Accelerating multi-objective neural architecture search by random-weight evaluation

Shengran Hu¹ · Ran Cheng¹ · Cheng He¹ · Zhichao Lu¹ · Jing Wang² · Miao Zhang²

Received: 21 July 2021 / Accepted: 11 November 2021 / Published online: 4 December 2021
© The Author(s) 2021

Abstract

For the goal of automated design of high-performance deep convolutional neural networks (CNNs), neural architecture search (NAS) methodology is becoming increasingly important for both academia and industries. Due to the costly stochastic gradient descent training of CNNs for performance evaluation, most existing NAS methods are computationally expensive for real-world deployments. To address this issue, we first introduce a new performance estimation metric, named random-weight evaluation (RWE) to quantify the quality of CNNs in a cost-efficient manner. Instead of fully training the entire CNN, the RWE only trains its last layer and leaves the remainders with randomly initialized weights, which results in a single network evaluation in seconds. Second, a complexity metric is adopted for multi-objective NAS to balance the model size and performance. Overall, our proposed method obtains a set of efficient models with state-of-the-art performance in two real-world search spaces. Then the results obtained on the CIFAR-10 dataset are transferred to the ImageNet dataset to validate the practicality of the proposed algorithm. Moreover, ablation studies on NAS-Bench-301 datasets reveal the effectiveness of the proposed RWE in estimating the performance compared to existing methods.

Keywords Neural architecture search · Efficient performance estimation · Multi-objective optimization · Evolutionary algorithm

Introduction

In recent years, deep convolutional neural networks (CNNs) have been widely studied and achieved astonishing per-

formance in different computer vision tasks. One crucial component among these studies is the design of dedicated architectures of neural networks, which significantly affects the performance and generalization ability of CNNs among various tasks [14,20,26]. Along with the architectural milestones, from the original AlexNet [20] to the ResNet [14], the performance of CNNs across extensive datasets and tasks keeps boosting. However, it still takes researchers enormous works to achieve these architectural advancements via trial-and-error tuning manually. Therefore, neural architecture search (NAS) has emerged as an alternative way to design CNN in an automated manner. Although NAS alleviates the laborious experiments by researchers, existing NAS algorithms still suffer from numerous computational overheads, leading to challenges in the real-world deployment [34,55].

The expensive evaluations of the performance of architectures contribute to dominant computational consumption in the NAS algorithm. Usually, a brute-force training of a network can cost days to weeks on a single GPU, varying from simple to complex datasets and tasks. Therefore, several approaches have been proposed to approximate the true performance with fewer computational costs and, as a result,

✉ Ran Cheng
ranchengcn@gmail.com
Shengran Hu
hu.shengran@outlook.com
Cheng He
chenghehust@gmail.com
Zhichao Lu
lu.zhichao@outlook.com
Jing Wang
wangjinger1218@163.com
Miao Zhang
zhangm-168@163.com

¹ Guangdong Key Laboratory of Brain-Inspired Intelligent Computation, Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China

² Shanghai Aircraft Design and Research Institute, Shanghai 200135, China

fewer fidelities. These works can be roughly divided into three categories.

The first category includes methods that reduce training budgets via decreasing the network sizes (e.g., the number of layers and channels), which are widely adopted in early NAS works [12,34,55]. Nevertheless, their effectiveness has not been systematically researched until recently [51,54], demonstrating that their effectiveness can be limited under inappropriate parameter settings. Moreover, these methods are computationally expensive due to the thorough training for every single network. Finally, these methods require that the CNN architectures in search spaces should be modular, i.e., the networks are constructed by repeatedly stacking modular blocks. For instance, several state-of-the-art search spaces in [22,27] do not follow the constraints above, and the extension of these methods to new search spaces is not trivial.

The second category is often known as the supernet-based method, which intends to avoid training every architecture from scratch [27,28,42]. This technique typically decouples NAS into two main stages to share weights during searching. In the first stage, it constructs a supernet that contains all possible architectures in the search space, such that each architecture becomes a subnet of the supernet. In the second stage, the search process begins, and each architecture inherits the weights from the supernet, and thus the evaluation of each architecture becomes a simple inference on the valid set. Although this technique can speed up the searching process, the construction of the supernet could be more time-consuming than a complete search [3]. Besides, the search spaces require substantial modifications to accommodate the construction of the supernet [23].

The third category consists of several studies known as the zero-cost proxies [1,32], which estimate the performance with a few mini-batches of forward/backward propagation. To be more specific, this category analyzes information such as gradients, activations, or magnitudes of parameters to achieve estimations for reducing the computational cost drastically. Notably, most of these techniques attempt to validate their effectiveness on several NAS-Bench datasets [8,38] which are public architecture datasets constructed by exhaustively evaluating search spaces. Nevertheless, they may perform well only on certain NAS-Bench datasets [1], or they are not validated on real-world search spaces [32].

On top of the performance estimation methods, a branch of works named predictor-based NAS [27,28] has been proposed to further improve the sampling efficiency. In these works, a regression model, i.e., a performance predictor, would be trained to fit the mapping from the architectures to the corresponding performance. After the establishment of the predictor, the estimated performance in the searching stage is achieved by the evaluations of the predictor instead of the expensive estimation methods, which improves the

sampling efficiency of NAS. The predictor can be built upon different performance estimation methods, e.g., based on the training with reduced budgets [40,46] and the evaluations of the supernet [27,28]. Also, several works explore different encoding methods for the architectures [24,33,49] and different machine learning models as the predictor [33,40,46].

In this work, we propose a Random-Weight Evaluation (RWE) approach. Compared with the existing methods, it is less expensive, more flexible, and the effectiveness is validated more solid. In detail, by training the last classification layer only and keeping others with randomly initialized weights, RWE saves orders of magnitudes computational resources compared with conventional methods. At the same time, RWE is conceptually flexible with any search space, and it does not need any modifications to the search space. Moreover, the effectiveness of RWE is validated by the searching on two modern real-world search spaces and some ablation studies on the NAS-Bench-301 dataset. We briefly summarize our main contributions below:

- We propose a novel performance estimation metric, namely RWE, for efficiently quantifying the quality of CNNs. RWE is highly efficient in computational costs compared with conventional methods, which reduces the wall-clock evaluation time from hours to seconds. Extensive experiments on both real-world search spaces and benchmark search space NAS-Bench further validate the effectiveness of RWE.
- Paired with a multi-objective evolutionary algorithm, our RWE based NAS algorithm can achieve a set of efficient networks in one run, where both the performance and efficiency of models are considered. For instance, the proposed algorithm achieves state-of-the-art performance on the CIFAR-10 dataset, resulting in the networks from the largest, with 2.98% Top-1 error and 1.5 M parameters, to the smallest, with 4.05% Top-1 error and 0.9 M parameters. The experiments of transferability on ImageNet further demonstrate the competitiveness of our method. With such competitive performance, the whole searching procedure only costs less than two hours on a single GPU card, making the algorithm highly practical in handling real-world applications.

The rest of this paper is organized as follows. In “Related works”, some related work on multi-objective NAS algorithms and the expressive power of randomly initialized convolution filters is introduced. Then we present our proposed approach in “Proposed approach”, including the detailed random-Weight evaluation, search strategy, and search space and encoding. Comparative studies are shown in “Experimental results” and the conclusions are drawn in “Conclusion”.

Related works

In this section, we briefly discuss two topics related to the technicalities of our approach, i.e., multi-objective NAS and randomly initialized convolution filters.

Multi-objective NAS

Single-objective optimization algorithms have dominated the early researches in NAS [23,34,55], which mainly propose architectures to maximize their performance on certain datasets and tasks. Though NAS algorithms have shown their practicality in solving benchmark tasks, they cannot meet the demands from deployment scenarios varying from GPU servers to edge devices [15]. Thus, NAS algorithms are expected to balance between multiple conflicting objectives, such as inference latency, memory footprint, and power consumption. Though recent attempts often convert multiple competing objectives into a single objective in a weighted sum manner [4,43], they may miss the global optima of the problem. As a result, multiple runs of the algorithm could be required in real-world applications, due to the difficulty of deciding the best coefficient for weighted sum. Also, the search strategies these works adopted are primarily based on gradient-based methods or reinforcement learning, and they cannot approximate the Pareto front in a single run.

There are also several works that adopt multi-objective evolutionary algorithms as search strategies for NAS [27–29]. The population-based strategies introduce the natural parallelism, which increases the practicality in large-scale applications, and the conflicting nature of multiple objectives is helpful to enhance the diversity of the population. Most of them aim to tradeoff between the performance and the complexity of networks [27,29], while some other works temp to exploit the performance among different datasets, similar to the concepts in multi-task learning [28]. Following successful practices, we adopt a classic multi-objective evolutionary algorithm, namely NSGA-II [5]. We aim to achieve a set of efficient architectures in one run, where the proposed performance metric and a complexity metric FLOPs are two conflicting objectives to be optimized.

Expressive power of randomly initialized convolution filters

The RWE is surprisingly powerful, inspired by the fact that the convolution filters are in terms of extracting the features for input images, even with randomly initialized weights [17]. It is indicated in [9,17] that, with proper architecture, the convolution filters with randomly initialized weights can be as competitive as the ones with fully trained weights on both visual and control tasks. Also, it is validated that the structure itself can introduce prior knowledge to be capable

of capturing the features for visual tasks [2,44]. Similarly, the local binary convolutional neural network achieves comparable performance to CNNs with fully trained convolution filters by learning a linear combination of randomly initialized convolution filters [18].

Some early works in the literature conceptually explore the potential of estimating the performance of networks from randomly initialized weights. In detail, Sax et al. mathematically proved that the convolutional lter with random weights still has its key functionality, which is frequency selectivity and translation invariance. These characteristics are utilized to rank shallow neural networks with different network configurations [37]. Rosenfeld and Tsotsos successfully predict the performance ranking on several widely used CNN architectures by only training a fraction of weights in convolutional filters [35].

Although previous works show the potential of randomly initialized convolution filters, those methods are not scalable to real-world applications. In this work, we randomly initialize and freeze the weights in convolutional kernels in CNN, only training for the last classification layer. Using the predictive performance after doing so as a performance metric, we demonstrate the scalability of our approach on complex datasets and modern CNN search spaces that contains deep yet powerful CNNs.

Proposed approach

The Multi-objective NAS problem for a target dataset $\mathcal{D} = \{\mathcal{D}_{trn}, \mathcal{D}_{vld}, \mathcal{D}_{tst}\}$ can be formulated as the following bilevel optimization problem [30],

$$\begin{aligned} & \underset{\alpha}{\text{minimize}} && f_1(\alpha; \mathbf{w}^*(\alpha)), f_2(\alpha), \dots, f_m(\alpha) \\ & \text{subject to} && \mathbf{w}^*(\alpha) \in \underset{\mathbf{w}}{\operatorname{argmin}} \mathcal{L}(\mathbf{w}; \alpha), \\ & && \alpha \in \Omega_\alpha, \quad \mathbf{w} \in \Omega_w \end{aligned}$$

where the upper level variable α defines an architecture in the search space Ω_α , and the lower level variable $\mathbf{w}(\alpha)$ represents the corresponding weights. $\mathcal{L}(\mathbf{w}; \alpha)$ is the loss function on the $\mathcal{D}_{trn}$ for the architecture α with weights $\mathbf{w}$. The first objective f_1 represents the classification error on $\mathcal{D}_{vld}$, which depends on both architectures and weights. Other objectives $f_2, \dots, f_m$ only depend on architectures, such as the number of parameters, floating-point operations (FLOPs), latencies, etc.

In our approach, we simplify the complex bilevel optimization by using the proposed performance metric RWE as a proxy of f_1 . In addition, we adopt the complexity metric FLOPs as the second objective f_2 to optimize. As a result,

the multi-objective formulation of this work becomes

$$\begin{aligned} & \underset{\alpha}{\text{minimize}} \quad \text{RWE}(\alpha), \text{FLOPs}(\alpha) \\ & \text{subject to} \quad \alpha \in \Omega_{\alpha}, \end{aligned}$$

where RWE and FLOPs represent the values of these metrics with respect to architecture α .

Random-weight evaluation

Algorithm 1: Random-Weight Evaluation

Input : An architecture α , a training and validation dataset $\mathcal{D}_{trn}, \mathcal{D}_{vld}$

Output: The performance metric RWE of α .

- 1 $net \leftarrow$ Decode the architecture α into CNN backbone;
- 2 Randomly initialize the net and a linear classifier $clsfr$;
- 3 Freeze the weights of net throughout the whole algorithm;
- 4 $features \leftarrow$ Infer $\mathcal{D}_{trn}$ on the net ;
- 5 Train the linear classifier $clsfr$ with the $features$ as input;
- 6 **foreach** $image, target \in \mathcal{D}_{vld}$ **do**
- 7 Infer $image$ on net with $clsfr$;
- 8 $prediction \leftarrow$ the label approved by $clsfr$;
- 9 Compare the $prediction$ with $target$ and record the result;
- 10 **end**
- 11 $RWE \leftarrow$ Calculate the classification error rate;
- 12 **return** RWE ;

As mentioned in “Expressive power of randomly initialized convolution filters”, randomly initialized convolution filters are surprisingly powerful in extracting the features from images, due to the frequency selectivity and translation invariance preserved with random weights [37]. Inspired by this amazing characteristic, this work attempts to judge the quality of the architectures based on the ability of architectures with random weights to extract “good” features. And we quantify the quality of the features by training a linear classifier that takes these features as input and calculates the classification error for that classifier.

We detail the proposed performance metric Random-Weight Evaluation (RWE) as follows. The overall procedure is shown in Algorithm 1. First, we decode the encoding of a candidate architecture α into a CNN backbone net , which refers to all layers before the last classification layer. Second, we initialize the net and a linear classifier $clsfr$ with random weights, the latter of which acts as the last classification layer in a complete CNN and its structure is identical for all candidate CNNs in the search space. Here, a modified version of the *Kaiming initialization* [13] is adopted to initialize the net (default setting in PyTorch). The weights in the backbone part will keep frozen throughout the algorithm. Third, we infer the training set $\mathcal{D}_{trn}$ on net and utilize the output *feature* to train $clsfr$. Finally, after assembling net and trained

$clsfr$ into a complete CNN, this CNN gets tested on the validation set $\mathcal{D}_{vld}$, the output error rate of which becomes the value of RWE.

Search strategy

We adopt a classic multi-objective evolutionary algorithm NSGA-II [5] in our approach, where the searching process is detailed as below.

First, we randomly initialize the population, the individual of which get evaluated with RWE and FLOPs as two objectives. Second, we apply the binary tournament selection to select the parents for offspring. Third, the two-point crossover and the polynomial mutation are applied to generate the offspring, followed by the evaluations of offspring. Finally, we apply the environment selection based on the non-dominated sorting and the crowding distance [5], and the process is repeated until reaching the max generation.

Search space and encoding

Our proposed RWE is conceptually flexible and can be applied to any search space. To validate the effectiveness of our algorithm on the real-world application, we experiment with two modern search spaces, the micro [55] and macro [47] search spaces, in our approach. As shown in Fig. 1 LEFT, these two search spaces are modular search spaces, in which two kinds of layer, the normal and reduction layers, are repeatedly stacked, forming the complete CNN. The former kind of layers keeps the resolution and the number of channels for input images, while the latter halves the resolution and doubles the number of channels. The main difference between the micro and macro spaces is the design of each layer and the way to stack them into a complete CNN.

Micro search space: In the micro search space [55], we search for both the normal and reduction layers, named the normal and reduction cells. All the normal cells share the same architecture in a CNN, in which the weights are different though, and the case for the reduction cells is the same. Typically, we scale networks using different repeating numbers (N) in searching and validation stages. The normal and reduction cells share the same template, except for the stride size in operators. In each kind of cell, we search for both of the connections between nodes and the operation applied on each connection, as shown in Fig. 1, middle.

Macro search space: In the macro search space [47], we search for only the normal layers, leaving the predefined reduction layers identical. Each normal layer in the macro search space is searched independently, where the repeating number in a phase (N) is equal to one. In the normal layers, only the connection patterns get searched, and the operation in each node is a predefined sequential operator,

Fig. 1 The micro [55] and macro [47] search spaces adopted in our approach. Left: overall network architecture. Middle and right: design of layers in micro and macro search space

including convolution operators, batch normalization layers, and activation functions. Figure 1, right shows an example for candidate connection patterns.

Experimental results

In this section, we first present the searching results of our proposed NAS algorithm on the micro and macro search spaces for a modern classification dataset CIFAR-10 [19]. Then, the ablation studies on NAS-Bench-301 [38] demonstrate the effectiveness of our evaluation method and the rationality of some design choices. Finally, the experiment on ImageNet [6], which is one of the most challenging classification benchmarks, shows the transferability of our architectures and illustrates the practicality for real-world applications.

Searching on CIFAR-10

In our approach, we search on a modern classification dataset CIFAR-10 [19], which contains ten categories and 60K 32×32 images. Conventionally, the dataset is split into a training set with 50K images and a test set with 10K images. Following common settings in NAS algorithms [29,34,55], we further split the training set (80–20%) in the searching stage to create the training and validation sets.

Here we introduce the detailed implementation and parameter settings in our NAS algorithm. In the searching stage, the population size is set to 20 and the max generation is set to 30. For RWE, the architectures in the micro search space have 10 initial channels and 5 layers, and the architectures in the macro search space have 32 initial channels.

Also, due to the randomness introduced in RWE, we adopt assemble learning technique [11] in the training of the linear classifier to stabilize the results. Specifically, there are five classifiers to be trained, each of which is only exposed to 4/5 features. We only introduce the normalization techniques in the preprocessing of the input images, without the data augmentation techniques introducing the randomness. SGD optimizer with an initial learning rate of 0.25 and a momentum of 0.9 is adopted, where the cosine annealing schedule [25] decays the learning rate to zero gradually. The batch size is set to 512 and the training iterations conduct for 30 epochs. The average CPU time for a single evaluation is approximately 10 s with a single Nvidia 2080Ti.

For the validation stage, we scale the architectures to the deployment settings, where the number of training epochs, layers, and channels increases. The architectures from the final Pareto front are selected to be trained from scratch, where the number of layers and initial channels is set to 20 and 34 for the micro search space, and the number of channels is set to 128 in all layers in the macro search space. We use the same SGD optimizer like the one in the training stage, except the initial learning rate is set to 0.025. The selected architectures are trained for 600 epochs with a batch size of 96. Also, the regularization techniques cutout [7] and scheduled path dropout [55] are introduced, where the cutout length and the drop out rate are set to 16 and 0.2. The settings are the same with state-of-the-art algorithms for a fair comparison [29].

The results of validation and the comparison to other state-of-the-art architecture are shown in Table 1. The representative architectures from the final Pareto front get compared to both hand-crafted and search-based architectures. In the experiments with the micro search space, the architecture

Table 1 The results of the proposed algorithm and other state-of-the-art methods on CIFAR-10

Architecture	Test error (%)	Params (M)	FLOPs (M)	Search cost (GPU days)	Search method
Wide ResNet [50]	4.17	36.5	–	–	Manual
DenseNet-BC [16]	3.47	25.6	–	–	Manual
BlockQNN [†] [53]	3.54	39.8	–	96	RL
SNAS [†] [48]	3.10	2.3	–	1.5	Gradient
NASNet-A [†] [55]	2.91	3.2	532	2000	RL
DARTS [†] [23]	2.76	3.3	547	4	Gradient
NSGA-Net [‡] + macro space [29]	3.85	3.3	1290	8	Evolution
Macro-L[†]	4.27	2.79	1074	0.14	Evolution
AE-CNN + E2EPP [40]	5.30	4.3	–	7	Evolution
Hier. evolution [22]	3.75	15.7	–	300	Evolution
AmoebaNet-A [†] [34]	2.77	3.3	533	3150	Evolution
NSGA-Net [†] [29]	2.75	3.3	535	4	Evolution
Micro-S[†]	4.05	0.9	203	0.05	Evolution
Micro-M[†]	3.37	1.2	249	0.05	Evolution
Micro-L[†]	2.98	1.5	340	0.05	Evolution

The results of proposed algorithm are indicated in bold

[‡]Results achieved by the same training setting with ours and reported in [29]

[†]Work that adopts the regularization technique cutout [7]

Fig. 2 The visualization of **Micro-L** and **Macro-L** architecture

with the lowest error rate (**Micro-L**) in our approach achieves a 2.98% Top-1 error rate with 340M FLOPs. With fewer FLOPs, it has competitive performance with state-of-the-art architectures. Also, **Micro-M**, **Micro-S** shows a different tradeoff between performance and complexity. Similar to the micro search space, the chosen architecture in the macro search space (**Macro-L**) has competitive performance and fewer FLOPs compared with the state-of-the-art. The visualization of the detailed structures of **Micro-L** in the micro space and **Macro-L** in the macro space are shown in Fig. 2.

Effectiveness of random-weight evaluation

To demonstrate the effectiveness of RWE, we conduct experiments on the NAS-Bench-301 dataset [38]. The dataset is constructed by the surrogate model trained by sampled architectures in search space, such that it covers the whole search space to help researchers to analyze their NAS algorithms. While other NAS-Bench datasets construct a toy search space for the convenience of studies [8], NAS-Bench-301 covers a real-world search space, which is the micro search space

adopted in our work. Thus, the ablation studies based on NAS-Bench-301 examine the behavior of our algorithm during the searching stage.

We evaluate the effectiveness of estimation strategies by calculating the Spearman correlation coefficient between the estimation and queried performance from NAS-Bench-301. The target individuals are from the union of the population of each generation and their offspring. The Spearman correlation coefficient, which ranged from $[-1, 1]$, is a nonparametric measure of rank correlation. The higher coefficient is, the ranking of two variables has a more similar rank to each other. The idea of experiment settings is that, during the optimization using the evolutionary algorithm, the only phase depending on the estimation strategy is the mating and survival selection, which happen in the union mentioned above. The higher the correlation coefficient is, the more reliable the estimation strategy is. Thus, the algorithm has more chances to choose good candidates from a set of architectures.

In the following experiments, we use the same search strategy as introduced in “Search strategy” and conduct the experiments in NAS-Bench-301 with 20 generations. The search space in NAS-Bench-301 is a subset of the Micro Space, where identical connections to a single node are not allowed in NAS-Bench-301. As a result, we add a fix operation in the search strategy, which randomly chooses another connection to avoid duplication when an invalid architecture is produced. Also, the results present the mean and the standard variation of five independent trials with different random seeds.

We first compare our estimation strategy RWE with the zero-cost proxies [1,32] and the training-based evaluation method [54,55]. For the zero-cost proxies, we choose the representative performance metrics *synflow*, *grasp*, and *fisher* from [1], and *jacob_conv* from [32]. For the training-based method, we train the network for 10 epochs with the number of initial channels and layers of 16 and 8. As shown in Fig. 3, the proposed RWE outperforms all zero-cost proxies after the initial stage of searching, ending up with a similar accuracy with the training-based method. The experiment shows the effectiveness of RWE is competitive to the one of the training-based methods while having much fewer computational overheads. Paired with the searching in the micro and macro spaces, it further shows that RWE performs well in the real-world search spaces.

We then investigate the effects of different initialization methods applied in our approach. The method we adopt in this paper is the default one in PyTorch, and we examine other four representative initialization methods, which are known as *Kaiming normal (uniform) initialization* [13] and *Xavier normal (uniform) initialization* [10]. As shown in Fig. 4, the initialization methods have minor impacts on the effectiveness of RWE, as we observe no significant different

Fig. 3 The Spearman correlation coefficient for different performance metrics on NAS-Bench-301

Fig. 4 The Spearman correlation coefficient for different initialization method on NAS-Bench-301

behaviors. The experiment shows that our approach is robust to different initialization methods.

Transferring to ImageNet

To validate the practicality of our output architectures, we experiment with the transferability of the architecture **Micro-L** from CIFAR-10 to ImageNet [6]. ImageNet dataset, which substantially shows its importance in real-world applications, contains more than one million images. With various resolutions, these images are unevenly distributed in 1K categories. The general idea of transferring is to scale the architectures with a greater number of channels but a smaller number of layers, which is introduced by some classic works in NAS [48,55]. More specifically, the architecture starts with three stem convolutional layers with stride 2, which downsample the resolution by eight times. Following, there are 14 layers and 48 initial channels, where the reduction cells appear on the fifth and ninth layers. Some common data augmentation techniques are also adopted, including the random resize, the random crop, the random horizon flip, and the color jitter. We train our model with the SGD optimizer with 250

Table 2 The results of the proposed algorithm and other state-of-the-art methods on ImageNet

Architecture	Test error (%)		Params (M)	FLOPs (M)
	Top-1	Top-5		
MobileNetV1 [15]	31.6	–	2.6	325
InceptionV1 [41]	30.2	10.1	6.6	1448
ShuffleNetV1 [52]	28.5	–	3.4	292
VGG [39]	28.5	9.9	138	–
MobileNetV2 [36]	28.0	9.0	3.4	300
ShuffleNetV2 1.5× [31]	27.4	–	–	299
NASNet-C [†] [55]	27.5	9.0	4.9	558
SNAS [†] [48]	27.3	9.2	4.3	533
EffNet [†] [45]	27.0	9.25	2.5	–
DARTS [†] [23]	26.7	8.7	4.7	574
AmoebaNet-B [†] [34]	26.0	8.5	5.3	555
PNAS [†] [21]	26.0	8.5	5.3	555
Micro-L[†]	27.6	9.4	3.7	363

The results of proposed algorithm are indicated in bold

[†]Methods that first get searched on CIFAR-10 and then get transferred to ImageNet

epochs, batch size of 1024, and resolution of 224×224 on 4 Nvidia Tesla V100 GPU. The initial learning rate is set to 0.5 and decays to 1×10^{-5} linearly. In addition, the warmup strategy is applied on the first five epochs, increasing the learning rate from 0 to 0.5 linearly. The label smooth technique with a smooth rate of 0.1 is also adopted. Table 2 shows the experimental results and the comparisons to state-of-the-art methods. It shows that our approach has outperformed the hand-crafted architectures and has performed competitively with state-of-the-art NAS algorithms.

Conclusion

This paper proposed a flexible performance metric Random-Weight Evaluation (RWE) to rapidly estimate the performance of CNNs. Inspired by the expressive power of randomly initialized convolution filters, RWE only trains the last classification layer and leaves the backbone with randomly initialized weights. As a result, RWE achieves a reliable estimation of architectures in seconds. We further integrated RWE with an evolutionary multi-objective algorithm, adopting the complexity metric as the second objective. The experimental results showed that our algorithm achieved a set of efficient networks with state-of-the-art performance on both micro and macro search spaces. The resulting architecture with 350M FLOPs achieved 2.98% Top-1 error in CIFAR-10 and 27.6% Top-1 error in ImageNet after transferring. Also, the careful ablation studies experiments on different performance metrics and initialization methods indicated the effectiveness of the proposed algorithm.

Acknowledgements This work was supported by the National Natural Science Foundation of China (nos. 61903178, 61906081, and U20A20306), the Shenzhen Science and Technology Program (no. RCBS20200714114817264), and the Guangdong Provincial Key Laboratory (no. 2020B121201001).

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abdelfattah MS, Mehrotra A, Dudziak Ł, Lane ND (2021) Zero-cost proxies for lightweight NAS. In: International conference on learning representations (ICLR 2021). Virtual only
2. Adebayo J, Gilmer J, Goodfellow IJ, Kim B (2018) Local explanation methods for deep neural networks lack sensitivity to parameter values. In: International conference on learning representations (ICLR 2018). Vancouver, Canada

3. Cai H, Gan C, Wang T, Zhang Z, Han S (2020) Once for all: train one network and specialize it for efficient deployment. In: International conference on learning representations (ICLR 2020). Virtual Only
4. Cai H, Zhu L, Han S (2019) ProxylessNAS: direct neural architecture search on target task and hardware. In: International conference on learning representations (ICLR 2019). New Orleans, United States
5. Deb K, Pratap A, Agarwal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Comput* 6(2):182–197
6. Deng J, Dong W, Socher R, Li L, Li K, Li F (2009) Imagenet: a large-scale hierarchical image database. In: IEEE computer society conference on computer vision and pattern recognition. (CVPR 2009). Miami Beach, United States, pp 248–255
7. DeVries T, Taylor GW (2017) Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*
8. Dong X, Yang Y (2020) NAS-Bench-201: extending the scope of reproducible neural architecture search. In: International conference on learning representations (ICLR 2020). Virtual Only
9. Gaier A, Ha D (2019) Weight agnostic neural networks. In: Advances in neural information processing systems, (NeurIPS 2019) vol 32. Vancouver, Canada, pp 5365–5379
10. Glorot X, Bengio Y (2010) Understanding the difficulty of training deep feedforward neural networks. In: International conference on artificial intelligence and statistics. (AISTATS 2010) Sardinia, Italy, pp 249–256
11. Hansen LK, Salamon P (1990) Neural network ensembles. *IEEE Trans Pattern Anal Mach Intell* 12(10):993–1001
12. He C, Tan H, Huang S, Cheng R (2021) Efficient evolutionary neural architecture search by modular inheritable crossover. *Swarm Evol Comput* 64:100894
13. He K, Zhang X, Ren S, Sun J (2015) Delving deep into rectifiers: surpassing human-level performance on imagenet classification. In: International conference on computer vision. (ICCV 2015) Santiago, Chile, pp 1026–1034
14. He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: IEEE computer society conference on computer vision and pattern recognition (CVPR 2016). Las Vegas, United States, pp 770–778
15. Howard AG, Zhu M, Chen B, Kalenichenko D, Wang W, Weyand T, Andreetto M, Adam H (2017) MobileNets: efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*
16. Huang G, Liu Z, Van Der Maaten L, Weinberger KQ (2017) Densely connected convolutional networks. In: IEEE computer society conference on computer vision and pattern recognition (CVPR 2017). Honolulu, Hawaii, pp 4700–4708
17. Jarrett K, Kavukcuoglu K, Ranzato M, LeCun Y (2009) What is the best multi-stage architecture for object recognition? In: International conference on computer vision. IEEE, pp 2146–2153
18. Juefei-Xu F, Naresh Boddeti V, Savvides M (2017) Local binary convolutional neural networks. In: IEEE computer society conference on computer vision and pattern recognition (CVPR 2017). Honolulu, Hawaii pp 19–28
19. Krizhevsky A, Hinton G et al (2009) Learning multiple layers of features from tiny images. University of Toronto, Tech. rep
20. Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems, (NeurIPS 2012) vol 25. Lake Tahoe, United States, pp 1097–1105
21. Liu C, Zoph B, Neumann M, Shlens J, Hua W, Li L.J., Fei-Fei L, Yuille A, Huang J, Murphy K (2018) Progressive neural architecture search. In: European conference on computer vision (ECCV 2018). Munich, Germany, pp 19–34
22. Liu H, Simonyan K, Vinyals O, Fernando C, Kavukcuoglu K (2018) Hierarchical representations for efficient architecture search. In: International conference on learning representations (ICLR 2018). Vancouver, Canada
23. Liu H, Simonyan K, Yang Y (2019) DARTS: differentiable architecture search. In: International conference on learning representations (ICLR 2019). New Orleans, United States
24. Liu Y, Tang Y, Sun Y (2021) Homogeneous architecture augmentation for neural predictor. In: International conference on computer vision (ICCV 2021). Virtual only
25. Loshchilov I, Hutter F (2016) SGDR: stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*
26. Lu Z, Deb K, Boddeti VN (2020) MUXConv: information multiplexing in convolutional neural networks. In: IEEE computer society conference on computer vision and pattern recognition (CVPR 2020). Virtual only, pp 12044–12053
27. Lu Z, Deb K, Goodman E, Banzhaf W, Boddeti VN (2020) NSGANetV2: evolutionary multi-objective surrogate-assisted neural architecture search. In: European conference on computer vision (ECCV 2020). Virtual only, pp 35–51
28. Lu Z, Sreekumar G, Goodman E, Banzhaf W, Deb K, Boddeti VN (2021) Neural architecture transfer. *IEEE Trans Pattern Anal Mach Intell* 43(09):2971–2989
29. Lu Z, Whalen I, Boddeti V, Dhebar Y, Deb K, Goodman E, Banzhaf W (2019) NSGA-net: neural architecture search using multi-objective genetic algorithm. In: Genetic and evolutionary computation conference (GECCO 2019). Prague, Czech Republic, pp 419–427
30. Lu Z, Whalen I, Dhebar Y, Deb K, Goodman E, Banzhaf W, Boddeti VN (2020) Multi-objective evolutionary design of deep convolutional neural networks for image classification. *IEEE Trans Evol Comput* 25(2):277–291
31. Ma N, Zhang X, Zheng HT, Sun J (2018) Shufflenet v2: practical guidelines for efficient CNN architecture design. In: European conference on computer vision (ECCV 2018). Munich, Germany, pp 116–131
32. Mellor J, Turner J, Storkey A, Crowley EJ (2021) Neural architecture search without training. In: International conference on machine learning. PMLR, pp 7588–7598
33. Ning X, Zheng Y, Zhao T, Wang Y, Yang H (2020) A generic graph-based neural architecture encoding scheme for predictor-based NAS. In: European conference on computer vision. Springer, pp 189–204
34. Real E, Aggarwal A, Huang Y, Le QV (2019) Regularized evolution for image classifier architecture search. *Proceed AAAI Conf Artif Intell* 33(01):4780–4789
35. Rosenfeld A, Tsotsos JK (2019) Intriguing properties of randomly weighted networks: generalizing while learning next to nothing. In: International conference on robotics and vision (ICRV 2019). Singapore, pp 9–16
36. Sandler M, Howard A, Zhu M, Zhmoginov A, Chen LC (2018) MobileNetV2: inverted residuals and linear bottlenecks. In: IEEE computer society conference on computer vision and pattern recognition (CVPR 2018). Salt Lake City, United States, pp 4510–4520
37. Saxe AM, Koh PW, Chen Z, Bhand M, Suresh B, Ng AY (2011) On random weights and unsupervised feature learning. In: International conference on machine learning (ICML 2011). Bellevue, United States
38. Siems J, Zimmer L, Zela A, Lukasik J, Keuper M, Hutter F (2020) NAS-Bench-301 and the case for surrogate benchmarks for neural architecture search. *arXiv preprint arXiv:2008.09777*
39. Simonyan K, Zisserman A (2015) Very deep convolutional networks for large-scale image recognition. In: International conference on learning representations (ICLR 2015). San Diego, United States

40. Sun Y, Wang H, Xue B, Jin Y, Yen GG, Zhang M (2020) Surrogate-assisted evolutionary deep learning using an end-to-end random forest-based performance predictor. *IEEE Trans Evol Comput* 24(2):350–364
41. Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A (2015) Going deeper with convolutions. In: *IEEE computer society conference on computer vision and pattern recognition (CVPR 2015)*. Boston, United States, pp 1–9
42. Tan H, Cheng R, Huang S, He C, Qiu C, Yang F, Luo P (2021) RelativeNAS: relative neural architecture search via slow-fast learning. *IEEE transactions on neural networks and learning systems*. pp 1–1
43. Tan M, Chen B, Pang R, Vasudevan V, Sandler M, Howard A, Le Q.V. (2019) MnasNet: platform-aware neural architecture search for mobile. In: *IEEE conference on computer vision and pattern recognition*. pp 2820–2828
44. Ulyanov D, Vedaldi A, Lempitsky V (2018) Deep image prior. In: *IEEE computer society conference on computer vision and pattern recognition (CVPR 2018)*. Salt Lake City, United States, pp 9446–9454
45. Wang B, Xue B, Zhang M (2021) Surrogate-assisted particle swarm optimization for evolving variable-length transferable blocks for image classification. *IEEE transactions on neural networks and learning systems*. pp 1–14. <https://doi.org/10.1109/TNNLS.2021.3054400>
46. Wen W, Liu H, Chen Y, Li H, Bender G, Kindermans PJ (2020) Neural predictor for neural architecture search. In: *European conference on computer vision*. Springer, pp 660–676
47. Xie L, Yuille A (2017) Genetic CNN. In: *International conference on computer vision (ICCV 2017)*. Venice, Italy
48. Xie S, Zheng H, Liu C, Lin L (2019) SNAS: stochastic neural architecture search. In: *International conference on learning representations (ICLR 2019)*. New Orleans, United States
49. Yan S, Zheng Y, Ao W, Zeng X, Zhang M (2020) Does unsupervised architecture representation learning help neural architecture search? *Adv Neural Inf Process Syst* 33: 12486–12498
50. Zagoruyko S, Komodakis N (2016) Wide residual networks. In: *British machine vision conference British machine vision conference (BMVC 2016)*. York, United Kingdom
51. Zela A, Klein A, Falkner S, Hutter F (2018) Towards automated deep learning: efficient joint neural architecture and hyperparameter search. *arXiv preprint [arXiv:1807.06906](https://arxiv.org/abs/1807.06906)*
52. Zhang X, Zhou X, Lin M, Sun J (2018) ShuffleNet: an extremely efficient convolutional neural network for mobile devices. In: *IEEE computer society conference on computer vision and pattern recognition (CVPR 2018)*. Salt Lake City, United States, pp 6848–6856
53. Zhong Z, Yang Z, Deng B, Yan J, Wu W, Shao J, Liu C (2021) BlockQNN: efficient block-wise neural network architecture generation. *IEEE Trans Pattern Anal Mach Intell* 43(7):2314–2328
54. Zhou D, Zhou X, Zhang W, Loy CC, Yi S, Zhang X, Ouyang W (2020) Econas: finding proxies for economical neural architecture search. In: *IEEE computer society conference on computer vision and pattern recognition (CVPR 2020)*. Virtual only, pp 11396–11404
55. Zoph B, Vasudevan V, Shlens J, Le QV (2018) Learning transferable architectures for scalable image recognition. In: *IEEE computer society conference on computer vision and pattern recognition (CVPR 2018)*, Salt Lake City, United States, pp 8697–8710

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.