

Fitting the Search Space of Weight-sharing NAS with Graph Convolutional Networks

Xin Chen¹, Lingxi Xie¹, Jun Wu², Longhui Wei¹, Yuhui Xu³, Qi Tian¹

¹ Huawei Cloud & AI

² Fudan University

³ Shanghai Jiao Tong University

chenxin061@gmail.com, 198808xc@gmail.com, wujun@fudan.edu.cn,
weilonghui1@huawei.com, yuhuixu@sjtu.edu.cn, tian.qi1@huawei.com

Abstract

Neural architecture search has attracted wide attentions in both academia and industry. To accelerate it, researchers proposed weight-sharing methods which first train a super-network to reuse computation among different operators, from which exponentially many sub-networks can be sampled and efficiently evaluated. These methods enjoy great advantages in terms of computational costs, but the sampled sub-networks are not guaranteed to be estimated precisely unless an individual training process is taken. This paper attributes such inaccuracy to the inevitable mismatch between assembled network layers, so that there is a random error term added to each estimation. We alleviate this issue by training a graph convolutional network to fit the performance of sampled sub-networks so that the impact of random errors becomes minimal. With this strategy, we achieve a higher rank correlation coefficient in the selected set of candidates, which consequently leads to better performance of the final architecture. In addition, our approach also enjoys the flexibility of being used under different hardware constraints, since the graph convolutional network has provided an efficient lookup table of the performance of architectures in the entire search space.

Neural architecture search (NAS) is an emerging research field of automated machine learning (AutoML), with the goal being exploring deep networks that have not been investigated by manual designs. Early NAS approaches (Zoph and Le 2017; Real et al. 2017) mostly sampled architectures from a large search space and evaluated them using an individual training-from-scratch process. Despite their ability in finding powerful architectures, the search process is often computationally expensive, *e.g.*, hundreds or even thousands of GPU-days are required even if efficient sampling strategies were used (Liu et al. 2018; Real et al. 2018).

To alleviate the computational burden, researchers started to consider reusing computation among differently sampled architectures (Cai et al. 2018). Going one step forward, an efficient framework named weight-sharing NAS (Pham et al. 2018; Liu, Simonyan, and Yang 2018; Xie et al. 2020) was proposed in which the search space is formulated into a **super-network**, an over-parameterized architecture which

contains the parameters of all architectures (often referred to as **sub-networks**) that can appear. With the super-network being pre-trained, each sampled sub-network can be estimated with reduced costs, so that the overall search process is largely accelerated, *e.g.*, by 4–5 orders of magnitudes if differentiable search is used (Chen et al. 2019; Xu et al. 2020).

One of the most important concerns of weight-sharing NAS lies in the accuracy of sub-network sampling. In other words, there is no guarantee if the estimated accuracy of the sub-networks can reflect the real performance of the corresponding architectures. In particular, given a search space, \mathcal{S} , and two architectures, \mathbb{M}_1 and \mathbb{M}_2 , we can either evaluate them by training them from scratch or train a super-network and then sample them from it. Then, how often will the relative performance (*i.e.*, whether \mathbb{M}_1 is better than \mathbb{M}_2) be consistent under the two evaluation methods? We perform experiments by sampling 8 architectures with similar hardware complexity from a search space containing 19 cells, each of which has 6 possibilities. The result is disappointing: the Kendall- τ coefficient between two rankings is merely 0.2143 (0 being random permutation). With such a low correlation, it is hard to guarantee that weight-sharing NAS can find a high-quality architecture in the search space.

To solve this problem, we make an assumption that the inaccuracy mostly comes from the randomness during training the super-network. We take a single-path training process (Guo et al. 2019; Chu et al. 2019b) as an example: in each training iteration, only one operator in each cell gets updated while others remain unchanged. For a candidate architecture being evaluated, its performance will potentially be high if it is closely related to one of the recently trained sub-networks. Therefore, for an arbitrary architecture, \mathbb{M} , we formulate the relationship between the real performance, $z_{\mathbb{M}}^*$, and the estimated one, $z_{\mathbb{M}}$, into a linear formula¹ of $z_{\mathbb{M}} = a \times z_{\mathbb{M}}^* + \epsilon_{\mathbb{M}} + b$, where a and b are constants and $\epsilon_{\mathbb{M}}$ is a zero-mean random variable associated to \mathbb{M} . The goal is to alleviate the impact of $\epsilon_{\mathbb{M}}$, since the remaining part will

¹A higher-order function may better fit the data distribution, but it also increases the risk of over-fitting. More importantly, we found no cues in NAS that suggests a higher-order relationship between the true and sampled accuracy, so using a linear assumption is a relatively safe choice.

not change the relative ranking of the sampled architectures.

Next, we assume that $z_{\mathbb{M}}^*$, though difficult (or expensive) to obtain, is a learnable function with respect to \mathbb{M} . Hence, we sample a set of training data, $\{(\mathbb{M}, z_{\mathbb{M}})\}$, and train a graph convolutional network $f(\cdot)$ by minimizing the average error of $|z_{\mathbb{M}} - f(\mathbb{M})|$. We expect that $f(\cdot)$, by seeing large amount of training pairs, $(\mathbb{M}, z_{\mathbb{M}})$, can minimize of the random noise, *i.e.*, $f(\mathbb{M}) \approx a \times z_{\mathbb{M}}^* + b$. Hence, $f(\mathbb{M})$ serves as an estimation of the performance of \mathbb{M} which is almost free to compute ($f(\cdot)$ is very light-weighted) meanwhile being more reliable than $z_{\mathbb{M}}$ since it produces the same ranking as $z_{\mathbb{M}}^*$.

We perform experiments on the search space defined by FairNAS (Chu et al. 2019b) which has 19 cells and 6 choices for each cell. Starting with the same super-network, our approach partitions the search space into three subspaces (with 7, 6, and 6 cells, respectively) due to the limitation of CPU memory (it is difficult to support more than 6^7 nodes in a graph). The iterative optimization takes 9 GPU-hours beyond super-network training which is faster than that of FairNAS which sampled 12.8K sub-networks directly. Our approach achieves a top-1 accuracy of 75.5% on ImageNet with 383M multi-add operations, or 75.6% with 393M multi-add operations, which consistently surpass the counterpart that uses random sampling. In addition, the graph convolutional network makes it easier to formulate the network performance in the entire search space into a lookup table, based on which a wide range of hardware constraints can be integrated towards more demands of architecture search.

Related Work

Neural architecture search aims to automate the network design process and discover architectures that perform better than hand-crafted ones (Elsken, Metzen, and Hutter 2019). In a general NAS pipeline, architectures are sampled from a pre-defined search space and evaluated on a specific task, *e.g.*, image classification. To reduce the number of sampled architectures and accelerate the search process, heuristic algorithms including evolutionary algorithms (EA) and reinforcement learning (RL) are adopted to guide the sampling process. Recently, some EA-based (Xie and Yuille 2017; Zoph and Le 2017; Real et al. 2018) and RL-based (Zoph et al. 2018; Liu et al. 2018; Tan et al. 2019) approaches have achieved state-of-the-art performance on a variety of computer vision and natural language processing tasks. However, these approaches share a common evaluation scheme that optimizes each sampled architecture from scratch, which results in a critical drawback of heavy computational overhead even on CIFAR10, a small proxy dataset.

To alleviate the computational burden, researchers proposed an efficient solution which trains an over-parameterized super-network to cover all architectures in the search space and reuses or shares network weights among multiple architectures. This was named one-shot NAS (Brock et al. 2017) or weight-sharing NAS (Pham et al. 2018). ENAS (Pham et al. 2018) proposed to share weights among child models and apply reinforcement learning to improve the efficiency of computations, which dramatically reduce the search time to less than half a day with a single GPU. Pushing one-shot NAS to a continuous parameter space, DARTS (Liu,

Simonyan, and Yang 2018) and its variants (Xu et al. 2020; Chen et al. 2019; Dong and Yang 2019; Xie et al. 2018; Bi et al. 2020) adopted a differentiable framework that assigned a set of architectural parameters aside from the parameters of the super-network and iteratively optimized them by gradient descent, where the importance of different candidates is determined by the value of these architectural parameters. The reduction in computational burden facilitated architecture search on large-scale proxy datasets, *e.g.*, ImageNet (Deng et al. 2009), with acceptable search cost. For example, ProxylessNAS (Cai, Zhu, and Han 2018) and FBNet (Wu et al. 2019) successfully searched architecture directly on ImageNet with differentiable schemes.

Despite the great success in accelerating NAS, one-shot or weight-sharing methods still suffer from a severe problem named ranking inconsistency, which refers to the fact that the estimation of a sub-network can be different when it is sampled from the super-network and when it is trained from scratch. Single-path one-shot NAS (Guo et al. 2019) used a uniformly sampling method to guarantee that all candidates are fully and equally trained, which is believed to be effective on alleviating the inconsistency. FairNAS (Chu et al. 2019a) paved one step further and proposed to train the super-network with a strict constraint on fairness and demonstrated a stronger correlation between the one-shot and stand-alone evaluation results.

Our Approach

Our work is based on the pipeline of weight-sharing NAS. Under this pipeline, $L \times O$ sets of parameters are trained to simulate the behavior of O^L architectures, where each architecture consists of L layers and each layer is chosen from O candidates. Although this scheme demonstrates superior efficiency, there is no guarantee that the performance of sub-networks can be accurately estimated in this way. We will delve deep into this problem in the following parts. More details about weight-sharing NAS for classification can be found in the Arxiv version (Chen et al. 2020).

Inaccuracy of Sub-network Sampling

Throughout this paper, we work on the search space defined by FairNAS (Chu et al. 2019b) which has 19 inverted residual cells (Sandler et al. 2018), each of which has 6 options differing from each other in the expansion ratio and kernel size. On the ImageNet dataset with 1,000 classes, we randomly sample 100 classes and use the corresponding subset to optimize the super-network. 90% of the training data is used to update the model parameters, and the remaining 10% is used for evaluating each of the sampled sub-networks. Here, we follow (Guo et al. 2019; Chu et al. 2019b) to directly feed each testing image into the sampled sub-network and obtain the classification accuracy. Evaluating each sub-network on the validation subset (around 13K images) takes an average of 4.48 seconds on an NVIDIA Tesla-V100 GPU.

The first and foremost observation is that the performance of sub-networks cannot be accurately evaluated in this manner. To show this, we randomly sample 8 architectures with a similar complexity (*i.e.*, 400M multi-adds), and evaluate

ID	FLOPs	Acc*	Acc ₁₅₀	Acc ₁₆₀
01	400M	85.78%	81.59%	81.20%
02	401M	85.76%	81.56%	81.41%
03	401M	85.59%	81.73%	81.55%
04	400M	85.48%	81.95%	81.67%
05	403M	85.32%	81.70%	81.37%
06	399M	85.28%	81.64%	81.15%
07	402M	84.98%	81.60%	81.58%
08	400M	84.60%	81.53%	81.46%
τ	N/A	—	0.2143	-0.1429




Figure 1: The inconsistency between the performance of sub-networks by sampled from the super-network and trained from scratch. We choose 8 architectures with similar FLOPs and use the super-networks after 150 and 160 epochs of training. Left: the table summarizing all classification accuracy over 100 classes, in which τ indicate the Kendall- τ coefficient between the ranking and the ground-truth. Right: visualizing the results on a 2D plane, which the horizontal and vertical axes denote the ground-truth and sampled accuracy, respectively.

their performance in two different ways, namely, sampling the corresponding sub-network from the well-trained super-network, or training the sub-network from scratch. Results are summarized in Figure 1. One can see that the ranking of the sampled accuracy can be very different from that of the ground-truth accuracy, *i.e.*, when each architecture undergoes a complete training process from scratch. That being said, the architecture with the highest sub-network sampling accuracy may not be the optimal solution. As a side note, FairNAS (Chu et al. 2019b) advocated for that sub-network sampling is good at performance estimation, but we point out that the seemingly accurate prediction was mostly due to the large difference between the sampled architectures (they sampled 13 architectures with FLOPs varying significantly, unlike the 8 architectures sampled by our work that are all close to 400M FLOPs).

We explain this phenomenon by noting that the single-path strategy of super-network training is highly random. Let \mathbb{M}_0 be an architecture to be evaluated using sub-network sampling. Since each layer of \mathbb{M}_0 shares the training process with other operators, it is probable that each layer of \mathbb{M}_0 gets updated in different training iterations. In this situation, when \mathbb{M}_0 is sampled as a sub-network, its layers may not ‘cooperate’ well with each other. In particular, if a layer does not get updated for a long time, its parameters are relatively ‘outdated’ and thus may incur a low recognition accuracy of \mathbb{M}_0 . On the other hand, if all layers of another architecture, \mathbb{M}_1 , happen to be updated sufficiently in the last few iterations, the recognition accuracy of \mathbb{M}_1 is potentially high. Nevertheless, this does not mean that \mathbb{M}_1 is better than \mathbb{M}_0 .

In this paper, we introduce a simple model to formulate the above randomness introduced by sub-network sampling. Let $z_{\mathbb{M}}^*$ be the ground-truth accuracy of \mathbb{M} , *e.g.*, when \mathbb{M} is trained individually, and $z_{\mathbb{M}}$ be the accuracy obtained by sub-network sampling. Note that $z_{\mathbb{M}}$ is related to the super-

network, \mathbb{S} , but we omit \mathbb{S} for simplicity. Our key assumption is that $z_{\mathbb{M}}$ can be written in a linear function of $z_{\mathbb{M}}^*$ added by a random perturbation:

$$z_{\mathbb{M}} = a \times z_{\mathbb{M}}^* + b + \epsilon_{\mathbb{M}}. \quad (1)$$

Here, a is a linear coefficient indicating the systematic error between $z_{\mathbb{M}}^*$ and $z_{\mathbb{M}}$, and b is a constant bias between $z_{\mathbb{M}}^*$ and $z_{\mathbb{M}}$. Most often, a is slightly smaller than 1.0 because the super-network is not sufficiently optimized compared to a complete training process. $\epsilon_{\mathbb{M}}$ is a random variable associated to \mathbb{M} . According to the above analysis, it is mostly determined by how the layers of \mathbb{M} are updated in the final iterations, so it is totally random and unavailable unless $z_{\mathbb{M}}^*$ is tested. As shown in Figure 1, the intensity of $\epsilon_{\mathbb{M}}$ can be large – for two architectures, \mathbb{M}_1 and \mathbb{M}_2 , it is possible that the difference between $\epsilon_{\mathbb{M}_1}$ and $\epsilon_{\mathbb{M}_2}$ is even larger than $a \times |z_{\mathbb{M}_1}^* - z_{\mathbb{M}_2}^*|$, which may alter the relative ranking between the two models.

Further Analysis on the Key Assumption

We perform a toy experiment to demonstrate the influence of the random error term. We adopt a small search space that contains 6 cells, each of which has 6 choices, *i.e.*, a search space with $6^6 = 46,656$ candidate architectures. The super-network is firstly trained for 150 epochs to guarantee a sufficient convergence degree and we denote the parameters of this super-network as θ_{150} . 10,000 candidate architectures (about 20% of the whole search space) are randomly sampled from the search space and evaluated with θ_{150} on the stand-alone validating set by sub-network sampling. We mark the selected architectures and their corresponding evaluation accuracy as $\{(\mathbb{M}_m, z_{\mathbb{M}_m})\}_{m=1}^{10000}$. After this process, the super-network is further trained for a few more iterations² with the same training data, which results with a set of slightly

²Here, we train the super-network for 12 iterations, *i.e.*, about one-tenth of a full epoch.

modified parameters, denoted as θ'_{150} . A similar process is performed with θ'_{150} to obtain a set of new evaluation results, $\{(\mathbb{M}_m, z'_{\mathbb{M}_m})\}_{m=1}^{10000}$.

After acquiring these two set of evaluation results, we calculate the Kendall- τ coefficient between them and the result is astonishing. The Kendall- τ coefficient is only 0.5470, which implies a big change on the ranking of selected architectures. This phenomenon strongly support our assumption that the random error term does matter and it can even alter the ranking of two architectures.

Alleviating Noise with Graph Convolutional Networks

In what follows, we train a model to estimate $a \times z_{\mathbb{M}}^* + b \equiv z_{\mathbb{M}} - \epsilon_{\mathbb{M}}$, *i.e.*, eliminating the noise, $\epsilon_{\mathbb{M}}$, from $z_{\mathbb{M}}$. Note that $\epsilon_{\mathbb{M}}$ is not predictable, but we assume that it is a zero-mean noise, *i.e.*, $\sum_{\mathbb{M} \in \mathcal{S}} \epsilon_{\mathbb{M}} = 0$, otherwise we can adjust the value of the bias, b , to achieve this goal. Then, we collect a number of pairs, $(\mathbb{M}_m, z_{\mathbb{M}_m})$, and train a function, $f(\mathbb{M}; \boldsymbol{\eta})$, to fit $z_{\mathbb{M}}$, where $\boldsymbol{\eta}$ are learnable parameters. Since \mathbb{M} is a structural data (the encoded architecture), we naturally choose a graph convolutional network (GCN) (Kipf and Welling 2017) to be the form of $f(\cdot)$ ³. The objective is written as:

$$\boldsymbol{\eta}^* = \arg \min_{\boldsymbol{\eta}} \mathbb{E}[|f(\mathbb{M}; \boldsymbol{\eta}) - z_{\mathbb{M}}|_1 \mid \mathbb{M} \in \mathcal{S}]. \quad (2)$$

Under the assumption that $\epsilon_{\mathbb{M}}$ is completely irrelevant to \mathbb{M} (*i.e.*, for any $\boldsymbol{\eta}$ and an arbitrary set of architectures, the correlation coefficient between the $f(\mathbb{M}; \boldsymbol{\eta})$ and $\epsilon_{\mathbb{M}}$ values is 0), we can derive that the optimal solution of Eqn (2) is $f(\mathbb{M}) \equiv a \times z_{\mathbb{M}}^* + b \equiv z_{\mathbb{M}} - \epsilon_{\mathbb{M}}$. In other words, the best \mathbb{M} 's that maximize $f(\mathbb{M})$ and $z_{\mathbb{M}}^*$ are the same, *i.e.*, the optimal architecture, \mathbb{M}^* .

The overall pipeline of using GCN for NAS is illustrated in Algorithm 1. Since the search space is very large, we cannot include all 6^{19} architectures in one graph, so we start with an initialized model, \mathbb{M}_0 , and apply an iterative process, each round of which updates a subset of layers of the current architecture. The idea that gradually optimizes the architecture is similar to that explored in PNAS (Liu et al. 2018). Thanks to the stable property of GCN, the choice of \mathbb{M}_0 barely impacts the final architecture. In this paper, we choose the ‘standard’ network in which all cells have a kernel size of 3 and an expansion ratio of 6, *i.e.*, the architecture is similar to MobileNet-v2 (Sandler et al. 2018).

Let there be T rounds of iteration in the search process. The t -th round, $t = 0, 1, \dots, T - 1$, allows a subset of layers, \mathcal{L}_t , to be searched, and keeps all others fixed. The optimal sub-architecture found in this process will replace the corresponding layers, advancing \mathbb{M}_t to \mathbb{M}_{t+1} , and the iterative process continues. Denote the subset of architectures that can appear in the t -th round by $\mathcal{S}_t \subset \mathcal{S}$. We construct a graph, $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$, where \mathcal{V}_t denotes the set of nodes (each node corresponds to a sub-network) and \mathcal{E}_t denotes the set of edges connecting nodes. For simplicity, we use the most

³We could replace the GCN with a MLP or an LSTM here, but previous literature has proved that GCN is one of the optimal choices (Shi et al. 2020).

Algorithm 1: Applying GCN for Weight-sharing NAS

Input : Search space \mathcal{S} , dataset \mathcal{D} , cell index set \mathcal{L} ;
Output : Optimal architecture \mathbb{M}^* ;

- 1 Split \mathcal{D} into $\mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{val}}$, train the super-network \mathbb{S} on $\mathcal{D}_{\text{train}}$;
- 2 Initialize \mathbb{M}_0 as a default architecture, $t \leftarrow 0$;
- 3 **repeat**
- 4 Sample a subset of active cells, \mathcal{L}_t , determine the subspace \mathcal{S}_t ;
- 5 Construct a graph $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$, in which $\mathcal{V}_t \equiv \mathcal{S}_t$;
- 6 Sample M architectures from \mathcal{S}_t and evaluate the performance to fill up part of \mathcal{G}_t ;
- 7 Train a GCN on \mathcal{G}_t , use the GCN to find top- K architectures in \mathcal{S}_t ;
- 8 Evaluate all K architectures to find the best one, \mathbb{M}_t^* ;
- 9 $\mathbb{M}_{t+1} \leftarrow \mathbb{M}_t^*$, $t \leftarrow t + 1$;
- 10 **until** $t = T$ **or** convergence;

Return : \mathbb{M}_t .

straightforward way of graph construction, *i.e.*, $\mathcal{V}_t \equiv \mathcal{S}_t$, and there exists a connection between two nodes if and only if the corresponding sub-networks have a Hamming distance of 1. We observe in experiments that introducing more edges (*e.g.*, using a Hamming threshold of 2) does not improve GCN prediction accuracy but considerably increases training complexity.

Note that our work is related to prior ones focusing on predicting network performance without actually training them from scratch (Luo et al. 2018; Friede et al. 2019). These approaches achieved promising results on some NAS benchmarks (Ying et al. 2019; Dong and Yang 2020), demonstrating that one can find top-ranked architectures by sampling a very small portion of the entire space. However, they lack the guarantee that the architectures in training, in particular those with very different topologies, have potential relationship to each other so that the trained model is indeed learning useful knowledge. Our work adds two assumptions that makes the prediction more reliable: (i) a well-trained super-network is provided so that the sampled sub-networks share the same set of network weights; (ii) only sub-networks with 1-cell difference are related to each other. There are also other methods apply GCN to predictor-based NAS (Shi et al. 2020; Wen et al. 2020; Ning et al. 2020). They mostly denotes one single architecture as a graph, while our method regards each single architecture as a vertex in the graph which represents the whole search space.

Below we describe the details of our approach. Step 0 is performed only once, and Steps 1–3 can be iteratively executed until convergence or a pre-defined number of rounds is achieved.

- **Step 0: Super-network Training**

We train the super-network in a single-path, one-shot process. During each iteration, a mini-batch is sampled from the training set, and a sub-network is constructed by randomly

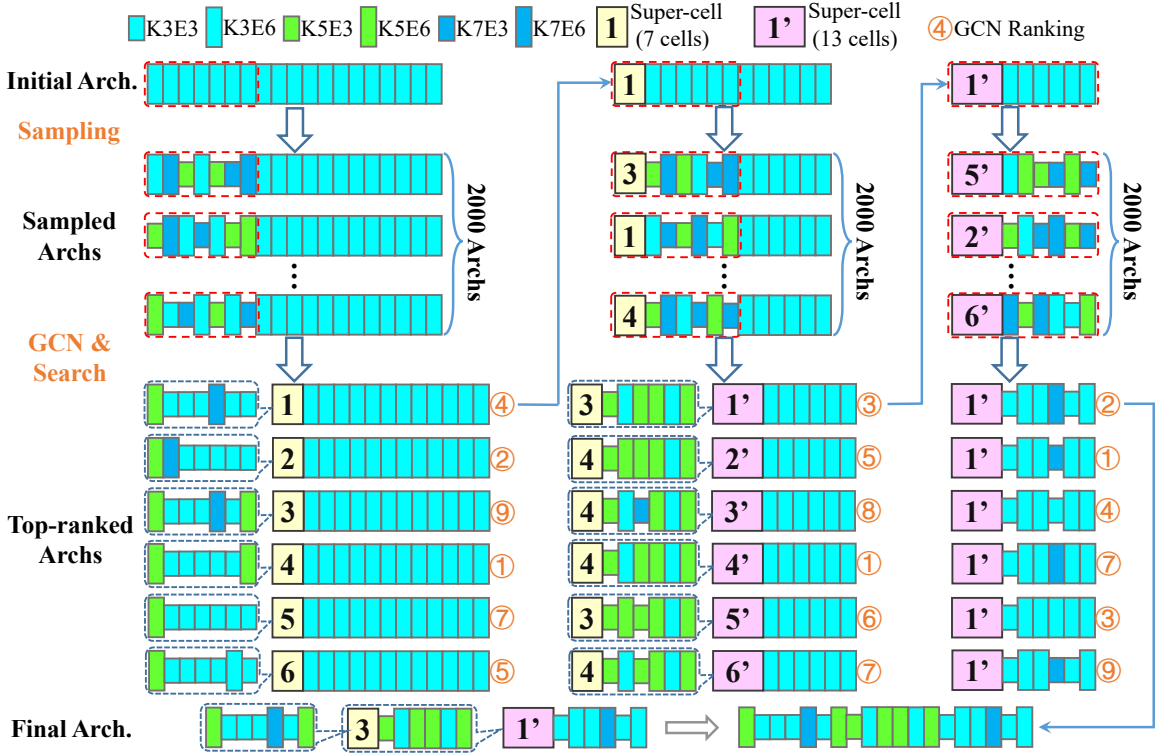


Figure 2: A typical process of how GCN assists to find the best architecture (best viewed in color). More detailed descriptions can be found in Section 10.

choosing a cell (building block) for each layer. Following FairNAS (Chu et al. 2019b), the candidate cells in each layer are always sampled with equal probability, regardless of the current status and how these cells contribute to network accuracy. This strategy is believed to improve the ‘fairness’ of sampling, *i.e.*, in the probabilistic viewpoint, all sub-networks have the equal chance of being sampled and optimized. Thus, the relative ranking among sub-networks is believed to be more accurate.

The length of the training stage is not very important, as we will show in the experimental section that the fitting degree of the super-network does not heavily impact the searched architecture. 150 epochs is often sufficient for super-network training, which takes around 13 hours on eight GPUs.

• **Step 1: Sub-network Sampling**

In the t -th round of iteration, given the set of layer indices to be searched, \mathcal{L}_t , we define a subspace of \mathcal{S} , denoted by \mathcal{S}_t . Let $\mathcal{V}_t = \mathcal{S}_t$, from which a small subset of sub-networks are uniformly sampled and evaluated. Let the number of sampled architectures be M , and the collected architecture-performance pairs be $\{(\mathbb{M}_m, z_{\mathbb{M}_m})\}_{m=1}^M$. Due to the potentially heavy computational overhead, M is often much smaller than the total number of sub-networks, $|\mathcal{V}_t|$.

• **Step 2: Training the GCN**

One of the key factors of GCN is named the adjacent matrix, \mathbf{A} , which defines how two nodes (*i.e.*, sub-networks) with an edge connection relate to each other and the way that we use the accuracy of one to predict that of another. For

simplicity, we simply define \mathbf{A} to be the similarity matrix, *i.e.*, the weight of each edge is determined by the similarity between two nodes it connects. The last factor of GCN is the feature representation of each node (sub-network). We simply follow the Gray coding scheme to encode each sub-network into a fixed-dimensional vector. More details about the graph construction can be found in the Arxiv version.

We follow a standard training procedure described in (Kipf and Welling 2017) to train the GCN, yet we modify the classification head of the network (applied to node features) into a regression module to predict sub-network accuracy. After GCN training is complete, we can predict the accuracy of each sub-network with its feature vector and the similarity matrix, \mathbf{A} , with a negligible cost.

• **Step 3: Updating the Optimal Architecture**

After the optimization is finished, the GCN provides a lookup table that requires negligible costs in estimating the performance of each sub-network in the space of \mathcal{S}_t . As we elaborate previously, although GCN is good at depicting the property of the overall search space, its prediction in each single sub-network is not always accurate. So, we enumerate over the search space, find a number of sub-networks with the best performance, and perform sub-network evaluation again to determine the best of them. Experiments show that this re-verification step, with negligible costs, brings a consistent accuracy gain of 0.1%–0.2%. Finally, the searched one is used to replace the corresponding layers of \mathbb{M}_t and \mathbb{M}_t becomes \mathbb{M}_{t+1} . If the search process does not terminate,

then we go back to Step 1 and repeat data collection, GCN training, and architecture update.

It is possible that the optimal sub-network found in \mathcal{S}_t does not lead to global optimality. To improve flexibility, we preserve K top-ranked sub-networks and feed them into the next round. To achieve this goal, the searched layers in this (t -th) round are not allowed to be searched in the next ($t + 1$ -st) round, *i.e.*, $\mathcal{L}_t \cap \mathcal{L}_{t+1} = \emptyset$. Hence, the next round can sample each of the K preserved sub-architectures as an entire – in other words, the cells in \mathcal{L}_t form a ‘super-cell’ in the next round, and each preserved architecture is a choice in the super-cell. We will see in experiments that this strategy improves the stability of search, *i.e.*, the chance of finding high-quality architectures is increased.

Experiments

All experiments are conducted on ILSVRC2012 (Rusakovsky et al. 2015). Additional hyper-parameter settings and implementation details can be found in the Arxiv version.

A Typical Search Procedure: How GCN Works?

Figure 2 shows a complete search process. The search starts with an initial model, in which all cells have a kernel size of 3 and an expansion ratio of 6, denoted by K3E6. After searching in the first subspace containing 7 cells, 6 top-ranked sub-architectures are preserved and sent into the next stage as a super-cell. This process continues two more times until the final architecture is obtained.

From this example, we provide some intermediate statistics to show how GCN assists in the architecture search process. In the first subspace, 2,000 sub-networks are sampled, with the best one (denoted by $\tilde{\mathbb{M}}_0$) reporting an accuracy of 81.25% (over the sampled 100 classes). However, after GCN is trained and 100 best sub-networks are chosen from its prediction, \mathbb{M}_0 does not appear in these 100 sub-networks. Yet, when these 100 sub-networks are sent into sub-network sampling, the best one (\mathbb{M}_0^*) reports an accuracy of 81.11% which is slightly lower than that of $\tilde{\mathbb{M}}_0$. But, this does not mean that $\tilde{\mathbb{M}}_0$ is indeed better – it is the noise introduced by sub-network sampling that surpasses the difference between their actual performance. To verify this, we perform another NAS process that does not involve using GCN for prediction but simply preserves the top-ranked architectures from sub-network sampling. This strategy reports a final accuracy (over 1,000 classes) of 75.2% (with 438M multi-adds), which is inferior to that (75.5%) with GCN applied.

Two more evidences verify the effectiveness of GCN. On the one hand, GCN is able to fit the data collected by sub-network sampling. In three search segments, using 200 left-out architecture-performance pairs for validation, the Kendall- τ coefficient (ranged in $[-1, 1]$ with 0 implying random guess) is always higher than 0.5, *i.e.*, the relative ranking of more than 75% pairs is consistent between sub-network sampling and GCN prediction. On the other hand, GCN prediction actually works better than sub-network sampling. We choose 8 architectures from the third subspace and evaluate their real performance by training them from scratch. The

Kendall- τ coefficient between the real performance and sub-network sampling is 0.2143 (still close to random guess), but the coefficient is improved to 0.6429 by GCN.

Furthermore, we perform linear regression on Acc_{150} , Acc_{160} and the GCN-predicted one, Acc_{GCN} , and re-plot corresponding dots and curves in the Arxiv version to show the results that GCN helps to improve the ranking consistency. The regression scores⁴ of $\text{Acc}_{150}^{\text{reg}}$, $\text{Acc}_{160}^{\text{reg}}$ and $\text{Acc}_{GCN}^{\text{reg}}$ are 0.07, 0.04 0.62, respectively. The large difference on regression scores between the super-network-evaluated accuracy and the GCN-predicted one is a proof that the random error term in our key assumption is largely suppressed.

Comparison to the State-of-the-Arts

The performance comparison on ImageNet classification with state-of-the-art architectures is listed in Table ?? . For a fair comparison, we do not list architectures that are either trained with additional data augmentations (*e.g.*, AutoAugment (Cubuk et al. 2019)) or equipped with extra architecture modifications (*e.g.*, MobileNet-v3 (Howard et al. 2019) introduced H-Swish and reported a top-1 accuracy of 75.2%). With a comparable amount of multi-add operations, the performance of our approach is on par with state-of-the-art methods that searched on a chain-styled search space. Our searched architecture enjoys a superior test accuracy than FairNAS (Chu et al. 2019b) due to a better fitting to the search space, which is explained in Section 10. Compared to those architectures searched on the DARTS-based search space, our discovered architecture achieves comparable performance while enjoys a smaller model size and multi-adds count. Notably, with nearly 20M less multi-adds, the performance of our searched architecture is comparable to that of MnasNet-A3, which was searched by sampling and evaluating architectures by training-from-scratch and cost thousands of GPU-days.

Diagnostic Studies

Relationship to the Overall Quality of the Super-network

The length of one-shot training is worth discussion. For example, training the super-network throughout 150 epochs achieves a 86.17% training accuracy, and increasing the number of epochs to 450 improves the accuracy to 94.67%, which potentially leads to a higher quality of sub-networks. However, such improvement does not necessarily cause the final architecture to be better, implying that the architecture-related random noise, $\epsilon_{\mathbb{M}}$, is still a major concern.

The Number of Sampled Architectures

The number of architectures sampled for each round is a critical factor related to the search cost since the graph construction and GCN training only take a few minutes. We have tested two different settings of $M = 2,000$ and $M = 5,000$. The results are listed in Table ?? and no obvious difference is observed on both Kendall- τ coefficients and the evaluation

⁴The regression score is a criterion to measure the determination of prediction. Larger regression score is better.

Architecture	Test Acc. (%)		Params (M)	$\times+$ (M)	Search Cost (GPU-days)	Search Method
	top-1	top-5				
Inception-v1 (Szegedy et al. 2015)	69.8	89.9	6.6	1448	-	manual
NASNet-A (Zoph et al. 2018)	74.0	91.6	5.3	564	1800	RL
AmoebaNet-C (Real et al. 2018)	75.7	92.4	6.4	570	3150	EA
DARTS (second order) (Liu, Simonyan, and Yang 2018)	73.3	91.3	4.7	574	4.0	GB
P-DARTS (CIFAR10) (Chen et al. 2019)	75.6	92.6	4.9	557	0.3	GB
PC-DARTS (ImageNet) (Xu et al. 2020)	75.8	92.7	5.3	597	3.8	GB
MnasNet-A3 [†] (Tan et al. 2019)	76.7	93.3	5.2	403	1667 [‡]	RL
ProxylessNAS (GPU) (Cai, Zhu, and Han 2018)	75.1	92.5	7.1	465	8.3	GB
FFairNAS-A (Chu et al. 2019b)	75.3	92.4	4.6	388	12	EA
One-Shot-NAS-GCN (ours)	75.5	92.7	4.4	383	4.7	GCN
One-Shot-NAS-GCN [†] (ours)	76.6	93.1	4.6	384	4.7	GCN

Table 1: Comparison with state-of-the-art architectures on ImageNet (mobile setting). [†]: SE module included, [‡]: estimated by (Cai, Zhu, and Han 2018), GB: gradient-based, BO: Bayesian Optimization.

# Archs	Sim.	τ_1	τ_2	τ_3	Test Acc. (%)
2,000	A	0.76	0.58	0.51	75.48
2,000	M	0.74	0.60	0.52	75.47
5,000	A	0.74	0.60	0.50	75.47
5,000	M	0.75	0.62	0.53	75.45

Table 2: Results for different number of architectures sampled for each round and similarity type. τ_t denotes the Kendall- τ coefficient of the t -th round. A: Assigned; M: Measured.

accuracy of the discovered architectures. It takes around 9 hours on a single GPU to finish the search process when $M = 2,000$, while the search cost is increased to about 21 GPU-hours when we use the $M = 5,000$ setting. Thus, we adopt the former setting for most of the experiments.

The Similarity Measurement in GCN

We have compared the influence of the way to define the similarity matrix and demonstrate the results in Table ?? . Due to the fact that both types of similarity definitions involve approximation, there is no obvious difference between them.

Preserving 1 vs. 6 Candidates

Preserving multiple candidates from the previous round is important because more candidates can increase the probability that the most promising sub-architectures are included in the final searched architecture. We conduct experiments to validate it. We run the search process for 3 times for both the settings of preserving 1 candidate and 6 candidates from the previous round. The classification accuracy of the former setting is $75.34\% \pm 0.08\%$, while the latter is $75.47\% \pm 0.01\%$, which supports our previous analysis.

Flexible NAS with Hardware Constraints

Thanks to the prediction ability of the trained GCN, we can easily build a lookup table that contains the predicted performance of all sub-networks in the search space (or subspace), upon which architectures that satisfy different hardware constraints can be selected. Without such a lookup table, the

target architecture must be included in the explored architectures, which is rigid and inconvenient when multiple architectures with different constraints are required. We pick multiple architectures with different multi-adds levels by filtering out the top-ranked architectures that meet the target constraint and selecting the top-1 according to the sampling and re-evaluation ranking. The selected architectures achieve 75.6%, 75.5% and 75.4% top-1 accuracies with 393M, 383M and 360M multi-adds, respectively. Under target hardware constraints, our searched architectures keep comparable performance with state-of-the-art architectures under similar constraints. In addition, one can freely generalize our approach to other types of hardware constraints, *e.g.*, network latency.

Conclusions

This paper introduces a novel idea that uses graph convolutional network to assist weight-sharing neural architecture search. The most important opinion is that there is an inevitable random noise between a well-trained super-network and the sub-networks sampled from it, and GCN, by averaging over the entire search space, can eliminate the error systematically and avoid the architecture search from falling into local minima. Experiments demonstrate the effectiveness and efficiency of our approach, in particular in the scenarios with additional hardware constraints.

Our research paves the way of projecting the architectures in the search space into another low-dimensional space. This is a new direction which may provide new insights to neural architecture search, but currently, there are still some important issues that remain uncovered. For example, it is well known that fine-tuning each sampled sub-network can improve the accuracy of estimation, but it requires considerable computation – this is a tradeoff between accuracy and efficiency. We can perform fine-tuning on a small number of sub-networks and assign the remaining ones to be evaluated by sub-network evaluation. The property of GCN in such a heterogeneous graph is worth investigating, which, as well as other topic, will be left for future work.

Acknowledgments

This work was supported by National Key R&D Program of China under Grant 2020YFA0711400, and the National Natural Science Foundation of China under Grant Nos. 61831018 and 61631017. This work was also supported by National Natural Science Foundation of China 61702130 Guangxi Natural Science Foundation 2020GXNSFAA159137. We thank Dr. Song Bai and Zhengsu Chen for their valuable suggestions.

References

- Bi, K.; Xie, L.; Chen, X.; Wei, L.; and Tian, Q. 2020. Gold-nas: Gradual, one-level, differentiable. *arXiv preprint arXiv:2007.03331* .
- Brock, A.; Lim, T.; Ritchie, J. M.; and Weston, N. 2017. SMASH: one-shot model architecture search through hypernetworks. *arXiv preprint arXiv:1708.05344* .
- Cai, H.; Chen, T.; Zhang, W.; Yu, Y.; and Wang, J. 2018. Efficient architecture search by network transformation. In *AAAI*.
- Cai, H.; Zhu, L.; and Han, S. 2018. ProxylessNAS: Direct Neural Architecture Search on Target Task and Hardware. *arXiv preprint arXiv:1812.00332* .
- Chen, X.; Xie, L.; Wu, J.; and Tian, Q. 2019. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *ICCV*.
- Chen, X.; Xie, L.; Wu, J.; Wei, L.; Xu, Y.; and Tian, Q. 2020. Fitting the search space of weight-sharing nas with graph convolutional networks. *arXiv preprint arXiv:2004.08423* .
- Chu, X.; Zhang, B.; Ma, H.; Xu, R.; Li, J.; and Li, Q. 2019a. Fast, Accurate and Lightweight Super-Resolution with Neural Architecture Search. *arXiv preprint arXiv:1901.07261* .
- Chu, X.; Zhang, B.; Xu, R.; and Li, J. 2019b. FairNAS: Rethinking evaluation fairness of weight sharing neural architecture search. *arXiv preprint arXiv:1907.01845* .
- Cubuk, E. D.; Zoph, B.; Mane, D.; Vasudevan, V.; and Le, Q. V. 2019. Autoaugment: Learning augmentation strategies from data. In *CVPR*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*.
- Dong, X.; and Yang, Y. 2019. Searching for A Robust Neural Architecture in Four GPU Hours. In *CVPR*.
- Dong, X.; and Yang, Y. 2020. NAS-Bench-102: Extending the Scope of Reproducible Neural Architecture Search. *arXiv preprint arXiv:2001.00326* .
- Elsken, T.; Metzen, J. H.; and Hutter, F. 2019. Neural Architecture Search: A Survey. *Journal of Machine Learning Research* 20(55): 1–21.
- Friede, D.; Lukasik, J.; Stuckenschmidt, H.; and Keuper, M. 2019. A Variational-Sequential Graph Autoencoder for Neural Architecture Performance Prediction. *arXiv preprint arXiv:1912.05317* .
- Guo, Z.; Zhang, X.; Mu, H.; Heng, W.; Liu, Z.; Wei, Y.; and Sun, J. 2019. Single path one-shot neural architecture search with uniform sampling. *arXiv preprint arXiv:1904.00420* .
- Howard, A.; Sandler, M.; Chu, G.; Chen, L.-C.; Chen, B.; Tan, M.; Wang, W.; Zhu, Y.; Pang, R.; Vasudevan, V.; et al. 2019. Searching for mobilenetv3. In *ICCV*.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- Liu, C.; Zoph, B.; Neumann, M.; Shlens, J.; Hua, W.; Li, L.-J.; Fei-Fei, L.; Yuille, A.; Huang, J.; and Murphy, K. 2018. Progressive neural architecture search. In *ECCV*.
- Liu, H.; Simonyan, K.; and Yang, Y. 2018. DARTS: Differentiable architecture search. *arXiv preprint arXiv:1806.09055* .
- Luo, R.; Tian, F.; Qin, T.; Chen, E.; and Liu, T.-Y. 2018. Neural architecture optimization. In *NeurIPS*.
- Ning, X.; Zheng, Y.; Zhao, T.; Wang, Y.; and Yang, H. 2020. A Generic Graph-based Neural Architecture Encoding Scheme for Predictor-based NAS. *arXiv preprint arXiv:2004.01899* .
- Pham, H.; Guan, M. Y.; Zoph, B.; Le, Q. V.; and Dean, J. 2018. Efficient neural architecture search via parameter sharing. In *ICML*.
- Real, E.; Aggarwal, A.; Huang, Y.; and Le, Q. V. 2018. Regularized evolution for image classifier architecture search. *arXiv preprint arXiv:1802.01548* .
- Real, E.; Moore, S.; Selle, A.; Saxena, S.; Suematsu, Y. L.; Tan, J.; Le, Q. V.; and Kurakin, A. 2017. Large-scale evolution of image classifiers. In *ICML*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. ImageNet large scale visual recognition challenge. *IJCV* 115(3): 211–252.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*.
- Shi, H.; Pi, R.; Xu, H.; Li, Z.; Kwok, J.; and Zhang, T. 2020. Bridging the Gap between Sample-based and One-shot Neural Architecture Search with BONAS. In *NeurIPS*.
- 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 *CVPR*.
- Tan, M.; Chen, B.; Pang, R.; Vasudevan, V.; Sandler, M.; Howard, A.; and Le, Q. V. 2019. Mnasnet: Platform-aware neural architecture search for mobile. In *CVPR*.
- Wen, W.; Liu, H.; Chen, Y.; Li, H.; Bender, G.; and Kindermans, P.-J. 2020. Neural predictor for neural architecture search. In *ECCV*.
- Wu, B.; Dai, X.; Zhang, P.; Wang, Y.; Sun, F.; Wu, Y.; Tian, Y.; Vajda, P.; Jia, Y.; and Keutzer, K. 2019. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *CVPR*.

Xie, L.; Chen, X.; Bi, K.; Wei, L.; Xu, Y.; Chen, Z.; Wang, L.; Xiao, A.; Chang, J.; Zhang, X.; et al. 2020. Weight-Sharing Neural Architecture Search: A Battle to Shrink the Optimization Gap. *arXiv preprint arXiv:2008.01475* .

Xie, L.; and Yuille, A. 2017. Genetic CNN. In *ICCV*.

Xie, S.; Zheng, H.; Liu, C.; and Lin, L. 2018. SNAS: Stochastic neural architecture search. *arXiv preprint arXiv:1812.09926* .

Xu, Y.; Xie, L.; Zhang, X.; Chen, X.; Qi, G.-J.; Tian, Q.; and Xiong, H. 2020. PC-DARTS: Partial Channel Connections for Memory-Efficient Architecture Search. In *ICLR*.

Ying, C.; Klein, A.; Christiansen, E.; Real, E.; Murphy, K.; and Hutter, F. 2019. NAS-Bench-101: Towards Reproducible Neural Architecture Search. In *ICML*.

Zoph, B.; and Le, Q. V. 2017. Neural architecture search with reinforcement learning. In *ICLR*.

Zoph, B.; Vasudevan, V.; Shlens, J.; and Le, Q. V. 2018. Learning transferable architectures for scalable image recognition. In *CVPR*.