



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Constrained deep neural network architecture search for IoT devices accounting for hardware calibration

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Scheidegger, F., Benini, L., Bekas, C., Malossi, C. (2019). Constrained deep neural network architecture search for IoT devices accounting for hardware calibration. 10010 NORTH TORREY PINES RD, LA JOLLA, CALIFORNIA 92037 USA : NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).

Availability:

This version is available at: <https://hdl.handle.net/11585/767350> since: 2020-07-29

Published:

DOI: <http://doi.org/>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Florian Scheidegger, Luca Benini, Costas Bekas, A. Cristiano I. Malossi (2019). Constrained deep neural network architecture search for IoT devices accounting for hardware calibration. In H. Wallach et al., Advances in Neural Information Processing Systems 32, Curran Associates, Inc, pag. 6056—6066.

The final published version is available online at: <https://papers.nips.cc/paper/8838-constrained-deep-neural-network-architecture-search-for-iot-devices-accounting-for-hardware-calibration>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Constrained deep neural network architecture search for IoT devices accounting hardware calibration^{* †}

Florian Scheidegger^{1,2}, Luca Benini^{1,3}, Costas Bekas², Cristiano Malossi²

¹ ETH Zürich, Säumerstrasse 101, 8092 Zürich, Switzerland

² IBM Research - Zürich, Säumerstrasse 4, 8803 Rüschlikon, Switzerland

³ Università di Bologna, Via Zamboni 33, 40126 Bologna, Italy

Abstract

Deep neural networks achieve outstanding results in challenging image classification tasks. However, the design of network topologies is a complex task and the research community makes a constant effort in discovering top-accuracy topologies, either manually or employing expensive architecture searches. In this work, we propose a unique narrow-space architecture search that focuses on delivering low-cost and fast executing networks that respect strict memory and time requirements typical of Internet-of-Things (IoT) near-sensor computing platforms. Our approach provides solutions with classification latencies below 10ms running on a \$35 device with 1GB RAM and 5.6GFLOPS peak performance. The narrow-space search of floating-point models improves the accuracy on CIFAR10 of an established IoT model from 70.64% to 74.87% respecting the same memory constraints. We further improve the accuracy to 82.07% by including 16-bit half types and we obtain the best accuracy of 83.45% by extending the search with model optimized IEEE 754 reduced types. To the best of our knowledge, we are the first that empirically demonstrate on over 3000 trained models that running with reduced precision pushes the Pareto optimal front by a wide margin. Under a given memory constraint, accuracy is improved by over 7% points for half and over 1% points further for running with the best model individual format.

Introduction

With an increasing number of published methods, data, models, new available deep learning frameworks, and hype of special purpose hardware accelerators that become more commercially available, the design of an economical viable artificial intelligence system becomes a formidable challenge. The availability of large scale datasets with known ground truth (Deng et al. 2009; Stallkamp et al. 2011; Krizhevsky and Hinton 2009) and widespread commercial availability of increased computational performance, usually achieved with graphics processing units (GPUs), enables the

current growth of deep learning and explains the large interest and the emergence of new businesses. Smart homes (Li et al. 2019), smart grids (Fenza, Gallo, and Loia 2019) and smart cities (Gaber et al. 2019) trigger a natural demand for the Internet of Things (IoT), which are products designed around low cost, low energy consumption and fast reaction times due to the inherent constraints given by the final application that typically demand for autonomy with long battery lifetimes or fast real-time operation. Experts estimate a number of around 30 billion IoT devices by 2020 (Nordrum 2016) many of which serve applications that profit from artificial intelligence deployment.

In this context, we propose an automatic way to design deep learning models satisfying user-given constraints that are specially tailored to match typical IoT requirements, such as inference latency bounds. Additionally, our approach is designed in a modular manner that allows future adaptations and specializations for novel network topology extensions to different IoT devices and reduced precision arithmetic. In summary, our main contributions are the following:

- We propose an end-to-end approach to synthesize models that satisfy IoT application and HW constraints.
- We propose a narrow-space architecture search algorithm to leverage knowledge from large reference models to generate a family of small and efficient models.
- We evaluate reduced precision formats for over 3000 models.
- We isolate IoT device characteristics and demonstrate how our concepts operate with analytical network properties and map to final platform specific metrics.

The remainder of the paper is organized as follows. Section describes the related work, Section introduces the core design procedures, Section details and merges a full synthesis workflow, Section states and discusses the obtained results, and Section concludes all findings.

Related work

Automated architecture search potentially discovers better models (Miikkulainen et al. 2019; Xie and Yuille 2017;

^{*}IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. Other product and service names might be trademarks of IBM or other companies.

[†]Published as a conference paper at NeurIPS 2019

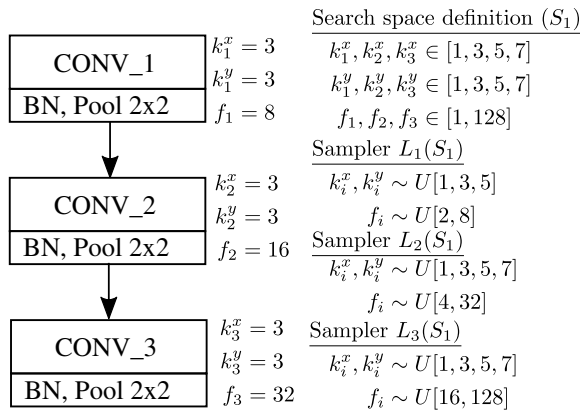


Figure 1: Simple three layer architecture with default configuration of search space with restricted sampling laws.

Zhong, Yan, and Liu 2017; Zoph and Le 2016; Zoph et al. 2018; Cai et al. 2018; Baker et al. 2016; Wistuba, Rawat, and Pedapati 2019). However, traditional approaches require a vast amount of computing resources or cause excessive execution times due to full training of candidate networks (Real et al. 2017). Early stopping based on learning curve predictors (Domhan, Springenberg, and Hutter 2015) or transferring learned weights improves the timings (Wistuba 2018). A method called train Train-less Accuracy Predictor for Architecture Search (TAPAS) demonstrates how to generalize architecture search results to new data without the need of training during the search process (Istrate et al. 2019). Architecture searches face the common challenge of defining the search space. Historically, it happened that new networks are independently developed by expert knowledge that outperform previously found networks generated by architectural search. In such cases, very expensive reconsiderations lead to follow up work to correctly account for a richer search space (Pham et al. 2018; Weng et al. 2019). Recent progress in the field, such as MnasNet (Tan et al. 2018) and FBNet (Wu et al. 2018) focus to tailor the search for smartphones by optimizing a multi-objective function including inference time. MnasNet trains a controller that adjusts to sample models that are more optimal according to the multi-objective. FBNet trains a supernet by a differentiable neural architecture search (DNAS) in a single step and claims to be $420\times$ faster since additional model training steps are avoided. In contrast to solving a joint optimization problem in one step, our proposed union of narrow-space searches follows a modular approach that separates the search process of finding architectures that strictly satisfy constraints from the training of candidate networks. That way, we can analyze ten-thousand architectures with zero training cost while only a small subset of suitable candidates are selected for training.

Compression, quantization and pruning techniques reduce heavy computational needs based on the inherent error resilience of deep neural networks (Rybalkin et al. 2017). Mobile nets (Howard et al. 2017a) or low-rank expansions (Jaderberg, Vedaldi, and Zisserman 2014) change the topol-




Figure 2: Statistics of number of parameters obtained when sampling up to one million networks from the base configuration space and when sampling 1000 networks from the restricted sampling laws.

ogy into layers that require fewer weights and cause reduced workload. Quantization studies the effect of using reduced precision floating point or fixed point formats (Hill et al. 2018; Loroeh et al. 2017), compression further tries to reduce the binary footprint of activation and weight maps (Cavigelli and Benini 2018), and pruning approaches avoid computation by enforcing sparsity (Ashiquzzaman et al. 2019). We use floatx, an IEEE 754 compliant reduced precision library (Flegar et al. 2019), to assess data format specific aspects of networks. The novelty of our work is that we jointly evaluate network topologies in combination with reduced precision.

Core design procedures

Architecture search

It's challenging to define a space S that produces enough variation and simultaneously reduces the probability of sampling suboptimal networks. We propose narrow-space architecture searches, where results are obtained over aggregation of n independent searches $S = \bigcup_{i=1}^n S_i$. Since a good search space should satisfy $S_r \subset S$ where $S_r = \{M_1, \dots, M_n\}$ is a set of reference models, we construct S by designing narrow-spaces that obey $M_i \in S_i$ in order to guarantee $S_r \subset S$. Instead of considering superpositions, we have specialized search spaces that produce simple sequence structure with residual bypass operations (ResNets (He et al. 2016a)) to even high fan-out and convergent structures such as they occur in the Inception module (Szegedy et al. 2015) or in DenseNets (Huang et al. 2017a). The aggregation allows extending results easily with a tailored narrow-space search for new reference architectures. Next, we define a set of distribution law configurations $L_1(S_i), \dots, L_k(S_i)$ that allow drawing samples




Figure 3: High correlations between the two analytical properties of network architectures.




Figure 4: The run time dependent latency is best correlated with the workload where different search space specific characteristics are present.

in a biased way such that models satisfy properties of interest. Figure 1 demonstrates with an example the advantages over a uniform distribution among valid networks. Consider a space of three-layer networks with allowed variations in kernel shapes in $\{1, 3, 5, 7\}$ and output channels in $[1, 128]$ leading to $|S| = 4^6 * 128^3 = 8.6 * 10^9$ network configurations.

Figure 2 shows the statistics over up to 10^6 samples compared against sampling only 1000 samples when using restricted samplers L_1, L_2 and L_3 . The restricted random laws enable to efficiently generate networks of interest in contrast to the uniform sampler that fails to deliver high sampling densities in certain regions. For example, only 132 out of 10^6 networks have less than 1000 parameters.

We define each narrow-space architecture search and its sampling laws according to the following design goals: first, the original model is included in the search space, second, only valid models are generated with a topology that resembles the original model, third, the main model-specific parameters are varied, fourth, the main way to generate small

and efficient models was achieved through lowering channel widths in convolutional layers, and fifth, all random laws follow a uniform distribution over available options where the lower and upper limits were used as way to bias the models to span several orders of magnitude targeting the range of parameter and flop counts that are relevant for IoT applications.

Precision analysis

The precision analysis evaluates model accuracies when models are running with reduced precision representations. To follow a general methodology, we perform the precision analysis on the backend device that has different execution capabilities than current or future targeted IoT devices. The methodology enforces to use emulated computation throughout the analysis to assess accuracy independent of the target hardware. Low precision can be applied to model parameters, to the computations performed by the models and to the activation maps that are passed between operators. In this work, we follow the extrinsic quantization approach (Loroch et al. 2017), where we enforce a precision caused by the reduced type $T_{w,t}$ of storage width $1 + w + t$ to be applied to all model parameters and all activation maps that are passed between operations. For the analysis, we follow the IEEE 754 standard (Zuras et al. 2008) that defines storage encoding, special cases (Nan, Inf), and rounding behavior of floating-point data. A sign s , an exponent e and the significand m represent a number $v = (-1)^s * 2^e * m$ where the exponent field width w and the trailing significant field width t limit dynamic range and precision. Types $T_{5,10}$ and $T_{8,23}$ correspond to standard formats *half* and *float*. Our experiments are based on a PyTorch (pyt) integration of the GPU quantization kernel based on the high performance floatx library (Flegar et al. 2019) that implements the type $T_{w,t}$. The fast realization of the precision analysis allows elaborating over 3'000 models with a full grid search of 214 types ($w \in [1, 8], t \in [1 - 23]$) on the full validation data.

Deployment and performance characterization

To evaluate model execution performance on the IoT target device we propose to perform a calibration to assess the execution speed of models of interest. Despite many choices of deep learning frameworks, ways of optimizing code depending on compilation or version of software and even several hardware platforms that accelerate deep learning models, we formulate the performance characterization general and as most decoupled from the topology architecture search and the precision analysis to ease later extensions. Performance measurements on the IoT device are affected by explicit and implicit settings. In this work, we demonstrate our search algorithm with performance measurements with the least amount of assumptions and requirements on the runtime. To that end, we selected a Raspberry-Pi 3(B+) as a representative IoT device. It features a Broadcom BCM2837B0, quad-core ARMv8 Cortex-A53 running at 1.4 GHz and the board is equipped with 1GB LPDDR2 memory (pi3). The Raspberry-Pi 3(B+) belongs to the general-purpose device category that is shipped with peripherals (WiFi, LAN, Bluetooth, and USB, HDMI), a full operating system (Raspbian,

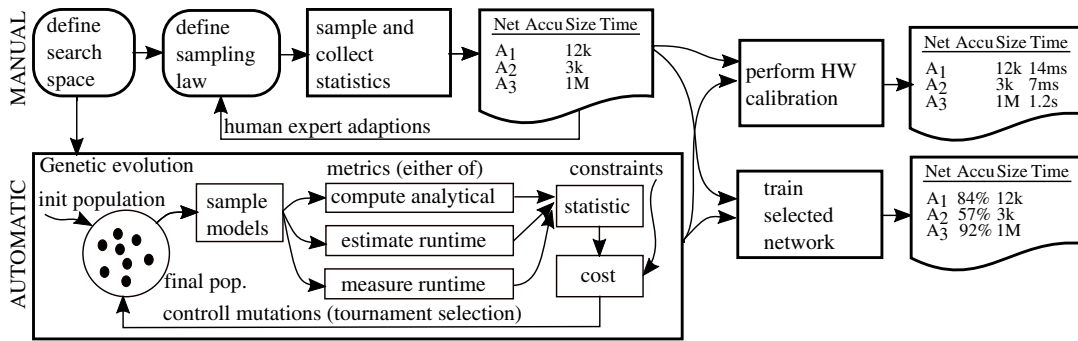


Figure 5: Manual and automatic workflow. First, sampling laws are defined to generate models of interest. Second, models are calibrated to check latency on the IoT device even if they are not yet trained. Third, models are trained to obtain their accuracy. Since training is the most expensive task, it is essential to reduce the amount of trained model to candidates of interest only.

a Linux distribution) available for a low cost of about \$35 per device (Mittal 2019). Throughout this work, we measure the model inference latency on the target device by averaging over 10 repetitions. We used a batch size of one to minimize latency and internal memory requirements. The latency study covers many relevant use cases, for example, the classification of sporadically arriving data in short time to prolong battery lifetime or frame processing of a video stream where the classification has to be completed before the next frame arrives.

For each model we consider two analytical properties, the number of trainable parameters and the workload measured as the number of floating-point operations required for inference. The calibration relates analytical properties to execution performance and allows to separate runtime metrics. Figure 3 and Figure 4 show high correlations between the number of parameters, the workload and the measured latency on the Raspberry-Pi 3(B+). Workload and parameters follow a similar scaling over five orders of magnitude with homogenous variations. The dynamic range of the latency spans more than two orders of magnitude with higher variations for larger models. However, due to the compute-bound nature of the kernels, the workload is the better latency time indicator than the number of parameters.

Fast cognitive design algorithms

In this section, we leverage the architecture search, the precision analysis, and the HW calibration to synthesize use case-specific solutions that satisfy given constraints. We address two tasks: First, the constraint search solves for the best model that satisfy given constraints. Second, the Pareto front elaboration provides insights into trade-offs over the full solution space. The two tasks are related. Solving the first task on a grid of constraints provides solutions to the second task while filtering the latter based on the given constraints allows returning to the former. Both tasks are solved in a manual and automated way as shown in Figure 5. In the manual task, the expert user defines the narrow-space search and for each space a list of sampling laws. Collected statistics over analytical network properties provide quick feedback to adapt the settings to cover the range of interest.




Figure 6: Manual defined sampling laws cover the full space.

Additionally, network run time metrics can be measured on the target device or estimated from calibration mea-

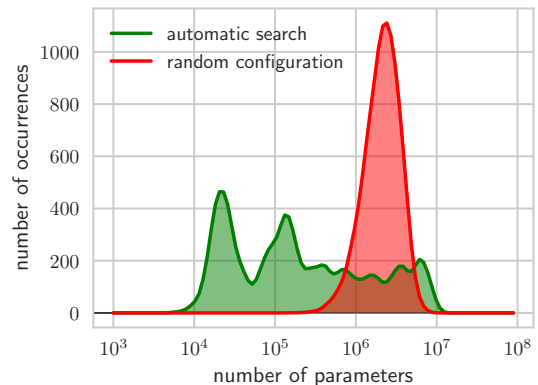


Figure 7: The automatic search finds configurations without human interaction and the distribution covers an higher dynamic range than just sampling uniformly.




Figure 8: Results of our architecture search compared against reference models. Each dot represents a model by its size and the obtained accuracy on the CIFAR-10 validation set. Our search finds results over five order of magnitudes and especially finds various models that are much smaller than out-of-the box available models. In the restricted IoT domain, our search delivers models that outperform the reference with a wide margin for fixed constraints.

surements. Next, depending on the task type, either a few candidate networks that satisfy constraints or a full wave of networks are selected for training. Large scale training takes the most time, each training job is of complexity $O(n_{train}C_{model}E)$, proportional to the amount of training data, the model complexity and the number of epochs the model is trained for.

We designed a genetic and clustering based algorithm to automatize the design of sampling laws. We define the valid space with a list of variables with absolute minimal and maximal ratings. A sampling law $L(S_i)$ is defined as an ordered set of uniform sampling laws $L = (U_x[l_x, h_x], \dots)$ with lower and upper limits l_x and h_x per variable x . The genetic algorithm automatically learns the search space specific sampling law limits $[l_x, h_x]$. The cost function is defined in a two step approach. First, the statistic $(\mu_m, \sigma_m) := E_m^n(L)$ is estimated by computing means and standard deviations over the metric m extracted from the n generated topologies. Second, cost is computed as $c((\mu_m, \sigma_m), (\tau_1, \tau_2)) := |\mu_m - \sigma_m - \tau_1| + |\mu_m + \sigma_m - \tau_2|$ in order that the high density range of the estimated distribution coincides with a given interval (τ_1, τ_2) . We avoided definitions based on single sided constraints like $\mu < \tau$ since such formulations might be either satisfied trivially (using the smallest network) or satisfied by undesirable laws having wide or narrow variations. We used the tournament selection variant of genetic algorithms (Goldberg and Deb 1991) and defined mutations by randomly adapting the sampling law of hyper-parameters l_x and h_x . We used an initial population of $n_{init} = 100$ and run the algorithm for

$n_{steps} = 900$ steps while using $n_{eval} = 10$ samples to estimate mean and standard deviation per configuration. This way, one search considers $(n_{init} + n_{steps}) * n_{eval} = 10'000$ networks. Since the final population might contain different sampling laws of similar quality, we perform spectral clustering (Stella and Shi 2003) to find $k = 10$ clusters with similar sampling laws. We assemble a list of the most different top- k laws by taking the best fitted law per cluster.

To elaborate the full search space with a Pareto optimal front, we split each decade into three intervals $[\tau, 2\tau, 5\tau, 10\tau]$ and define a grid for $\tau = 10^3, 10^4, 10^5, 10^6$ spanning five orders of magnitude.

We run the genetic search algorithm multiple times by setting the target bounds (τ_1, τ_2) in a sliding window manner over consecutive values from the defined grid. Finally, we accumulate results from 12 genetic searches each found 10 sampling laws, where we sampled each law $n_{val} = 100$ times to obtain the statistic of 12'000 network architectures per narrow-space search. Figure 6 and Figure 7 show results for manual and automatic sampled networks. Even though the manual search allows to nicely cover the region of interest, human expertise is required to correctly define the parameters of the laws L_1 up to L_6 . The naive sampling approach in the full search space produces a narrow distribution and is highly skewed towards larger networks. In contrast, the genetic algorithm was able to equalize the distribution and provides samples that cover much higher dynamic ranges, especially extending the scale for smaller networks without manually restricting the architecture.




Figure 9: Left: zoomed view of direct comparison, manual and automatic search perform equally well. Middle: manual and automatic search results. In the manual case clusters are visible while the automatic search was able to sample more homogeneously. Right: results for one narrow-space search with marked clusters matching Figure 6 and Figure 7.

Results

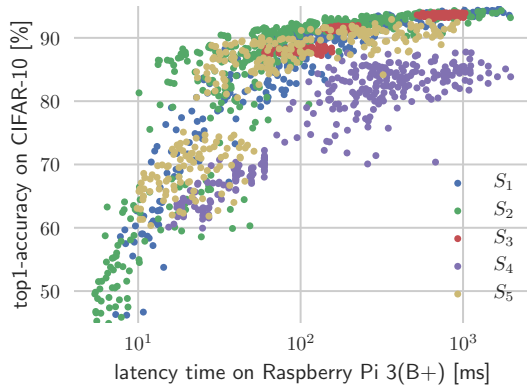


Figure 10: Final result showing the achievable tradeoffs between on the IoT device measured model latency and the model accuracy. Our search is able to deliver models that run below 10ms on the Raspberry Pi 3(B+) which we consider as representative cost limited IoT device.

To study our algorithm we run full design space explorations on the well established CIFAR-10 (Krizhevsky and Hinton 2009) classification task and compare our results with those obtained with established reference models. Figure 8 shows the trade-off between the model size and the obtained accuracy including manual and automatic generated results of the aggregate search spaces. The Pareto optimal front follows a smooth curve that saturates towards the best accuracy obtainable for large models. The number of parameters is logarithmic and the accuracy linearly scaled. Even very small models with less than 1000 parameters can achieve above 45% of accuracy. The accuracy increase per decade of added parameters is in the order of 30%, 15%, 3%

and $< 2\%$ points and diminishes very quickly. This effect allows constructing models that consist of multiple orders of magnitude fewer parameters and provides economical interesting solutions when IoT devices are powerful enough to process data in real-time. We compare our results with three sources of reference models, a) with traditional reference models, b) with ProbeNets (Scheidegger et al. 2019) that are designed to be small and fast and, c) with models that were designed and run on the parallel ultra-low power (PULP) platform (Conti et al. 2016). Traditional models include 30 reference topologies including variants of VGG (Simonyan and Zisserman 2014), ResNets (He et al. 2016b), GoogleNet (Szegedy et al. 2016), MobileNets (Howard et al. 2017b) dual path nets (DPNs) (Chen et al. 2017) and DenseNets (Huang et al. 2017b) where most of them (28/30) exceed 1M parameters. ProbeNets are originally introduced to characterize the classification difficulty and are by design considerably smaller (Scheidegger et al. 2019). They act as reference points for manual designed networks that cover the relevant lower tail in terms of parameters. In the IoT relevant domain ($< 10M$ parameters) our search outperforms all the listed reference models. The top three fronts in Figure 8 show the results of the precision analysis. For each trained model we evaluated the effect of running models with all configurations of type $T_{w,t}$ and we extract and plot the Pareto-optimal front. We considered three cases, running all models with half precision, running all models with the type T_{4_3} which is the best choice for types of 8-bit length and running each model with its individual best trade-off type $T_{w,t}$. We empirically demonstrate that running with reduced precision pushes the Pareto optimal front. Under a given memory constraint, accuracy improves by over 7% points for half and over 1% points further for running with the model individual format. Figure 9 shows details about manual and automatic searches both leading to very similar results. The right figure shows results obtained for one narrow-space search, where manually defined sampling laws lead to clusters. The




Figure 11: We demonstrate scalability of our approach by applying our search to three constraints on 13 datasets. Best models per dataset and constraint are connected with a line.

automatic search was able to homogeneously cover a similar range. Figure 10 shows inference times when the same set of models is executed on the Raspberry Pi 3(B+). Similarly, towards the small model end of the scale, given additional time for the latency results in dominant accuracy gains, however towards the traditional high accuracy domain, even slight accuracy improvements are only achieved with even more complex models that cause long evaluation times. Figure 11 demonstrates the scalability of our approach. We applied our search for three constraints $\tau = 10^3, 10^4, 10^5$ on thirteen datasets (Scheidegger et al. 2019) where we spend a training effort of ten architectures per dataset and constraint. The lines connect the best per constraint and dataset performing architectures.

Conclusion

We studied the solution of synthesizing deep neural networks that are eligible candidates to efficiently run on IoT devices. We propose a narrow-space search approach to quickly leverage knowledge from existing architectures that is modular enough to be further adapted to new design patterns. Manually and automatically designed sampling laws allows generating various models with the number of parameters covering multiple orders of magnitude. We demonstrate that reduced precision improves top1 accuracy by over 8% points for constraint weight memory in the IoT relevant domain. A strong correlation between model size and latency enables to create small enough models that provide superior inference response latencies below 10ms on a \$35 edge device.

Acknowledgments This work was funded by the the European Union’s H2020 research and innovation programme under grant agreement No 732631, project OPRECOMP.

References

[Ashiquzzaman et al. 2019] Ashiquzzaman, A.; Ma, L. V.; Kim, S.; Lee, D.; Um, T.; and Kim, J. 2019. Compacting deep neural net-

works for light weight iot scada based applications with node pruning. In *2019 International Conference on Artificial Intelligence in Information and Communication (ICAIC)*, 082–085.

[Baker et al. 2016] Baker, B.; Gupta, O.; Naik, N.; and Raskar, R. 2016. Designing neural network architectures using reinforcement learning. *CoRR* abs/1611.02167.

[Cai et al. 2018] Cai, H.; Chen, T.; Zhang, W.; Yu, Y.; and Wang, J. 2018. Efficient architecture search by network transformation. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

[Cavigelli and Benini 2018] Cavigelli, L., and Benini, L. 2018. Extended bit-plane compression for convolutional neural network accelerators. *CoRR* abs/1810.03979.

[Chen et al. 2017] Chen, Y.; Li, J.; Xiao, H.; Jin, X.; Yan, S.; and Feng, J. 2017. Dual path networks. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc. 4467–4475.

[Conti et al. 2016] Conti, F.; Rossi, D.; Pullini, A.; Loi, I.; and Benini, L. 2016. Pulp: A ultra-low power parallel accelerator for energy-efficient and flexible embedded vision. *Journal of Signal Processing Systems* 84(3):339–354.

[Deng et al. 2009] Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE CVPR*, 248–255.

[Domhan, Springenberg, and Hutter 2015] Domhan, T.; Springenberg, J. T.; and Hutter, F. 2015. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.

[Fenza, Gallo, and Loia 2019] Fenza, G.; Gallo, M.; and Loia, V. 2019. Drift-aware methodology for anomaly detection in smart grid. *IEEE Access* 7:9645–9657.

[Flegar et al. 2019] Flegar, G.; Scheidegger, F.; Novakovic, V.; Mariani, G.; Tomas, A.; Malossi, C.; and Quintana-Ortí, E. 2019. Float x: A c++-library for customized floating-point arithmetic. submitted.

[Gaber et al. 2019] Gaber, M. M.; Aneiba, A.; Basurra, S.; Batty, O.; Elmisery, A. M.; Kovalchuk, Y.; and Rehman, M. H. U. 2019. Internet of things and data mining: From applications to techniques

- and systems. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 9(3):e1292.
- [Goldberg and Deb 1991] Goldberg, D. E., and Deb, K. 1991. A comparative analysis of selection schemes used in genetic algorithms. In *Foundations of genetic algorithms*, volume 1. Elsevier. 69–93.
- [He et al. 2016a] He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016a. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [He et al. 2016b] He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016b. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- [Hill et al. 2018] Hill, P.; Zamirai, B.; Lu, S.; Chao, Y.; Laurenzano, M.; Samadi, M.; Papaefthymiou, M. C.; Mahlke, S. A.; Wensch, T. F.; Deng, J.; Tang, L.; and Mars, J. 2018. Rethinking numerical representations for deep neural networks. *CoRR* abs/1808.02513.
- [Howard et al. 2017a] Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017a. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR* abs/1704.04861.
- [Howard et al. 2017b] Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017b. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR* abs/1704.04861.
- [Huang et al. 2017a] Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017a. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4700–4708.
- [Huang et al. 2017b] Huang, G.; Liu, Z.; van der Maaten, L.; and Weinberger, K. Q. 2017b. Densely connected convolutional networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [Istrate et al. 2019] Istrate, R.; Scheidegger, F.; Mariani, G.; Nikolopoulos, D. S.; Bekas, C.; and Malossi, A. C. I. 2019. Tapas: Train-less accuracy predictor for architecture search.
- [Jaderberg, Vedaldi, and Zisserman 2014] Jaderberg, M.; Vedaldi, A.; and Zisserman, A. 2014. Speeding up convolutional neural networks with low rank expansions. *CoRR* abs/1405.3866.
- [Krizhevsky and Hinton 2009] Krizhevsky, A., and Hinton, G. 2009. Learning multiple layers of features from tiny images.
- [Li et al. 2019] Li, W.; Logenthiran, T.; Phan, V.-T.; and Woo, W. L. 2019. A novel smart energy theft system (sets) for iot based smart home. *IEEE Internet of Things Journal*.
- [Loroch et al. 2017] Loroch, D. M.; Pfreundt, F.-J.; Wehn, N.; and Keuper, J. 2017. Tensorquant: A simulation toolbox for deep neural network quantization. In *Proceedings of the Machine Learning on HPC Environments, MLHPC'17*, 1:1–1:8. New York, NY, USA: ACM.
- [Miikkulainen et al. 2019] Miikkulainen, R.; Liang, J.; Meyerson, E.; Rawal, A.; Fink, D.; Francon, O.; Raju, B.; Shahrzad, H.; Navruzyan, A.; Duffy, N.; and Hodjat, B. 2019. Chapter 15 - evolving deep neural networks. In Kozma, R.; Alippi, C.; Choe, Y.; and Morabito, F. C., eds., *Artificial Intelligence in the Age of Neural Networks and Brain Computing*. Academic Press. 293 – 312.
- [Mittal 2019] Mittal, S. 2019. A survey on optimized implementation of deep learning models on the nvidia jetson platform. *Journal of Systems Architecture*.
- [Nordrum 2016] Nordrum, A. 2016. The internet of fewer things [news]. *IEEE Spectrum* 53(10):12–13.
- [Pham et al. 2018] Pham, H.; Guan, M. Y.; Zoph, B.; Le, Q. V.; and Dean, J. 2018. Efficient neural architecture search via parameter sharing. *CoRR* abs/1802.03268.
- [pi3] Raspberry pi 3 model b+ product description. <https://www.raspberrypi.org/products/raspberry-pi-3-model-b-plus/>. Accessed: 2019-05-14.
- [pyt] Pytorch. <https://pytorch.org/>. Accessed: 2019-05-22.
- [Real et al. 2017] Real, E.; Moore, S.; Selle, A.; Saxena, S.; Sue-matsu, Y. L.; Tan, J.; Le, Q. V.; and Kurakin, A. 2017. Large-scale evolution of image classifiers. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2902–2911. JMLR. org.
- [Rybalkin et al. 2017] Rybalkin, V.; Wehn, N.; Yousefi, M. R.; and Stricker, D. 2017. Hardware architecture of bidirectional long short-term memory neural network for optical character recognition. In *Proceedings of the Conference on Design, Automation & Test in Europe*, 1394–1399. European Design and Automation Association.
- [Scheidegger et al. 2019] Scheidegger, F.; Istrate, R.; Mariani, G.; Benini, L.; Bekas, C.; and Malossi, C. 2019. Efficient image dataset classification difficulty estimation for predicting deep-learning accuracy. submitted.
- [Simonyan and Zisserman 2014] Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- [Stallkamp et al. 2011] Stallkamp, J.; Schlipsing, M.; Salmen, J.; and Igel, C. 2011. The german traffic sign recognition benchmark: A multi-class classification competition. In *The 2011 International Joint Conference on Neural Networks*, 1453–1460.
- [Stella and Shi 2003] Stella, X. Y., and Shi, J. 2003. Multiclass spectral clustering. In *null*, 313. IEEE.
- [Szegedy et al. 2015] Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9.
- [Szegedy et al. 2016] Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [Tan et al. 2018] Tan, M.; Chen, B.; Pang, R.; Vasudevan, V.; and Le, Q. V. 2018. Mnasnet: Platform-aware neural architecture search for mobile. *CoRR* abs/1807.11626.
- [Weng et al. 2019] Weng, Y.; Zhou, T.; Liu, L.; and Xia, C. 2019. Automatic convolutional neural architecture search for image classification under different scenes. *IEEE Access* 7:38495–38506.
- [Wistuba, Rawat, and Pedapati 2019] Wistuba, M.; Rawat, A.; and Pedapati, T. 2019. A survey on neural architecture search. *arXiv preprint arXiv:1905.01392*.
- [Wistuba 2018] Wistuba, M. 2018. Deep learning architecture search by neuro-cell-based evolution with function-preserving mutations. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 243–258. Springer.
- [Wu et al. 2018] Wu, B.; Dai, X.; Zhang, P.; Wang, Y.; Sun, F.; Wu, Y.; Tian, Y.; Vajda, P.; Jia, Y.; and Keutzer, K. 2018. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. *CoRR* abs/1812.03443.
- [Xie and Yuille 2017] Xie, L., and Yuille, A. 2017. Genetic cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, 1379–1388.

- [Zhong, Yan, and Liu 2017] Zhong, Z.; Yan, J.; and Liu, C. 2017. Practical network blocks design with q-learning. *CoRR* abs/1708.05552.
- [Zoph and Le 2016] Zoph, B., and Le, Q. V. 2016. Neural architecture search with reinforcement learning. *CoRR* abs/1611.01578.
- [Zoph et al. 2018] Zoph, B.; Vasudevan, V.; Shlens, J.; and Le, Q. V. 2018. Learning transferable architectures for scalable image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [Zuras et al. 2008] Zuras, D.; Cowlishaw, M.; Aiken, A.; Applegate, M.; Bailey, D.; Bass, S.; Bhandarkar, D.; Bhat, M.; Bindel, D.; Boldo, S.; et al. 2008. Ieee standard for floating-point arithmetic. *IEEE Std 754-2008* 1–70.