadratic to linear while preserving functionality. In addition, it avoids hardware-inefficient operations like softmax, making it more suitable for hardware deployment (Figure 4).

However, ReLU linear attention alone has limited capacity due to the lack of local information extraction and multi-



Figure 1: Latency/Throughput vs. Performance. All performance results are obtained with the single model and singlescale inference. The GPU latency/throughput results are obtained on one edge GPU (Jetson AGX Orin) and one cloud GPU (A100) using TensorRT and fp16. EfficientViT consistently achieves a remarkable boost in speed on diverse hardware platforms while providing the same/higher performances on Cityscapes, ADE20K, and ImageNet than prior segmentation/classification models.

scale learning ability. Therefore, we propose to enhance ReLU linear attention with convolution and introduce the multi-scale linear attention module to address the capacity limitation of ReLU linear attention. Specifically, we aggregate nearby tokens with small-kernel convolutions to generate multi-scale tokens. We perform ReLU linear attention on multi-scale tokens (Figure 2) to combine the global receptive field with multi-scale learning. We also insert depthwise convolutions into FFN layers to further improve the local feature extraction capacity.

We extensively evaluate EfficientViT on two popular high-resolution dense prediction tasks: semantic segmentation and super-resolution. EfficientViT provides significant performance boosts over prior SOTA high-resolution dense prediction models. More importantly, EfficientViT does not involve hardware-inefficient operations, so our #FLOPs reduction can easily translate to latency reduction on hardware devices (Figure 1).

In addition to these conventional high-resolution dense prediction tasks, we apply EfficientViT to Segment Anything [13], an emerging promptable segmentation task that allows zero-shot transfer to many vision tasks. EfficientViT achieves  $84 \times$  acceleration on A100 GPU than ViT-Huge [13] while maintaining comparable zero-shot image segmentation quality. We summarize our contributions as follows:

- We introduce a new multi-scale linear attention module for efficient high-resolution dense prediction. It achieves the global receptive field and multi-scale learning while maintaining good efficiency on hardware. To the best of our knowledge, our work is the first to demonstrate the effectiveness of linear attention for high-resolution dense prediction.
- We design EfficientViT, a new family of high-resolution vision models, based on the proposed multi-scale linear

attention module.

• Our model demonstrates remarkable speedup on semantic segmentation, super-resolution, Segment Anything, and ImageNet classification on diverse hardware platforms (mobile CPU, edge GPU, and cloud GPU) over prior SOTA models.

# 2. Method

This section first introduces the multi-scale linear attention module. Unlike prior works, our multi-scale linear attention simultaneously achieves the global receptive field and multi-scale learning with only hardware-efficient operations. Then, based on the multi-scale linear attention, we present a new family of vision transformer models named EfficientViT for high-resolution dense prediction.

#### 2.1. Multi-Scale Linear Attention

Our multi-scale linear attention balances two crucial aspects of efficient high-resolution dense prediction, i.e., performance and efficiency. Specifically, the global receptive field and multi-scale learning are essential from the performance perspective. Previous SOTA high-resolution dense prediction models provide strong performances by enabling these features but fail to provide good efficiency. Our module tackles this issue by trading slight capacity loss for significant efficiency improvements.

An illustration of the proposed multi-scale linear attention module is provided in Figure 2 (right). In particular, we propose to use ReLU linear attention [12] to enable the global receptive field instead of the heavy softmax attention [8]. While ReLU linear attention [12] and other linear attention modules [14, 15, 16, 17] have been explored in other domains, it has never been successfully applied to highresolution dense prediction. To the best of our knowledge,



Figure 2: EfficientViT's Building Block (left) and Multi-Scale Linear Attention (right). *Left*: EfficientViT's building block consists of a multi-scale linear attention module and an FFN with depthwise convolution (FFN+DWConv). Multi-scale linear attention is responsible for capturing context information, while FFN+DWConv captures local information. *Right*: After getting Q/K/V tokens via the linear projection layer, we generate multi-scale tokens by aggregating nearby tokens via lightweight small-kernel convolutions. ReLU linear attention is applied to multi-scale tokens, and the outputs are concatenated and fed to the final linear projection layer for feature fusing.



Figure 3: **Softmax Attention vs. ReLU Linear Attention.** Unlike softmax attention, ReLU linear attention cannot produce sharp attention distributions due to a lack of the non-linear similarity function. Thus, its local information extraction ability is weaker than the softmax attention.

EfficientViT is the first work demonstrating ReLU linear attention's effectiveness in high-resolution dense prediction. In addition, our work introduces novel designs to address its capacity limitation.

**Enable Global Receptive Field with ReLU Linear Attention.** Given input  $x \in \mathbb{R}^{N \times f}$ , the generalized form of softmax attention can be written as:

$$O_{i} = \sum_{j=1}^{N} \frac{Sim(Q_{i}, K_{j})}{\sum_{j=1}^{N} Sim(Q_{i}, K_{j})} V_{j},$$
(1)

where  $Q = xW_Q$ ,  $K = xW_K$ ,  $V = xW_V$  and  $W_Q/W_K/W_V \in \mathbb{R}^{f \times d}$  is the learnable linear projection matrix.  $O_i$  represents the i-th row of matrix O.  $Sim(\cdot, \cdot)$  is the similarity function. When using the similarity function  $Sim(Q, K) = \exp(\frac{QK^T}{\sqrt{d}})$ , Eq. (1) becomes the original softmax attention [8].

Apart from  $\exp(\frac{QK^T}{\sqrt{d}})$ , we can use other similarity functions. In this work, we use ReLU linear attention [12] to achieve both the global receptive field and linear computational complexity. In ReLU linear attention, the similarity



Figure 4: Latency Comparison Between Softmax Attention and ReLU Linear Attention. ReLU linear attention is  $3.3-4.5 \times$  faster than softmax attention with similar computation, thanks to removing hardware-unfriendly operations (e.g., softmax). Latency is measured on the Qualcomm Snapdragon 855 CPU with TensorFlow-Lite, batch size 1, and fp32.

function is defined as

$$Sim(Q, K) = \operatorname{ReLU}(Q)\operatorname{ReLU}(K)^{T}.$$
(2)

With  $Sim(Q, K) = \text{ReLU}(Q)\text{ReLU}(K)^T$ , Eq. (1) can be rewritten as:

$$O_{i} = \sum_{j=1}^{N} \frac{\text{ReLU}(Q_{i})\text{ReLU}(K_{j})^{T}}{\sum_{j=1}^{N} \text{ReLU}(Q_{i})\text{ReLU}(K_{j})^{T}} V_{j}$$
$$= \frac{\sum_{j=1}^{N} (\text{ReLU}(Q_{i})\text{ReLU}(K_{j})^{T})V_{j}}{\text{ReLU}(Q_{i})\sum_{j=1}^{N} \text{ReLU}(K_{j})^{T}}.$$

Then, we can leverage the associative property of matrix multiplication to reduce the computational complexity and memory footprint from quadratic to linear without changing its functionality:

$$O_{i} = \frac{\sum_{j=1}^{N} [\text{ReLU}(Q_{i})\text{ReLU}(K_{j})^{T}]V_{j}}{\text{ReLU}(Q_{i})\sum_{j=1}^{N} \text{ReLU}(K_{j})^{T}}$$
$$= \frac{\sum_{j=1}^{N} \text{ReLU}(Q_{i})[(\text{ReLU}(K_{j})^{T}V_{j})]}{\text{ReLU}(Q_{i})\sum_{j=1}^{N} \text{ReLU}(K_{j})^{T}}$$
$$= \frac{\text{ReLU}(Q_{i})(\sum_{j=1}^{N} \text{ReLU}(K_{j})^{T}V_{j})}{\text{ReLU}(Q_{i})(\sum_{j=1}^{N} \text{ReLU}(K_{j})^{T})}.$$
(3)

As shown in Eq. (3), we only need to compute  $(\sum_{j=1}^{N} \operatorname{ReLU}(K_j)^T V_j) \in \mathbb{R}^{d \times d}$  and  $(\sum_{j=1}^{N} \operatorname{ReLU}(K_j)^T) \in \mathbb{R}^{d \times 1}$  once, then can reuse them for each query, thereby only requires  $\mathcal{O}(N)$  computational cost and  $\mathcal{O}(N)$  memory.

Another key merit of ReLU linear attention is that it does not involve hardware-unfriendly operations like softmax, making it more efficient on hardware. For example, Figure 4 shows the latency comparison between softmax attention and ReLU linear attention. With similar computation, ReLU linear attention is significantly faster than softmax attention on the mobile CPU.

Address ReLU Linear Attention's Limitations. Although ReLU linear attention is superior to softmax attention in terms of computational complexity and hardware latency, ReLU linear attention has limitations. Figure 3 visualizes the attention maps of softmax attention and ReLU linear attention. Because of the lack of the non-linear similarity function, ReLU linear attention cannot generate concentrated attention maps, making it weak at capturing local information.

To mitigate its limitation, we propose to enhance ReLU linear attention with convolution. Specifically, we insert a depthwise convolution in each FFN layer. An overview of the resulting building block is illustrated in Figure 2 (left), where the ReLU linear attention captures context information and the FFN+DWConv captures local information.

Furthermore, we propose to aggregate the information from nearby Q/K/V tokens to get multi-scale tokens to enhance the multi-scale learning ability of ReLU linear attention. This information aggregation process is independent for each Q, K, and V in each head. We only use small-kernel depthwise-separable convolutions [18] for information aggregation to avoid hurting hardware efficiency. In the practical implementation, independently executing these aggregation operations is inefficient on GPU. Therefore, we take advantage of the group convolution to reduce the number of total operations. Specifically, all DWConvs are fused into a single DWConv while all 1x1 Convs are combined into a single 1x1 group convolution (Figure 2 right) where the number of groups is  $3 \times$ #heads and the number of channels in each group is d. After getting multi-scale tokens, we perform ReLU linear attention upon them to extract multiscale global features. Finally, we concatenate the features along the head dimension and feed them to the final linear projection layer to fuse the features.

#### 2.2. EfficientViT Architecture

We build a new family of vision transformer models based on the proposed multi-scale linear attention module. The core building block (denoted as 'EfficientViT Module') is illustrated in Figure 2 (left). The macro architecture of EfficientViT is demonstrated in Figure 5. We use the standard backbone-head/encoder-decoder architecture design.

• **Backbone.** The backbone of EfficientViT also follows the standard design, which consists of the input stem and four stages with gradually decreased feature map size and gradually increased channel number. We insert the Effi-



Figure 5: **Macro Architecture of EfficientViT.** We adopt the standard backbone-head/encoder-decoder design. We insert our EfficientViT modules in Stages 3 and 4 in the backbone. Following the common practice, we feed the features from the last three stages (P2, P3, and P4) to the head. We use addition to fuse these features for simplicity and efficiency. We adopt a simple head design that consists of several MBConv blocks and output layers.

cientViT module in Stages 3 and 4. For downsampling, we use an MBConv with stride 2.

• Head. P2, P3, and P4 denote the outputs of Stages 2, 3, and 4, forming a pyramid of feature maps. For simplicity and efficiency, we use 1x1 convolution and standard upsampling operation (e.g., bilinear/bicubic upsampling) to match their spatial and channel size and fuse them via addition. Since our backbone already has a strong context information extraction capacity, we adopt a simple head design that comprises several MBConv blocks and the output layers (i.e., prediction and upsample). In the experiments, we empirically find this simple head design is sufficient for achieving SOTA performances.

In addition to dense prediction, our model can be applied to other vision tasks, such as image classification, by combining the backbone with task-specific heads.

Following the same macro architecture, we design a series of models with different sizes to satisfy various efficiency constraints. We name these models EfficientViT-B0, EfficientViT-B1, EfficientViT-B2, and EfficientViT-B3, respectively. In addition, we designed the EfficientViT-L series for the cloud platforms. Detailed configurations of these models are provided in our official GitHub repository<sup>1</sup>.

## **3. Experiments**

# 3.1. Setups

**Datasets.** We evaluate the effectiveness of EfficientViT on three representative high-resolution dense prediction tasks, including semantic segmentation, super-resolution, and Segment Anything.

For semantic segmentation, we use two popular benchmark datasets: Cityscapes [24] and ADE20K [25]. In addition, we evaluate EfficientViT under two settings for

Table 1: **Ablation Study.** The mIoU and MACs are measured on Cityscapes with 1024x2048 input resolution. We rescale the width of the models so that they have the same MACs. Multi-scale learning and the global receptive field are essential for obtaining good semantic segmentation performance.

Compo	onents	mIoU ↑	Params 1	MACs ↓
Multi-scale	Global att.			•
		68.1	0.7M	4.4G
$\checkmark$		72.3	0.7M	4.4G
	$\checkmark$	72.2	0.7M	4.4G
$\checkmark$	$\checkmark$	74.5	0.7M	4.4G

super-resolution: lightweight super-resolution (SR) and high-resolution SR. We train models on DIV2K [26] for lightweight SR and test on BSD100 [27]. For high-resolution SR, we train models on the first 3000 training images of FFHQ [28] and test on the first 500 validation images of FFHQ<sup>2</sup>.

Apart from dense prediction, we also study the effectiveness of EfficientViT for image classification using the ImageNet dataset [29].

**Latency Measurement.** We measure the mobile latency on Qualcomm Snapdragon 8Gen1 CPU with Tensorflow-Lite<sup>3</sup>, batch size 1 and fp32. We use TensorRT<sup>4</sup> and fp16 to measure the latency on edge GPU and cloud GPU. The data transfer time is included in the reported latency/throughput results.

**Implementation Details.** We implement our models using Pytorch [30] and train them on GPUs. We use the

https://github.com/mit-han-lab/efficientvit

<sup>&</sup>lt;sup>2</sup>https://rb.gy/7jela

<sup>&</sup>lt;sup>3</sup>https://www.tensorflow.org/lite

<sup>&</sup>lt;sup>4</sup>https://docs.nvidia.com/deeplearning/tensorrt/

Models	Top1 Acc↑	Top5 Acc↑	Params ↓	MACs↓	Later Nano(bs1)	ncy↓ Orin(bs1)	Throughput ↑ A100 (image/s)
CoAtNet-0 [19]	81.6	-	25M	4.2G	95.8ms	4.5ms	3011
ConvNeXt-T [20]	82.1	-	29M	4.5G	87.9ms	3.8ms	3303
EfficientViT-B2 (r256)	82.7	96.1	24M	2.1G	58.5ms	2.8ms	5325
Swin-B [21]	83.5	-	88M	15G	240ms	6.0ms	2236
CoAtNet-1 [19]	83.3	-	42M	8.4G	171ms	8.3ms	1512
ConvNeXt-S [20]	83.1	-	50M	8.7G	146ms	6.5ms	2081
EfficientViT-B3 (r224)	83.5	96.4	49M	4.0G	101ms	4.4ms	3797
CoAtNet-2 [19]	84.1	-	75M	16G	254ms	10.3ms	1174
ConvNeXt-B [20]	83.8	-	89M	15G	211ms	7.8ms	1579
EfficientViT-B3 (r288)	84.2	96.7	49M	6.5G	141ms	5.6ms	2372
CoAtNet-3 [19]	84.5	-	168M	35G	-	15.4ms	642
ConvNeXt-L [20]	84.3	-	198M	34G	-	11.5ms	1032
EfficientNetV2-S [22]	83.9	-	22M	8.8G	-	4.3ms	2869
EfficientViT-L1 (r224)	84.5	96.9	53M	5.3G	-	2.6ms	6207
EfficientNetV2-M [22]	85.2	-	54M	25G	-	9.2ms	1160
FasterViT-4 [23]	85.4	97.3	425M	37G	-	13.0ms	1382
EfficientViT-L2 (r288)	85.6	97.4	64M	11G	-	4.3ms	3102
FasterViT-6 [23]	85.8	97.4	1360M	142G	-	-	594
EfficientNetV2-L [22]	85.7	-	120M	53G	-	-	696
EfficientViT-L2 (r384)	86.0	97.5	64M	20G	-	-	1784

Table 2: Backbone Performance on ImageNet Classification. 'r224' means the input resolution is 224x224. 'bs1' represents that the latency is measured with batch size 1.

AdamW optimizer with cosine learning rate decay for training our models. For multi-scale linear attention, we use a two-branch design for the best trade-off between performance and efficiency, where 5x5 nearby tokens are aggregated to generate multi-scale tokens.

For semantic segmentation experiments, we use the mean Intersection over Union (mIoU) as our evaluation metric. The backbone is initialized with weights pretrained on ImageNet and the head is initialized randomly, following the common practice.

For super-resolution, we use PSNR and SSIM on the Y channel as the evaluation metrics, same as previous work [31]. The models are trained with random initialization.

#### **3.2.** Ablation Study

**Effectiveness of EfficientViT Module.** We conduct ablation study experiments on Cityscapes to study the effectiveness of two key design components of our EfficientViT module, i.e., multi-scale learning and global attention. To eliminate the impact of pre-training, we train all models from random initialization. In addition, we rescale the width of the models so that they have the same #MACs. The results are summarized in Table 1. We can see that removing either global attention or multi-scale learning will significantly hurt the performances. It shows that all of them are essential for achieving a better trade-off between performance and efficiency.

**Backbone Performance on ImageNet.** To understand the effectiveness of EfficientViT's backbone in image classification, we train our models on ImageNet following the standard training strategy. We summarize the results and compare our models with SOTA image classification models in Table 2.

Though EfficientViT is designed for high-resolution dense prediction, it achieves highly competitive performances on ImageNet classification. In particular, EfficientViT-L2-r384 obtains 86.0 top1 accuracy on ImageNet, providing +0.3 accuracy gain over EfficientNetV2-L and 2.6x speedup on A100 GPU.

## 3.3. Semantic Segmentation

**Cityscapes.** Table 3 reports the comparison between EfficientViT and SOTA semantic segmentation models on Cityscapes. EfficientViT achieves remarkable efficiency improvements over prior SOTA semantic segmentation models without sacrificing performances. Specifically, compared with SegFormer, EfficientViT obtains up to 13x #MACs saving and up to 8.8x latency reduction on the edge GPU (Jetson AGX Orin) with higher mIoU. Compared

Table 3: Comparison with SOTA Semantic Segmentation Models on Cityscapes. The input resolution is 1024x2048 for all models. Models with similar mIoU are grouped for efficiency comparison.

Models	mIoU ↑	Params ↓	$MACs\downarrow$	Latency↓ Nano(bs1) Orin(bs1)		Throughput ↑ A100(image/s)
DeepLabV3plus-Mbv2 [32]	75.2	15M	555G	0.28s	83.5ms	102
EfficientViT-B0	75.7	0.7M	4.4G		<b>9.9ms</b>	263
SegFormer-B1 [7]	78.5	14M	244G	5.6s	146ms	49
SegNeXt-T [9]	79.8	4.3M	51G	2.2s	93.2ms	95
EfficientViT-B1	80.5	4.8M	25G	<b>0.82s</b>	<b>24.3ms</b>	<b>175</b>
SegFormer-B3 [7]	81.7	47M	963G	-	407ms	18
SegNeXt-S [9]	81.3	14M	125G	3.4s	127ms	70
EfficientViT-B2	82.1	15M	74G	<b>1.7s</b>	<b>46.5ms</b>	<b>112</b>
SegFormer-B5 [7]	82.4	85M	1460G	-	638ms	12
SegNeXt-B [9]	82.6	28M	276G	-	228ms	41
EfficientViT-B3	83.0	40M	179G	-	81.8ms	70
EfficientViT-L1	82.7	40M	282G	-	<b>45.9ms</b>	122
SegNeXt-L [9]	83.2	49M	578G		374ms	26
EfficientViT-L2	83.2	53M	396G		<b>60.0ms</b>	<b>102</b>

Table 4: Comparison with SOTA Semantic Segmentation Models on ADE20K. The shorter side of the image is resized to 512, following the common practice.

Models	mIoU ↑	Params ↓	MACs↓	Later Nano(bs1)	ncy↓ Orin(bs1)	Throughput ↑ A100(image/s)
SegFormer-B1 [7]	42.2	14M	16G	389ms	12.3ms	542
SegNeXt-T [9]	41.1	4.3M	6.6G	281ms	12.4ms	842
EfficientViT-B1	42.8	4.8M	3.1G	110ms	4.0ms	1142
SegNeXt-S [9]	44.3	14M	16G	428ms	17.2ms	592
EfficientViT-B2	45.9	15M	9.1G	212ms	7.3ms	846
Mask2Former [33]	47.7	47M	74G	-	-	-
MaskFormer [34]	46.7	42M	55G	-	-	-
SegFormer-B2 [7]	46.5	28M	62G	920ms	24.3ms	345
SegNeXt-B [9]	48.5	28M	35G	806ms	32.9ms	347
EfficientViT-B3	49.0	39M	22G	411ms	12.5ms	555
EfficientViT-L1	49.2	40M	36G	-	7.2ms	947
SegFormer-B4 [7]	50.3	64M	96G	-	44.9ms	212
EfficientViT-L2	50.7	51M	45G	-	9.0ms	758

with SegNeXt, EfficientViT provides up to 2.0x MACs reduction and 3.8x speedup on the edge GPU (Jetson AGX Orin) while maintaining higher mIoU. On A100 GPU, EfficientViT delivers up to 3.9x higher throughput than Seg-NeXt and 10.2x higher throughput than SegFormer while achieving the same or higher mIoU. Having similar computational cost, EfficientViT also yields significant performance gains over previous SOTA models. For example, EfficientViT-B3 delivers +4.5 mIoU gain over SegFormer-B1 with lower MACs.

In addition to the quantitative results, we visualize EfficientViT and the baseline models qualitatively on

Cityscapes. The results are shown in Figure 6. We can find that EfficientViT can better recognize boundaries and small objects than the baseline models while achieving lower latency on GPU.

**ADE20K.** Table 4 summarizes the comparison between EfficientViT and SOTA semantic segmentation models on ADE20K. Like Cityscapes, we can see that EfficientViT also achieves significant efficiency improvements on ADE20K. For example, with +0.6 mIoU gain, EfficientViT-B1 provides 5.2x MACs reduction and up to

Model	F	FHQ (512x	$512 \rightarrow 1024 \mathrm{x10}$	)24)	BSD100 (160x240 $\rightarrow$ 320x480)			
	$PSNR\uparrow$	$\mathbf{SSIM} \uparrow$	A100(bs1) $\downarrow$	Speedup $\uparrow$	PSNR $\uparrow$	$\text{SSIM} \uparrow$	A100(bs1)↓	Speedup $\uparrow$
Restormer [35]	43.43	0.9806	92.0ms	1x	32.31	0.9021	15.1ms	1x
SwinIR [31]	43.49	0.9807	61.2ms	1.5x	32.31	0.9012	9.7ms	1.6x
VapSR [36]	-	-	-	-	32.27	0.9011	4.8ms	3.1x
BSRN [37]	-	-	-	-	32.24	0.9006	4.5ms	3.4x
EfficientViT w0.75	43.54	0.9809	14.3ms	6.4x	32.31	0.9016	2.8ms	5.4x
EfficientViT	43.58	0.9810	17.8ms	5.2x	32.33	0.9019	3.2ms	4.7x

Table 5: Comparison with SOTA super-resolution models.

Table 6: Zero-Shot Image Segmentation Results on COCO val2017. We prompt the SAM models with the ground-truth object bounding box annotations to predict the segmentation masks.

Model		mIoU ↑				MACs	Throughput $\uparrow$
	all	large	medium	small		1.11 100 \$	A100(image/s)
MobileSAM [38] NanoSAM [39]	72.8 70.6	80.4 79.6	75.9 73.8	65.8 62.4		-	297 744
EfficientViT-L0 EfficientViT-L1	74.5 <b>75.2</b>	81.4 <b>81.8</b>	77.2 <b>78.1</b>	68.2 <b>68.9</b>	31M 44M	35G 49G	<b>1009</b> 815



Figure 6: Qualitative results on Cityscapes.

3.5x GPU latency reduction than SegFormer-B1. With +1.6 mIoU gain, EfficientViT-B2 requires 1.8x fewer computational costs and runs 2.4x faster on Jetson AGX Orin GPU than SegNeXt-S.

#### 3.4. Super-Resolution

Table 5 presents the comparison of EfficientViT with SOTA ViT-based SR methods (SwinIR [31] and Restormer [35]) and SOTA CNN-based SR methods (VapSR [36] and BSRN [37]). EfficientViT provides a better latency-performance trade-off than all compared methods.

On lightweight SR, EfficientViT provides up to 0.09dB gain in PSNR on BSD100 while maintaining the same or

lower GPU latency compared with SOTA CNN-based SR methods. Compared with SOTA ViT-based SR methods, EfficientViT provides up to  $5.4 \times$  speedup on GPU and maintains the same PSNR on BSD100.

On high-resolution SR, the advantage of EfficientViT over previous ViT-based SR methods becomes more significant. Compared with Restormer, EfficientViT achieves up to  $6.4 \times$  speedup on GPU and provides 0.11dB gain in PSNR on FFHQ.

#### 3.5. Segment Anything

For simplicity, we use the image embedding extracted by the pretrained SAM-ViT-H [13] as the training target to train our image encoders. We use the same prompt encoder and mask decoder as SAM-ViT-H [13].

In table 6, we compare EfficientViT with prior efficient Segment Anything models, including MobileSAM and NanoSAM. EfficientViT obtains higher zero-shot image segmentation results on COCO val2017 and runs faster than MobileSAM and NanoSAM.

In addition, we visualize the outputs of EfficientViT Segment Anything models in Figures 7, 8, and 9. We can see that EfficientViT can achieve similar zero-shot image segmentation quality as ViT-H [13] with significantly higher throughput on A100 GPU.

## 4. Related Work

High-Resolution Dense Prediction. Dense prediction targets producing predictions for each pixel given the in-



Figure 7: Visualization of Segment Anything Model with Bounding Box as Prompt.



Figure 8: Visualization of Segment Anything Model with Point as Prompt.

put image. It can be viewed as an extension of image classification from per-image prediction to per-pixel predictions. Extensive studies have been done to improve the performance of CNN-based high-resolution dense prediction models [1, 2, 3, 4, 5, 6].

In addition, there are also some works targeting improving the efficiency of high-resolution dense prediction models [40, 41, 42, 43]. While these models provide good efficiency, their performances are far behind SOTA highresolution dense prediction models.

Compared to these works, our models provide a better trade-off between performance and efficiency by enabling a global receptive field and multi-scale learning with lightweight operations.

**Efficient Vision Transformer.** While ViT provides impressive performances in the high-computation region, it is usually inferior to previous efficient CNNs [44, 45, 46, 47] when targeting the low-computation region. To close the gap, MobileViT [48] proposes to combine the strength of CNN and ViT by replacing local processing in convolu-

tions with global processing using transformers. Mobile-Former [49] proposes to parallelize MobileNet and Transformer with a two-way bridge in between for feature fusing. NASViT [50] proposes to leverage neural architecture search to search for efficient ViT architectures.

However, these models mainly focus on image classification and still rely on softmax attention with quadratic computational complexity, thus unsuitable for high-resolution dense prediction.

Efficient Deep Learning. Our work is also related to efficient deep learning, which aims at improving the efficiency of deep neural networks so that we can deploy them on hardware platforms with limited resources, such as mobile phones and IoT devices. Typical technologies in efficient deep learning include network pruning [51, 52, 53], quantization [54], efficient model architecture design [18, 55], and training techniques [56, 57, 58]. In addition to manual designs, many recent works use AutoML techniques [59, 60, 61] to automatically design [46], prune [62] and quantize [63] neural networks.



Figure 9: Automatically Generate Masks with Segment Anything Models.

## 5. Conclusion

In this work, we studied efficient architecture design for high-resolution dense prediction. We introduced a lightweight multi-scale attention module that simultaneously achieves a global receptive field, and multi-scale learning with lightweight and hardware-efficient operations, thus providing significant speedup on diverse hardware devices without performance loss than SOTA highresolution dense prediction models. For future work, we will explore applying EfficientViT to other vision tasks and further scaling up our EfficientViT models.

## Acknowledgments

We thank MIT-IBM Watson AI Lab, MIT AI Hardware Program, Amazon and MIT Science Hub, Qualcomm Innovation Fellowship, National Science Foundation for supporting this research.

## References

- Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017. 1, 9
- [2] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. Unet: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III* 18, pages 234–241. Springer, 2015. 1, 9
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 1, 9

- [4] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2881–2890, 2017. 1, 9
- [5] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Objectcontextual representations for semantic segmentation. In *European conference on computer vision*, pages 173–190. Springer, 2020. 1, 9
- [6] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions* on pattern analysis and machine intelligence, 43(10):3349– 3364, 2020. 1, 9
- [7] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. Advances in Neural Information Processing Systems, 34, 2021. 1, 7
- [8] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 1, 2, 3
- [9] Meng-Hao Guo, Cheng-Ze Lu, Qibin Hou, Zheng-Ning Liu, Ming-Ming Cheng, and Shi min Hu. Segnext: Rethinking convolutional attention design for semantic segmentation. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, Advances in Neural Information Processing Systems, 2022. 1, 7
- [10] Xiaohan Ding, Xiangyu Zhang, Yizhuang Zhou, Jungong Han, Guiguang Ding, and Jian Sun. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In *Proceedings of the IEEE conference on computer vision and pattern* recognition, 2022. 1
- [11] Yihan Wang, Muyang Li, Han Cai, Wei-Ming Chen, and Song Han. Lite pose: Efficient architecture design for 2d human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2022. 1

- [12] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pages 5156–5165. PMLR, 2020. 1, 2, 3
- [13] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. arXiv preprint arXiv:2304.02643, 2023. 2, 8
- [14] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, and Judy Hoffman. Hydra attention: Efficient attention with many heads. In *Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 35–49. Springer, 2023. 2
- [15] Krzysztof Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. arXiv preprint arXiv:2009.14794, 2020. 2
- [16] Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Efficient attention: Attention with linear complexities. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 3531– 3539, 2021. 2
- [17] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. arXiv preprint arXiv:2006.04768, 2020. 2
- [18] 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. 4, 9
- [19] Zihang Dai, Hanxiao Liu, Quoc V Le, and Mingxing Tan. Coatnet: Marrying convolution and attention for all data sizes. Advances in Neural Information Processing Systems, 34:3965–3977, 2021. 6
- [20] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In Proceedings of the IEEE conference on computer vision and pattern recognition, 2022. 6
- [21] 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. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 10012–10022, 2021. 6
- [22] Mingxing Tan and Quoc Le. Efficientnetv2: Smaller models and faster training. In *International Conference on Machine Learning*, pages 10096–10106. PMLR, 2021. 6
- [23] Ali Hatamizadeh, Greg Heinrich, Hongxu Yin, Andrew Tao, Jose M Alvarez, Jan Kautz, and Pavlo Molchanov. Fastervit: Fast vision transformers with hierarchical attention. arXiv preprint arXiv:2306.06189, 2023. 6
- [24] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe

Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 5

- [25] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017. 5
- [26] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 126–135, 2017. 5
- [27] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001. 5
- [28] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 4401–4410, 2019. 5
- [29] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In CVPR, 2009. 5
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems, 32, 2019.
- [31] Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1833–1844, 2021. 6, 8
- [32] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision* (ECCV), pages 801–818, 2018. 7
- [33] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1290–1299, 2022. 7
- [34] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Perpixel classification is not all you need for semantic segmentation. Advances in Neural Information Processing Systems, 34, 2021. 7
- [35] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image

restoration. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 5728– 5739, 2022. 8

- [36] Lin Zhou, Haoming Cai, Jinjin Gu, Zheyuan Li, Yingqi Liu, Xiangyu Chen, Yu Qiao, and Chao Dong. Efficient image super-resolution using vast-receptive-field attention. arXiv preprint arXiv:2210.05960, 2022. 8
- [37] Zheyuan Li, Yingqi Liu, Xiangyu Chen, Haoming Cai, Jinjin Gu, Yu Qiao, and Chao Dong. Blueprint separable residual network for efficient image super-resolution. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 833–843, 2022. 8
- [38] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. arXiv preprint arXiv:2306.14289, 2023.
- [39] NVIDIA-AI-IOT. Nanosam. https://github.com/ NVIDIA-AI-IOT/nanosam, 2023. 8
- [40] Hengshuang Zhao, Xiaojuan Qi, Xiaoyong Shen, Jianping Shi, and Jiaya Jia. Icnet for real-time semantic segmentation on high-resolution images. In *Proceedings of the European conference on computer vision (ECCV)*, pages 405– 420, 2018. 9
- [41] Rudra PK Poudel, Stephan Liwicki, and Roberto Cipolla. Fast-scnn: Fast semantic segmentation network. arXiv preprint arXiv:1902.04502, 2019. 9
- [42] Hanchao Li, Pengfei Xiong, Haoqiang Fan, and Jian Sun. Dfanet: Deep feature aggregation for real-time semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9522–9531, 2019. 9
- [43] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 325–341, 2018. 9
- [44] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 9
- [45] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *ICCV*, 2019. 9
- [46] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once for all: Train one network and specialize it for efficient deployment. In *ICLR*, 2020. 9
- [47] Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, and Chang Xu. Ghostnet: More features from cheap operations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1580–1589, 2020. 9
- [48] Sachin Mehta and Mohammad Rastegari. Mobilevit: Lightweight, general-purpose, and mobile-friendly vision transformer. In *International Conference on Learning Representations*, 2022. 9

- [49] Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Xiaoyi Dong, Lu Yuan, and Zicheng Liu. Mobileformer: Bridging mobilenet and transformer. In Proceedings of the IEEE conference on computer vision and pattern recognition, 2022. 9
- [50] Chengyue Gong, Dilin Wang, Meng Li, Xinlei Chen, Zhicheng Yan, Yuandong Tian, qiang liu, and Vikas Chandra. NASVit: Neural architecture search for efficient vision transformers with gradient conflict aware supernet training. In *International Conference on Learning Representations*, 2022. 9
- [51] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *NeurIPS*, 2015. 9
- [52] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *ICCV*, 2017.9
- [53] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *ICCV*, 2017. 9
- [54] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *ICLR*, 2016. 9
- [55] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In ECCV, 2018. 9
- [56] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531, 2015. 9
- [57] Han Cai, Chuang Gan, Ji Lin, and Song Han. Network augmentation for tiny deep learning. arXiv preprint arXiv:2110.08890, 2021. 9
- [58] Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. Tinytl: Reduce memory, not parameters for efficient on-device learning. Advances in Neural Information Processing Systems, 33:11285–11297, 2020. 9
- [59] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *ICLR*, 2017. 9
- [60] Han Cai, Tianyao Chen, Weinan Zhang, Yong Yu, and Jun Wang. Efficient architecture search by network transformation. In AAAI, 2018. 9
- [61] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *ICLR*, 2019. 9
- [62] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *ECCV*, 2018. 9
- [63] Tianzhe Wang, Kuan Wang, Han Cai, Ji Lin, Zhijian Liu, Hanrui Wang, Yujun Lin, and Song Han. Apq: Joint search for network architecture, pruning and quantization policy. In *CVPR*, 2020. 9