

Binarized Neural Architecture Search

Hanlin Chen,¹ Li'an Zhuo,¹ Baochang Zhang,^{1*}
 Xiawu Zheng,² Jianzhuang Liu,⁴ David Doermann,³ Rongrong Ji²

¹Beihang University, ²Xiamen University, ³University at Buffalo

⁴Shenzhen Institutes of Advanced Technology, University of Chinese Academy of Sciences
 {hlchen, bczhang}@buaa.edu.cn

Abstract

Neural architecture search (NAS) can have a significant impact in computer vision by automatically designing optimal neural network architectures for various tasks. A variant, binarized neural architecture search (BNAS), with a search space of binarized convolutions, can produce extremely compressed models. Unfortunately, this area remains largely unexplored. BNAS is more challenging than NAS due to the learning inefficiency caused by optimization requirements and the huge architecture space. To address these issues, we introduce channel sampling and operation space reduction into a differentiable NAS to significantly reduce the cost of searching. This is accomplished through a performance-based strategy used to abandon less potential operations. Two optimization methods for binarized neural networks are used to validate the effectiveness of our BNAS. Extensive experiments demonstrate that the proposed BNAS achieves a performance comparable to NAS on both CIFAR and ImageNet databases. An accuracy of 96.53% vs. 97.22% is achieved on the CIFAR-10 dataset, but with a significantly compressed model, and a 40% faster search than the state-of-the-art PC-DARTS.

Introduction

Neural architecture search (NAS) have attracted great attention with remarkable performance in various deep learning tasks. Impressive results have been shown for reinforcement learning (RL) based methods (Zoph et al. 2018; Zoph and Le 2016), for example, which train and evaluate more than 20,000 neural networks across 500 GPUs over 4 days. Recent methods like differentiable architecture search (DARTs) reduce the search time by formulating the task in a differentiable manner (Liu, Simonyan, and Yang 2018). DARTS relaxes the search space to be continuous, so that the architecture can be optimized with respect to its validation set performance by gradient descent, which provides a fast solution for effective network architecture search. To reduce the redundancy in the network space, partially-connected DARTs (PC-DARTs) was recently introduced to perform

a more efficient search without compromising the performance of DARTS (Xu et al. 2019).

Although the network optimized by DARTS or its variants has a smaller model size than traditional light models, the searched network still suffers from an inefficient inference process due to the complicated architectures generated by multiple stacked full-precision convolution operations. Consequently, the adaptation of the searched network to an embedded device is still computationally expensive and inefficient. Clearly the problem requires further exploration to overcome these challenges.

One way to address these challenges is to transfer the NAS to a binarized neural architecture search (BNAS), by exploring the advantages of binarized neural networks (BNNs) on memory saving and computational cost reduction. Binarized filters have been used in traditional convolutional neural networks (CNNs) to compress deep models (Rastegari et al. 2016a; Courbariaux et al. 2016; Courbariaux, Bengio, and David 2015; Xu, Boddeti, and Savvides 2016), showing up to 58-time speedup and 32-time memory saving. In (Xu, Boddeti, and Savvides 2016), the XNOR network is presented where both the weights and inputs attached to the convolution are approximated with binary values. This results in an efficient implementation of convolutional operations by reconstructing the unbinarized filters with a single scaling factor. In (Gu et al. 2019), a projection convolutional neural network (PCNN) is proposed to realize BNNs based on a simple back propagation algorithm. In our BNAS framework, we re-implement XNOR and PCNN for the effectiveness validation. We show that the BNNs obtained by BNAS can outperform conventional models by a large margin. It is a significant contribution in the field of BNNs, considering that the performance of conventional BNNs are not yet comparable with their corresponding full-precision models in terms of accuracy.

The search process of our BNAS consists of two steps. One is the operation potential ordering based on partially-connected DARTs (PC-DARTs) (Xu et al. 2019) which serves as a baseline for our BNAS. It is further sped up with a second operation reduction step guided by a performance-based strategy. In the operation reduction step, we prune one operation at each iteration from one-half of the operations

*Baochang Zhang is the corresponding author.
 Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

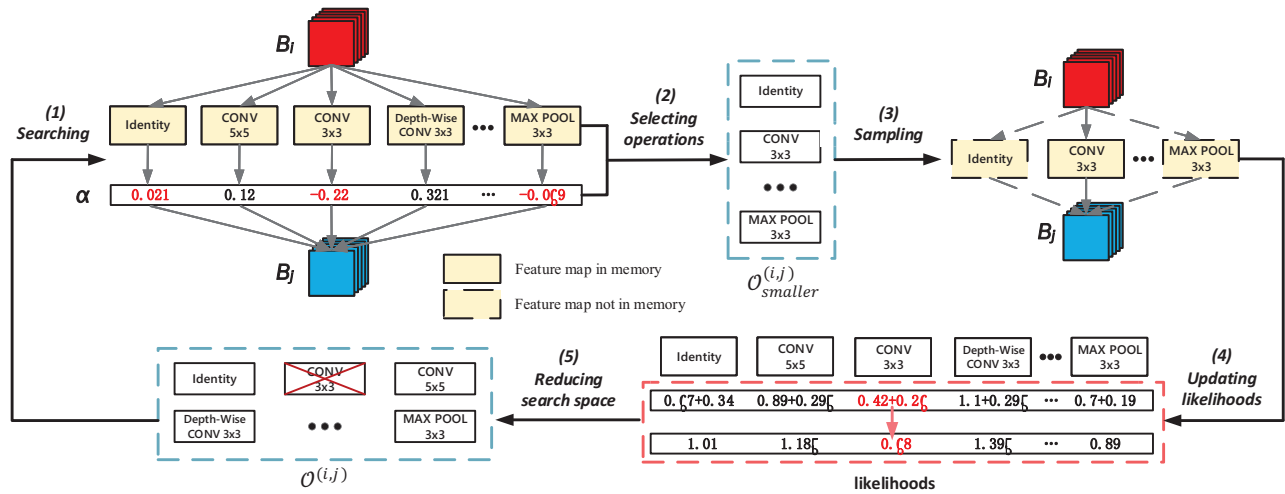


Figure 1: The main steps of our BNAS: (1) Search an architecture based on $\mathcal{O}^{(i,j)}$ using PC-DARTS. (2) Select half the operations with less potential from $\mathcal{O}^{(i,j)}$ for each edge, resulting in $\mathcal{O}_{smaller}^{(i,j)}$. (3) Select an architecture by sampling (without replacement) one operation from $\mathcal{O}_{smaller}^{(i,j)}$ for every edge, and then train the selected architecture. (4) Update the operation selection likelihood $s(o_k^{(i,j)})$ based on the accuracy obtained from the selected architecture on the validation data. (5) Abandon the operation with the minimal selection likelihood from the search space $\{\mathcal{O}^{(i,j)}\}$ for every edge.

with less potential as calculated by PC-DARTS. As such, the optimization of the two steps becomes faster and faster because the search space is reduced due to the operation pruning. We can take advantage of the differential framework of DARTS where the search and performance evaluation are in the same setting. We also enrich the search strategy of DARTS. Not only is the gradient used to determine which operation is better, but the proposed performance evaluation is included for further reduction of the search space. In this way BNAS is fast and well built. The contributions of our paper include:

- BNAS is developed based on a new search algorithm which solves the BNNs optimization and architecture search in a unified framework.
- The search space is greatly reduced through a performance-based strategy used to abandon operations with less potential, which improves the search efficiency by 40%.
- Extensive experiments demonstrate that the proposed algorithm achieves much better performance than other light models on CIFAR-10 and ImageNet.

Related Work

Thanks to the rapid development of deep learning, significant gains in performance have been realized in a wide range of computer vision tasks, most of which are manually designed network architectures (Krizhevsky, Sutskever, and Hinton 2012; Simonyan and Zisserman 2014; He et al. 2016; Huang et al. 2017). Recently, the new approach called neural architecture search (NAS) has been attracting increased attention. The goal is to find automatic ways of designing neural architectures to replace conventional hand-crafted

ones. Existing NAS approaches need to explore a very large search space and can be roughly divided into three type of approaches: evolution-based, reinforcement-learning-based and one-shot-based.

In order to implement the architecture search within a short period of time, researchers try to reduce the cost of evaluating each searched candidate. Early efforts include sharing weights between searched and newly generated networks (Cai et al. 2018a). Later, this method was generalized into a more elegant framework named one-shot architecture search (Brock et al. 2017; Cai, Zhu, and Han 2018; Liu, Simonyan, and Yang 2018; Pham et al. 2018; Xie et al. 2018). In these approaches, an over-parameterized network or super network covering all candidate operations is trained only once, and the final architecture is obtained by sampling from this super network. For example, Brock et al. (Brock et al. 2017) trained the over-parameterized network using a HyperNet (Ha, Dai, and V. Le 2016), and Pham et al. (Pham et al. 2018) proposed to share parameters among child models to avoid retraining each candidate from scratch. The paper (Liu et al. 2017) is based on DARTS, which introduces a differentiable framework and thus combines the search and evaluation stages into one. Despite its simplicity, researchers have found some of its drawbacks and proposed a few improved approaches over DARTS (Cai, Zhu, and Han 2018; Xie et al. 2018; Chen et al. 2019).

Unlike previous methods, we study BNAS based on efficient operation reduction. We prune one operation at each iteration from one-half of the operations with smaller weights calculated by PC-DARTS, and the search becomes faster and faster in the optimization.

Binarized Neural Architecture Search

In this section, we first describe the search space in a general form, where the computation procedure for an architecture (or a cell in it) is represented as a directed acyclic graph. We then review the baseline PC-DARTS (Xu et al. 2019), which improves the memory efficiency, but is not enough for BNAS. Finally, an operation sampling and a performance-based search strategy are proposed to effectively reduce the search space. Our BNAS framework is shown in Fig. 1 and additional details of which are described in the rest of this section.

Search Space

Following Zoph et al. (2018); Real et al. (2018); Liu et al. (2018a;b), we search for a computation cell as the building block of the final architecture. A network consists of a pre-defined number of cells (Zoph and Le 2016), which can be either normal cells or reduction cells. Each cell takes the outputs of the two previous cells as input. A cell is a fully-connected directed acyclic graph (DAG) of M nodes, *i.e.*, $\{B_1, B_2, \dots, B_M\}$, as illustrated in Fig. 2(a). Each node B_i takes its dependent nodes as input, and generates an output through a sum operation $B_j = \sum_{i < j} o^{(i,j)}(B_i)$. Here each node is a specific tensor (*e.g.*, a feature map in convolutional neural networks) and each directed edge (i, j) between B_i and B_j denotes an operation $o^{(i,j)}(\cdot)$, which is sampled from $\mathcal{O}^{(i,j)} = \{o_1^{(i,j)}, \dots, o_K^{(i,j)}\}$. Note that the constraint $i < j$ ensures there are no cycles in a cell. Each cell takes the outputs of two dependent cells as input, and we define the two input nodes of a cell as B_{-1} and B_0 for simplicity. Following (Liu, Simonyan, and Yang 2018), the set of the operations \mathcal{O} consists of $K = 8$ operations. They include 3×3 max pooling, no connection (zero), 3×3 average pooling, skip connection (identity), 3×3 dilated convolution with rate 2, 5×5 dilated convolution with rate 2, 3×3 depth-wise separable convolution, and 5×5 depth-wise separable convolution, as illustrated in Fig. 2(b). The search space of a cell is constructed by the operations of all the edges, denoted as $\{\mathcal{O}^{(i,j)}\}$.

Unlike conventional convolutions, our BNAS is achieved by transforming all the convolutions in \mathcal{O} to binarized convolutions. We denote the full-precision and binarized kernels as X and \hat{X} respectively. A convolution operation in \mathcal{O} is represented as $B_j = B_i \otimes \hat{X}$ as shown in Fig. 2(b), where \otimes denotes convolution. To build BNAS, one key step is how to binarize the kernels from X to \hat{X} , which can be implemented based on state-of-the-art BNNs, such as XNOR or PCNN. As we know, the optimization of BNNs is more challenging than that of conventional CNNs (Gu et al. 2019; Rastegari et al. 2016b), which adds an additional burden to NAS. To solve it, we introduce channel sampling and operation space reduction into differentiable NAS to significantly reduce the cost of GPU hours, leading to an efficient BNAS.

PC-DARTS

The core idea of PC-DARTS is to take advantage of partial channel connections to improve memory efficiency. Tak-

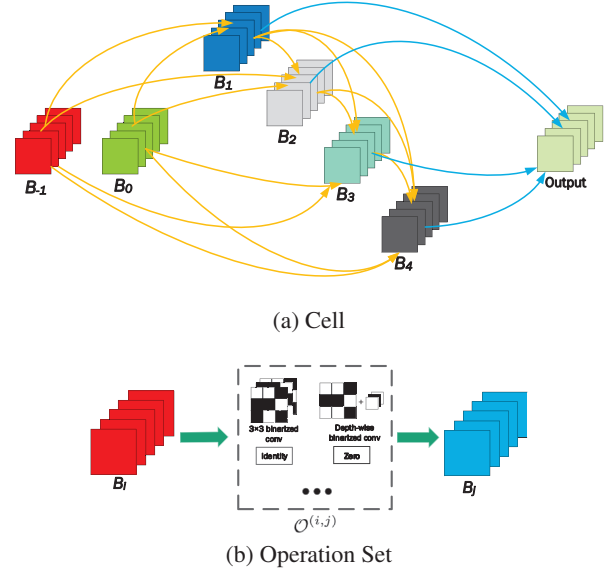


Figure 2: (a) A cell contains 7 nodes, two input nodes B_{-1} and B_0 , four intermediate nodes B_1, B_2, B_3, B_4 that apply sampled operations on the input nodes and upper nodes, and an output node that concatenates the outputs of the four intermediate nodes. (b) The set of operations $\mathcal{O}^{(i,j)}$ between B_i and B_j , including binarized convolutions.

ing the connection from B_i to B_j for example, this involves defining a channel sampling mask $S^{(i,j)}$, which assigns 1 to selected channels and 0 to masked ones. The selected channels are sent to a mixed computation of $|\mathcal{O}^{(i,j)}|$ operations, while the masked ones bypass these operations. They are directly copied to the output, which is formulated as:

$$\begin{aligned}
 & f^{(i,j)}(B_i, S^{(i,j)}) \\
 &= \sum_{o_k^{(i,j)} \in \mathcal{O}^{(i,j)}} \frac{\exp\{\alpha_{o_k^{(i,j)}}\}}{\sum_{o_{k'}^{(i,j)} \in \mathcal{O}^{(i,j)}} \exp\{\alpha_{o_{k'}^{(i,j)}}\}} \cdot o_k^{(i,j)}(S^{(i,j)} * B_i) \\
 &+ (1 - S^{(i,j)}) * B_i,
 \end{aligned} \tag{1}$$

where $S^{(i,j)} * B_i$ and $(1 - S^{(i,j)}) * B_i$ denote the selected and masked channels, respectively, and $\alpha_{o_k^{(i,j)}}$ is the parameter of operation $o_k^{(i,j)}$ between B_i and B_j .

PC-DARTS sets the proportion of selected channels to $1/C$ by regarding C as a hyper-parameter. In this case, the computation cost can also be reduced by C times. However, the size of the whole search space is $2 \times K^{|\mathcal{E}_{\mathcal{M}}|}$, where $\mathcal{E}_{\mathcal{M}}$ is the set of possible edges with M intermediate nodes in the fully-connected DAG, and the "2" comes from the two types of cells. In our case with $M = 4$, together with the two input nodes, the total number of cell structures in the search space is $2 \times 8^{2+3+4+5} = 2 \times 8^{14}$. This is an extremely large space to search for a binarized neural architectures which need more time than a full-precision NAS. Therefore, efficient optimization strategies for BNAS are required.

Sampling for BNAS

For BNAS, PC-DARTS is still time and memory consuming because of the large search space, although it is already faster than most of existing NAS methods. We introduce another approach to increasing memory efficiency by reducing the search space $\{\mathcal{O}^{(i,j)}\}$. According to $\{\alpha_{o_k^{(i,j)}}\}$, we can select half the operations with less potential from $\mathcal{O}^{(i,j)}$ for each edge, resulting in $\mathcal{O}_{smaller}^{(i,j)}$. We then sample an operation from $\mathcal{O}_{smaller}^{(i,j)}$ for each edge guided by a performance-based strategy proposed in the next section in order to reduce the search space. We follow the rule of sampling without replacement $K/2$ times. Here sampling without replacement means that after one operation is sampled randomly from $\mathcal{O}_{smaller}^{(i,j)}$, this operation is removed from $\mathcal{O}_{smaller}^{(i,j)}$. For convenience of description, the $K/2$ operations in each edge are transformed to a one-hot indicator vector. In other words we sample only one operation according to the performance-based strategy, which effectively reduces the memory cost compared with PC-DARTS (Xu et al. 2019).

Performance-based Strategy for BNAS

Reinforcement learning is inefficient in the architecture search due to the delayed rewards in network training, i.e., the evaluation of a structure is usually done after the network training converges. On the other hand, we can perform the evaluation of a cell when training the network. Inspired by (Ying et al. 2019), we use a performance-based strategy to boost the search efficiency by a large margin. Ying *et al.* (Ying et al. 2019) did a series of experiments showing that in the early stage of training, the validation accuracy ranking of different network architectures is not a reliable indicator of the final architecture quality. However, we observe that the experiment results actually suggest a nice property that if an architecture performs badly in the beginning of training, there is little hope that it can be part of the final optimal model. As the training progresses, this observation shows less uncertainty. Based on this observation, we derive a simple yet effective operation abandoning process. During training, along with the increasing epochs, we progressively abandon the worst performing operation in each edge.

To this end, we randomly sample one operation from the $K/2$ operations in $\mathcal{O}_{smaller}^{(i,j)}$ for every edge, then obtain the validation accuracy by training the sampled network for one epoch, and finally assign this accuracy to all the sampled operations. These three steps are performed $K/2$ times by sampling without replacement, leading to each operation having exactly one accuracy for every edge.

We repeat it T times. Thus each operation for every edge has T accuracies $\{y_{k,1}^{(i,j)}, y_{k,2}^{(i,j)}, \dots, y_{k,T}^{(i,j)}\}$. Then we define the selection likelihood of the k th operation in $\mathcal{O}_{smaller}^{(i,j)}$ for each edge as:

$$s_{smaller}(o_k^{(i,j)}) = \frac{\exp\{\bar{y}_k^{(i,j)}\}}{\sum_m \exp\{\bar{y}_m^{(i,j)}\}}, \quad (2)$$

where $\bar{y}_k^{(i,j)} = \frac{1}{T} \sum_t y_{k,t}^{(i,j)}$. And the selection likelihoods of

the other operations not in $\mathcal{O}_{smaller}^{(i,j)}$ are defined as:

$$s_{larger}(o_k^{(i,j)}) = \frac{1}{2} (\max_{o_k^{(i,j)}} \{s_{smaller}(o_k^{(i,j)})\} + \frac{1}{\lceil K/2 \rceil} \sum_{o_k^{(i,j)}} s_{smaller}(o_k^{(i,j)})), \quad (3)$$

where $\lceil K/2 \rceil$ denotes the smallest integer $\geq K/2$. The reason to use it is because K can be an odd integer during iteration in the proposed Algorithm 1. Eq. 3 is an estimation for the rest operations using a value balanced between the maximum and average of $s_{smaller}(o_k^{(i,j)})$. Then, $s(o_k^{(i,j)})$ is updated by:

$$s(o_k^{(i,j)}) \leftarrow \frac{1}{2} s(o_k^{(i,j)}) + q_k^{(i,j)} s_{smaller}(o_k^{(i,j)}) + (1 - q_k^{(i,j)}) s_{larger}(o_k^{(i,j)}), \quad (4)$$

where $q_k^{(i,j)}$ is a mask, which is 1 for the operations in $\mathcal{O}_{smaller}^{(i,j)}$ and 0 for the others.

Finally, we abandon the operation with the minimal selection likelihood for each edge. Such that the search space size is significantly reduced from $2 \times |\mathcal{O}^{(i,j)}|^{14}$ to $2 \times (|\mathcal{O}^{(i,j)}| - 1)^{14}$. We have:

$$\mathcal{O}^{(i,j)} \leftarrow \mathcal{O}^{(i,j)} - \{\arg \min_{o_k^{(i,j)}} s(o_k^{(i,j)})\}. \quad (5)$$

The optimal structure is obtained when there is only one operation left in each edge. Our performance-based search algorithm is presented in Algorithm 1. Note that in line 1, PC-DARTS is performed for L epochs as the warm-up to find an initial architecture, and line 14 is used to update the architecture parameters $\alpha_{o_k^{(i,j)}}$ for all the edges due to the reduction of the search space $\{\mathcal{O}^{(i,j)}\}$.

Optimization for BNAS

In this paper, the binarized kernel weights are computed based on XNOR (Rastegari et al. 2016b) or PCNN (Gu et al. 2019). Both methods are easily implemented in our BNAS framework, and the source code will be publicly available soon.

Binarizing CNNs, to the best of our knowledge, shares the same implementation framework. Without loss of generality, at layer l , let D_i^l be the direction of a full-precision kernel X_i^l , and A^l be the shared amplitude. For the binarized kernel \hat{X}_i^l corresponding to X_i^l , we have $\hat{X}_i^l = A^l \odot D_i^l$, where \odot denotes the element-wise multiplication between two matrices. We then employ an amplitude loss function to reconstruct the full-precision kernels as:

$$L_A = \frac{\theta}{2} \sum_{i,l} \|X_i^l - \hat{X}_i^l\|^2 = \frac{\theta}{2} \sum_{i,l} \|X_i^l - A^l \odot D_i^l\|^2, \quad (6)$$

where $D_i^l = \text{sign}(X_i^l)$. The element-wise multiplication combines the binarized kernels and the amplitude matrices to approximate the full-precision kernels. The amplitudes A^l

Algorithm 1: Performance-Based Search

Input: Training data, Validation data, Searching hyper-graph: \mathcal{G} , $K = 8$, $s(o_k^{(i,j)}) = 0$ for all edges;
Output: Optimal structure α ;

- 1 Search an architecture for L epochs based on $\mathcal{O}^{(i,j)}$ using PC-DARTS;
- 2 **while** ($K > 1$) **do**
- 3 Select $\mathcal{O}_{smaller}^{(i,j)}$ consisting of $\lceil K/2 \rceil$ operations with smallest $\alpha_{o_k}^{(i,j)}$ from $\mathcal{O}^{(i,j)}$ for every edge;
- 4 **for** $t = 1, \dots, T$ epoch **do**
- 5 $\mathcal{O}'_{smaller}^{(i,j)} \leftarrow \mathcal{O}_{smaller}^{(i,j)}$;
- 6 **for** $e = 1, \dots, \lceil K/2 \rceil$ epoch **do**
- 7 Select an architecture by sampling (without replacement) one operation from $\mathcal{O}'_{smaller}^{(i,j)}$ for every edge;
- 8 Train the selected architecture and get the accuracy on the validation data;
- 9 Assign this accuracy to all the sampled operations;
- 10 **end**
- 11 **end**
- 12 Update $s(o_k^{(i,j)})$ using Eq. 4;
- 13 Update the search space $\{\mathcal{O}^{(i,j)}\}$ using Eq. 5;
- 14 Search the architecture for V epochs based on $\mathcal{O}^{(i,j)}$ using PC-DARTS;
- 15 $K = K - 1$;
- 16 **end**

are solved in different BNNs, such as (Gu et al. 2019) and (Rastegari et al. 2016b). The complete loss function L for BNAS is defined as:

$$L = L_S + L_A, \quad (7)$$

where L_S is the conventional loss function, e.g., cross-entropy.

Experiments

In this section, we compare our BNAS with state-of-the-art NAS methods, and also compare the BNNs obtained by our BNAS based on XNOR (Rastegari et al. 2016b) and PCNN (Gu et al. 2019).

Experiment Protocol

In these experiments, we first search neural architectures on an over-parameterized network on CIFAR-10, and then evaluate the best architecture with a stacked deeper network on the same data set. Then we further perform experiments to search architectures directly on ImageNet. We run the experiment multiple times and find that the resulting architectures only show slight variation in performance, which demonstrates the stability of the proposed method.

We use the same datasets and evaluation metrics as existing NAS works (Liu, Simonyan, and Yang 2018; Cai et al. 2018b; Zoph et al. 2018; Liu et al. 2018). First, most experiments are conducted on CIFAR-10 (Krizhevsky, Hinton, and others 2009), which has 50K training images and

10K testing images with resolution 32×32 and from 10 classes. The color intensities of all images are normalized to $[-1, +1]$. During architecture search, the 50K training samples of CIFAR-10 is divided into two subsets of equal size, one for training the network weights and the other for finding the architecture hyper-parameters. When reducing the search space, we randomly select 5K images from the training set as a validation set (used in line 8 of Algorithm 1). To further evaluate the generalization capability, we stack the discovered optimal cells on CIFAR-10 into a deeper network, and then evaluate the classification accuracy on ILSVRC 2012 ImageNet (Russakovsky et al. 2015), which consists of 1,000 classes with 1.28M training images and 50K validation images.

In the search process, we consider a total of 6 cells in the network, where the reduction cell is inserted in the second and the fourth layers, and the others are normal cells. There are $M = 4$ intermediate nodes in each cell. Our experiments follow PC-DARTS. We set the hyper-parameter C in PC-DARTS to 2 for CIFAR-10 so only $1/2$ features are sampled for each edge. The batch size is set to 128 during the search of an architecture for $L = 5$ epochs based on $\mathcal{O}^{(i,j)}$ (line 1 in Algorithm 1). Note for $5 \leq L \leq 10$, the larger L has little effect on the final performance, but will cost more search time. We freeze the network hyper-parameters such as α , and only allow the network parameters such as filter weights to be tuned in the first 3 epochs. Then in the next 2 epochs, we train both the network hyper-parameters and the network parameters. This is to provide an initialization for the network parameters and thus alleviates the drawback of parameterized operations compared with free parameter operations. We also set $T = 3$ (line 4 in Algorithm 1) and $V = 1$ (line 14), so the network is trained less than 60 epochs, with a larger batch size of 400 (due to few operation samplings) during reducing the search space. The initial number of channels is 16. We use SGD with momentum to optimize the network weights, with an initial learning rate of 0.025 (annealed down to zero following a cosine schedule), a momentum of 0.9, and a weight decay of 5×10^{-4} . The learning rate for finding the hyper-parameters is set to 0.01.

After search, in the architecture evaluation step, our experimental setting is similar to (Liu, Simonyan, and Yang 2018; Zoph et al. 2018; Pham et al. 2018). A larger network of 20 cells (18 normal cells and 2 reduction cells) is trained on CIFAR-10 for 600 epochs with a batch size of 96 and an additional regularization cutout (DeVries and Taylor 2017). The initial number of channels is 36. We use the SGD optimizer with an initial learning rate of 0.025 (annealed down to zero following a cosine schedule without restart), a momentum of 0.9, a weight decay of 3×10^{-4} and a gradient clipping at 5. When stacking the cells to evaluate on ImageNet, the evaluation stage follows that of DARTS, which starts with three convolution layers of stride 2 to reduce the input image resolution from 224×224 to 28×28 . 14 cells (12 normal cells and 2 reduction cells) are stacked after these three layers, with the initial channel number being 64. The network is trained from scratch for 250 epochs using a batch size of 512. We use the SGD optimizer with a momentum of 0.9, an initial learning rate of 0.05 (decayed down to

Architecture	Test Error (%)	# Params (M)	Search Cost (GPU days)	Search Method
ResNet-18 (He et al. 2016)	3.53	11.1 (32 bits)	-	Manual
WRN-22 (Zagoruyko and Komodakis 2016)	4.25	4.33 (32 bits)	-	Manual
DenseNet (Huang et al. 2017)	4.77	1.0 (32 bits)	-	Manual
SENet (Hu, Shen, and Sun 2018)	4.05	11.2 (32 bits)	-	Manual
ResNet-18 (XNOR)	6.69	11.17 (1 bit)	-	Manual
ResNet-18 (PCNN)	5.63	11.17 (1 bit)	-	Manual
WRN22 (PCNN) (Gu et al. 2019)	5.69	4.29 (1 bit)	-	Manual
Network in (McDonnell 2018)	6.13	4.30 (1 bit)	-	Manual
NASNet-A (Zoph et al. 2018)	2.65	3.3 (32 bits)	1800	RL
AmoebaNet-A (Real et al. 2018)	3.34	3.2 (32 bits)	3150	Evolution
PNAS (Liu et al. 2018)	3.41	3.2 (32 bits)	225	SMBO
ENAS (Pham et al. 2018)	2.89	4.6 (32 bits)	0.5	RL
Path-level NAS (Cai et al. 2018b)	3.64	3.2 (32 bits)	8.3	RL
DARTS(first order) (Liu, Simonyan, and Yang 2018)	2.94	3.1 (32 bits)	1.5	Gradient-based
DARTS(second order) (Liu, Simonyan, and Yang 2018)	2.83	3.4 (32 bits)	4	Gradient-based
PC-DARTS	2.78	3.5 (32 bits)	0.15	Gradient-based
BNAS (full-precision)	2.84	3.3 (32 bits)	0.08	Performance-based
BNAS (XNOR)	5.71	2.3 (1 bit)	0.104	Performance-based
BNAS (XNOR, larger)	4.88	3.5 (1 bit)	0.104	Performance-based
BNAS (PCNN)	3.94	2.6 (1 bit)	0.09375	Performance-based
BNAS (PCNN, larger)	3.47	4.6 (1 bit)	0.09375	Performance-based

Table 1: Test error rates for human-designed full-precision networks, human-designed binarized networks, full-precision networks obtained by NAS, and networks obtained by our BNAS on CIFAR-10. Note that the parameters are 1 bit in binarized networks, and are 32 bits in full-precision networks. For fair comparison, we select the architectures by NAS with similar parameters ($< 5M$). In addition, we also train an optimal architecture in a larger setting, *i.e.*, with more initial channels (44 in XNOR or 48 in PCNN).

zero following a cosine schedule), and a weight decay of 3×10^{-5} . Additional enhancements are adopted including label smoothing and an auxiliary loss tower during training. All the experiments and models are implemented in PyTorch (Paszke et al. 2017).

Results on CIFAR-10

We compare our method with both manually designed networks and networks searched by NAS. The manually designed networks include ResNet (He et al. 2016), Wide ResNet (WRN) (Zagoruyko and Komodakis 2016), DenseNet (Huang et al. 2017) and SENet (Hu, Shen, and Sun 2018). For the networks obtained by NAS, we classify them according to different search methods, such as RL (NASNet (Zoph et al. 2018), ENAS (Pham et al. 2018), and Path-level NAS (Cai et al. 2018b)), evolutionary algorithms (AmoebaNet (Real et al. 2018)), Sequential Model Based Optimization (SMBO) (PNAS (Liu et al. 2018)), and gradient-based methods (DARTS (Liu, Simonyan, and Yang 2018) and PC-DARTS (Xu et al. 2019)).

The results for different architectures on CIFAR-10 are summarized in Tab. 1. Using BNAS, we search for two binarized networks based on XNOR (Rastegari et al. 2016b) and PCNN (Gu et al. 2019). In addition, we also train a larger XNOR variant with 44 initial channels and a larger PCNN variant with 48 initial channels. We can see that the test errors of the binarized networks obtained by our BNAS are comparable to or smaller than those of the full-precision hu-

man designed networks, and are significantly smaller than those of the other binarized networks.

Compared with the full-precision networks obtained by other NAS methods, the binarized networks by our BNAS have comparable test errors but with much more compressed models. Note that the numbers of parameters of all these searched networks are less than 5M, but the binarized networks only need 1 bit to save one parameter, while the full-precision networks need 32 bits. In terms of search efficiency, compared with the previous fastest PC-DARTS, our BNAS is 40% faster (tested on our platform (NVIDIA GTX TITAN Xp)). We attribute our superior results to the proposed way of solving the problem with the novel scheme of search space reduction.

Our BNAS method can also be used to search full-precision networks. In Tab. 1, BNAS (full-precision) and PC-DARTS perform equally well, but BNAS is 47% faster. Both the binarized methods XNOR and PCNN in our BNAS perform well, which shows the generalization of BNAS. Fig. 3 and Fig. 4 show the best cells searched by BNAS based on XNOR and PCNN, respectively.

We also use PC-DARTS to perform a binarized architecture search based on PCNN on CIFAR10, resulting in a network denoted as PC-DARTS (PCNN). Compared with PC-DARTS (PCNN), BNAS (PCNN) achieves a better performance (95.12% vs. 96.06% in test accuracy) with less search time (0.18 vs. 0.09375 GPU days). The reason for this may be because the performance based strategy can help find bet-

Architecture	Accuracy (%)		Params (M)	Search Cost (GPU days)	Search Method
	Top1	Top5			
ResNet-18 (Gu et al. 2019)	69.3	89.2	11.17 (32 bits)	-	Manual
MobileNetV1 (Howard et al. 2017)	70.6	89.5	4.2 (32 bits)	-	Manual
ResNet-18 (PCNN) (Gu et al. 2019)	63.5	85.1	11.17 (1 bit)	-	Manual
NASNet-A (Zoph et al. 2018)	74.0	91.6	5.3 (32 bits)	1800	RL
AmoebaNet-A (Real et al. 2018)	74.5	92.0	5.1 (32 bits)	3150	Evolution
AmoebaNet-C (Real et al. 2018)	75.7	92.4	6.4 (32 bits)	3150	Evolution
PNAS (Liu et al. 2018)	74.2	91.9	5.1 (32 bits)	225	SMBO
DARTS (Liu, Simonyan, and Yang 2018)	73.1	91.0	4.9 (32 bits)	4	Gradient-based
PC-DARTS (Xu et al. 2019)	75.8	92.7	5.3 (32 bits)	3.8	Gradient-based
BNAS (PCNN)	71.3	90.3	6.2 (1 bit)	2.6	Performance-based

Table 2: Comparison with the state-of-the-art image classification methods on ImageNet. BNAS and PC-DARTS are obtained directly by NAS and BNAS on ImageNet, others are searched on CIFAR-10 and then directly transferred to ImageNet.

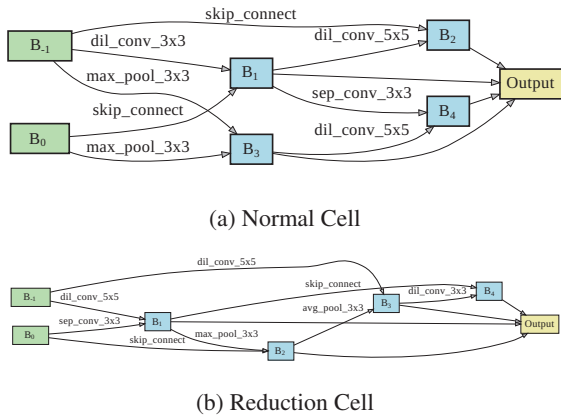


Figure 3: Detailed structures of the best cells discovered on CIFAR-10 using BNAS based on XNOR. In the normal cell, the stride of the operations on 2 input nodes is 1, and in the reduction cell, the stride is 2.

ter operations for recognition.

Results on ImageNet

We further compare the state-of-the-art image classification methods on ImageNet. All the searched networks are obtained directly by NAS and BNAS on ImageNet by stacking the cells. Our binarized network is based on PCNNs. From the results in Tab. 2, we have the following observations: (1) BNAS (PCNN) performs better than human-designed binarized networks (71.3% vs. 63.5%) and has far fewer parameters (6.1M vs. 11.17M). (2) BNAS (PCNN) has a performance similar to the human-designed full-precision networks (71.3% vs. 70.6%), with a much more highly compressed model. (3) Compared with the full-precision networks obtained by other NAS methods, BNAS (PCNN) has little performance drop, but is fastest in terms of search efficiency (0.09375 vs. 0.15 GPU days) and is a much more highly compressed model due to the binarization of the network. The above results show the excellent transferability of

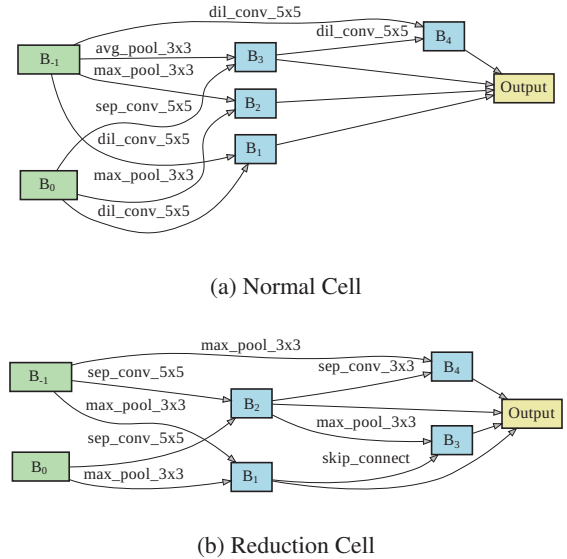


Figure 4: Detailed structures of the best cells discovered on CIFAR-10 using BNAS based on PCNN. In the normal cell, the stride of the operations on 2 input nodes is 1, and in the reduction cell, the stride is 2.

our BNAS method.

Conclusion

In this paper, we have proposed BNAS, the first binarized neural architecture search algorithm, which effectively reduces the search time by pruning the search space in early training stages. It is faster than the previous most efficient search method PC-DARTS. The binarized networks searched by BNAS can achieve excellent accuracies on CIFAR-10 and ImageNet. They perform comparable to the full-precision networks obtained by other NAS methods, but with much compressed models.

Acknowledgements

The work was supported in part by National Natural Science Foundation of China under Grants 61672079, 61473086, 61773117, 614730867. This work is supported by Shenzhen Science and Technology Program KQTD2016112515134654. Baochang Zhang is also with Shenzhen Academy of Aerospace Technology, Shenzhen 100083, China.

References

- Brock, A.; Lim, T.; Ritchie, J. M.; and Weston, N. 2017. Smash: one-shot model architecture search through hyper-networks. *arXiv*.
- Cai, H.; Chen, T.; Zhang, W.; Yu, Y.; and Wang, J. 2018a. Efficient architecture search by network transformation. In *Proc. of AAAI*.
- Cai, H.; Yang, J.; Zhang, W.; Han, S.; and Yu, Y. 2018b. Path-level network transformation for efficient architecture search. *arXiv*.
- Cai, H.; Zhu, L.; and Han, S. 2018. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv*.
- Chen, X.; Xie, L.; Wu, J.; and Tian, Q. 2019. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. *arXiv*.
- Courbariaux, M.; Bengio, Y.; and David, J.-P. 2015. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Proc. of NIPS*.
- Courbariaux, M.; Hubara, I.; Soudry, D.; El-Yaniv, R.; and Bengio, Y. 2016. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv*.
- DeVries, T., and Taylor, G. W. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv*.
- Gu, J.; Li, C.; Zhang, B.; Han, J.; Cao, X.; Liu, J.; and Doermann, D. 2019. Projection convolutional neural networks for 1-bit cnns via discrete back propagation. In *Proc. of AAAI*.
- Ha, D.; Dai, A.; and V. Le, Q. 2016. Hypernetworks. *arXiv*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proc. of CVPR*.
- Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv*.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proc. of CVPR*.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proc. of CVPR*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. Technical report, Cite-seer.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Proc. of NIPS*.
- Liu, H.; Simonyan, K.; Vinyals, O.; Fernando, C.; and Kavukcuoglu, K. 2017. Hierarchical representations for efficient architecture search. *arXiv*.
- 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 *Proc. of ECCV*.
- Liu, H.; Simonyan, K.; and Yang, Y. 2018. Darts: Differentiable architecture search. *arXiv*.
- McDonnell, M. D. 2018. Training wide residual networks for deployment using a single bit for each weight. *arXiv*.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch. In *Proc. of NIPS*.
- Pham, H.; Guan, M. Y.; Zoph, B.; Le, Q. V.; and Dean, J. 2018. Efficient neural architecture search via parameter sharing. *arXiv*.
- Rastegari, M.; Ordonez, V.; Redmon, J.; and Farhadi, A. 2016a. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Proc. of ECCV*.
- Rastegari, M.; Ordonez, V.; Redmon, J.; and Farhadi, A. 2016b. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Proc. of ECCV*.
- Real, E.; Aggarwal, A.; Huang, Y.; and Le, Q. V. 2018. Regularized evolution for image classifier architecture search. *arXiv*.
- 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. *International Journal of Computer Vision*.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv*.
- Xie, S.; Zheng, H.; Liu, C.; and Lin, L. 2018. Snas: stochastic neural architecture search. *arXiv*.
- Xu, Y.; Xie, L.; Zhang, X.; Chen, X.; Qi, G.-J.; Tian, Q.; and Xiong, H. 2019. Partial channel connections for memory-efficient differentiable architecture search. *arXiv*.
- Xu, J. F.; Boddeti, V. N.; and Savvides, M. 2016. Local binary convolutional neural networks. In *Proc. of CVPR*.
- Ying, C.; Klein, A.; Real, E.; Christiansen, E.; Murphy, K.; and Hutter, F. 2019. Nas-bench-101: Towards reproducible neural architecture search. *arXiv*.
- Zagoruyko, S., and Komodakis, N. 2016. Wide residual networks. In *Proc. of BMVC*.
- Zoph, B., and Le, Q. V. 2016. Neural architecture search with reinforcement learning. *arXiv*.
- Zoph, B.; Vasudevan, V.; Shlens, J.; and Le, Q. V. 2018. Learning transferable architectures for scalable image recognition. In *Proc. of CVPR*.