

Searching for Efficient Neural Architectures for On-Device ML on Edge TPUs

Berkin Akin

bakin@google.com

Anton Spiridonov

tohaspiridonov@google.com

Hao Xu

iamxuhao@google.com

Suyog Gupta

suyoggupta@google.com

Zhuo Wang

zhuowang@google.com

Ping Zhou

zhouping@google.com

Yun Long

longy@google.com

Marie White

mariewhite@google.com

Yanqi Zhou

yanqiz@google.com

Abstract

On-device ML accelerators are becoming a standard in modern mobile system-on-chips (SoC). Neural architecture search (NAS) comes to the rescue for efficiently utilizing the high compute throughput offered by these accelerators. However, existing NAS frameworks have several practical limitations in scaling to multiple tasks and different target platforms. In this work, we provide a two-pronged approach to this challenge: (i) a NAS-enabling infrastructure that decouples model cost evaluation, search space design, and the NAS algorithm to rapidly target various on-device ML tasks, and (ii) search spaces crafted from group convolution based inverted bottleneck (IBN) variants that provide flexible quality/performance trade-offs on ML accelerators, complementing the existing full and depthwise convolution based IBNs. Using this approach we target a state-of-the-art mobile platform, Google Tensor SoC, and demonstrate neural architectures¹ that improve the quality-performance pareto frontier for various computer vision (classification, detection, segmentation) as well as natural language processing tasks.

1. Introduction

Due to the diminishing returns in the performance gains with the technology scaling in the post-Moore era, specialized ML accelerators became an essential component in most of the modern mobile system-on-chip (SoC) platforms to serve the needs of the real-time on-device ML workloads. Specialized ML accelerators (such as TPUs [2], NPUs [1, 18]) provide a substantial peak computation throughput, however the neural networks can extract optimal performance only when they are co-designed for the

¹<https://github.com/tensorflow/models/tree/master/official/projects/edgetpu/>

Model Performance on Google Tensor TPU




Figure 1. Proposed MobileNetEdgeTPUV2 models achieve higher ImageNet top-1 accuracy at lower latency when running on Google Tensor’s TPU [2] compared to MobileNetEdgeTPU, EfficientNet, FBNet, MobilenetV3. All models are quantized unless noted otherwise.

underlying hardware architecture.

There has been significant effort in hand crafting optimized neural architectures for specific target platforms [16, 24]. However, the increased complexity of the neural models and the variety of the target platforms gave rise to automated neural architecture search (NAS/AutoML) approaches [6–8, 22, 26, 34]. Although there are a variety of NAS frameworks, there are practical scalability limitations when it comes to designing neural architectures for different task domains and/or target platforms.

Either using NAS [14, 29, 31] or through manual design [24, 30], inverted bottleneck (IBN) layers have been predominant in building computer vision models. Although conventional IBNs that use depthwise convolutions have been very successful for mobile CPUs, prior work highlighted the use of full convolutions can significantly

improve the model’s accuracy-latency trade-off [14, 31]. Moreover, using full convolutions in IBNs allow *fusing* the pointwise expansion with the main full convolutions that can enable further latency optimizations on ML accelerators such as Edge TPUs [4, 14]. However, fused-IBNs can have substantially high computational and memory requirements for spatially narrow, channel-wise deep tensor shapes that are typical in the later stages of vision models, limiting their use throughout the model and leaving the depthwise-IBN as the only alternative.

We make an observation that a key factor in the high hardware efficiency of full convolutions on ML accelerators is the increased data reuse due to channel-wise convolutions. Depthwise separable convolutions remove the channel-wise convolution dimension which reduces the overall parameter and operation count but at the same time leading to extremely low hardware utilizations. We propose group convolution (GC) based IBNs, where the channel-wise convolution is still performed but limited to within each group. This allows GC-IBN variants to reach hardware utilization levels similar to the full convolution based IBNs but with much fewer parameter/operation counts.

Moreover, to address the practical limitations of NAS frameworks, we built a scalable infrastructure which decouples the search space design, platform cost evaluation (e.g. latency, energy) and the NAS algorithm. We provide cost evaluation as a gRPC [3] service where multi-trial or one-shot NAS clients can plug into, either for directly evaluating a search candidate (multi-trial) or to build learned cost models (one-shot).

As a concrete case study, we use the proposed infrastructure on the search spaces including the proposed GC-based IBNs and target the Edge TPU ML accelerator in the Google Tensor mobile SoC [2] for the on-device ML tasks identified by MLPerf Mobile Inference suite [5] which includes image classification, object detection, semantic segmentation and natural language processing. For each task, we demonstrate that the models designed by using our framework significantly improves the accuracy-performance pareto-frontier through the latency and energy measurements from Pixel 6 devices.

2. Related Work

Neural Architecture Search (NAS) was proposed to automate the design of neural network architectures, often aiming to improve model quality given a cost metric [6–8, 22, 26, 34]. Since evaluating a candidate model’s quality requires expensive training jobs in a *multi-trial* NAS [26], one-shot approaches with weight sharing in a super-network are proposed [6, 7, 22]. In this work we are not proposing a new NAS algorithm. Rather we are making an observation that various NAS methods and their implementations come with either algorithmic or practical limi-

tations/benefits. We build an infrastructure which can interface with various NAS clients and exercise it for rapid development targeting a state-of-the-art platform for multiple on-device ML tasks from different domains.

Inverted bottleneck blocks (IBN) have been used extensively in building computer vision models [16, 17, 24, 29–31]. Conventionally the use of depthwise separable convolutions along with separate pointwise expansion and projection has shown to be very effective for mobile CPUs [16, 17, 24]. Recent work also showed that using full convolutions where expansion and the $K \times K$ kernel is fused can be very efficient on ML accelerators [4, 14, 31].

Group convolutions (GC) were originally intended for model parallelism across GPUs in AlexNet [19], yet they were also used as part of the IBN blocks [29, 30, 32] to improve model quality. Recent FBNet [29] use GC in pointwise convolutions while keeping depthwise convolutions. ResNext [30] divides the ResNet bottleneck blocks into groups, while ShuffleNet also uses shuffle operations to add cross-group feature exchange. In this work, we propose flexible GC based IBN variants that use GC as the $K \times K$ kernel and optionally keep the pointwise full convolutions. We exploit the flexibility of GC to implement the expansion/projection as part of the $K \times K$ GC kernel and achieve fused GC IBNs similar to fused full convolution versions. We demonstrate that using GC IBNs opens up the search space between depthwise and full convolution based IBNs, and create unique opportunities for efficient execution on ML accelerators.

3. Neural Architecture Search Infrastructure

In this section, we introduce a scalable infrastructure we built to perform neural architecture search for optimizing various models on a dedicated ML accelerator (Edge TPU). There are two major challenges to address with this infrastructure. First, different from optimizing models for CPUs, performance and power metrics of a model are harder to predict directly from the number of operations/parameters on ML accelerators. With the software managed memory hierarchies of ML accelerators, achieved performance highly depend on how compiler maps the neural networks on the hardware. Therefore, we need a way to collect accurate performance and power evaluations (PPE) for guiding the search. Second, since we target diverse applications, the framework should unify the search space description and exploration flow and scale to different model domains.

3.1. Performance Power Evaluation (PPE) Service

Figure 2 shows the components of the PPE service. The server integrates the Edge TPU compiler, a cycle-accurate simulator, an analytical performance model for fast yet less accurate model simulation and a power estimator. Clients can send independent estimation requests to the server via




Figure 2. PPE service for model power/performance evaluation. Multiple clients can have many workers serviced by several server replicas.




Figure 3. Comparing on-device latency measurement with the PPE service results.

gRPC [3] interface for a candidate model. The server is scaled to thousands of machines, so several requests can be served at the same time to serve the needs for highly parallel NAS.

Figure 3 presents the correlation of on-device latency versus the latency from PPE service for randomly selected real use-case models. We observe that PPE latency is in general lower compared to the on-device latency, due to simulator’s optimistic assumptions on system resources such as DRAM bandwidth. However, there is a very strong linear correlation between the PPE results and the real-device measurements ($R^2 = 0.99$). This leads to correct relative ranking of the models which is the most critical in NAS.

3.2. NAS Integration

PPE service can be integrated with different NAS back-ends. In this work, we have utilized both one-shot and multi-trial NAS algorithms. For one-shot NAS, we used a weight sharing method based on TuNAS framework [6]. We have leveraged the one-shot NAS for classification and detection tasks. However, mostly due to practical challenges of searching for end-to-end models (backbone and

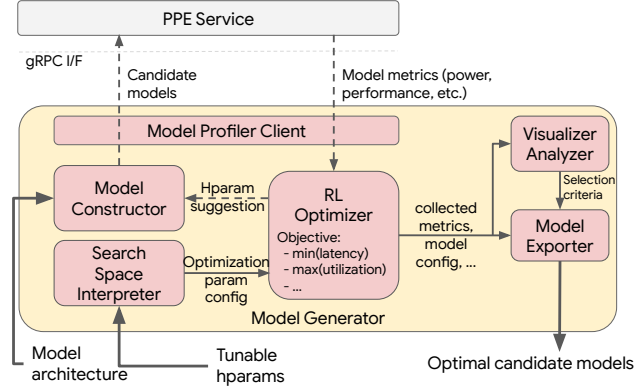


Figure 4. Model generator framework.

head) and compatibility of the neural operations (e.g. in transformer blocks) with weight sharing, we used multi-trial NAS for segmentation and NLP tasks. However, note that this is not a fundamental limitation but rather an implementation decision for rapid deployment targeting a state-of-the-art platform.

Model Generator. For the multi-trial approach, we have designed a model generator framework which is platform and task domain agnostic. Figure 4 is an overview of the framework. A user-friendly interface allows defining flexible search spaces by specifying model architecture topology and searchable parameters. A reinforcement learning (RL) back-end takes in searchable parameters and provides trial suggestions in iterations. The suggested candidate models are constructed and fed to PPE server as estimation requests. The PPE server responses with the model metrics are used to train the Vizier-based [13] RL agent for optimizing the back-end to make the next iteration of suggestions towards an optimization goal (latency, power, model size, etc.). A visualization and analysis tool is also integrated for assisting the selection of the best candidate models from the pool. The user can then export the selected models in preferred formats for further evaluation (e.g. training) and deployment. We use model generator as (i) a multi-trial NAS agent and (ii) an inverted bottleneck based neural block analyzer (see Section 4).

4. Neural Architecture Search Space

4.1. Inverted Bottlenecks (IBN)

Inverted bottleneck layers, commonly abbreviated as IBNs, have been a predominant building block in state-of-the-art computer vision models for mobile platforms [16, 24, 31]. The concept of a (inverted) bottleneck have also been extended to design of edge-device friendly NLP models [25]. As shown in Figure 5, a conventional IBN features a point-wise (1×1) convolution that expands the




Figure 5. IBN using depthwise convolution [24] (Depthwise-IBN).




Figure 6. IBN using full convolution for the fused expansion and main kernel [14, 31] (Fused-IBN).

input channel dimension to a larger value before applying a $K \times K$ depthwise convolution on the spatial dimensions. Finally, another point-wise convolution is used to project the expanded channel dimension to the desired final value.

IBNs are originally designed for mobile CPUs to reduce the overall operation count in FLOPS (floating-point operations), the number of trainable parameters and improve hardware efficiency. The separation of convolutions along the channel and spatial dimension serves this goal compared to performing a full convolution at the expanded channel dimension. However, also observed by prior work [14, 31], not all FLOPS have the same efficiency, especially on mobile ML accelerators, where a regular convolution may run $3\times$ as fast on Edge TPUs than a depthwise convolution even with $7\times$ as many FLOPS.

Motivated by this observation, Fused-IBN variants as shown in Figure 6 uses a regular full convolution instead of a separate pointwise expansion and a depthwise convolution kernel. Neural architecture search spaces augmented with the Fused-IBN were shown to improve model quality/latency trade-off for object detection [31] and image classification [14] tasks.

Although the Fused-IBN variant can provide an efficient alternative to Depthwise-IBN, we observed that Fused-IBN were primarily used in the early layers of the vision models where the channel dimension is relatively shallower. As the channel dimension gets deeper and the spatial dimensions get narrower, Fused-IBN uses a large amount of FLOPS and parameters which substantially increases the latency cost.

4.2. Group Convolution Based Inverted Bottlenecks

Group convolutions (GC) divide their input/output feature maps along the channel dimension into groups where




Figure 7. A $K \times K$ group convolution with g groups represented as a series of regular convolutions.

channel-wise convolutions are limited to within each group [19]. A group convolution operation can be represented with a series of full convolutions applied to the groups of input and output tensors as shown in Figure 7. GC can be considered as a generalized convolution representation such that when $g = 1$ a GC becomes a regular full convolution and when $Z_i = Z_o = g$ a GC degenerates into a depthwise convolution. Therefore, one can consider the number of groups in a GC as a *knob* to tune the number of parameters and operations of the convolution. This property of GC makes it a versatile tool that can be used in crafting IBN blocks. To this end, we propose GC based IBN variants to fill the gap in the neural architecture search spaces constructed solely from Depthwise and Fused IBNs.

A generalized form of GC-based IBN is provided in Figure 8. Firstly, GC can be used simply as a replacement of the depthwise convolution of a Depthwise-IBN to increase the total trainable parameters. However, in contrast to a depthwise convolution, GC does not constrain its input and output channel dimensions to be the same size. This allows performing a part of the channel expansion/projection using the pointwise convolutions and the remaining part by the GC kernel. For example, a total channel expansion of $m\times$, can be split into $n\times$ on pointwise convolution and $p\times$ on the GC such that $Z_{e'} = Z_i \times n$ and $Z_{o'} = Z_{e'} \times p$ where $m = n \times p$ (reverse can be applied to the projection side). Moreover, the entire expansion/projection can also be performed by the GC in which case the pointwise expansion/projection becomes ineffectual and can be eliminated (e.g. $n = 1$). This instance can be considered as a Fused-IBN where the $K \times K$ convolution is replaced with a GC (Figure 9). Due to this property we will refer to this special instance as a GC-IBN. GC-IBN provides advantages similar to the Fused-IBN as the pointwise expansion is fused into the GC kernel, yet it is more flexible thanks to the group count knob. Moreover, since the pointwise projection is kept, it provides a cross-group convolution. This allows us to avoid commonly used but hardware-unfriendly channel shuffle operations [29, 32]. Note that a dual of this block, where the projection is fused also exists. However, we did




Figure 8. A generalized IBN using group convolution as the main kernel. GC can implement part of the expansion/projection since there is no constraint such that $Z_{e'} = Z_o'$.




Figure 9. A special instance of Figure 8: IBN using GC for the fused pointwise expansion with the main kernel (GC-IBN). (A dual block also exists where the projection is fused.)

not include it in our search space due to its inferior performance on Edge TPUs.

4.3. Hardware Utilization Trade-offs

ML accelerator architectures commonly use wide single-instruction multiple-data (SIMD) execution units to extract the highest processing throughput. However, often times feeding these wide execution units from the memory system becomes the real bottleneck.

Depthwise convolutions require significantly lower number of parameters to mitigate the memory requirements. However, they fall short in utilizing the wide SIMD units of ML accelerators [14]. An overlooked key insight related to the low utilization is the lack of the activation operand reuse in depthwise convolutions. Every input feature map element fetched from the memory is only used once when computing the output feature maps in a depthwise convolution. This puts a heavy pressure on the activation fetch bandwidth requirements, and leads to low utilization of the compute units. We make the observation that in a group convolution operation, every input feature map element fetched from the memory is reused for computing the output feature maps within its group. This is significant since this means that we can amplify the activation operand data reuse by controlling the group size as needed by the SIMD width of the hardware while requiring fewer parameters than a full convolution.

To concretely demonstrate the computational characteristics of GC-IBNs, we leveraged the model generator as a *neural block analyzer* (Section 3) to generate neural nets solely based on IBN variants that can run on the Edge TPU accelerator. In Figure 10, first we observe that GC-

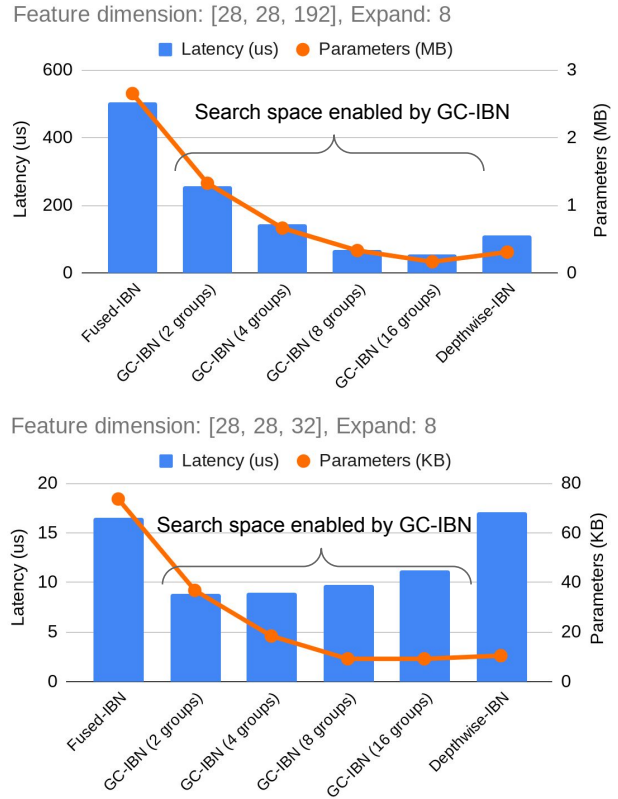


Figure 10. Executing different IBN variants for two different input sizes on Pixel 6 Tensor SoC. All IBNs use 3×3 kernel size and int8 data-type.

IBN blocks can provide $4 \times$ the trainable parameters and number of operations while having $0.5 \times$ the latency cost of Depthwise-IBNs. This indeed demonstrates the importance of the data reuse in reaching high hardware utilization. We also observe that GC-IBNs hardware utilization can be closer to the Fused-IBN blocks especially with smaller number of groups (hence larger group sizes). With smaller group sizes we start to lose data reuse and hit diminishing returns. Finally, we observe that the latency vs. trainable parameter count trade-offs are highly dependent on the tensor shapes and choosing the optimal IBN variant and its configuration (e.g., group size) is not a straight-forward task which calls for an automated exploration methodology using a neural architecture search (NAS).

With the inclusion of the proposed IBN variants the neural architecture search space becomes extremely large. Although choosing the optimal blocks that will maximize the model quality with a given latency target is a very difficult task to perform by hand, we observe that some block choices can be inherently sub-optimal for certain places in the neural network topology. For example, in the later stages of the neural network with the growth of the channel

dimension and the reduction of the spatial dimension, parameter reuse drops significantly and Fused-IBNs become much less efficient. We leverage the neural block analyzer to alleviate the search space size and filter out such choices by carefully analyzing the IBN’s performance characteristics.

5. Edge TPU Optimized Models

In this section, we use the proposed infrastructure on the search spaces including the proposed GC-based IBNs and target the Edge TPU ML accelerator in the Google Tensor mobile SoC [2] for the on-device ML tasks identified by MLPerf Mobile Inference suite [5] which includes image classification, object detection, semantic segmentation and natural language processing.

5.1. Image Classification

We start from a model topology similar to MobileNet/MobileDets due to their efficiency for mobile platforms including TPUs which is our primary target [4, 14, 31]. In the search space, we include the IBN variants described in Section 4 including Depthwise, Fused and GC-IBNs. We include residual skip connections over the IBN blocks with unit stride but omit them for the blocks that use stride > 1. We also omit swish non-linearity and the squeeze-and-excite blocks which are known to be less efficient on edge ML accelerators. As mentioned previously, we fine-tune the search space by filtering the IBN variants with consistently sub-optimal performance characteristics at certain blocks of the model topology instead of including all variants globally. Furthermore, considering the wide SIMD engines of ML accelerators, we pick a minimum group size of 32 and omit GC-IBNs with smaller group sizes (i.e. larger group counts) based on the neural block analyzer (Section 4).

Using this search space, we target the Edge TPU ML accelerator in the Google Pixel 6 Tensor SoC. We search for 5 different models with progressively increasing latency budgets which are named as Tiny, XS, S, M, L variants of MobileNetEdgeTPUV2. Accuracy vs. latency trade-offs provided by these models after post-training quantization to int8 datatype are provided in Figure 1 in comparison to other state-of-the-art (SOTA) mobile models. We observe that MobileNetEdgeTPUV2 model family outperforms even the prior SOTA MobileNetEdgeTPU models that are optimized for the Edge TPUs.

Our primary optimization target is the TPU accelerator, however our search space includes operations that also run well on mobile CPUs. Moreover, we implement GC using functionally equivalent series of commonly used ML primitives (slice, full convolution, concatenation) as shown in Figure 7, so that various platform compilers can efficiently support them since the native GC support may be missing.

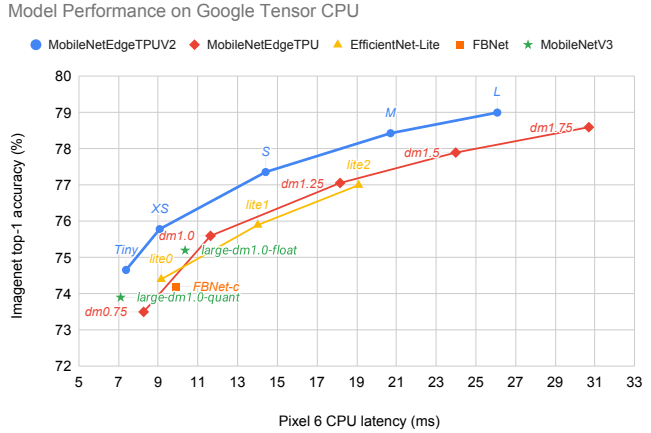


Figure 11. MobileNetEdgeTPUV2 models also demonstrate better accuracy-latency trade-off on Google Tensor CPU.

Also GC-IBNs tend to have fewer operations than Fused-IBNs and mobile CPUs show a stronger correlation between the number of operations in the neural network and latency. As a result, in Figure 11 we observe that when executed on the Google Tensor CPU, MobileNetEdgeTPUV2 family also outperform other SOTA models.

5.2. Semantic Segmentation

Many vision models consist of two components, the base feature extractor for understanding general features of the image, and the head for understanding domain-specific features, such as semantic segmentation. For feature extraction, we start from a model topology similar to EfficientNet [27] with IBN variants described in Section 4. We use a MobileNetEdgeTPUV2 classification model coupled with the DeepLabv3 [9] segmentation head as our baseline model and find that it improves the quality of on-device segmentation.

To further improve the segmentation model quality, we use the bidirectional feature pyramid network (BiFPN) [28] as the segmentation head, which performs weighted fusion of different features extracted by the feature extractor. Using NAS we find the optimal configuration of blocks in both the feature extractor and the BiFPN head. Specifically, we search for the kernel size from {3, 5, 7} for each IBN layer, and we also search for the expansion ratio from {3, 6} for each block except for the first one, which has the default expansion ratio of 1. In addition, we apply a channel multiplier that is among {1/2, 1/4, 1, 3/4, 2} to scale the model up and down. In the BiFPN head, we search over the number of repeats and minimum feature level, which produces trade-offs between accuracy and latency. The resulting models, named AutoSeg-EdgeTPU, produce even higher-quality segmentation results, while also running faster (Figure 12).

The final layers of the segmentation model contribute

Semantic segmentation on Pixel 6 Edge TPU

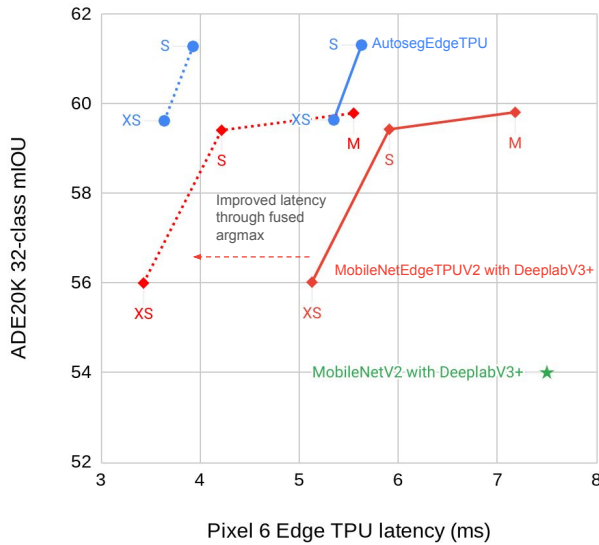


Figure 12. Segmentation model performance on Edge TPU.

significantly to the overall latency, mainly due to the operations involved in generating a high resolution segmentation map. To optimize the latency on TPU, we introduce an approximate method for generating the high resolution segmentation map that reduces the memory requirement and provides a nearly 1.5x speedup, without significantly impacting the segmentation quality.

5.3. Object Detection

Modern one stage object detection architectures typically produce one or more feature maps from the input image and use either an anchorless detection head such as the CenterNet [33] or anchor based detection head such as the SSD [23]. In the past, the classic process to design architectures for either of these two types of object detectors requires choosing a backbone network such as the MobileNets [16, 24] for low latency applications or ResNets [15] for accurate applications.

There are various methods to fuse together the feature maps from different endpoints of the backbone, such as FPN [20] which iteratively includes more low level information into the feature map as it upsamples the top feature map.

We notice that most of the classic object detection architectures allocate more than 70% of the total budget to the backbone area of the network while limiting the feature map fusion to less than 30%. We want to explore if rebalancing such allocation could lead to a better detection architecture. Also, recent NAS works such as MnasFPN [8] introduces a non-trivial connection pattern to fuse feature maps from

different endpoints in the backbone network. We want to utilize the success in such connection patterns as we design our object detection architecture.

With this in mind, we have created the Spaghetti Search Space, aptly named for the spaghetti-like connections between architecture blocks. For the COCO Object Detection Task [21], the search space consists of a stem node and 12 main blocks, each with the choice of between 2-4 layers. 6 blocks form the backbone whilst the other 6 form the head. Blocks in the head use the MnasFPN [8] connection pattern. Each layer may consist of depthwise separable convolutions [10], Inverted Bottleneck blocks (IBN) [24] and Grouped Convolution based IBNs (GC-IBN).

As seen in Figure 13, models found with this search space outperform MobileDet-EdgeTPU [31], current state-of-the-art detection models targeting Edge TPU platform. To verify the usefulness of GC-IBNs we remove GC-IBNs from the search space, and observe that the optimal models perform similarly to MobileDet-EdgeTPU, which demonstrates that the proposed search space provides more efficient options compared to the existing depthwise and full convolution based IBNs.

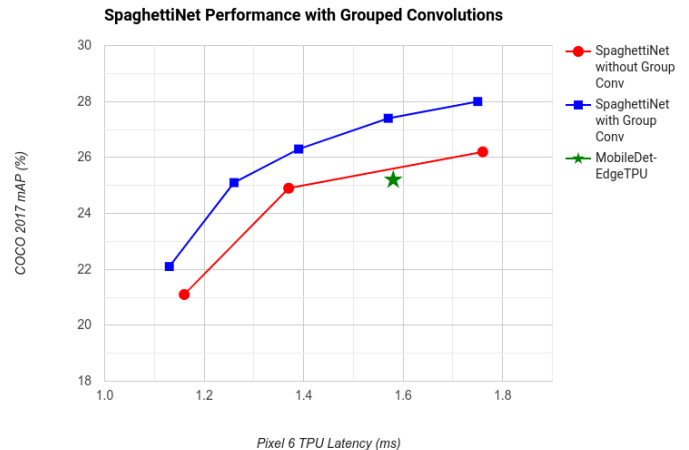


Figure 13. The performance of SpaghetNet models compared to MobileDet-EdgeTPU. When GC-IBN blocks are incorporated into the Spaghetti search space, it achieves +2.2% mAP more than MobileDet-EdgeTPU at the same latency.

5.4. Natural Language Processing

Deploying low-latency, high-quality transformer based language models on-device is highly desirable, and can potentially benefit multiple applications such as automatic speech recognition (ASR), translation, sentence auto-completion, and even some vision tasks [12]. While we mainly focused on vision tasks so far, the NAS infrastructure is domain agnostic, and can be easily extended to applications beyond vision, such as BERT [11] variant of language models.

Due to the limitations on weight sharing support in TuNAS for transformers, we simply use a multi-trial approach exploiting the flexibility of the proposed NAS infrastructure. Named as Mobilebert-EdgeTPU, we set up our NLP model architecture search space based on MobileBERT [25] and leverage the proposed NAS framework to find models with up to 2x better Edge TPU hardware utilization. With higher utilization, we are able to bring larger and more accurate models on chip, and meanwhile the models can still outperform the baseline MobileBERT latency. To complement the model generator for multi-trial NAS, we developed a customized knowledge distillation based training pipeline to quickly assess the generated model’s quality without full training during search. The final model is fully trained. As shown in figure 14, the quantized MobileBERT-EdgeTPU models establish a new pareto-frontier for the question answering tasks and also exceed the accuracy of the float $BERT_{base}$ [11] model, a 400MB+ model in float32 precision which is too large to run on edge devices.




Figure 14. Performance of MobileBERT-EdgeTPU models on the SQuAD v1.1 dataset.

As an alternative to the quant models, we also provide a set of Edge TPU friendly float models, as shown in figure 14. Notably, the float MobileBERT-EdgeTPU-M model yields accuracy that is even comparable to the $BERT_{large}$ [11], which has 1.3GB model size in float32 precision. Quantization now becomes an optional optimization rather than a prerequisite, which can greatly benefit use cases where quantization is infeasible or introduce large accuracy deterioration, and potentially reduce the time-to-market.

5.5. Energy Efficiency

As the energy consumption is critical for on-device ML use cases, we also setup an energy measurement harness and benchmark our models. Our benchmarking setup uses the nominal device settings and runs the models at 30 in-




Figure 15. On-device energy per inference measurements when running the models on Edge TPU.

ferences per second. We first measure the average power while the TPU is idle. Then, we subtract the idle power from the average power measured when the TPU is running the model for 100 inferences to find the active TPU power consumption. Finally, this power consumption rate is multiplied by the model latency to find the energy consumed per inference. Figure 15 demonstrates that the energy efficiency trends are similar to the latency measurements. This is expected since the efficient utilization of the hardware not only improves performance but also minimizes the use of inefficient operations and reduces the energy consumption. For the other targeted tasks we observe similar trends but for brevity we only report the large image classification results.

6. Conclusion

In this work we target optimizing various on-device ML tasks on edge ML accelerators. We propose flexible inverted-bottleneck (IBN) variants using group convolutions (GC) and design search spaces including these blocks. GC based IBNs opens up the search space between depth-wise and full convolution based IBNs, and create unique opportunities for efficient execution on ML accelerators. To easily find optimized models for various on-device ML tasks, we propose a scalable NAS-enabling infrastructure that decouples cost evaluation, neural search space design. Using this infrastructure with the proposed search spaces and targeting a state-of-the-art mobile SoC platform Google Tensor TPU, we demonstrate significant improvements in quality, latency and energy metrics for mobile ML tasks including computer vision (classification, detection, segmentation) and natural language processing (NLP).

References

- [1] Apple unleashes ml. <https://www.apple.com/newsroom/2020/11/apple-unleashes-ml/>. 1
- [2] Google tensor is a milestone for machine learning. <https://blog.google/products/pixel/introducing-google-tensor/>. Accessed: Oct 19, 2021. 1, 2, 6
- [3] A high performance, open source universal rpc framework. <https://grpc.io/>. 2, 3
- [4] Introducing the next generation of on-device vision models: Mobilenetv3 and mobilenetdedgetpu. <https://ai.googleblog.com/2019/11/introducing-next-generation-on-device.html>. 2, 6
- [5] Mlperf mobile inference benchmark. <https://mlcommons.org/en/inference-mobile-11/>. 2, 6
- [6] Gabriel Bender, Hanxiao Liu, Bo Chen, Grace Chu, Shuyang Cheng, Pieter-Jan Kindermans, and Quoc Le. Can weight sharing outperform random architecture search? an investigation with tunas. *CoRR*, abs/2008.06120, 2020. 1, 2, 3
- [7] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware, 2019. 1, 2
- [8] Bo Chen, Golnaz Ghiasi, Hanxiao Liu, Tsung-Yi Lin, Dmitry Kalenichenko, Hartwig Adam, and Quoc V. Le. Mnasfpn: Learning latency-aware pyramid architecture for object detection on mobile devices. *CoRR*, abs/1912.01106, 2019. 1, 2, 7
- [9] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation, 2017. 6
- [10] François Chollet. Xception: Deep learning with depthwise separable convolutions. *CoRR*, abs/1610.02357, 2016. 7
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 7, 8
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua 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. 7
- [13] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Elliot Karro, and D. Sculley, editors. *Google Vizier: A Service for Black-Box Optimization*, 2017. 3
- [14] Suyog Gupta and Berkin Akin. Accelerator-aware neural network design using automl, 2020. 1, 2, 4, 5, 6
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. 7
- [16] 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. 1, 2, 3, 7
- [17] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017. 2
- [18] Jun-Woo Jang, Sehwan Lee, Dongyoung Kim, Hyunsun Park, Ali Shafiee Ardestani, Yeongjae Choi, Channoh Kim, Yoojin Kim, Hyeongseok Yu, Hamzah Abdel-Aziz, Jun-Seok Park, Heonsoo Lee, Dongwoo Lee, Myeong Woo Kim, Hanwoong Jung, Heewoo Nam, Dongguen Lim, Seungwon Lee, Joon-Ho Song, Suknam Kwon, Joseph Hassoun, SukHwan Lim, and Changkyu Choi. Sparsity-aware and reconfigurable npu architecture for samsung flagship mobile soc. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, pages 15–28, 2021. 1
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Commun. ACM*, 60(6):84–90, May 2017. 2, 4
- [20] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. *CoRR*, abs/1612.03144, 2016. 7
- [21] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. 7
- [22] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search, 2019. 1, 2
- [23] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. *CoRR*, abs/1512.02325, 2015. 7
- [24] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation. *CoRR*, abs/1801.04381, 2018. 1, 2, 3, 4, 7
- [25] Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. Mobilebert: a compact task-agnostic bert for resource-limited devices. *arXiv preprint arXiv:2004.02984*, 2020. 3, 8
- [26] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile, 2019. 1, 2
- [27] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *ICML*, 2019. 6
- [28] Mingxing Tan, Ruoming Pang, and Quoc V. Le. Efficientdet: Scalable and efficient object detection, 2020. 6
- [29] 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, 2019. 1, 2, 4

- [30] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks, 2017. [1](#), [2](#)
- [31] Yunyang Xiong, Hanxiao Liu, Suyog Gupta, Berkin Akin, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Vikas Singh, and Bo Chen. Mobiledets: Searching for object detection architectures for mobile accelerators. *CoRR*, abs/2004.14525, 2020. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [32] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices, 2017. [2](#), [4](#)
- [33] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *CoRR*, abs/1904.07850, 2019. [7](#)
- [34] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition, 2018. [1](#), [2](#)