

# MobileNetV4: Universal Models for the Mobile Ecosystem

Danfeng Qin<sup>†‡</sup>, Chas Leichner<sup>†‡</sup>, Manolis Delakis<sup>†</sup>, Marco Fornoni, Shixin Luo, Fan Yang, Weijun Wang, Colby Banbury, Chengxi Ye, Berkin Akin, Vaibhav Aggarwal, Tenghui Zhu, Daniele Moro, and Andrew Howard<sup>†§</sup>

Google

**Abstract.** We present the latest generation of MobileNets: MobileNetV4 (MNv4). They feature universally-efficient architecture designs for mobile devices. We introduce the Universal Inverted Bottleneck (UIB) search block, a unified and flexible structure that merges Inverted Bottleneck (IB), ConvNext, Feed Forward Network (FFN), and a novel Extra Depth-wise (ExtraDW) variant. Alongside UIB, we present Mobile MQA, an attention block for mobile accelerators, delivering a significant 39% speedup. An optimized neural architecture search (NAS) recipe is also introduced which improves MNv4 search effectiveness. The integration of UIB, Mobile MQA and the refined NAS recipe results in a new suite of MNv4 models that are mostly Pareto optimal across mobile CPUs, DSPs, GPUs, as well as accelerators like Apple Neural Engine and Google Pixel EdgeTPU. This performance uniformity is not found in any other models tested. We introduce performance modeling and analysis techniques to explain how this performance is achieved. Finally, to further boost accuracy, we introduce a novel distillation technique. Enhanced by this technique, our MNv4-Hybrid-Large model delivers 87% ImageNet-1K accuracy, with a Pixel 8 EdgeTPU runtime of 3.8ms.

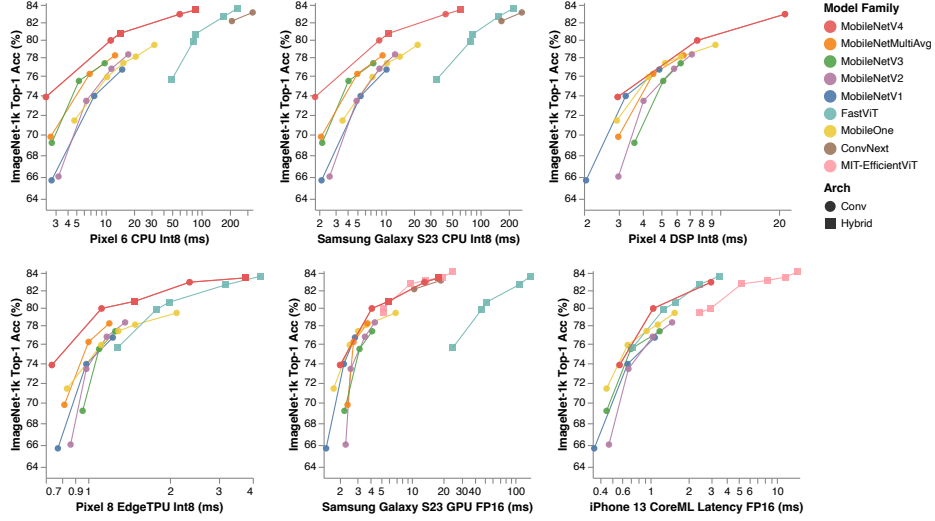
## 1 Introduction

Efficient on-device neural networks not only enable fast, real-time and interactive experiences, but also avoid streaming private data through the public internet. However, the computational constraints of mobile devices pose the significant challenge of balancing accuracy and efficiency. To this end, we introduce UIB and Mobile MQA, two innovative building blocks integrated via a refined NAS recipe to create a series of universally mostly-Pareto-optimal mobile models.<sup>1</sup> Additionally, we present a distillation technique that further improves efficiency.

Our Universal Inverted Bottleneck (UIB) block improves the Inverted Bottleneck block [38] by incorporating two optional depthwise convolutions [21]. Despite its simplicity, UIB unifies prominent micro-architectures - Inverted Bot-

<sup>0</sup> †Equal primary contribution. ‡Project Lead. §Senior Lead.

<sup>1</sup> All MobileNetV4 models are available at <https://github.com/tensorflow/models/blob/master/official/vision/modeling/backbones/mobilenet.py>



**Fig. 1: MNv4 Models are Universally Mostly Pareto Optimal:** MNv4 performs strongly compared to leading efficient models across diverse hardware. All models were trained on ImageNet-1k solely. MobileNetV1-V3 were retrained with updated recipes. Most models were optimized for one device, but MNv4 is Pareto optimal across most devices. Hybrid models and ConvNext are DSP-incompatible. Due to PyTorch-to-TF Lite export tool limitations, EfficientViTs [14] [15] are not benchmarked on CPUs and EdgeTPU. MNv4-Hybrid models were excluded from CoreML evaluation due to the lack of PyTorch implementation of Mobile MQA.

tleneck (IB), ConvNext [34], and FFN [13] - and introduces the Extra Depthwise (ExtraDW) IB block. UIB offers flexibility in spatial and channel mixing, the option to extend the receptive field, and enhanced computational efficiency.

Our optimized Mobile MQA block achieves over a 39% inference speedup on mobile accelerators with respect to Multi-Head Attention [47].

Our two-phase NAS approach, separating coarse and fine-grained searches, significantly boosts search efficiency and facilitates the creation of models that are significantly larger than previous state-of-the-art models [43]. Additionally, incorporating an offline distillation dataset reduces noise in NAS reward measurements, resulting in improved model quality.

By integrating UIB, MQA, and an improved NAS recipe, we present the MNv4 suite of models which achieve mostly Pareto optimal performance across diverse hardware platforms, including CPUs, DSPs, GPUs, and accelerators. Our models range from the extremely compact MNv4-Conv-S to the MNv4-Hybrid-L high-end variant that establishes a new reference for mobile model accuracy. MNv4-Conv-S achieves 73.8% top-1 ImageNet-1K accuracy with 3.8M parameters, 0.2G MACs and 2.4 ms of Pixel 6 CPU latency. MNv4-Hybrid-L gets 83.4% top-1 within 3.8 ms on Pixel 8 EdgeTPU. Our novel distillation recipe mixes datasets with different augmentations and adds balanced in-class data, enhancing generalization and increasing accuracy. With these techniques, MNv4-Hybrid-L achieves a 87% top-1 accuracy on ImageNet-1K: 0.5% less than its teacher, despite having 39x less MACs.

## 2 Related Work

Optimizing models for both accuracy and efficiency is a well studied problem.

**Mobile Convolutional Networks:** Key work includes MobileNetV1 [22] with depthwise-separable convolutions for better efficiency, MobileNetV2 [38] introducing linear bottlenecks and inverted residuals, GhostNet [17] increasing the relative frequency of depthwise convolutions, MnasNet [42] integrating lightweight attention in bottlenecks, and MobileOne [46] adding and re-parameterizing linear branches in inverted bottlenecks at inference time.

**Efficient Hybrid Networks:** This research combines convolutions and attention. MobileViT [35] merges CNN strengths with ViT [13] through global attention blocks. GhostNetV2 [44] uses FC layers to capture long-range dependencies. MobileFormer [7] parallelizes a MobileNet and a Transformer with a two-way bridge in between for feature fusing. FastViT [45] adds attention to the last stage with large convolutional kernels instead of early stage self-attention.


**Efficient Attention:** Research has focused on enhancing MHSA [47] efficiency. EfficientViT [14] and MobileViTv2 [36] introduce self-attention approximations for linear complexity with minor accuracy impacts. EfficientFormerV2 [29] downsamples Q, K, and V for efficiency, while CMT [16] and NextViT [28] downsample only K and V.

**Hardware-aware Neural Architecture Search (NAS):** Another common technique is to automate the model design process using hardware-aware Neural Architecture Search (NAS). NetAdapt [52] uses empirical latency tables to optimize the accuracy of a model under a target latency constraint. MnasNet [42] also uses latency tables, but applies reinforcement learning to do hardware-aware NAS. FBNet [50] accelerates multi-task hardware-aware search via differentiable NAS. MobileNetV3 [20] is tuned to mobile phone CPUs through a combination of hardware-aware NAS, the NetAdapt algorithm, and architecture advances. MobileNet MultiHardware [9] optimizes a single model for multiple hardware targets. Once-for-all [6] separates training and search for efficiency.

## 3 Hardware-Independent Pareto Efficiency

**The Roofline Model:** For a model to be universally efficient, it must perform well on hardware targets with vastly different bottlenecks that limit the model’s performance. These bottlenecks are largely determined by the hardware’s peak computational throughput and its peak memory bandwidth.

To this end, we use the Roofline Model [49] which estimates the performance of a given workload and predicts whether it is memory-bottlenecked or compute-bottlenecked. In short, it abstracts away specific hardware details and only considers a workload’s operational intensity ( $\text{LayerMACs}_i / (\text{WeightBytes}_i + \text{ActivationBytes}_i)$ ) *vs.* the theoretical limits of the hardware’s processor and memory system. Memory and compute operations happen roughly in parallel, so the slower of the two approximately determines the latency bottleneck. To apply the Roofline Model to neural networks with layers indexed by  $i$ , we can



**Fig. 2: Ridge Points and Latency/Accuracy Trade-Offs:** In the roofline performance model, the ridge point summarizes the relationship between memory bandwidth and MACs. If memory bandwidth is constant, high-compute hardware (accelerators) have a higher ridge point than low-compute hardware (CPUs). MobileNetV4 is mostly Pareto-optimal from a ridge point of 0 to 500 MACs/byte. These analytically-derived (Eq. (1)) charts reflect the real hardware measurements in Fig. 1. Appendix F contains further analysis of this relationship.

calculate the model inference latency, ModelTime, as follows:

$$\text{ModelTime} = \sum_i \max(\text{MACTime}_i, \text{MemTime}_i) \quad (1)$$

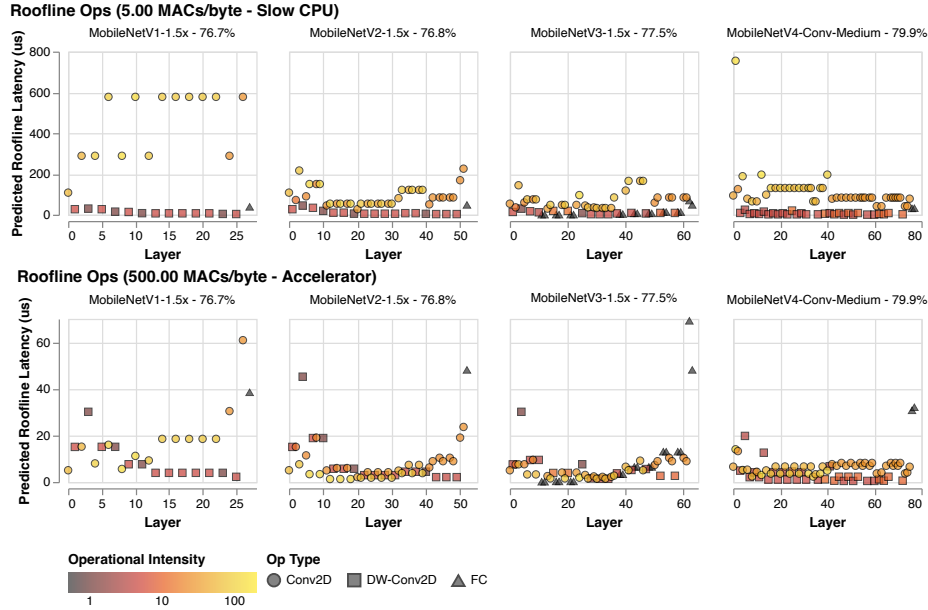
$$\text{MACTime}_i = \frac{\text{LayerMACs}_i}{\text{PeakMACs}}, \quad \text{MemTime}_i = \frac{\text{WeightBytes}_i + \text{ActivationBytes}_i}{\text{PeakMemBW}}$$

In the roofline model, hardware behavior is summarized by the *Ridge Point* (RP)—the ratio of a hardware’s PeakMACs to PeakMemBW *i.e.* the minimum operational intensity required to achieve maximum performance.<sup>2</sup> In order to optimize for hardware with a wide range of bottlenecks, as seen in Fig. 2 and Fig. 3, we analyze our algorithms’ latency while sweeping the RP from its lowest expected value (0 MAC/byte) to its highest expected value (500 MACs/byte)—see Appendix F for more details. Roofline Models only depend on the ratio of data transfer to compute, so all hardware with the same RP will rank workloads the same by latency.<sup>3</sup> This means that swept-RP roofline analysis (see next paragraph) applies to future hardware and software if the RP of the new targets is contained in the swept range.

**Ridge Point Sweep Analysis:** As seen in Fig. 2 and Fig. 3, the roofline model sheds light on how MobileNetV4 models achieve hardware-independent mostly-Pareto-optimal performance against other convolutional MobileNets. On low-RP hardware (e.g. CPUs), models are more likely to be compute-bound than memory-bound. So, to improve latency, you minimize the total number of MACs even at the cost of increased memory complexity (MobileNetV3Large-1.5x). Data movement is the bottleneck on high-RP hardware, so MACs do not meaningfully slow down the model but can increase model capacity (MobileNetV1-1.5x). So models optimized for low-RPs run slowly at high-RPs because memory-intensive

<sup>2</sup> The common practice of using a model’s total MACs to proxy latency is the same as targeting a roofline model with a Ridge Point (RP) = 0. This is equivalent to infinite bytes per MAC so,  $\forall i, \text{MemTime}_i = 0$  and  $\text{ModelTime} = \sum_i \text{MACTime}_i$ .

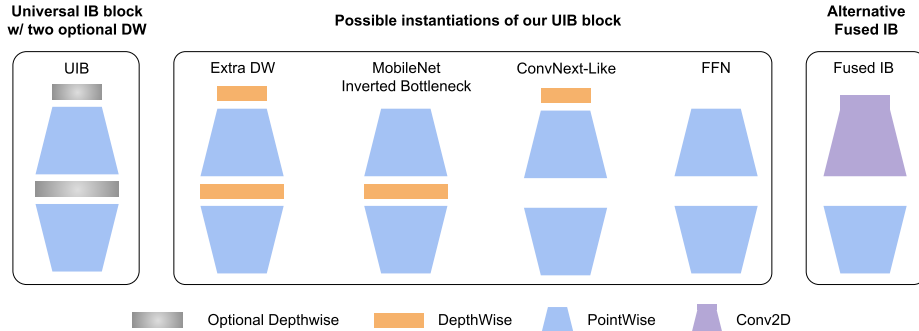
<sup>3</sup> The Roofline Model assumes that software implementation has no impact on workload performance. This means techniques with complex memory access (e.g. pruning) perform much better on a Roofline Model than on a real device.



**Fig. 3: Op Cost vs. Ridge Point:** Each sub-chart displays the roofline latency (Eq. (1)) of a network’s ops. Networks start on the left. Large Conv2Ds are expensive on low ridge point (RP) hardware (*top row*), but add cheap model capacity on high-RP hardware (*bottom row*). FC layers and DW-Conv2Ds are cheap at low RPs and expensive at high RPs. MobileNetV4 balances MAC-intensive Conv2D layers and memory-intensive FC layers where they contribute most to the network—the beginning and end, respectively. Full sweeps and data for all MobileNetV4-Conv models are in Appendix F.

and low-MAC fully-connected (FC) layers are bottlenecked on memory bandwidth and can’t take advantage of the high available PeakMACs.

**MobileNetV4 Design:** MobileNetV4 balances investing MACs and memory bandwidth where they will provide the maximum return for the cost, paying particular attention to the start and end of the network. At the beginning of the network, MobileNetV4 uses large and expensive initial layers to substantially improve the models’ capacity and downstream accuracy. These initial layers are dominated by a high number of MACs, so they are only expensive on low-RP hardware. At the end of the network, all MobileNetV4 variants use the same size final FC layers to maximize accuracy, even though this causes smaller MNV4 variants to suffer higher FC latency on high-RP hardware. Since large initial Conv layers are expensive on low-RP hardware but not high-RP hardware while the final FC layers are expensive on high-RP hardware but not low-RP hardware, MobileNetV4 models will never see both slowdowns at the same time. In other words, MNv4 models are able to use expensive layers that disproportionately improve accuracy but do not suffer the simultaneous combined costs of the layers, resulting in mostly Pareto-optimal performance at all ridge points.



**Fig. 4:** Universal Inverted Bottleneck (UIB) blocks.

## 4 Universal Inverted Bottlenecks

With an established foundation of roofline modeling and operational intensity, we proceed to discuss our architectural blocks. First is the Universal Inverted Bottleneck (UIB) Block, a building block for efficient network design that can adapt to a variety of optimization targets while remaining simple enough to use with Neural Architecture Search (NAS). Figure 4 shows the UIB block structure.

UIB extends the MobileNet Inverted Bottleneck (IB) block (introduced in MobileNetV2 [38]), which has become the standard building block for efficient networks [13, 20, 34, 43]. We introduce an optional DW before the expansion layer and also make the DW between the expansion and projection layer optional. The NAS procedure selects which DW ops to include, resulting in novel architectures. Despite the simplicity of this modification, our new building block unifies important existing blocks: the original IB block, ConvNext block, and the FFN block in ViT. Additionally, UIB introduces a novel variant: the Extra DepthWise IB (ExtraDW) block. The NAS SuperNet size is manageable because the pointwise expansion and projection components of each block are shared between instantiations and the depthwise ops are searchable options. In a SuperNet-based NAS algorithm, this approach shares >95% of the parameters between instantiations so NAS remains efficient. We further use FusedIBs to improve the efficiency: A  $k \times k$  FusedIB is a  $k \times k$  Conv2D into a  $1 \times 1$  Conv2D [1]. FusedIBs are used in all MNV4 model stems (Appendix D, Tabs. 11 - 15).

**UIB Instantiations:** The two optional depthwise convolutions in the UIB block have four possible instantiations (Fig. 4), resulting in different tradeoffs.

**MobileNet Inverted Bottleneck (IB)** performs spatial mixing on the expanded features’ activations for greater model capacity at increased cost.

**ConvNext-Like** allows for a cheaper spatial mixing with larger kernel size by performing the spatial mixing before the expansion.

**ExtraDW** inexpensively increases the network depth and receptive field, combining the benefits of ConvNext-Like and IB.<sup>4</sup>

**FFN** is a stack of two  $1 \times 1$  pointwise convolutions (PW) with activation and

<sup>4</sup> ExtraDW could be seen as a MobileNetV1-style factorization of two standard convolutional blocks.

**Table 1:** Comparison between searches using Inverted Bottleneck blocks, ConvNext-Like blocks, and full UIB blocks.

Block	Top-1	GMACs	MParams	P8 EdgeTPU
<b>UIB</b>	<b>83.3%</b>	<b>6.2</b>	<b>33.0</b>	<b>2.68 ms</b>
CN	83.2%	6.9	35.1	2.69 ms
IB	82.3%	6.1	32.4	2.61 ms

normalization layers in between. PW is very accelerator-friendly but works best with other blocks.

At each network stage, UIB provides flexibility to: (1) Strike an ad-hoc spatial and channel mixing tradeoff. (2) Enlarge the receptive field as needed. (3) Maximize the computational utilization. Tab. 1 shows impact on accuracy and latency across three searches.

## 5 Mobile MQA

In this section we present Mobile MQA, a novel accelerator-optimized attention block which speeds up attention by >39%.

**Importance of Operational Intensity:** Vision model research has largely focused on improving efficiency by reducing MACs. Since accelerators greatly increase computational capabilities without proportionally increasing memory bandwidth, many models are bottlenecked by memory access and solely minimizing MACs will not improve performance. Instead we must consider the Operational Intensity—the ratio of arithmetic operations to memory access.

**MQA is efficient in hybrid models:** MHSA [47] projects the queries, keys, and values into multiple spaces to capture different aspects of the information. Multi-Query Attention (MQA) [39] simplifies this by sharing keys and values across all heads. While large language models require multiple query heads, they can share a single head for keys and values without sacrificing accuracy [8] [27]. When the number of batched tokens is small compared to the feature dimensions, sharing one head across keys and values reduces memory bandwidth requirements—significantly improving Operational Intensity. In hybrid mobile vision models, the tokens are often small compared to features because attention is only used in the low-resolution later stages with high feature dimensions and because batch size one operation is common. Our experiments confirm MQA’s advantage in hybrid models. As shown in Tab. 2 MQA achieves >39% acceleration on EdgeTPUs and Samsung S23 GPU with negligible quality loss (-0.03%) compared to MHSA. MQA also reduces MACs and model parameters by >25%. To our knowledge, we are the first to use MQA for mobile vision. Furthermore, we introduce an additional Einsum optimization (see Appendix G), specifically tailored for accelerated inference on hardware accelerators.

**Incorporate asymmetric spatial down-sampling:** Drawing inspiration from MQA, which utilizes asymmetric computation across queries, keys, and values, we add Spatial Reduction Attention (SRA) [48] to our optimized MQA block to downscale key and value resolution while retaining high-resolution queries. This strategy is motivated by the observed correlation between spatially adjacent

**Table 2: MQA Impact:** MNv4-Conv-L base model. Attention blocks are added to the last stage. Percentage improvements only consider attention block latency *vs.* MHSA.

model	Top-1 Acc(%)	MACs (G)	Params (M)	EdgeTPU Pixel 7	Pixel 8	Samsung S23 GPU
base model	84.88	6.0	30.9	4.31 ms	2.35 ms	13.15 ms
+3 MHSA	85.27	6.7	36.0	9.69 ms	2.76 ms	16.46 ms
+3 MQA	85.24 (-0.03%)	6.5 (-28.6%)	34.7 (-25.5%)	5.16 ms (-84.2%)	2.60 ms (-39.0%)	15.10ms (-41.1%)

**Table 3: Impact of Downsampling on Mobile MQA:** MNv4-Hybrid-M base model on Samsung S23. Stride-2 down-sampling is applied at penultimate 16x16 stage.

down-sampling on KV	Top-1 Acc	MACs (G)	CPU (ms)	GPU (ms)
No	80.77	1.285	15.8	7.4
Yes	80.71	1.245	12.8	5.9
Efficiency Gain	-	+3%	+23%	+25%

tokens in hybrid models attributed to spatial mixing convolution filters in early layers. Unlike [48], our method replaces AvgPooling with a stride-2 3x3 DW for spatial reduction—a cost-effective way to boost model capacity.

**Mobile MQA** Here we present our Mobile MQA block:

$$\text{Mobile\_MQA}(\mathbf{X}) = \text{Concat}(\text{attention}_1, \dots, \text{attention}_n) \mathbf{W}^O$$

$$\text{where } \text{attention}_j = \text{softmax} \left( \frac{(\mathbf{XW}^{Q_j})(SR(\mathbf{X})\mathbf{W}^K)^T}{\sqrt{d_k}} \right) (SR(\mathbf{X})\mathbf{W}^V) \quad (2)$$

where  $SR$  denotes either spatial reduction, our stride-2 DW, or, if spatial reduction isn’t used, the identity function. As shown in Tab. 3, asymmetric spatial down-sampling adds >20% efficiency with minimal accuracy loss (-0.06%).

## 6 Design of MNv4 Models

**Our Design Philosophy: Simplicity Meets Efficiency.** In developing the latest MobileNets, our core goal was Pareto optimality across diverse mobile platforms. To achieve this, we started by conducting extensive correlation analyses on existing models and hardware. Through empirical examination, we found a set of components and parameters that ensure high correlations between cost models (the prediction of cost of latency) across various devices while approaching the Pareto frontier in performance. Our investigation unveiled critical insights:

*Multi-path efficiency concerns:* Group convolutions [55] and similar multi-path designs, despite lower MAC counts, can be less efficient due to memory access complexity.

*Hardware support matters:* Advanced modules like Squeeze and Excite (SE) [23], GELU [18], and LayerNorm [2] are not well supported on DSPs, with LayerNorm also lagging behind BatchNorm [25], and SE is slow on accelerators.

*The Power of Simplicity:* Conventional components – depthwise and pointwise convolutions, ReLU [37], BatchNorm, and simple attention (e.g., MHSA) –



**Table 4:** Comparison between one-stage and two-stage searches, highlighting accuracy improvements and latency reduction on Pixel 6 EdgeTPU.

Search Method	Top-1 Acc (Val)	Top-1 Acc (Train)	Pixel 6 EdgeTPU (ms)
One-stage	81.26	74.64	3.85
Two-stage	81.48 (+ <b>0.22</b> )	78.24 (+ <b>3.60</b> )	3.67 ( <b>-4.68%</b> )

demonstrate superior efficiency and hardware compatibility.

Based on these findings, we established a set of design principles:

- **Standard Components:** We prioritize widely supported elements for seamless deployment and hardware efficiency.
- **Flexible UIB Blocks:** Our searchable UIB block lets NAS tune spatial and channel mixing, adjust receptive fields, and improve hardware utilization.
- **Employ Straightforward Attention:** Our Mobile MQA mechanism prioritizes simplicity for optimal performance.

These principles allow MobileNetV4 to be mostly Pareto-optimal on all hardware evaluated. In the following, we detail our refined NAS recipe for UIB model search, outline specific search configurations for various MNv4-Conv model sizes, and explain the construction of hybrid models.

## 6.1 Refining NAS for Enhanced Architectures

To effectively instantiate the UIB blocks, we adopt TuNAS [4] with tailored enhancements for improved performance. We use per-size searches and search spaces instead of using fixed scaling rules such as in EfficientNet [43].

**Enhanced Search Strategy:** Our approach mitigates TuNAS’s bias towards smaller filters and expansion factors, attributed to parameter sharing, by implementing a two-stage search. This strategy addresses the variance in parameter counts between UIB’s depthwise layers and other search options.

*Coarse-Grained Search:* Initially, we focus on determining optimal filter sizes while maintaining fixed parameters: an inverted bottleneck block with a default expansion factor of 4 and a 3x3 depthwise kernel.

*Fine-Grained Search:* Building on the initial search’s outcomes, we search the configuration of UIB’s two depthwise layers (including their presence and kernel size of either 3x3 or 5x5), keeping the expansion factor constant at 4.

Tab. 4 demonstrates the enhanced efficiency and model quality achieved through our two-stage search compared to a conventional one-stage search, where a unified search space was explored in a single TuNAS pass.

**Enhancing TuNAS with Robust Training:** The success of TuNAS hinges on accurately evaluating architecture quality, crucial for reward calculation and policy learning. Originally, TuNAS leveraged ImageNet-1k for training the SuperNet, but ImageNet performance is notably affected by data augmentation, regularization, and hyper-parameter choices. Given TuNAS’s evolving architecture samples, finding a stable set of hyper-parameters is challenging.

We address this with an offline distillation dataset, eliminating the need for extra augmentations and reducing sensitivity to regularization and optimization

**Table 5:** Performance Boost from JFT Distillation: NAS Training on ImageNet-1k vs. JFT Data. Highlights efficiency improvements and slight accuracy differences.

NAS Dataset	Top-1 Acc (Val/Train)	MACs	Params	Pixel 4 GPU	Pixel 6 CPU
ImageNet	82.4 / 72.9	7.2G	43.5M	59.2ms	70.4ms
JFT distill	82.3 / 74.0	6.2G	34.4M	51.0ms	67.3ms
Gain	-0.1 / +1.1	+ <b>13.9%</b>	+ <b>20.9%</b>	+ <b>13.9%</b>	+ <b>4.4%</b>

settings. The JFT distillation dataset, as detailed in Sec. 8, serves as our training set for TuNAS, with notable improvements shown in Tab. 5. Acknowledging that depth-scaling surpasses width-scaling in extended training sessions [3], we extend TuNAS training to 750 epochs, yielding deeper, higher-quality models.

## 6.2 Optimization of MNv4 Models

We constructed MNv4-Conv models from NAS-optimized UIB blocks, tailoring them for specific resource constraints. More details are given in Appendix A. In line with other hybrid models, we found that adding attention to the last stages of convolution models is most effective. In MNv4-Hybrid models, we interlace Mobile MQA blocks with UIB blocks for enhanced performance. For comprehensive model specifications, refer to Appendix D.

## 7 Results

In this section, we demonstrate the mostly Pareto-optimal performance of MobileNetV4 (MNv4) on ImageNet-1K classification and COCO object detection.

### 7.1 ImageNet classification

**Experimental Setup:** To assess model architecture performance, we train exclusively with the ImageNet-1k [12] training split and measure Top-1 accuracy on its validation split. Our latency analysis includes a representative selection of mobile hardware, including ARM Cortex CPUs (Pixel 6, Samsung S23), Qualcomm Hexagon DSP (Pixel 4), ARM Mali GPU (Pixel 7), Qualcomm Snapdragon (S23 GPU), Apple Neural Engine, and Google EdgeTPU. Our complete training recipe is detailed in the Appendix C.

We benchmark our models against the leading efficient models, including hybrid (MiT-EfficientViT [14], FastViT [45], NextViT [28]) and convolutional models (MobileOne [46], ConvNext [34], and previous MobileNets [21] [38] [20]) based on their reported Top-1 Accuracies and our latency evaluations. We used modern training recipes to improve MobileNetV1-V3 accuracy: a +3.4% (to 74.0%) for V1, +1.4% (to 73.4%) for V2, and +0.3% (to 75.5%) for V3. These new figures are used throughout the paper to isolate architectural advancements.

**Results:** Our results, seen in Fig. 1 and Tab. 6, demonstrate that MNv4 models are mostly Pareto-optimal across a range of accuracies and mobile targets, including CPUs, DSPs, GPUs, and accelerators like the Apple Neural Engine

**Table 6: Classification results on ImageNet-1K [12], along with on-device benchmarks.** Median latency is reported. A – indicates that we did not benchmark a model due to missing corresponding model file for a platform. *Failed* indicates that the model is not supported by the platform. Dividers denote approximate latency classes.

Model	Top-1	Latency (ms)								
		Params (M)	MACs (G)	Pixel 6 CPU	Pixel 8 EdgeTPU	iPhone 13 CoreML	Pixel 4 Hexagon	Pixel 7 GPU	Samsung S23 CPU	S23 GPU
MobileNet-V2-0.5x [38]	66	2.0	0.1	2.4	0.7	0.5	2.9	8.3	1.8	1.9
MobileNet-V3L-0.5x [20]	69.2	2.7	0.1	2.4	0.8	0.45	3.5	9.9	2.0	2.1
MobileOne-S0 [46]	71.4	2.1	0.3	4.2	0.7	0.5	2.9	10.7	3.3	1.7
MobileNet-V2 [38]	73.4	3.5	0.3	5.0	0.7	0.7	3.9	13.6	4.1	2.5
<b>MNv4-Conv-S</b>	<b>73.8</b>	<b>3.8</b>	<b>0.2</b>	<b>2.4</b>	<b>0.7</b>	<b>0.6</b>	<b>2.4</b>	<b>8.4</b>	<b>1.8</b>	<b>2.0</b>
MobileNet-V1 [22]	74.0	4.2	0.6	6.1	0.8	0.7	3.2	13.0	4.6	2.1
FastViT-T8 <sup>†</sup> [45]	75.6	3.6	0.7	49.3	1.3	0.7	<i>Failed</i>	40.7	43.6	24.7
MobileNet-V2-1.5x [38]	76.8	6.8	0.7	9.3	0.9	1.0	5.6	16.4	7.3	3.3
MultiHardware-MAX-1.5x [9]	77.9	8.9	0.8	9.8	1.0	-	5.7	23.2	-	4.1
MultiHardware-AVG-1.5x [9]	78.2	10.0	1.0	12.0	1.1	-	6.1	20.3	-	4.5
MobileNet-V2-2.0x [38]	78.4	11.2	1.1	13.9	1.1	1.5	6.9	19.1	10.6	4.2
MobileOne-S4 [46]	79.4	14.8	1.5	26.7	1.7	1.5	9.0	28.6	19.4	5.9
FastViT-S12 <sup>†</sup> [45]	79.8	8.8	1.8	83.0	1.8	1.6	<i>Failed</i>	75.0	69.2	47.0
MIT-EfficientViT-B1-r224 [14]	79.4	9.1	0.5	-	-	2.4	-	-	18.1	5.0
<b>MNv4-Conv-M</b>	<b>79.9</b>	<b>9.2</b>	<b>1.0</b>	<b>11.4</b>	<b>1.1</b>	<b>1.1</b>	<b>7.3</b>	<b>18.1</b>	<b>8.6</b>	<b>4.1</b>
FastViT-SA12 [45]	80.6	10.9	1.9	86.5	2.0	1.6	<i>Failed</i>	79.6	69.5	52.1
<b>MNv4-Hybrid-M</b>	<b>80.7</b>	<b>10.5</b>	<b>1.2</b>	<b>14.3</b>	<b>1.5</b>	-	<i>Failed</i>	<b>17.9</b>	<b>10.8</b>	<b>5.9</b>
FastViT-SA24 [45]	82.6	20.6	3.8	171.6	3.2	2.4	<i>Failed</i>	131.9	136.3	107.5
MIT-EfficientViT-B2-r256 [14]	82.7	24.0	2.1	-	-	5.4	-	-	64.9	9.5
<b>MNv4-Conv-L</b>	<b>82.9</b>	<b>31</b>	<b>5.9</b>	<b>59.9</b>	<b>2.4</b>	<b>3.0</b>	<b>20.8</b>	<b>37.6</b>	<b>43.0</b>	<b>13.2</b>
ConvNext-S [34]	83.1	50	8.7	314.9	3.7	-	<i>Failed</i>	45.2	243.9	18.5
NextViT-B [28]	83.2	44.8	8.3	-	-	-	-	-	-	-
<b>MNv4-Hybrid-L</b>	<b>83.4</b>	<b>35.9</b>	<b>7.2</b>	<b>87.6</b>	<b>3.8</b>	-	<i>Failed</i>	<b>61.3</b>	<b>61.8</b>	<b>18.1</b>
MIT-EfficientViT-B3-r224 [14]	83.5	49.0	4.0	-	-	12.2	-	-	125.9	18.4
FastViT-SA36 [45]	83.6	30.4	5.6	241.6	4.3	-	<i>Failed</i>	186.5	206.3	138.1

and Google EdgeTPU.

MNv4 performs notably well on CPU—roughly 2x faster than MobileNetV3 and substantially faster than iso-accuracy models. On EdgeTPUs, MNv4 models are as accurate as MobileNetV3 and 2x as fast. MNv4-Conv-M is >50% faster than MobileOne-S4 and FastViT-S12 and has +1.5% more Top-1 accuracy than MobileNetV2 at comparable latency. On S23 GPU and iPhone 13 CoreML (ANE), MNv4 is mostly at the Pareto front. MIT-EfficientViT—which has the closest performance on S23 GPU—has >2x the latency as MNv4 on CoreML at the same accuracy. FastViT—optimized for Apple Neural Engine—is 2nd on CoreML but has >5x the latency of MNv4 on S23 GPU. While some models, such as EfficientViT, reach the same accuracy with fewer MACs, MobileNetV4 models are optimized for high accuracy and minimal latency on the most hardware possible. Increasing MACs often decreases memory bandwidth and op complexity which is often more important for achieving this goal. Like many hybrid models, MNv4-hybrid models are not compatible with DSPs. MNv4-Conv models remain the top performers on DSP, emphasizing the compatibility and efficiency across diverse hardware provided by our UIB block, NAS recipe, and search spaces. MNv4-Hybrid performs well on CPUs and accelerators which demonstrates the broad efficiency of Mobile MQA.

**Table 7:** Object detection results on the COCO-17 [33] Val. set. The width-multiplier is reported next to the MobileNet backbones that were scaled-up.

Backbone	COCO Val AP	MACs (G)	Params (M)	Pixel 6 CPU latency (ms)
EfficientFormer L1 [30]	29.5	6.54	12.77	84.3
MobileNet v1 @ 1.5 [21]	31.0	6.68	9.05	66.4
<b>MNv4-Conv-M</b>	<b>32.6</b>	<b>5.06</b>	<b>9.79</b>	<b>51.3</b>
MobileNet Multi-AVG @ 1.5 [9]	32.7	5.42	9.51	58.1
MobileNet v2 @ 2.0 [38]	32.9	5.81	10.15	66.4
MobileNet v3 Large [20] @ 2.0	33.2	4.99	17.92	59.9
<b>MNv4-Hybrid-M</b>	<b>34.0</b>	<b>5.62</b>	<b>11.15</b>	<b>60.5</b>

Mobile models should perform well on diverse hardware, but we show that many models fail to meet this requirement. MobileNetV3 performs well on CPUs but not on EdgeTPU, DSPs, and GPUs. FastViT performs well on ANE but not on CPUs and GPUs. EfficientViT has good performance on GPUs but not on ANE. In contrast, MNv4-Conv models achieves mostly-Pareto-optimal performance across CPUs, GPUs, DSPs, the Apple Neural Engine, and Google EdgeTPUs. This versatility ensures MNv4-Conv models can be easily deployed across the mobile ecosystem and sets a new benchmark for mobile model universality.

## 7.2 COCO Object Detection

**Experimental Setup:** We evaluate the effectiveness of MNv4 backbones for object detection tasks on the COCO 17 [33] dataset. We compare MNv4 medium backbones against SOTA backbones with a MAC count. For each backbone, we build a detector using the RetinaNet [32] framework. We attach a 256-d FPN [31] decoder to the P3 - P7 endpoints, as well as a 256-d prediction head with 4 convolutional layers. As usual for mobile detectors, we use depth-separable convolutions for an efficient FPN decoder and box prediction head. We train all models on COCO 17 [33] for 600-epochs. Images are resized to  $384px$  and augmented with random horizontal flip, random scale, and Randaug [10]. We exclude Shear and Rotate from Randaug, as those deteriorate small-object detection AP. The models are trained with a 2048 batch size, Adam [26], and a 0.00003 L2 weight decay, plus a cosine LR schedule with 24 epochs warm-up. The learning rate is tuned per-model. For all baselines, filter multipliers are tuned to similar MACs. Following classification, MobileNetV4 backbones are trained using a 0.2 stochastic drop [24]. MobileNet baselines were from Tensorflow Model Garden [19] implementation. EfficientFormer was reimplemented in Tensorflow.

**Results:** Results are reported in Tab. 7. Parameters, MACs and benchmarks are computed using the entire detector at the  $384px$  input resolution. The MNv4-Conv-M detector achieves 32.6% AP, similar to MobileNetMultiAvg and MobileNetV2. However, this model is 12% faster than MobileNetMultiAvg and 23% faster than MobileNetV2 on Pixel 6 CPU. The MNv4-Hybrid-M detector gets +1.6% AP over MNv4-Conv-M while running 18% slower on Pixel 6 CPU. This demonstrates the effectiveness of MNv4 hybrid models on tasks like object detection.

## 8 Enhanced distillation recipe

Complementing architectural innovation, distillation is a powerful tool for enhancing machine learning efficiency. This is particularly true for mobile models where distillation can greatly increase accuracy without increasing latency. Building upon the Patient Teacher distillation baseline [5], we introduce two novel techniques to further boost performance.

**Dynamic Dataset Mixing:** Data augmentation is crucial for distillation performance. While prior methods rely on a fixed augmentation sequence, we find that dynamically mixing multiple datasets with diverse augmentation strategies improves distillation. Our experiments use three distillation datasets:

$\mathcal{D}_1$  : Inception Crop [41] followed by RandAugment [11] l2m9 applied to 500 ImageNet-1k replicas.

$\mathcal{D}_2$  : Inception Crop followed by extreme Mixup [54] applied to 1000 ImageNet-1k replicas (mirroring the Patient Teacher approach).

$\mathcal{D}_1 + \mathcal{D}_2$  : A dynamic mixture of  $\mathcal{D}_1$  and  $\mathcal{D}_2$  during training.

Our results (Tab. 8) show that  $\mathcal{D}_2$  outperforms  $\mathcal{D}_1$  (84.1% vs. 83.8% student accuracy), but a dynamic mixture of the two ( $\mathcal{D}_1 + \mathcal{D}_2$ ) elevates accuracy to 84.4% (+0.3%). This suggests that mixing expands the augmented image space, increases difficulty and diversity, and leads to improved student performance.

**JFT Data Augmentation:** To increase training data volume, we add in-domain, class-balanced data by resampling the JFT-300M [40] dataset to 130K images per class (130M total). Following Noisy Student [51] and using EfficientNet-B0 trained on ImageNet-1K, we select images with a relevance threshold above 0.3. For classes with abundant data, we choose the top 130K images; for rare classes, we replicate images for balance. This dataset is replicated 10x. Due to JFT’s complexity, we apply weaker augmentations (Inception Crop + RandAugment l2m5). This is dataset  $\mathcal{D}_3$ . Tab. 8 shows that using solely  $\mathcal{D}_3$  drops accuracy by 2%. However, combining ImageNet and JFT data ( $\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3$ ) raises accuracy by +0.6%. The additional data improves generalization.

**Our distillation recipe:** Our combined distillation recipe dynamically mixes datasets  $\mathcal{D}_1$ ,  $\mathcal{D}_2$ , and  $\mathcal{D}_3$  for diverse augmentations and leverages class-balanced JFT data. As shown in Tab. 8 and Tab. 9, our method improves top-1 accuracy >0.8% over the previous SOTA [5]. Training an MNv4-Conv-L student model for 2000 epochs yields 85.9% top-1 accuracy. Our approach is effective: the student has 15x fewer parameters and 48x fewer MACs than its teacher (EfficientNet-L2), but is only 1.6% less accurate. MNv4-Conv-Hybrid reaches 87.0% top-1 accuracy by combining this distillation with pretraining on JFT. More details of our distillation recipe can be found in Appendix H.

## 9 Conclusion

In this paper, we presented MobileNetV4, a series of universal, efficient models that run efficiently across the mobile ecosystem. Multiple advances make

**Table 8:** Distillation results using MNv4-Conv-L as student, highlighting gains over SOTA and marking our contributions explicitly.

Dataset	Data source	Augmentations	Mixing Ratio	Top-1 Acc (Val/Train)	
				400 epochs	2000 epochs
$\mathcal{D}_1$	1000× ImageNet-1k	Inception Crop & RandAug l2m9	-	83.8/86.6	-
$\mathcal{D}_2$ (SOTA [5])	1000× ImageNet-1k	Inception Crop & Extreme Mixup	-	84.1/85.6	-
$\mathcal{D}_3$	10× JFT subset	Inception Crop & RandAug l2m5	-	81.8/84.1	-
<b>Ours: <math>\mathcal{D}_1 + \mathcal{D}_2</math></b>			1:1	<b>84.4/85.0 (+0.3)</b>	-
<b>Ours: <math>\mathcal{D}_2 + \mathcal{D}_3</math></b>			1:1	<b>84.7/82.7 (+0.6)</b>	-
<b>Ours: <math>\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3</math></b>			1:1:2	<b>84.9/82.6 (+0.8)</b>	85.9/85.5 (+1.8)

**Table 9: Top-1 Accuracy Comparison Across Training Approaches:** This table contrasts baseline ImageNet-1k training, SOTA distillation, and our distillation.

Model	IN-1k Only Only	SOTA Distill [5]	Our Distill	Our Gain Over IN-1k / SOTA
MNv4-Conv-S	73.8	-	<b>75.5</b>	+1.7 / -
MNv4-Conv-M	79.9	81.5	<b>82.7</b>	+2.8 / +1.2
MNv4-Hybrid-M	80.7	82.7	<b>83.7</b>	+3.0 / +1.0
MNv4-Conv-L	82.9	84.4	<b>85.9</b>	+3.0 / +1.5
MNv4-Hybrid-L	83.4	85.7	<b>86.6</b>	+3.2 / +0.9

MobileNetV4 mostly-Pareto-optimal on all mobile CPUs, GPUs, DSPs and specialized accelerators, a characteristic not found in any other models tested. We introduced the Universal Inverted Bottleneck and Mobile MQA layers and combined them with improved NAS recipes. With these and a novel, SOTA distillation approach, we achieve 87% ImageNet-1K accuracy at 3.8ms Pixel 8 EdgeTPU latency, setting a new state-of-the-art. Finally, we introduced a framework for understanding model universality on heterogeneous devices. We hope the novel contributions and analysis further spur advances in mobile computer vision.

## Acknowledgements

We appreciate Tammo Spalink, Yeqing Li, Sage Stevens, Bob Muniz, Liviu Panait, David Wood, Lynn Nguyen, and Lucas Beyer for their support while developing this work. We also thank the MLPerf Mobile working group for their feedback and collaboration.<sup>5</sup> We particularly thank Ross Wightman for his reimplementation of MobileNetV4 for `timm` and training recipe improvements.<sup>6</sup>

<sup>5</sup> MobileNetV4-Conv-L was used for v4.0 of <https://mlcommons.org/benchmarks/inference-mobile/>

<sup>6</sup> MobileNetV4 models are available in `timm` at <https://huggingface.co/collections/timm/mobilenetv4-pretrained-weights-6669c22cda4db4244def9637>

## References

1. Berkin Akin, Suyog Gupta, Yun Long, Anton Spiridonov, Zhuo Wang, Marie White, Hao Xu, Ping Zhou, and Yanqi Zhou. Searching for efficient neural architectures for on-device ML on edge tpus. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19-20, 2022*, pages 2666–2675. IEEE, 2022.
2. Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
3. Irwan Bello, William Fedus, Xianzhi Du, Ekin Dogus Cubuk, Aravind Srinivas, Tsung-Yi Lin, Jonathon Shlens, and Barret Zoph. Revisiting resnets: Improved training and scaling strategies. *Advances in Neural Information Processing Systems*, 34:22614–22627, 2021.
4. Gabriel Bender, Hanxiao Liu, Bo Chen, Grace Chu, Shuyang Cheng, Pieter-Jan Kindermans, and Quoc V. Le. Can weight sharing outperform random architecture search? an investigation with tunas. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
5. Lucas Beyer, Xiaohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov. Knowledge distillation: A good teacher is patient and consistent. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10925–10934, 2022.
6. Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791*, 2019.
7. Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Xiaoyi Dong, Lu Yuan, and Zicheng Liu. Mobile-former: Bridging mobilenet and transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5270–5279, 2022.
8. Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
9. Grace Chu, Okan Arikan, Gabriel Bender, Weijun Wang, Achille Brighton, Pieter-Jan Kindermans, Hanxiao Liu, Berkin Akin, Suyog Gupta, and Andrew Howard. Discovering multi-hardware mobile models via architecture search. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2021, virtual, June 19-25, 2021*, pages 3022–3031. Computer Vision Foundation / IEEE, 2021.
10. Ekin Dogus Cubuk, Barret Zoph, Jonathon Shlens, and Quoc Le. Randaugment: Practical automated data augmentation with a reduced search space. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
11. Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020.
12. J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, June 2009.

13. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
14. Han Cai et al. Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction. In *ICCV*, 2023.
15. Xinyu Liu et al. Efficientvit: Memory efficient vision transformer with cascaded group attention. In *CVPR*, 2023.
16. Jianyuan Guo, Kai Han, Han Wu, Yehui Tang, Xinghao Chen, Yunhe Wang, and Chang Xu. Cmt: Convolutional neural networks meet vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12175–12185, 2022.
17. 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 (CVPR)*, June 2020.
18. Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
19. Xianzhi Du Yeqing Li Abdullah Rashwan Le Hou Pengchong Jin Fan Yang Frederick Liu Jaeyoun Kim Hongkun Yu, Chen Chen and Jing Li. TensorFlow Model Garden. <https://github.com/tensorflow/models>, 2020.
20. Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for mobilenetv3. *CoRR*, abs/1905.02244, 2019.
21. 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. *CoRR*, abs/1704.04861, 2017.
22. 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.
23. Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
24. Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 646–661. Springer, 2016.
25. Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015.
26. Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015.
27. Ariel N Lee, Cole J Hunter, and Nataniel Ruiz. Platypus: Quick, cheap, and powerful refinement of llms. *arXiv preprint arXiv:2308.07317*, 2023.
28. Jiashi Li, Xin Xia, Wei Li, Huixia Li, Xing Wang, Xuefeng Xiao, Rui Wang, Min Zheng, and Xin Pan. Next-vit: Next generation vision transformer for efficient deployment in realistic industrial scenarios. *arXiv preprint arXiv:2207.05501*, 2022.
29. Yanyu Li, Ju Hu, Yang Wen, Georgios Evangelidis, Kamyar Salahi, Yanzhi Wang, Sergey Tulyakov, and Jian Ren. Rethinking vision transformers for mobilenet



- size and speed. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16889–16900, 2023.
30. Yanyu Li, Geng Yuan, Yang Wen, Ju Hu, Georgios Evangelidis, Sergey Tulyakov, Yanzhi Wang, and Jian Ren. Efficientformer: Vision transformers at mobilenet speed. In *NeurIPS*, 2022.
  31. Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 936–944. IEEE Computer Society, 2017.
  32. Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2999–3007. IEEE Computer Society, 2017.
  33. Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer, 2014.
  34. Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 11966–11976. IEEE, 2022.
  35. Sachin Mehta and Mohammad Rastegari. Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer. *arXiv preprint arXiv:2110.02178*, 2021.
  36. Sachin Mehta and Mohammad Rastegari. Separable self-attention for mobile vision transformers. *arXiv preprint arXiv:2206.02680*, 2022.
  37. Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.
  38. Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *CoRR*, abs/1801.04381, 2018.
  39. Noam Shazeer. Fast transformer decoding: One write-head is all you need. *arXiv preprint arXiv:1911.02150*, 2019.
  40. Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 843–852, 2017.
  41. Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
  42. Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. *CoRR*, abs/1807.11626, 2018.
  43. Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
  44. Yehui Tang, Kai Han, Jianyuan Guo, Chang Xu, Chao Xu, and Yunhe Wang. Ghostnetv2: Enhance cheap operation with long-range attention. *Advances in Neural Information Processing Systems*, 35:9969–9982, 2022.

45. Pavan Kumar Anasosalu Vasu, James Gabriel, Jeff Zhu, Oncel Tuzel, and Anurag Ranjan. Fastvit: A fast hybrid vision transformer using structural reparameterization. *arXiv preprint arXiv:2303.14189*, 2023.
46. Pavan Kumar Anasosalu Vasu, James Gabriel, Jeff Zhu, Oncel Tuzel, and Anurag Ranjan. Mobileone: An improved one millisecond mobile backbone. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7907–7917, 2023.
47. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
48. Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 568–578, 2021.
49. Samuel Williams, Andrew Waterman, and David Patterson. Roofline: an insightful visual performance model for multicore architectures. *Communications of the ACM*, 52(4):65–76, 2009.
50. Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. *CoRR*, abs/1812.03443, 2018.
51. Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10687–10698, 2020.
52. Tien-Ju Yang, Andrew G. Howard, Bo Chen, Xiao Zhang, Alec Go, Vivienne Sze, and Hartwig Adam. Netadapt: Platform-aware neural network adaptation for mobile applications. *CoRR*, abs/1804.03230, 2018.
53. Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019.
54. Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
55. Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. *CoRR*, abs/1707.01083, 2017.