

# E-DNAS: Differentiable Neural Architecture Search for Embedded Systems

Javier García López  
FICOSA ADAS S.L.U  
08232 Barcelona, Spain  
Email: jgarcia@iri.upc.edu

Antonio Agudo and Francesc Moreno-Noguer  
Institut de Robòtica i Informàtica Industrial, CSIC UPC  
08028, Barcelona, Spain  
Email: {aagudo, fmoreno}@iri.upc.edu

**Abstract**—Designing optimal and light weight networks to fit in resource-limited platforms like mobiles, DSPs or GPUs is a challenging problem with a wide range of interesting applications, *e.g.* in embedded systems for autonomous driving. While most approaches are based on manual hyperparameter tuning, there exist a new line of research, the so-called NAS (Neural Architecture Search) methods, that aim to optimize several metrics during the design process, including memory requirements of the network, number of FLOPs, number of MACs (Multiply-ACcumulate operations) or inference latency. However, while NAS methods have shown very promising results, they are still significantly time and cost consuming.

In this work we introduce E-DNAS, a differentiable architecture search method, which improves the efficiency of NAS methods in designing light-weight networks for the task of image classification. Concretely, E-DNAS computes, in a differentiable manner, the optimal size of a number of meta-kernels that capture patterns of the input data at different resolutions. We also leverage on the additive property of convolution operations to merge several kernels with different compatible sizes into a single one, reducing thus the number of operations and the time required to estimate the optimal configuration. We evaluate our approach on several datasets to perform classification. We report results in terms of the SoC (System on Chips) metric, typically used in the Texas Instruments TDA2x families for autonomous driving applications. The results show that our approach allows designing low latency architectures significantly faster than state-of-the-art.

**Index Terms**—Deep Learning, Neural Architecture Search, Convolutional Meta Kernels.

## I. INTRODUCTION

Designing light Deep Neural Networks (DNNs) and doing it in an efficient manner, are two of the main challenges faced in industries like the automotive, which typically need to deal with resource-constrained platforms. This has been addressed in recent works, like SqueezeNet [1] or MNet [2], focused on optimizing the design of neural networks to alleviate their computational cost without losing performance. Most these studies, however, are based on the optimization of "indirect metrics", such as the number of Multiply-ACcumulate operations (MACs) or the number of architecture parameters, which might not be good approximations to the "direct metrics" like energy consumption or latency. As discussed in [3, 4], the relationship between these direct and indirect metrics can be highly non-linear and platform-dependent. Another drawback of [1] and [2] is that they require manual approaches and

prior expertise, limiting thus their applicability and design efficiency.

The design method has been automatized by the so-called Neural Architecture Search (NAS) [5, 6, 7] approaches. These techniques aim to automatically design light and accurate DNNs by optimizing over a search space defined by all possible operations of the target architecture. This optimization is carried on using either reinforcement learning [5, 6] or evolutionary computing [7].

While NAS-based approaches provide state-of-the-art results in classification tasks for small datasets like CIFAR, they are very computationally and time demanding. There have been attempts to speed up the search process using weight prediction techniques or weight sharing across multiple architectures [8]. Unfortunately, the improvement is still far from providing solutions that can scale to large datasets like ImageNet due to the prohibitive time and resources required.

In this paper we introduce E-DNAS, a differentiable NAS approach that optimizes the direct metrics of an embedded platform, yielding accurate and low-latency DNNs that can be deployed in memory-constrained platforms. The presented research builds upon three main ideas. First, we apply a depth-aware convolution over the input image to compute high-resolution feature maps. Second, we propose a parallel architecture search pipeline that operates on these feature maps and learns the optimal size and parameters of the convolution kernels. This optimization process is ruled by a multi-objective differentiable loss function that combines classification accuracy and minimal latency, a direct metric. And third, we boost the architecture search velocity through a novel block that connects the learned meta-kernels during training. This block is shown in Figure 1 and aims to update the learned meta-kernel (from *feature map 1*) on each iteration with the result of the weighted sum of that kernel and a second one being learned in parallel (the one from *feature map 2*). We show that this training information exchange on each iteration speeds up the search for the optimal kernels.

We demonstrate remarkable results in terms of search-time and classification accuracy compared to other state-of-the-art NAS methods and comparable to other recent breakthroughs like [9] or [10], which are more oriented for mobile devices rather than to be integrated into embedded systems, that typically have less flexible architectures.




FIG. 1: **General overview of E-DNAS.** Our approach has two main building blocks: a depth-aware convolution with a high resolution  $11 \times 11$  kernel followed by pairwise learning of meta-kernels with loopy flow of information on each iteration between training paths.

## II. RELATED WORK

Deep Learning is revolutionizing many technological areas, but it has some important constraints or limitations that need to be overcome in order to obtain its full potential. Some of these are the large amount of hardware resources (memory e.g.) needed to run some deep learning applications and also the manual network and parameter configuration traditionally done by experts to obtain an optimal DNN for a particular application.

In this direction, some works have focused on reducing the network size and optimizing its hyperparameter by pruning weights from DNNs, like [11] or [12]. Although these approaches could be effective for some applications, manually selecting the redundant weights and using unstructured sparse filters does not necessarily mean a real advantage in real platforms. Based on a similar idea, some recently published papers propose a method to design networks that can evolve during the design process based on some feedback in order to obtain the optimal number and type of layers for a specific application. These so called NAS (neural architecture search) approaches have recently achieved better performance than hand-crafted models by automating the architecture design.

Some NAS approaches like [5, 6, 13] apply the concept of reinforcement learning for finding the best neural architecture. These approaches propose a framework with a recurrent neural network (RNN) as a controller from which child architectures will be extracted and trained to get their accuracy. Based on this accuracy, the reward signal for the controller will be calculated and fed back to it, so that on next iteration the controller will give higher probabilities to architectures that receive higher accuracies (controller learns to improve its search over time), [6], [5]. Although this reward-based approaches showed really good results in providing efficient network architectures to be executed on mobile platforms, it still had one big disadvantage, which is the extremely long training time needed (e.g. [5] requires 2000 GPU days in the

ImageNet or CIFAR-10 dataset or approach proposed in [7] takes 3150 GPU days).

Lately, a faster version of the NAS has appeared, which can get to an optimal network design quicker by using gradient-based optimizations, like DARTS[9]. The differentiable neural architectural search (DNAS) propose to relax the search space to be continuous so that the architecture can be optimized with respect to its validation set through gradient descent. These techniques achieve a big efficiency improvement reducing drastically the cost of architecture finding in comparison to the non-differentiable approaches (NAS).

Although methods like DARTS have given good results in terms of accuracy and searching time compared to NAS, they still face some weak points, such as the still relatively long time needed for the architecture finding. Together with this, DNAS approaches such as DARTS [9] have proofed not to be practical to be used in large datasets. State-of-the-art works such as [10] or [14] have also exploited a similar approach as the one proposed in this work, making use of the additive property of the convolution to merge the searched operations and reduce the number of parameters in the DNN architecture. The main contribution of this work is the reduction of the search time through the self-designed *feedback block* defined in Section III-A2. Moreover, this work extends the DNAS method not only to general-purpose computing platforms like mobile devices but also to embedded platforms such as DSP, which, as above commented, are more restrictive and less flexible and are designed for single pre-defined functions.

## III. METHOD

In this work we propose a methodology for automatic neural architecture design to be executed on embedded platforms. We demonstrate state-of-the-art results on feature extraction and object detection tasks, as presented in Table I. We present a DNAS approach that aims to find optimal neural architecture with low latency to be executed on memory-constrained system on chips (SoCs), such as the one used for the experiments in this work.

The presented pipeline has two main steps:

- High resolution feature extraction through depthwise convolution using big dimension convolutional kernels.
- Pairwise neural architecture cross-search for the calculated feature maps on previous step.

In this work we regard network MACs and FLOPs as the proxy of the computation consumption.

### A. Formulation

1) *Convolutional filters:* As it was demonstrated in AlexNet [15], each convolutional kernel is responsible to capture a local image pattern. The larger the convolutional kernel is, the higher resolution patterns it tends to detect at the cost of more parameters and computations.

In particular, there is an important idea proposed in Mix-Conv work [14] that we exploit here and consists in having multiple kernels with different sizes in a single convolution operation to allow the network to capture different types of

features from the input images. Based on this, we present a two-step pipeline in which: first a large convolutional kernel is applied over the input image to capture high resolution patterns, and second several learnable kernels with different sizes are applied on the calculated feature maps to learn different of patterns on the input data.

In order to reduce the number of operations and, hence, the resulting network size, there are two considerations that shall be taken into account:

- The first step proposed in this paper suggests a separable *depthwise* convolution with a  $11 \times 11$  kernel applied on the input image that leads to a reduced parameter size and computational cost, compared to the traditional convolution operation, [2, 14, 16, 17].
- The filters applied on resulting feature maps after a first  $11 \times 11$  convolution are a sum of  $3 \times 3$ ,  $5 \times 5$  and  $7 \times 7$  filters. This work exploits the additivity property of convolution: if several 2D kernels with compatible sizes operate on the same input with the same stride to produce outputs of the same resolution, and their outputs are summed up, these kernels are finally added on the corresponding position to obtain the equivalent filter which will produce the same output, [18].

The main difference between the traditional convolution operation and the mentioned separable depthwise convolution over an input image (or tensor) is the number of steps in which this operation is applied.

In this context, the additivity property is applicable because the sizes of the filter or kernels are compatible, which means, smaller ones can be "contained" in bigger ones (with same center). That is:

$$\mathbf{I} * \mathbf{K}_1 + \mathbf{I} * \mathbf{K}_2 = \mathbf{I} * (\mathbf{K}_1 \oplus \mathbf{K}_2), \quad (1)$$

where  $\mathbf{I}$  is the input feature map,  $\mathbf{K}_1$  and  $\mathbf{K}_2$  are two 2D kernels with compatible sizes, and  $\oplus$  is the element-wise addition of the kernel or filter parameters on the corresponding positions, [10, 18].

The application of the additivity property is also valid for the following batch normalization (BN) so that each single BN applied after each convolution from Eq. (1) produces the same output, as the summation of each single convolution and BN with added bias, [18]:

$$\mathbf{O} = \mathbf{I} * \left( \frac{\gamma_1}{\sigma_1} \mathbf{K}_{3 \times 3} \oplus \frac{\gamma_2}{\sigma_2} \mathbf{K}_{3 \times 1} \oplus \frac{\gamma_3}{\sigma_3} \mathbf{K}_{1 \times 3} \right) + b, \quad (2)$$

where  $\mathbf{O}$  represents the output feature map,  $\mathbf{I}$  is the input data or feature map generated by the previous layer,  $\sigma$  is batch standard deviation and  $\gamma$  and  $b$  are the BN parameters to be learned. The input  $\mathbf{I}$  may need to be appropriately padded depending on the resolutions.

2) *Feedback-block*: One of the contributions of this paper is the addition of one *feedback-block* into the training pipeline of each feature map to update the learned convolutional filters or kernels on each iteration, see Figure 2. The implementation




FIG. 2: Example of a summed convolutional kernel (E), resulting of summing  $1 \times 1$  kernel (A),  $3 \times 3$  (B),  $5 \times 5$  (C) and  $7 \times 7$  kernel (D).

of this *feedback-block* is just the weighted sum of the learned meta-kernels being trained in parallel:

$$\mathbf{K}'_1 = \mathbf{K}'_2 = \beta_1 * \mathbf{K}_1 + \beta_2 * \mathbf{K}_2, \quad (3)$$

where  $\mathbf{K}'_1$  and  $\mathbf{K}'_2$  are the two meta-kernels candidates being learned,  $\mathbf{K}_1$  and  $\mathbf{K}_2$  are the kernels before the update and  $\beta_1$  and  $\beta_2$  are the weights for these kernels. These weights are calculated according to the loss calculated on each training "path".

We next compute the weights, in such a way that they will be close to one for small losses in the forward pass:

$$\begin{aligned} \beta_1 &= \frac{\tanh \frac{1}{L_1}}{\tanh \frac{1}{L_1} + \tanh \frac{1}{L_2}} \\ \beta_2 &= \frac{\tanh \frac{1}{L_2}}{\tanh \frac{1}{L_1} + \tanh \frac{1}{L_2}} \\ &\text{with } \beta_1 + \beta_2 = 1. \end{aligned} \quad (4)$$

where  $L_1$  and  $L_2$  represent the value of the loss function on two parallel network candidates being searched (illustrated in Fig. 1). Through this implementation, on each iteration the closer kernel to the "expected" one has more influence.

After training of each image, all learned kernels are encoded into one, following a similar approach as detailed in Eq. (4):

$$\begin{aligned} \mathbf{K} &= \sum_{j=1}^{j=N} \Gamma_j * \mathbf{K}_j \\ &\text{with } \Gamma_j = \frac{\beta_j}{\sum_i \beta_i}. \end{aligned} \quad (5)$$

## B. Search space

Following [9], we define the search space of each output  $x^{(i)}$  (e.g. feature map in convolutional networks) as the combination of operations  $o^{(i,j)}$  applied on inputs  $x^{(i)}$ , assuming the inputs as the outputs of the previous two layers:

$$x^{(j)} = \sum_{i < j} o^{(i,j)}(x^{(i)}). \quad (6)$$

The gradient-based NAS methodologies ([9], [19]) relax the categorical choice to a softmax to make it continuous. Let  $\mathcal{O}$  be a set of candidate operations (e.g. max pooling, convolutions) where each operation represents some function

$o(-)$  to be applied on the input  $x^{(i)}$ , a particular operation can be represented as:

$$\bar{o}^{(ij)}(x) = \sum_{o \in O} \frac{\exp(\alpha_o^{(ij)})}{\sum_{o' \in O} \exp(\alpha_{o'}^{(ij)})} o(x). \quad (7)$$

As demonstrated in [9], the task of architecture search reduces to learning a set of continuous variables  $\alpha = \{\alpha^{(ij)}\}$ .

The relaxation of the categorical choice presented in Eq. (7) can also be defined as follows, using the additivity property of convolutions. Based on this, each operation can be calculated as  $I * K$ , where  $I$  is the input of the operation and  $K^{(i)}$  is the kernel to be learned:

$$o(x) = I * K^{(i)}. \quad (8)$$

$$\bar{o}^{(ij)}(x) = \sum_{o \in O} \frac{\exp(\alpha_o^{(ij)})}{\sum_{o' \in O} \exp(\alpha_{o'}^{(ij)})} (I * K^{(i)}). \quad (9)$$

After relaxation of the search space, the proposed search network algorithm aims to learn jointly the architecture  $\alpha$  and the weights  $w$ .

### C. Multi-objective loss function

Based on the equation (9), the goal of the presented DNAS approach shall be calculating the weights  $w$  that minimize the validation loss:

$$\min_{\alpha} = L_{val}(w(\alpha), \alpha). \quad (10)$$

These weights  $w$  are learned during forward- and backward-pass and represent typical connection weights but include also the  $\beta$  values defined in Eq.3 and Eq. 4, which aim to update the kernel's value after each training iteration to speed up the search stage. In order to let this method generate models adaptively depending on the target embedded platform, we propose to include one more term to the global loss function to be minimized during training that attends to the latency of the network candidate.

The proposed loss function has a term to observe the latency of the proposed architecture. As demonstrated in different works such as [3] or [20], extracting indirect metrics like MACs or number of weights might not be good proxies for the resource consumption of a network since networks with fewer number of MACs can be slower when executed on embedded targets:

$$L_{LAT}(\alpha) = \sum_l LAT(b_l)^{(\alpha)}. \quad (11)$$

where  $b_l^{(\alpha)}$  denotes the block at layer- $l$  from the network architecture candidate  $\alpha$ , [20].

In the proposed implementation, we use a latency look up table model to estimate the global latency of the network candidate based on the runtime of each operator, similar as proposed in [20]. The latency lookup table has been created by checking the runtime of multiple operators on the target platform.

### D. The search algorithm

The formulation of the network search problem that is solved through the proposed method can be expressed as follows:

$$\max_{a_i} Accuracy(a_i) \quad (12)$$

constrained by  $LAT(a_i) \leq Budget, a = 1, \dots, N$ .

where  $a_i$  is the sampled network from the search space,  $LAT(a_i)$  is the latency on the platform of the sampled network and  $Budget$  is the predefined latency budget.

The problem presented in equation (12) is solved iteratively by the presented method by minimizing on each iteration the following loss function:

$$L(a, w_a) = CE(a, w_a) + \beta L_{LAT}(a), \quad (13)$$

where  $CE(a, w_a)$  is the cross-entropy loss of the network candidate  $a$  with weights  $w_a$  and  $LAT(a)$  is the measured latency of network candidate  $a$  in microseconds, [20]. Typical NAS approaches like [5], [17] or [6] are based on the iteratively training of sampled architectures candidates from the search space on a small proxy dataset through some epochs to be then transferred to the target dataset after training. In the end, the objective of these NAS methodologies is finding the network weights  $w$  and optimizing the network candidate  $a \in A$  (being  $A$  the search space), same as in the presented work but the needed resources and time to train thousand of network architectures before reaching the optimal solution make them some times infeasible.

Motivated by this problem, DNAS like [9], [19], [20] or [10] have become more popular lately. In this research we adopt a different paradigm of solving the same problem based on DNAS, as explained in Section (I).

In the presented work we relax the categorical choice of a particular filter or kernel in the target architecture by formulating the sampling process in the search stage, similar as proposed in [21] and [20], and define the probability of sampling the  $i$ -th kernel candidate  $K^i$ .

$$P(\bar{K} == K^i) = softmax(\alpha^{(i)}) = \frac{\exp(\alpha^{(i)})}{\sum_{j=0}^N \exp(\alpha^{(j)})}. \quad (14)$$

where  $\bar{K}$  is the sampled kernel during the search stage. Following this we reformulate the equation (9) and focus on making the Eq. (14) differentiable so that the loss function (13) can be optimized through stochastic gradient descent (SGD) approach, [22], [23], [20].

The objective function in Eq. (14) is already differentiable with respect to the weights of the kernels but not to the architecture parameters  $\alpha$  due to the sampling process. To solve this, we follow a similar approach as in NAS related works [17], [24], [19], [21]. We adopt the Gumbel Softmax function [25] to rewrite the equation (14):

$$P(\bar{K} == K^i) = \frac{\exp((\log(\theta_i) + g_i)/\tau)}{\sum_{j=0}^N \exp((\log(\theta_j) + g_j)/\tau)}, \quad (15)$$

where  $g_i \sim \text{Gumbel}(0,1)$  represents a white noise function that follows the Gumbel distribution between zero and one,  $\tau$  represents the temperature parameter of the Gumbel Softmax function, [25], which makes the discrete sampling probability function in Eq. (14) become continuous as  $\tau$  approximates to one. Lastly  $\theta_i$  represents the class probabilities calculated in Eq. 14, [10].

Once the loss function is differentiable, the SGD method to optimize function (13) is applied, so that on each iteration the network architecture weights  $w_i$  and probability parameters  $\alpha$  are updated based on the partial derivative of the loss function with respect to  $w$  and  $\alpha$ , respectively.

The search process is now equivalent to training the stochastic network after generating all kernel candidates from the  $11 \times 11$  meta kernel, similar to [10]. During training, the value of the loss function  $L(a, w_a)$  in Eq. (13) is calculated. Together with it,  $\partial L/\partial w_a$  and  $\partial L/\partial \alpha$  are computed to update weights and architecture probability parameters on each iteration. Through this, we train each operator's weight and update the sampling probability for each operator, respectively, so that when training finishes, we can obtain the optimal network architectures with the best kernels from the learned  $\alpha$  parameters.

As will be shown in the experiment section, the proposed approach works faster than RL and typical NAS methodologies and provides acceptable results for high resolution input images.

---

**Algorithm 1:** The search architecture methodology

---

**Result:** Find weights  $w_i$  and architecture probability parameters  $\alpha$  to optimize the global loss function (13), given a defined search space with a combination of operations  $\bar{o}^{(i,j)}$ , defined in Eq. (9), a latency budget and an input dataset.

random initialization of  $\alpha$  parameters

**while not converge do**

    Similar to [10], we generate the kernel candidates.

    Calculate Loss through Eq. (13).

    Calculate  $\partial L/\partial w_a$  and  $\partial L/\partial \alpha$ .

    Update weights and architecture probability parameters  $\alpha$ .

    Update Kernels using Eq. 3 and Eq. 5.

**end**

Extract more optimal architecture from learned  $\alpha$  parameters.

---

## IV. EXPERIMENTS

In this section we aim to demonstrate the performance and efficiency of the proposed method comparing the results obtained on different datasets.

We have conducted experiments on the commonly used ImageNet benchmark [26] and PascalVOC [27] datasets applying several important training techniques, as detailed in Section IV-A.

These results can be seen in Table I. In this section, a comparison of the proposed method with state-of-the-art is also presented.

### A. Implementation details

We have trained the models using 8 GPU NVIDIA Tesla V100. As proposed in [35] we have applied several implementation tricks to improve the training process, such as the following:

- Randomly sample an image and decode it into 32-bit floating point raw pixel values in [0,255].
- Random crop of a rectangular region.
- Horizontal flip with 0.5 probability.
- Normalize RGB channels.
- Scale hue, saturation and brightness coefficients.

Due to the importance of the learning rate in the training process, in the conducted experiments we have applied a learning rate warm up (use small learning rate at the beginning and then switch back to the initial learning rate when training process is stable) followed by decaying cosine learning rate to improve the training process, as commented in [35] and proposed by [36].

In contrast to the typical exponentially decaying learning rate used  $l_r = l_{r0} * e^{-Kt}$  in this work we have applied the following formula:

$$l_r = \frac{1}{2} \left( 1 + \cos \frac{b\pi}{B} \right) l_{r0} \quad (16)$$

$$w^* = w + l_r \frac{\partial L}{\partial w},$$

where  $B$  is the total number of batches,  $b$  is the actual batch during training,  $w^*$  the updated weight,  $w$  the weight before update and  $l_{r0}$  is the initial learning rate (in our case 0.65).

### B. Target platform

The embedded platform targeted to check the effectiveness of the proposed method is the TDA2 system on chip which can accelerate deep neural network layers using the C66x DSP cores together with the Texas Instruments Deep Learning suite (TIDL) to convert the trained network from floating point to fixed point and to enable the inference of the network on the embedded platform.

The mentioned TDA2 hardware has ARM Cortex-A15 cores running at up to 1.5 gigahertz, a dual-core DSP C66x processors that are capable of running deep learning inference, together with one embedded vision engine subsystem (EVE). It can run 16 times 16-bit (enough resolution for deep learning applications) MAC operations per cycle reaching up to 20.8 GMACs/s, [37].

To check the effectiveness of the proposed method we have mainly attended to MAC and FLOPs in order to compare the results with the theoretical performance of the target platform. To calculate this, based on the hardware specifications mentioned above the SoC has two C66x DSP cores running at 1000 MHz frequenz ( $2 * 32GMACs$ ), other one EVE core running at 900 MHz ( $16 * 900MMACs$ ) and other 2 ARM

TABLE I: ImageNet classification performance compared with other state-of-the-art methods. The proposed approach in this paper demonstrates a good Top-1 accuracy with less number of parameters and FLOPs. The number of parameters, FLOPs and Top-1 accuracy metrics presented in this table for the rest of the methodologies have been directly extracted from their respective papers. As it can be seen, the proposed method E-DNAS achieves similar accuracy results compared to other state-of-the-art methods, such as [10] and [14] in less time, as presented in Figure 3.

Model	Search Method	Search Space	Search Dataset	# Params(M)	FLOPs(M)	acc(%)
<i>MNetV2</i> [28]	manual	-	-	3.4	300	72.0
CondenseNet(G=C=8) [29]	manual	-	-	4.8	529	73.8
<i>EfficientNet-B0</i> [30]	manual	-	-	5.3	390	76.3
<i>NASNet-A</i> [5]	RL	cell	CIFAR-10	5.3	564	74.0
<i>PNASNet</i> [31]	SMBO	cell	CIFAR-10	5.1	588	74.2
<i>DARTS</i> [9]	gradient	cell	CIFAR-10	4.7	574	73.3
<i>PDARTS</i> [32]	gradient	cell	CIFAR-10	4.9	557	75.6
<i>GDAS</i> [21]	gradient	cell	CIFAR-10	4.4	497	72.5
<i>MnasNet</i> [17]	RL	stage-wise	ImageNet	3.9	312	75.2
<i>Single-Path NAS</i> [33]	gradient	layer-wise	ImageNet	4.3	365	75.0
<i>ProxylessNAS-R</i> [24]	RL	layer-wise	ImageNet	4.1	320	74.6
<i>ProxylessNAS-G</i> [24]	gradient	layer-wise	ImageNet	-	-	74.2
<i>FBNet</i> [20]	gradient	layer-wise	ImageNet	5.5	375	74.9
<i>MNetV3 Large</i> [34]	RL	layer-wise	ImageNet	5.4	219	75.2
<i>MNetV3 Small</i> [34]	RL	layer-wise	ImageNet	2.9	66	67.4
<i>MixNet</i> [14]	RL	kernel-wise	ImageNet	5.0	360	77.0
<i>MetaKernels</i> [10]	gradient	kernel-wise	ImageNet	7.2	357	77.0
<b>Ours</b>	<b>gradient</b>	<b>parallel kernel-wise</b>	<b>ImageNet</b>	<b>5.9</b>	<b>365</b>	<b>76.9</b>

TABLE II: Comparison of the obtained results on the Pascal VOC2007 test set. As suggested by other works like [10], the VOC2007 trainval and VOC2012 trainval are combined as the training data and the PascalVOC2007 dataset is used as test-set. We demonstrate a better classification accuracy than other similar methods such as [10] or [34].

Model	mAP	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	bike	person	plant	sheep	sofa
<i>MNetV2</i> [28]	75.8	84.5	83.4	76.1	68.3	58.7	78.9	84.8	86.5	54.4	80.7	70.9	84.0	85.0	83.6	76.8	48.7	78.7	72.8
<i>MNetV3-S</i> [34]	69.3	77.4	76.9	67.0	62.0	43.7	76.3	79.1	82.1	47.2	75.4	65.2	78.4	81.0	79.7	72.5	39.4	68.7	67.1
<i>MNetV3-L</i>	76.7	84.1	84.1	77.0	69.9	75.9	84.8	85.1	88.1	56.3	84.8	64.8	84.3	87.9	84.7	77.9	46.3	80.5	73.9
<i>metaKernel</i> [10]	77.3	86.1	84.8	76.8	68.6	59.2	83.6	86.3	87.1	56.9	85.2	67.2	86.6	87.2	86.0	77.7	49.1	80.8	74.5
<b>Ours</b>	<b>77.8</b>	<b>85.8</b>	<b>84.6</b>	<b>74.9</b>	<b>70.1</b>	<b>57.4</b>	<b>63.4</b>	<b>87.8</b>	<b>88.1</b>	<b>58.7</b>	<b>83.9</b>	<b>72.1</b>	<b>86.2</b>	<b>86.3</b>	<b>86.8</b>	<b>87.1</b>	<b>50.6</b>	<b>79.6</b>	<b>74.9</b>

Cortex A15 cores running at 1500 MHz ( $2 \times 8 \times 1500 \text{MMAC}s$ ). This results in 105 GMACs as theoretical performance, which means 210 GDLOPs, assuming DLOPs as 8-bit arithmetic or conditional operation (Multiply/Add/Compare). It can be assumed than  $1 \text{MAC} = 2 \text{DLOPs}$ .

For the experiments done on the mentioned TI and presented in Table III, the trained networks using the proposed method in this paper and ImageNet benchmark were converted from floating-point to fixed point to be then executed on the DSP dual core of the above mentioned hardware. In the floating-point to fixed-point conversion an accuracy loss of around 3 % could be extracted from the results.

### C. Results

The experimental results and comparison with other state-of-the-art methods are presented in Table I. It is there presented that our method achieves a good Top-1 accuracy better than several methods with a lower number of parameters and FLOPs.

For the first variant of this method in which input data passes first through a convolution with  $11 \times 11$  kernel we find

that a better accuracy with a slight increment of the number of FLOPs can be achieved by increasing the size of this first filter to  $13 \times 13$  until  $17 \times 17$ . Beyond that, the rate between accuracy and number of parameters decreases. After several tests, like showed in Table II, we have empirically seen that the mentioned  $11 \times 11$  kernel size gives the best trade-off between accuracy, number of operations and simplicity in the implementation.

Larger kernel sizes increase the model size with more parameters and also more operations and for this reason, using bigger kernels in the initial step of the pipeline would have led to a bigger network, not so suitable for embedded targets.

With regard to the search process speed, the experiments show that this proposal achieves an optimal architecture faster than other DNAS works, as it can be seen in Figure 3. Our experiment results are summarized in Table 3 where we compare our method with state-of-the-art efficient models both designed automatically and manually. In the case of the MnasNet, this paper does not disclose the exact search cost (in terms of GPU-hours or days) so in this paper we have

TABLE III: Results on ImageNet Benchmark comparing extracted multiply-accumulate operations from different methods and ours. The estimated inference latency on the described TI platform based on the calculated MACs is 38ms.

Model	# Params (M)	MACs (M)	Time (ms)
MNet [2]	4.2	569	75
NasNet-A [5]	5.3	564	183
<b>Ours</b>	<b>5.9</b>	<b>535</b>	<b>38</b>

assumed the prediction made for the search cost in MnasNet by [24] and [10]. Our experiment results are summarized in Table 3 where we compare our method with state-of-the-art efficient models both designed automatically and manually. In the case of the MnasNet, this paper does not disclose the exact search cost (in terms of GPU-hours or days) so in this paper we have assumed the prediction made for the search cost in MnasNet by [24] and [10].

Our experiment results are summarized in Table 3 where we compare our method with state-of-the-art efficient models both designed automatically and manually. In the case of the MnasNet, this paper does not disclose the exact search cost (in terms of GPU-hours or days) so in this paper we have assumed the prediction made for the search cost in MnasNet by [24] and [10].

## V. CONCLUSION

In this work we present a network search approach to design light and optimal DNNs reducing the searching time. We propose a two-step pipeline that learns different meta-kernel sizes, able to treat different resolution patterns. We propose a pairwise searching with circular feedback on each iteration to speed up the process by updating the target weights and network parameters iteratively with, not only the loss calculated during its training, but also with the loss calculated on the parallel kernel being learned.

We demonstrate that our method provides good results in terms of accuracy and searching speed compared to other methods like [17] or [9] under similar computation resource constraints.

## ACKNOWLEDGMENTS

This work was supported by the AGAUR Scholarship of the "Doctorats Industrials" program and by the Spanish government under projects HuMoUR TIN2017-90086-R and María de Maeztu Seal of Excellence MDM-2016-0656.

## REFERENCES

[1] Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, Kurt Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5mb model size," *International Conference on Learning Representations (ICLR)*, 2017.

[2] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam,

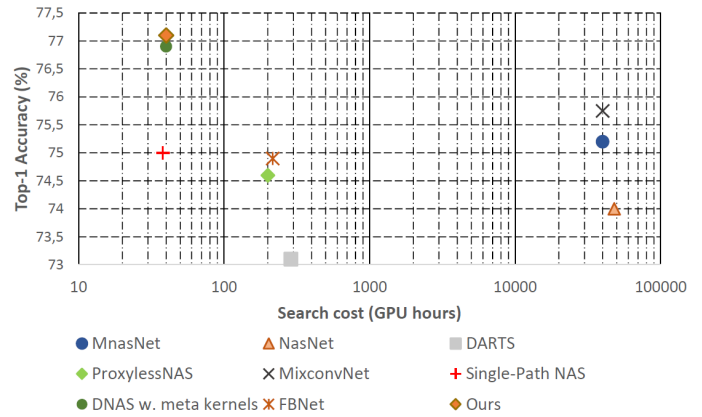


FIG. 3: Comparison of several NAS and DNAS methods in terms of search cost. The data has been directly obtained from their papers. As also commented in [20], the search cost for MnasNet is estimated according to the description in [17]. The search cost of PNAS [31] is estimated based on the results claimed on that work that their method is 8x faster than NAS [5].

"Mobilenets: Efficient convolutional neural networks for mobile vision applications," *ArXiv*, vol. abs/1704.04861, 2017.

[3] Tien-Ju Yang, Andrew Howard, Bo Chen, Xiao Zhang, Alec Go, Mark Sandler, Vivienne Sze, Hartwig Adam, "Netadapt: Platform-aware neural network adaptation for mobile applications," *European Conference on Computer Vision (ECCV)*, 2018.

[4] Tien-Ju Yang, Yu-Hsin Chen, Vivienne Sze, "Designing energy-efficient convolutional neural networks using energy-aware pruning," *Computer Vision and Pattern Recognition (CVPR)*, 2017.

[5] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, Quoc V. Le, "Learning transferable architectures for scalable image recognition," *International conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[6] Barret Zoph, Quoc V. Le, "Neural architecture search with reinforcement learning," *International Conference on Learning Representation (ICLR)*, 2016.

[7] Esteban Real, Alok Aggarwal, Yanping Huang, Quoc V Le, "Regularized evolution for image classifier architecture search," *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.

[8] Thomas Elsken, Jan-Hendrik Metzen, and Frank Hutter., "Simple and efficient architecture search for convolutional neural networks." *International Conference on Learning Representations (ICLR)*, 2018.

[9] Hanxiao Liu, Karen Simonyan, Yiming Yang, "Darts: Differentiable architecture search," *International Conference on Learning Representation (ICLR)*, 2019.

[10] Shoufa Chen, Yunpeng Chen, Shuicheng Yan, Jiashi Feng, "Efficient differentiable neural architecture search with meta kernels," *Computer Vision and Pattern Recognition (CVPR)*, 2019.

[11] Song Han, Jeff Pool, John Tran, William J. Dally,

- “Learning both weights and connections for efficient neural networks,” *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [12] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, Jan Kautz, “Pruning convolutional neural networks for resource efficient inference,” *International Conference on Learning Representation (ICLR)*, 2017.
- [13] Y. He and S. Han, “Adc: Automated deep compression and acceleration with reinforcement learning,” *ArXiv*, vol. abs/1802.03494, 2018.
- [14] Mingxing Tan, Quoc V. Le, “Mixconv: Mixed depthwise convolutional kernels,” *British Machine Vision Conference (BMVC)*, 2019.
- [15] Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton, “Imagenet classification with deep convolutional neural networks,” 2015. [Online]. Available: [http://vision.stanford.edu/teaching/cs231b\\_spring1415/slides/alexnet\\_tugce\\_kyunghee.pdf](http://vision.stanford.edu/teaching/cs231b_spring1415/slides/alexnet_tugce_kyunghee.pdf)
- [16] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, Jian Sun, “Shufflenet v2: Practical guidelines for efficient cnn architecture design,” *European Conference on Computer Vision (ECCV)*, 2018.
- [17] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, Quoc V. Le, “Mnasnet: Platform-aware neural architecture search for mobile,” *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [18] Xiaohan Ding, Yuchen Guo, Guiguang Ding, Jungong Han, “Acnet: Strengthening the kernel skeletons for powerful cnn via asymmetric convolution blocks,” *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [19] Guilin Li, Xing Zhang, Zitong Wang, Zhenguo Li, Tong Zhang, “Stacnas: Towards stable and consistent differentiable neural architecture search,” *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [20] 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,” *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [21] Xuanyi Dong, Yi Yang, “Searching for a robust neural architecture in four gpu hours,” *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [22] Cui, Xiaodong and Zhang, Wei and Tüske, Zoltán and Picheny, Michael, “Evolutionary stochastic gradient descent for optimization of deep neural networks,” *Advances in Neural Information Processing Systems (NIPS)*, 2018.
- [23] Sebastian Ruder, “An overview of gradient descent optimization algorithms,” 2017. [Online]. Available: <https://arxiv.org/pdf/1609.04747.pdf>
- [24] Han Cai, Ligeng Zhu, Song Han, “Proxyless nas: Direct neural architecture search on target task and hardware,” *International Conference on Learning Representation (ICLR)*, 2019.
- [25] Eric Jang, Shixiang Gu, Ben Poole, “Categorical reparameterization with gumbel-softmax,” *International Conference on Learning Representation (ICLR)*, 2017.
- [26] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” *Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [27] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman, “The pascal visual object classes (voc) challenge,” *International Journal of Computer Vision (IJCV)*, 2010.
- [28] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, Liang-Chieh Chen, “Mobilenetv2: Inverted residuals and linear bottlenecks,” *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [29] Gao Huang, Shichen Liu, Laurens van der Maaten, Kilian Q. Weinberger, “Condensenet: An efficient densenet using learned group convolutions,” *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [30] Mingxing Tan, Quoc V. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” *International Conference on Machine Learning (ICML)*, 2019.
- [31] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, Kevin Murphy, “Progressive neural architecture search,” *European Conference Computer Vision (ECCV)*, 2018.
- [32] Xin Chen, Lingxi Xie, Jun Wu, Qi Tian, “Progressive differentiable architecture search: Bridging the depth gap between search and evaluation,” 2019. [Online]. Available: <https://arxiv.org/abs/1904.12760>
- [33] Dimitrios Stamoulis, Ruizhou Ding, Di Wang, Dimitrios Lymberopoulos, Bodhi Priyantha, Jie Liu, Diana Marculescu, “Single-path nas: Designing hardware-efficient convnets in less than 4 hours,” 2019. [Online]. Available: <https://arxiv.org/abs/1904.02877>
- [34] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, Hartwig Adam, “Searching for mobilenetv3,” *International Conference on Computer Vision (ICCV)*, 2019.
- [35] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, Mu Li, “Bag of tricks for image classification with convolutional neural networks,” *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [36] TIllya Loshchilov, Frank Hutter, “Sgdr: Stochastic gradient descent with warm restarts,” *International Conference on Learning Representations (ICLR)*, 2017.
- [37] TEXAS INSTRUMENTS, “Texas instruments tidl,” 2018. [Online]. Available: <http://www.ti.com/tool/SITARA-MACHINE-LEARNING>