

MobileDets: Searching for Object Detection Architectures for Mobile Accelerators

Yunyang Xiong*
University of Wisconsin-Madison
yxiong43@wisc.edu

Hanxiao Liu*
Google
hanxiaol@google.com

Suyog Gupta
Google
suyoggupta@google.com

Berkin Akin, Gabriel Bender, Yongzhe Wang, Pieter-Jan Kindermans, Mingxing Tan
Google
{bakin, gbender, yongzhe, pikinder, tanmingxing}@google.com

Vikas Singh
University of Wisconsin-Madison
vsingh@biostat.wisc.edu

Bo Chen
Google
bochen@google.com

Abstract

Inverted bottleneck layers, which are built upon depthwise convolutions, have been the predominant building blocks in state-of-the-art object detection models on mobile devices. In this work, we investigate the optimality of this design pattern over a broad range of mobile accelerators by revisiting the usefulness of regular convolutions. We discover that regular convolutions are a potent component to boost the latency-accuracy trade-off for object detection on accelerators, provided that they are placed strategically in the network via neural architecture search. By incorporating regular convolutions in the search space and directly optimizing the network architectures for object detection, we obtain a family of object detection models, MobileDets, that achieve state-of-the-art results across mobile accelerators. On the COCO object detection task, MobileDets outperform MobileNetV3+SSDLite by 1.7 mAP at comparable mobile CPU inference latencies. MobileDets also outperform MobileNetV2+SSDLite by 1.9 mAP on mobile CPUs, 3.7 mAP on Google EdgeTPU, 3.4 mAP on Qualcomm Hexagon DSP and 2.7 mAP on Nvidia Jetson GPU without increasing latency. Moreover, MobileDets are comparable with the state-of-the-art MnasFPN on mobile CPUs even without using the feature pyramid, and achieve better mAP scores on both EdgeTPUs and DSPs with up to 2× speedup. Code and models are available in the TensorFlow Object Detection API [16]: https://github.com/tensorflow/models/tree/master/research/object_detection.

*Equal contribution.

1. Introduction

In many computer vision applications it can be observed that higher capacity networks lead to superior performance [29, 44, 15, 24]. However, they are often more resource-consuming. This makes it challenging to find models with the right quality-compute trade-off for deployment on edge devices with limited inference budgets.

A lot of effort has been devoted to the manual design of lightweight neural architectures for edge devices [14, 27, 42, 39]. Unfortunately, relying on human expertise is time-consuming and can be sub-optimal. This problem is made worse by the speed at which new hardware platforms are released. In many cases, these newer platforms have differing performance characteristics which make a previously developed model sub-optimal.

To address the need for automated tuning of neural network architectures, many methods have been proposed. In particular, neural architecture search (NAS) methods [5, 32, 30, 13, 10] have demonstrated a superior ability in finding models that are not only accurate but also efficient on a specific hardware platform.

Despite many advancements in NAS algorithms [43, 24, 1, 19, 5, 32, 40, 30, 13, 10], it is remarkable that inverted bottlenecks (IBN) [27] remain the predominant building block in state-of-the-art mobile models. IBN-only search spaces have also been the go-to setup in a majority of the related NAS publications [30, 5, 32, 13]. IBN layers rely heavily on depthwise and depthwise separable convolutions [28]. The resulting models have relatively low FLOPS and parameter counts, and can be executed efficiently on CPUs.

However, the advantage of depthwise convolutions for

mobile inference is less clear for hardware accelerators such as DSPs or EdgeTPUs which are becoming increasingly popular on mobile devices. For example, for certain tensor shapes and kernel dimensions, a regular convolution can run $3\times$ as fast as the depthwise variation on an EdgeTPU despite having $7\times$ as many FLOPS. This observation leads us to question the exclusive use of IBN-only search spaces in most current state-of-the-art mobile architectures.

Our work seeks to rethink the use of IBN-only search space on modern mobile accelerators. To this end, we propose the MobileDet search space family, which includes not only IBNs but also flexible full convolution sequences motivated by the structure of tensor decomposition [33, 6, 21]. Using the task of object detection as an example (one of the most popular mobile vision applications), we show the MobileDet search space family enables NAS methods to identify models with substantial better latency-accuracy trade-offs on mobile CPUs, DSPs, EdgeTPUs and edge GPUs.

To evaluate our proposed MobileDet search space, we perform latency-aware NAS for object detection, targeting a diverse set of mobile platforms. Experimental results show that by using our MobileDet search space family and directly searching on detection tasks, we can consistently improve the performance across all hardware platforms. By leveraging full convolutions at selected positions in the network, our method outperforms IBN-only models by a significant margin. Our searched models, MobileDets, outperform MobileNetV2 classification backbone by 1.9mAP on CPU, 3.7mAP on EdgeTPU, 3.4mAP on DSP, and 2.7mAP on edge GPU at comparable inference latencies. MobileDets also outperform the state-of-the-art MobileNetV3 classification backbone by 1.7mAP at similar CPU inference efficiency. Further, the searched models achieved comparable performance with the state-of-the-art mobile CPU detector, MnasFPN [7], without leveraging the NAS-FPN head which may complicate the deployment. On both EdgeTPUs and DSPs, MobileDets are more accurate than MnasFPN while being more than twice as fast.

Our main contributions can be summarized as follows:

- Unlike many existing works which exclusively focused on IBN layers for mobile applications, we propose an augmented search space family with building blocks based on regular convolutions. We show that NAS methods can substantially benefit from this enlarged search space to achieve better latency-accuracy trade-off on a variety of mobile devices.
- We deliver MobileDets, a set of mobile object detection models with state-of-the-art quality-latency trade-offs on multiple hardware platforms, including mobile CPUs, EdgeTPUs, DSPs and edge GPUs. Code and models will be released to benefit a wide range of on-device object detection applications.

2. Related Work

2.1. Mobile Object Detection

Object detection is a classic computer vision challenge where the goal is to learn to identify objects of interest in images. Existing object detectors can be divided into two categories: two-stage detectors and one-stage single shot detectors. For two-stage detectors, including Faster RCNN [26], R-FCN [9] and ThunderNet [23], region proposals must be generated first before the detector can make any subsequent predictions. Two-stage detectors are not efficient in terms of inference time due to this multi-stage nature. On the other hand, one-stage single shot detectors, such as SSD [20], SSDLite [27], YOLO [25], SqueezeDet [38] and Pelee [35], require only one single pass through the network to predict all the bounding boxes, making them ideal candidates for efficient inference on edge devices. We therefore focus on one-stage detectors in this work.

SSDLite [27] is an efficient variant of SSD that has become one of the most popular lightweight detection heads. It is well suited for use cases on mobile devices. Efficient backbones, such as MobileNetV2 [27] and MobileNetV3 [13], are paired with SSDLite to achieve state-of-the-art mobile detection results. Both models will be used as baselines to demonstrate the effectiveness of our proposed search spaces over different mobile accelerators.

2.2. Mobile Neural Architecture Search (NAS)

NetAdapt [41] and AMC [12] were among the first attempts to utilize latency-aware search to finetune the number of channels of a pre-trained model. MnasNet [31] and MobileNetV3 [13] extended this idea to find resource-efficient architectures within the NAS framework. With a combination of techniques, MobileNetV3 delivered state-of-the-art architectures on mobile CPU. As a complementary direction, there are many recent efforts aiming to improve the search efficiency of NAS [3, 1, 22, 19, 5, 37, 4].

2.3. NAS for Mobile Object Detection

A large majority of the NAS literature [32, 30, 13] focuses on classification and only re-purposes the learned feature extractor as the backbone for object detection without further searches. Recently, multiple papers [8, 34, 7] have shown that better latency-accuracy trade-offs are obtained by searching directly for object detection models.

One strong detection-specific NAS baseline for mobile detection models is MnasFPN [7], which searches for the feature pyramid head with a mobile-friendly search space that heavily utilizes depthwise separable convolutions. Several factors limit its generalization towards mobile accelerators: (1) so far both depthwise convolutions and feature pyramids are less optimized on these platforms, and (2) MnasFPN does not search for the backbone, which is a




Figure 1: Platform-aware NAS and our MobileDet search space work synergistically to boost object detection performance on accelerators. SSDLite object detection performance on Pixel-4 DSPs with different backbone designs: manually-designed MobileNetV2, searched with IBN-only search space, and searched with the proposed MobileDet space (with both IBNs and full conv-based building blocks). Layers are visualized as vertical bars where color indicates layer type and length indicates expansion ratio. C_4 and C_5 mark the feature inputs to the SSDLite head. While conducting platform-aware NAS in an IBN-only search space achieves a 1.6 mAP boost over the handcrafted baseline, searching within the MobileDet space brings another 1.6 mAP gain.

bottleneck for latency. By comparison, our work relies on SSD heads and proposes a new search space for the backbone based on full-convolutions, which are more amenable to mobile acceleration. While it is challenging to develop a generic search space family that spans a set of diverse and dynamically-evolving mobile platforms, we take a first step towards this goal, starting from the most common platforms such as mobile CPUs, DSPs and EdgeTPUs.

3. Revisiting Full Convolutions for Mobile Search Spaces

In this section, we first explain why IBN layers may not be sufficient to handle mobile accelerators beyond mobile CPUs. We then propose new building blocks based on regular convolutions to enrich our search space, and discuss the connections between these building blocks and Tucker/CP decompositions [33, 6].

Are IBNs all we need? The layout of an Inverted Bottleneck (IBN) is illustrated in Figure 2. IBNs are designed to reduce the number of parameters and FLOPS, and leverage depthwise and pointwise (1×1) convolutional kernels to achieve high efficiency on mobile CPUs. However, not all FLOPS are the same, especially for modern mobile accelerators such as EdgeTPU and DSPs. For example, a regular convolution may run $3 \times$ as fast on EdgeTPUs than its depthwise variation even with $7 \times$ as many FLOPS. The observation indicates that the widely used IBN-only search space can be suboptimal for modern mobile accelerators. This motivated us to propose new building blocks by revisiting regular (full) convolutions to enrich IBN-only search spaces for mobile accelerators. Specifically, we propose two flexible layers to perform channel expansion and compression, respectively, which are detailed below.




Figure 2: Inverted bottleneck layer: 1×1 pointwise convolution transforms the input channels from C_1 to $s \times C_1$ with input expansion ratio $s > 1$, then $K \times K$ depthwise convolution transforms the input channels from $s \times C_1$ to $s \times C_1$, and the last 1×1 pointwise convolution transforms the channels from $s \times C_1$ to C_2 . The highlighted C_1 , s , K , C_2 in IBN layer are searchable.




Figure 3: Fused inverted bottleneck layer: $K \times K$ regular convolution transforms the input channels from C_1 to $s \times C_1$ with input expansion ratio $s > 1$, and the last 1×1 pointwise convolution transforms the channels from $s \times C_1$ to C_2 . The highlighted C_1 , K , s , C_2 in the fused inverted bottleneck layer are searchable.




Figure 4: Tucker layer: 1×1 pointwise convolution transforms the input channels C_1 to $s \times C_1$ with input compression ratio $s < 1$, then $K \times K$ regular convolution transforms the input channels from $s \times C_1$ to $e \times C_2$ with output compression ratio $e < 1$, and the last 1×1 pointwise convolution transforms the channels from $e \times C_2$ to C_2 . The highlighted C_1 , s , K , e , C_2 in Tucker layer are searchable.

3.1. Fused Inverted Bottleneck Layers (Expansion)

The depthwise-separable convolution [28] is a critical element of an inverted bottleneck [14] (Figure 2). The idea behind the depthwise-separable convolution is to replace an “expensive” full convolution with a combination of a depthwise convolution (for spatial dimension) and a 1×1 pointwise convolution (for channel dimension). However, the notion of expensiveness was largely defined based on FLOPS or the number of parameters, which are not necessarily correlated with the inference efficiency on modern mobile accelerators. To incorporate regular convolutions, we propose to modify an IBN layer by fusing together its first 1×1 convolution and its subsequent $K \times K$ depthwise convolution into a single $K \times K$ regular convolution (Figure 3). Like a standard inverted bottleneck layer, the initial convolution in our fused inverted bottleneck increases the number of filters by a factor $s > 1$. The expansion ratio of this layer will be determined by the NAS algorithm.

3.2. Tucker Convolution Layers (Compression)

Bottleneck layers were introduced in ResNet [11] reduce the cost of large convolutions over high-dimensional feature maps. A bottleneck layer with compression ratio $s < 1$ consists of a 1×1 convolution with C_1 input filters and $s \cdot C_1$ output filters, followed by a $K \times K$ convolution with $s \cdot C_1$ output filters, followed by a 1×1 convolution with C_2 output filters. We generalize these bottlenecks (Figure 4) by allowing the initial 1×1 convolution to potentially have a different number of output filters than the $K \times K$ convolution and let the NAS algorithm to decide the best configurations.

We refer to these new building blocks as *Tucker convolution layers* because of their connections with Tucker decomposition. (See supplement for details.)

4. Architecture Search Method

Search Algorithm. Our proposed search spaces are complementary to any neural architecture search algorithms. In our experiments, we employ TuNAS [2] for its scalability and its reliable improvement over random baselines. TuNAS constructs a one-shot model that encompasses all architectural choices in a given search space, as well as a controller whose goal is to pick an architecture that optimizes a platform-aware reward function. The one-shot model and the controller are trained together during a search. In each step, the controller samples a random architecture from a multinomial distribution that spans over the choices, then the portion of the one-shot model’s weights associated with the sampled architecture are updated, and finally a reward is computed for the sampled architecture, which is used to update the controller. The update is given by applying REINFORCE algorithm [36] on the following reward function:

$$R(M) = mAP(M) + \tau \left| \frac{c(M)}{c_0} - 1 \right|$$

where $mAP(M)$ denotes the mAP of an architecture M , $c(M)$ is the inference cost (in this case, latency), c_0 is the given cost budget, and $\tau < 0$ is a hyper-parameter that balances the importance of accuracy and inference cost. The search quality tends to be insensitive to τ , as shown in [2].

Cost Models. We train a *cost model*, $c(\cdot)$ – a linear regression model whose features are composed of, for each layer, an indicator of the cross-product between input/output channel sizes and layer type. This model has high fidelity across platforms ($r^2 \geq 0.99$). The linear cost model is related to previously proposed methods based on lookup tables [41, 12, 30], but only requires us to benchmark the latencies of randomly selected models within the search space, and does not require us to measure the costs of individual network operations such as convolutions.

Since $R(M)$ is computed at every update step, efficiency is key. During search, we estimate $mAP(M)$ based on a small mini-batch for efficiency, and use the regression model as a surrogate for on-device latency $c(M)$. To collect training data for the cost model, we randomly sample several thousand network architectures from our search space and benchmark each one on device. This is done only once per hardware and prior to search, eliminating the need for direct communication between server-class ML hardware and mobile phones. For final evaluation, the found architectures are benchmarked on the actual hardware instead of the cost model.

5. Experiments

We use the COCO dataset for our object detection experiments. We report mean average precision (mAP) for object detection and the real latency of searched models on accelerators with image size 320×320 . We conduct our experiments in two stages: architecture search for optimal architecture and architecture evaluation via retraining the found architecture from scratch.

5.1. Implementation Details

Standalone Training. We use 320×320 image size for both training and evaluation. The training is carried out using 32 Cloud TPU v2 cores. For fair comparison with existing models, we use standard preprocessing in the TensorFlow object detection API without additional enhancements such as drop-block or auto-augment. We use SGD with momentum 0.9 and weight decay 5×10^{-4} . The learning rate is warmed up in the first 2000 steps and then follows cosine decay. All models are trained from scratch without any ImageNet pre-trained checkpoint. We consider two different training schedules:

- *Short-schedule*: Each model is trained for 50K steps with batch size 1024 and an initial learning rate of 4.0.
- *Long-schedule*: Each model is trained for 400K steps with batch size 512 and an initial learning rate of 0.8.

The short schedule is about $4\times$ faster than the long schedule but would result in slightly inferior quality. Unless otherwise specified, we use the short schedule for ablation studies and the long schedule for the final results in Table 1.

Architecture Search. To avoid overfitting the true validation dataset, we split out 10% of the COCO training data to evaluate the models and compute rewards during search. Hyperparameters for training the shared weights follow the short schedule for standalone training. As for reinforcement learning, we use Adam optimizer with an initial learning rate of 5×10^{-3} , $\beta = (0, 0.999)$ and $\epsilon = 10^{-8}$. We search for 50K steps for ablation studies and search for 100K steps to obtain the best candidates in the main results table.

5.2. Latency Benchmarking

We report on-device latencies for all of our main results. We benchmark using TF-Lite for CPU, EdgeTPU and DSP, which relies on NNAPI to delegate computations to accelerators. For all benchmarks, we use single-thread and a batch size of 1. In Pixel 1 CPU, we use only a single large core. For Pixel 4 EdgeTPU and DSP, the models are fake-quantized [17] as required. The GPU models are optimized and benchmarked using TensorRT 7.1 converted from an intermediate ONNX format.

5.3. Search Space Definitions

The overall layout of our search space resembles that of ProxylessNAS [5] and TuNAS [2]. We consider three variants with increasing sizes:

- *IBN*: The smallest search space that contains IBN layers only. Each layer may choose from kernel sizes (3, 5) and expansion factors (4, 8).
- *IBN+Fused*: An enlarged search space that not only contains all the IBN variants above, but also Fused convolution layers in addition with searchable kernel sizes (3, 5) and expansion factors (4, 8).
- *IBN+Fused+Tucker*: A further enlarged search space that contains Tucker (compression) layers in addition. Each Tucker layer allows searchable input and output compression ratios within (0.25, 0.75).

For all search spaces variants above, we also search for the output channel sizes of each layer among the options of (0.5, 0.625, 0.75, 1.0, 1.25, 1.5, 2.0) times a *base channel size* (rounded to multiples of 8 to be more hardware-friendly). Layers in the same block share the same base

channel size, though they can end up with different expansion factors. The base channel sizes for all the blocks (from stem to head) are 32-16-32-48-96-96-160-192-192. The multipliers and base channel sizes are designed to approximately subsume several representative architectures in the literature, such as MobileNetV2 and MnasNet.

Hardware-Specific Adaptations. The aforementioned search spaces are slightly adapted depending on the target hardware. Specifically: **(a)** all building blocks are augmented with Squeeze-and-Excitation blocks and h-Swish activation functions (to replace ReLU-6) when targeting CPUs. This is necessary to obtain a fair comparison against the MobileNetV3+SSDLite baseline, which also includes these operations. Neither primitive is well supported on EdgeTPUs or DSPs; **(b)** when targeting DSPs, we exclude 5×5 convolutions from the search space, since they are highly inefficient due to hardware/software constraints.

5.4. Search Space Ablation

With a perfect architecture search algorithm, the largest search space is guaranteed to outperform the smaller ones because it subsumes the solutions of the latter. This is not necessarily the case in practice, however, as the algorithm may end up with sub-optimal solutions, especially when the search space is large [7]. In this work, a search space is considered useful if it enables NAS methods to identify *sufficiently good* architectures even if they are not optimal. In the following, we evaluate the usefulness of different search spaces by pairing them with TuNAS [2], a scalable latency-aware NAS implementation.

CPU. Figure 5 shows the NAS results for Pixel-1 CPUs. As expected, MobileNetV3+SSDLite is a strong baseline as the efficiency of its backbone has been heavily optimized for the same hardware platform over the classification task on ImageNet. We also note that the presence of regular convolutions does not offer clear advantages in this particular case, as IBN-only is already strong under FLOPS/CPU latency. Nevertheless, conducting domain-specific architecture search w.r.t. the object detection task offers non trivial gains on COCO (+1mAP in the range of 150-200ms).

EdgeTPU. Figure 6 shows the NAS results when targeting Pixel-4 EdgeTPUs. Conducting hardware-aware architecture search with any of the three search spaces significantly improves the overall quality. This is largely due to the fact that the baseline architecture (MobileNetV2)¹ is heavily optimized towards CPU latency, which is strongly correlated with FLOPS/MAdds but not well calibrated with the EdgeTPU latency. Notably, while IBN-only still offers the best accuracy-MAdds trade-off (middle plot), having regular convolutions in the search space (either IBN+Fused

¹MobileNetV3 is not well supported on EdgeTPUs due to h-swish and squeeze-and-excite blocks.




Figure 5: NAS results on Pixel-1 CPU using different search space variants.

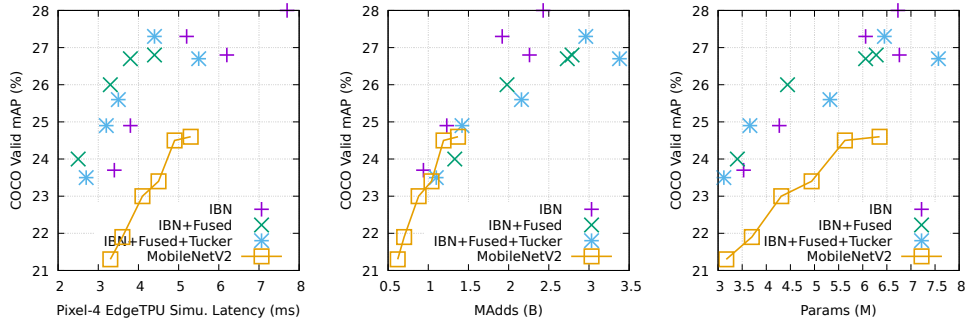


Figure 6: NAS results on Pixel-4 EdgeTPU using different search space variants.




Figure 7: NAS results on Pixel-4 DSP using different search space variants.

or IBN+Fused+Tucker) offers clear further advantages in terms of accuracy-latency trade-off. The results demonstrate the usefulness of full convolutions on EdgeTPUs.

DSP. Figure 7 shows the search results for Pixel-4 DSPs. Similar to EdgeTPUs, it is evident that the inclusion of regular convolutions in the search space leads to substantial mAP improvement under comparable inference latency.

5.5. Main Results

We compare our architectures obtained via latency-aware NAS against state-of-the-art mobile detection models on COCO [18]. For each target hardware platform, we report results obtained by searching over each of the three

search space variants. Results are presented in Table 1. On CPUs, we find that searching directly on a detection task allows us to improve mAP by 1.7mAP over MobileNetV3, but do not find evidence that our proposed fused IBNs and Tucker bottlenecks are required to obtain high quality models. On DSPs and EdgeTPUs, however, we find that our new primitives allow us to significantly improve the trade-off between model speed and accuracy.

On mobile CPUs, MobileDet outperforms MobileNetV3+SSDLite [13], a strong baseline based on the state-of-the-art image classification backbone, by 1.7 mAP at comparable latency. The result confirms the effectiveness of detection-specific NAS. The models also

Model/Search Space	Target hardware	mAP (%)		Latency (ms)			MAdds (B)	Params (M)
		Valid	Test	CPU	EdgeTPU	DSP		
MobileNetV2 ^{◇‡}		–	22.1	162	8.4	11.3	0.80	4.3
MobileNetV2 (ours) [◇]		22.2	21.8	129	6.5	9.2	0.62	3.17
MobileNetV2 × 1.5 (ours) [◇]		25.7	25.3	225	9.0	12.1	1.37	6.35
MobileNetV3 ^{‡‡}		–	22.0	119	*	*	0.51	3.22
MobileNetV3 (ours) [†]		22.2	21.8	108	*	*	0.46	4.03
MobileNetV3 × 1.2 (ours) [†]		23.6	23.1	138	*	*	0.65	5.58
MnasFPN (ours) [◁]		25.6	26.1	185	18.5	25.1	0.92	2.50
MnasFPN (ours) × 0.7 [◁]		24.3	23.8	120	16.4	23.4	0.53	1.29
IBN+Fused+Tucker [†]	CPU	24.2	23.7	122	*	*	0.51	3.85
IBN+Fused [†]		23.0	22.7	107	*	*	0.39	3.57
IBN [†]		23.9	23.4	113	*	*	0.45	4.21
IBN+Fused+Tucker	EdgeTPU	25.7	25.5	248	6.9	10.8	1.53	4.20
IBN+Fused		26.0	25.4	272	6.8	9.9	1.76	4.79
IBN		25.1	24.7	185	7.4	10.4	0.97	4.17
IBN+Fused+Tucker [◇]	DSP	28.9	28.5	420	8.6	12.3	2.82	7.16
IBN+Fused [◇]		29.1	28.5	469	8.6	11.9	3.22	9.15
IBN [◇]		27.3	26.9	259	8.7	12.2	1.43	4.85

Table 1: Test AP scores are based on COCO test-dev. †: Augmented with Squeeze-Excite and h-Swish (CPU-only); *: Not well supported by the hardware platform; ◇: 3×3 kernel size only (DSP-friendly); ◁: Augmented with NAS-FPN head; ‡: Endpoint C4 located after the 1×1 expansion in IBN.

achieved competitive results with MnasFPN, the state-of-the-art detector for mobile CPUs, without leveraging the NAS-FPN head which may complicate the deployment process. It is also interesting to note that the incorporation of full convolutions is quality-neutral over mobile CPUs, indicating that IBNs are indeed promising building blocks for this particular hardware platform.

On EdgeTPUs, MobileDet outperforms MobileNetV2+SSDLite by 3.7 mAP on COCO test-dev at comparable latency. We attribute the gains to both task-specific search (w.r.t. COCO and EdgeTPU) and the presence of full convolutions. Specifically, IBN+Fused+Tucker leads to 0.8 mAP improvement together with 7% latency reduction as compared to the IBN-only search space.

On DSPs, MobileDet achieves 28.5 mAP on COCO with 12.3 ms latency, outperforming MobileNetV2+SSDLite (×1.5) by 3.2 mAP at comparable latencies. The same model also outperforms MnasFPN by 2.4 mAP with more than 2× speed-up. Again it is worth noticing that including full convolutions in the search space clearly improved the architecture search results from 26.9 mAP @ 12.2 ms to 28.5 mAP @ 11.9ms.

These results support the fact that mobile search spaces have traditionally targeted CPU, which heavily favors separable convolutions, and that we need to rethink this decision for networks targeting other hardware accelerators.

Cross-Hardware Transferability of Searched Architectures. Fig. 8 compares MobileDets (obtained by targeting




Figure 8: Transferability of architectures (searched for different target platforms) across hardware platforms. For each given architecture, we report the original model and its scaled version with channel multiplier 1.5×.

at different accelerators) w.r.t. different hardware platforms. Our results indicate that architectures searched on EdgeTPUs and DSPs are mutually transferable. In fact, both searched architectures extensively leveraged regular convolutions. On the other hand, architectures specialized w.r.t. EdgeTPUs or DSPs (which tend to be FLOPS-intensive) do not transfer well to mobile CPUs.

Generalization to New Hardware. The proposed search space family was specifically developed for CPUs, EdgeTPUs and DSPs. It remains an interesting question whether the search space can extrapolate to new hardware. To answer this, we conduct architecture search for NVIDIA Jetson GPUs as a “holdout” device. We follow the DSP setup in Table 1, except that the number of filters are rounded to the multiples of 16. Results are reported in Table 2. The searched model within our expanded search space achieves +3.7mAP boost over MobileNetV2 while being faster. The

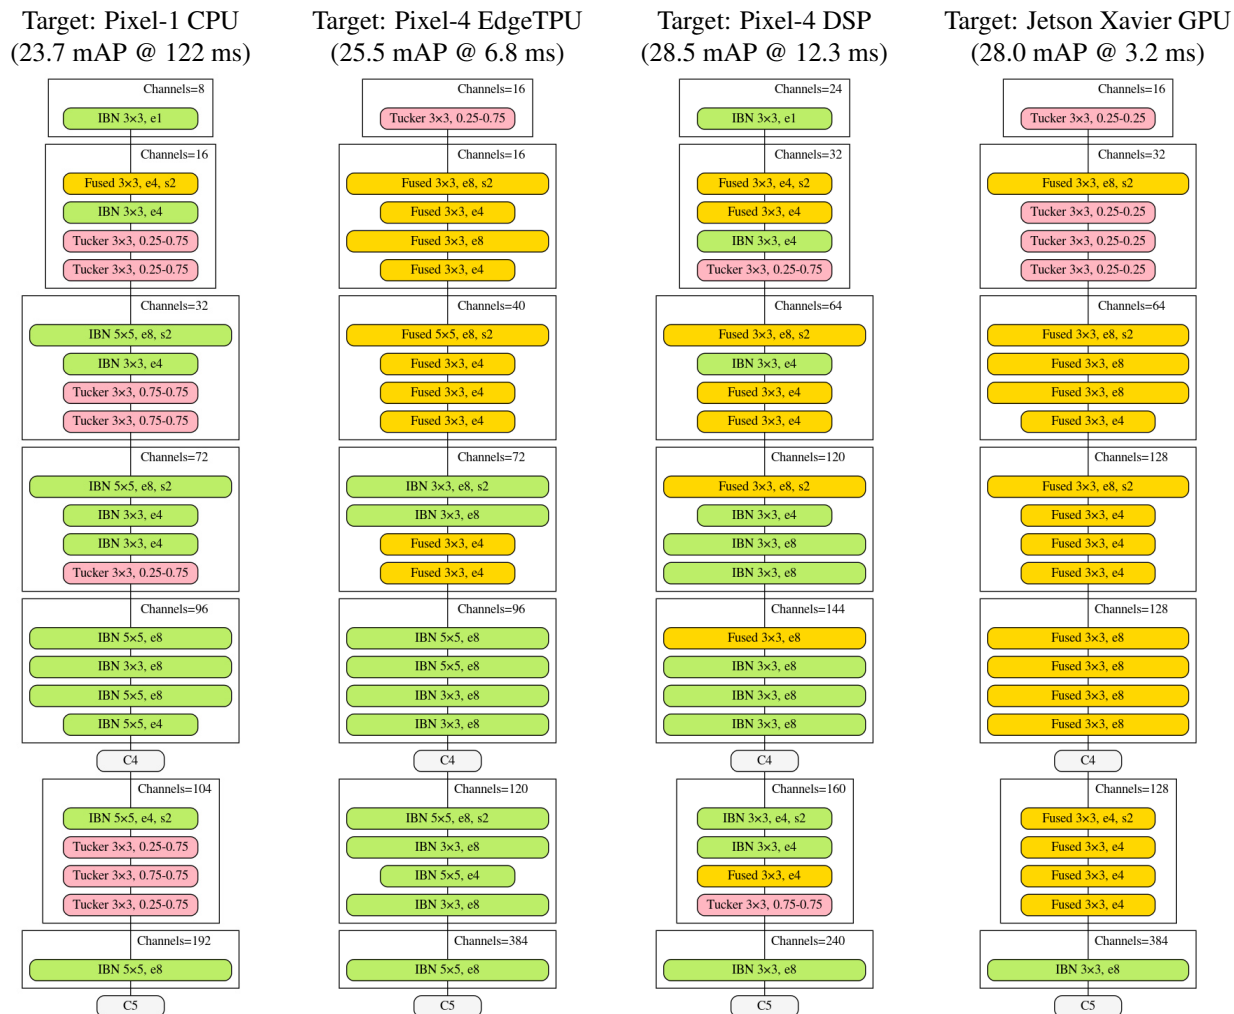


Figure 9: Best architectures searched in the IBN+Fused+Tucker space w.r.t. different mobile accelerators. Endpoints C4 and C5 are used by the SSDLite head. In the figures above, e refers to the expansion factor, s refers to the stride, and “Tucker $3 \times 3, 0.25-0.75$ ” refers to a Tucker layer with kernel size 3×3 , input compression ratio 0.25 and output compression ratio 0.75.

Model	mAP (%)		Jetson Xavier FP16 Latency (ms)
	Valid	Test	
MobileNetV2 (ours)	22.2	21.8	2.6
MobileNetV2 \times 1.5 (ours)	25.7	25.3	3.4
Searched (IBN+Fused+Tucker)	27.6	28.0	3.2

Table 2: Generalization of the MobileDet search space family to unseen hardware (edge GPU).

results confirm that the proposed search space family is generic enough to handle a new hardware platform.

Architecture visualizations. Fig. 9 illustrates our searched object detection architectures, MobileDets, by targeting at different mobile hardware platforms using our largest search space. One interesting observation is that MobileDets use regular convolutions extensively on EdgeTPU and DSP, especially in the early stage of the network where depthwise convolutions tend to be less efficient. The results

confirm that IBN-only search space is not optimal for these accelerators.

6. Conclusion

In this work, we question the predominant design pattern of using depthwise inverted bottlenecks as the only building block for vision models on the edge. Using the object detection task as a case study, we revisit the usefulness of regular convolutions over a variety of mobile accelerators, including mobile CPUs, EdgeTPUs, DSPs and edge GPUs. Our results reveal that full convolutions can substantially improve the accuracy-latency trade-off on several accelerators when placed at the right positions in the network, which can be efficiently identified via neural architecture search. The resulting architectures, MobileDets, achieve superior detection results over a variety of hardware platforms, significantly outperforming the prior art by a large margin.

References

- [1] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *International Conference on Machine Learning*, pages 550–559, 2018. 1, 2
- [2] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14323–14332, 2020. 4, 5
- [3] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Smash: one-shot model architecture search through hypernetworks. *arXiv preprint arXiv:1708.05344*, 2017. 2
- [4] Han Cai, Chuang Gan, and Song Han. Once for all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791*, 2019. 2
- [5] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019. 1, 2, 5
- [6] J Douglas Carroll and Jih-Jie Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of “eckart-young” decomposition. *Psychometrika*, 35(3):283–319, 1970. 2, 3
- [7] Bo Chen, Golnaz Ghiasi, Hanxiao Liu, Tsung-Yi Lin, Dmitry Kalenichenko, Hartwig Adams, and Quoc V Le. MnasFPN: Learning latency-aware pyramid architecture for object detection on mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2, 5
- [8] Yukang Chen, Tong Yang, Xiangyu Zhang, Gaofeng Meng, Xinyu Xiao, and Jian Sun. Detnas: Backbone search for object detection. In *Advances in Neural Information Processing Systems*, pages 6638–6648, 2019. 2
- [9] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *Advances in neural information processing systems*, pages 379–387, 2016. 2
- [10] Xiaoliang Dai, Peizhao Zhang, Bichen Wu, Hongxu Yin, Fei Sun, Yanghan Wang, Marat Dukhan, Yunqing Hu, Yiming Wu, Yangqing Jia, et al. Chamnet: Towards efficient network design through platform-aware model adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11398–11407, 2019. 1
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 4
- [12] 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 *Proceedings of the European Conference on Computer Vision*, pages 784–800, 2018. 2, 4
- [13] 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 *Proceedings of the IEEE International Conference on Computer Vision*, pages 1314–1324, 2019. 1, 2, 6
- [14] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 1, 4
- [15] 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. 1
- [16] Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, et al. Speed/accuracy trade-offs for modern convolutional object detectors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7310–7311, 2017. 1
- [17] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2704–2713, 2018. 5
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [19] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. 1, 2
- [20] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 2
- [21] Ronak Mehta, Rudrasis Chakraborty, Yunyang Xiong, and Vikas Singh. Scaling recurrent models via orthogonal approximations in tensor trains. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10571–10579, 2019. 2
- [22] Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. *arXiv preprint arXiv:1802.03268*, 2018. 2
- [23] Zheng Qin, Zeming Li, Zhaoning Zhang, Yiping Bao, Gang Yu, Yuxing Peng, and Jian Sun. Thundernet: Towards real-time generic object detection on mobile devices. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6718–6727, 2019. 2
- [24] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pages 4780–4789, 2019. 1
- [25] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016. 2
- [26] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region

- proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 2
- [27] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 1, 2
- [28] Laurent Sifre. Rigid-motion scattering for image classification. *Ph.D. thesis section 6.2*, 2014. 1, 4
- [29] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*, 2017. 1
- [30] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2820–2828, 2019. 1, 2, 4
- [31] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. In *The IEEE Conference on Computer Vision and Pattern Recognition*, June 2019. 2
- [32] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019. 1, 2
- [33] Ledyard R Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966. 2, 3
- [34] Ning Wang, Yang Gao, Hao Chen, Peng Wang, Zhi Tian, and Chunhua Shen. Nas-fcos: Fast neural architecture search for object detection. *arXiv preprint arXiv:1906.04423*, 2019. 2
- [35] Robert J Wang, Xiang Li, and Charles X Ling. Pelee: A real-time object detection system on mobile devices. In *Advances in Neural Information Processing Systems*, pages 1963–1972, 2018. 2
- [36] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992. 4
- [37] 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. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10734–10742, 2019. 2
- [38] Bichen Wu, Forrest Iandola, Peter H Jin, and Kurt Keutzer. Squeezedet: Unified, small, low power fully convolutional neural networks for real-time object detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 129–137, 2017. 2
- [39] Yunyang Xiong, Hyunwoo J Kim, and Varsha Hedau. Antnets: Mobile convolutional neural networks for resource efficient image classification. *arXiv preprint arXiv:1904.03775*, 2019. 1
- [40] Yunyang Xiong, Ronak Mehta, and Vikas Singh. Resource constrained neural network architecture search: Will a submodularity assumption help? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1901–1910, 2019. 1
- [41] Tien-Ju Yang, Andrew Howard, Bo Chen, Xiao Zhang, Alec Go, Mark Sandler, Vivienne Sze, and Hartwig Adam. Nektadap: Platform-aware neural network adaptation for mobile applications. In *Proceedings of the European Conference on Computer Vision*, pages 285–300, 2018. 2, 4
- [42] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6848–6856, 2018. 1
- [43] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. 1
- [44] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8697–8710, 2018. 1