

Evolving Search Space for Neural Architecture Search

Yuanzheng Ci¹, Chen Lin², Ming Sun³, Boyu Chen¹, Hongwen Zhang⁴, Wanli Ouyang¹

¹The University of Sydney, ²University of Oxford, ³SenseTime Research, ⁴CASIA

{yuanzheng.ci, boyu.chen, wanli.ouyang}@sydney.edu.au, chen.lin@eng.ox.ac.uk,
sunming1@sensetime.com, hongwen.zhang@cripac.ia.ac.cn

Abstract

Automation of neural architecture design has been a coveted alternative to human experts. Various search methods have been proposed aiming to find the optimal architecture in the search space. One would expect the search results to improve when the search space grows larger since it would potentially contain more performant candidates. Surprisingly, we observe that enlarging search space is unbeneficial or even detrimental to existing NAS methods such as DARTS, ProxylessNAS, and SPOS. This counterintuitive phenomenon suggests that enabling existing methods to large search space regimes is non-trivial. However, this problem is less discussed in the literature.

We present a Neural Search-space Evolution (NSE) scheme, the first neural architecture search scheme designed especially for large space neural architecture search problems. The necessity of a well-designed search space with constrained size is a tacit consent in existing methods, and our NSE aims at minimizing such necessity. Specifically, the NSE starts with a search space subset, then evolves the search space by repeating two steps: 1) search an optimized space from the search space subset, 2) refill this subset from a large pool of operations that are not traversed. We further extend the flexibility of obtainable architectures by introducing a learnable multi-branch setting. With the proposed method, we achieve 77.3% top-1 retrain accuracy on ImageNet with 333M FLOPs, which yielded a state-of-the-art performance among previous auto-generated architectures that do not involve knowledge distillation or weight pruning. When the latency constraint is adopted, our result also performs better than the previous best-performing mobile models with a 77.9% Top-1 retrain accuracy. Code is available at https://github.com/orashi/NSE_NAS.

1. Introduction

Deep neural networks are prevailing in myriad fields of real-world applications. The emergence of Neural Archi-

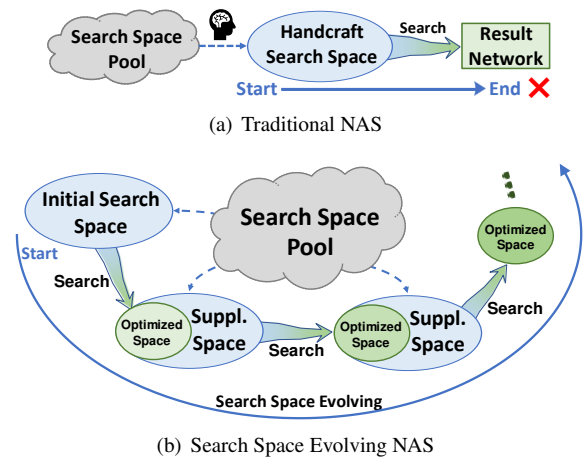


Figure 1. Comparison of search schemes. (a) Traditional pipeline. (b) Our proposed search space evolving pipeline.

ture Search (NAS) has brought up a possibility to automate the customization of deep neural network architectures for specific applications. Researchers have investigated Reinforcement Learning (RL) and Evolutionary Algorithm (EA) based methods [58, 29, 30, 55] to achieve the automation of architecture design. Weight sharing based methods [21, 7, 2, 9, 22, 56, 19, 8, 33] that can substantially reduce the computational cost have been proposed and became one of the off-the-shelf approaches of NAS research. These methods successfully yielded promising results that have surpassed human-designed architectures [41].

In addition to the search method, another key component of NAS is the search space. When compared with early NAS works [59, 29], the quality of the search space has been improved along with the development of search algorithms [27]. It has been observed that the improvement of search space design imposed a positive effect on the performance of many existing works [27, 45]. In particular, the research community has devoted multiple efforts to search space design, from selecting well-suited operations based on the prior knowledge [39, 14] to utilizing channel-level fine-grained model search [40, 35, 14, 49, 6, 1] over

a smaller set of operations. Recent methods often limit the number of candidate operations in each layer to less than ten (excluding decisions on activation functions or SE modules [16]). However, such search space improvements fall back into the paradigm of expertise design, which is a counter-march of automated learning of architectures.

A natural question now could be considered is: *can we construct a huge search space which is a super-set of a formerly mentioned space and obtain superior results from it?* If the answer is YES, this approach would solve the problem of search space design by simply using the largest search space one can construct. However, it is less discussed by NAS literature. Yu *et al.* [51] and Zhang *et al.* [53] provide results suggesting that, under the traditional NAS pipeline as shown in Figure 1(a), simply enlarging the search space could be detrimental to the final result.

To look into this issue in detail, we set up a search space consists of 27 distinct operations and then tested 4 reasonably fast NAS algorithms, including DARTS [21], Proxyless [7], SPOS [14] and One-Shot [2]. We show that all of these methods do not hold the behavior of obtaining better search results with a larger search space. Furthermore, some have their search cost prohibitively high or failed to converge, while others perform poorly even with their training epochs increased (effectiveness of longer training schedule is suggested by [54, 3]). Another relevant technique is search space simplification, which can also assist neural architecture search [54, 17, 13], we will show that such technique is not enough to help NAS algorithm exploit large space effectively.

In this work, we aim at large space neural architecture search by proposing a Neural Search-space Evolution (NSE) scheme. Instead of directly confronting the negative impacts derived from a large search space, the NSE starts with a search space subset, which is a reasonably sized random subset of the full search space, and search an optimized space from this subset, then refill this subset and repeat the search-then-refill steps as shown in Figure 1(b) to traverse the whole space progressively. The resulted NSE enables ever-evolving search space for Neural Architecture Search.

NSE progressively explores extra operation candidates while retaining past knowledge. The search process is constructed as an iterative process to traverse the pending unseen operation candidates. During the iterative process, instead of keeping a single architecture as the intermediate result, we combine all architectures on the Pareto front found by a supernet trained with One-Shot [2] method to obtain an optimized search space, which will be inherited to the next round of search. By maintaining an optimized search space as knowledge, the search process can always proceed with new candidate operations added to the pending list, which means we can always add newly proposed operations in CNN literature that are less verified yet potentially effi-

cient for specific tasks.

To effectively exploit more complex architectures, we further adopt the proposed paradigm to the multi-branch scheme, which has orders of magnitude more distinct structures when compared to its single-branch counterpart. Compared with previous single-branch schemes like DARTS that only allow one operation to be selected, the multi-branch scheme allows multiple operations to be selected adaptively. By constructing a probabilistic model for the multi-branch scheme, we can retrieve the fitness of every candidate operation. Operations with fitness lower than a certain threshold will be dropped from the search space so that the complexity of the search space is progressively reduced and the degree of co-adaptation for the rest of the possible path combinations could also be enhanced.

We conduct experiments on ImageNet [31] with two resource constraints, *i.e.* FLOPs and Latency. For these two constraints, NAS under our NSE scheme effectively exploited the potential of a very large search space and secured a continual performance increment in the iterative process, leading to state-of-the-art results.

In summary, the key contributions of this paper are summarized as follows:

- We propose NSE, the first neural architecture search scheme that designed especially for large space neural architecture search problems, which empowers NAS to minimize the necessity of dedicated search space design. The inheritance property of the evolving process keeps knowledge derived from previous search space while improving that knowledge with new operations added into the current search space.
- We propose a probabilistic modeling of operation-wise fitness for the multi-branch scheme, which makes it feasible to gradually simplify the multi-branch scheme search space as shared weights converge through one-shot training. Such an annealing paradigm gradually simplifies the complexity of the sub-task and helps the remaining shared weights to be learned better [54].

2. Related Work

2.1. NAS Algorithm Design

Network architecture search algorithms are essential for finding good architectures. These algorithms are based on Bayesian optimization, Reinforcement Learning (RL), Genetic Algorithm (GA), Weight sharing, and One-shot. There are Bayesian optimization methods that formulate the NAS as a hyperparameter optimization problem and allow to search non-fixed length neural architectures [26, 4]. RL-based NAS [58, 59, 55] adopted RL to learn generating the best architecture. GA-based methods [30, 29] use

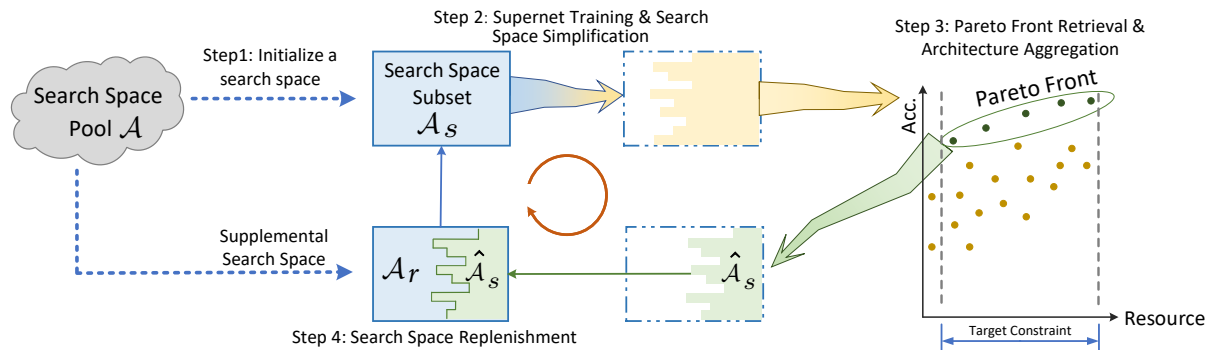


Figure 2. Search space update scheme for proposed approach. A search space subset is retained to incorporate new unseen search space while preserving existing knowledge.

GA to generate candidate architectures and are also popular with impressive results. Weight sharing approaches [2, 21, 7, 14, 10, 23] utilize a supernet that shares weights among different architectures. One-shot NAS [2, 14] by directly training the supernet with dropout, the supernet with shared weights are utilized for predicting model score predictor. However, when new candidate operations are available, these NAS approaches have to restart the search. In contrast, under our NSE scheme, a subset of search space is maintained to efficiently accommodate new candidate operations and inherit the knowledge of previously searched candidate operations. Recent works [49, 6, 1, 25] also integrated knowledge distillation or weight pruning into the NAS pipeline for instant performance. Nevertheless, these techniques can still be conducted independently with considerable gain [44, 5] and are orthogonal to our focus.

2.2. Search Space Design

The NAS search space is identified to have posed a non-negligible impact on search results [51, 53]. Meanwhile, search space design has been improved together with the NAS algorithms [27]. Zoph *et al.* [58] adopted design space with naive building blocks and jump connections. Many approaches repeat the same building block they searched via NAS to construct the network. [59, 55, 29, 21]. For better flexibility, many latter approaches enable searching different operations for different blocks but have to constrain the number of candidate operations for controlling the search space size [37, 7, 42, 14, 13, 52]. Orthogonal works that search for fine-grained model adjustments and training configs are also proposed [50, 48, 25, 40, 12]. Recent works [42, 14, 7, 37, 52, 15, 39] also include new candidate operations or module designs such as hand-crafted multi-branch cells [39], tree-structure [7], shuffle operation [24], Squeeze-and-Excitation (SE) module [16] and swish activation [28]. Our NSE scheme enables neural architecture search with ever-evolving new candidate operations.

3. Method

Problem Formulation. As the neural architecture commonly uses a feed-forward structure, we represent an overall search space pool \mathcal{A} as a direct acyclic graph (DAG) of L layers, $\bigcup_{l=1}^L \mathbf{E}_l$, where \mathbf{E}_l represents available operations (e.g., 3×3 convolution, pooling, or identity) in the l -th layer of DAG. We denote a neural network within the search space as $a = \bigcup_{l=1}^L e_l$, where $e_l \subseteq \mathbf{E}_l$.

Noticing that multi-branch based networks like Inception [36] and ResNeXt [43] were important inventions to the CNN literature, we include this scheme into our search space evolution. For a network architecture in the *multi-branch* scheme, a layer consists of multiple operations $\{op_n\}$ selected from N operation candidates, i.e. $\mathbf{e} = \mathbf{o}_g = \{op_n | g_n = 1, n \in \{1, \dots, N\}\}$, where \mathbf{g} denotes a specific set of operation configuration $\{g_n\}$ and the binary gates $g_n \in \{0, 1\}$ denotes whether the n -th operation is selected or not. In this case, the number of selected operations within \mathbf{o}_g is $\sum_{n=1}^N g_n$, and the total amount of possible operation configurations (i.e. combinations) is 2^N .

The goal of the proposed pipeline is to explore an extremely large search space by evolving through the search space subsets and looking for the best-performing architecture $a^* \subseteq \mathcal{A}$ at the same time.

3.1. Overview

We formulate our search space evolving NAS pipeline as an adaptive process that is capable of exploring from a stream of search space subset replenishment, as depicted in Figure 2. The pipeline of NSE scheme is as follows:

Step 1. Initially, a search space subset \mathcal{A}_s is sampled from the overall search space pool \mathcal{A} . This is achieved by randomly sampling K candidate operations for each layer. In practice, the search space subset \mathcal{A}_s is much smaller than the whole set \mathcal{A} to avoid the large search space dilemma we mentioned. It consists of candidate operations currently being considered for NAS algorithms to search; we believe this search space subset, limited in size, is naturally easier

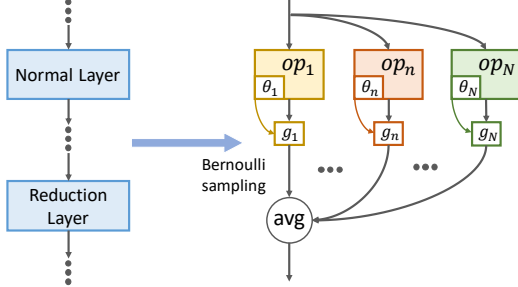


Figure 3. Forwarding scheme for multi-branch paths.

for a search algorithm to handle.

Step 2. Supernet training (Section 3.2) and search space simplification (Section 3.3.1) are alternately conducted within search space subset \mathcal{A}_s . If the fitness of an operation is found to be lower than a certain threshold during the supernet training, then this operation will be dropped from the search space subset.

Step 3. The optimized search space subset is obtained by sampling architectures from the search space subset \mathcal{A}_s and evaluating them (Section 3.3.2). Specifically, the fitness of these sampled architectures is evaluated by validation accuracy to retrieve a Pareto front. We aggregate architectures at the Pareto front and obtain the optimized search space subset $\hat{\mathcal{A}}_s$, which is the smallest network search space containing all architectures in the Pareto Front.

Step 4. A new search space \mathcal{A}'_s is constructed (Section 3.3.3). \mathcal{A}'_s is a combination of optimized search space subset $\hat{\mathcal{A}}_s$ from Step 3 and a supplemental subset of search space \mathcal{A}_r . \mathcal{A}_r is a replenishment sampled from \mathcal{A} to ensure that there are still K candidate operations in each layer. The traversed search space is excluded from \mathcal{A}_r so that the same operation in a layer will not be sampled twice. In this way, the accumulated knowledge in $\hat{\mathcal{A}}_s$ is inherited. Let $\mathcal{A}_s = \mathcal{A}'_s$, re-initialize all weights, then go to Step 2. When there is not enough operation remained in \mathcal{A} , the loop ends, and the Pareto front for \mathcal{A}_s in Step 2 is used as the final result of the architecture search.

3.2. Supernet Training

3.2.1 Forward Path in Multi-Branch Scheme

Given an input feature map \mathbf{X} , the output of a branch op_n can be written as $g_n op_n(\mathbf{X})$. As illustrated in Figure 3, the output of a multi-branch layer under the configuration of $\mathbf{g} = \{g_n\}$ can be defined as follows:

$$\mathbf{o}_{\mathbf{g}}(\mathbf{X}) = \frac{1}{\sum_{n \in \mathbf{K}} g_n} \sum_{n \in \mathbf{K}} g_n op_n(\mathbf{X}), \quad (1)$$

where \mathbf{K} denotes the set of index for K candidate operations. Notably, an identity operation op_{id} is additionally indexed in \mathbf{K} for normal layers with g_{id} always equals 1.

3.2.2 Weight Update

With weight sharing [2, 21, 7, 14], every subgraph, *i.e.* architecture denoted by a , inherits its weights denoted by $W_{\mathcal{A}_s}(a)$ from the weights of the supernet denoted by $W_{\mathcal{A}_s}$, where $a \subseteq \mathcal{A}_s$. The optimization of shared weights $W_{\mathcal{A}_s}$ can be expressed as follows:

$$W_{\mathcal{A}_s} = \arg \min_{W_{\mathcal{A}_s}} \mathbb{E}_{a \sim \mathbb{U}_a} [\mathcal{L}_{CE}(\mathcal{N}(a, W_{\mathcal{A}_s}(a)))], \quad (2)$$

where $\mathbb{E}[\cdot]$ denotes expectation, $\mathcal{L}_{CE}(\cdot)$ denotes the cross-entropy loss, and $\mathcal{N}(a, W_{\mathcal{A}_s}(a))$ denotes the network with architecture a and parameters $W_{\mathcal{A}_s}(a)$. The minimization over expectation $\mathbb{E}[\cdot]$ in Eq. (2) is implemented by sampling architecture a from the space \mathcal{A}_s and then updating the corresponding weights $W_{\mathcal{A}_s}(a)$ using stochastic gradient descent. Multiple works [10, 51] points out that the supernet needs to be trained evenly so that the shared weights can be good for predicting accuracy. Therefore, we sample every possible architecture equally, *i.e.* architectures will be sampled from a uniform distribution \mathbb{U}_a where every branch is sampled with $g_n \sim \text{Bernoulli}(0.5)$.

3.3. Search Space Evolution

Each iteration of search space evolution involves three steps as shown in Figure 2, here we provide detailed illustrations for each step.

3.3.1 Search Space Simplification

The multi-branch scheme we adopted enlarge the search space by more than 90 orders of magnitudes (see appendix for detailed numbers). To effectively search for the optimal architecture, we propose to progressively simplify the search space subset with an operation-wise fitness while training the supernet. Specifically, we adopt learnable fitness indicators $\Theta = \{\theta_n^l\}$ to predict the fitness of an operation and guide the simplification of the search space subset, they are assigned to every operation candidate individually as shown in Figure 3.

Simplifying Search Space with Lock and Rehearse. The n -th path with its corresponding indicator θ_n below a certain threshold will be dropped from supernet, except that: 1) the path is the last remained path in a reduction cell, or 2) the operation of this path is inherited from the previous iteration. The second case of exceptions keeps the inherited operations even if their predicted fitness is below the threshold. It also allows us to re-evaluate previous Pareto optimal architectures on the current supernet, thus protecting the inherited knowledge from being underestimated. This is inspired by the knowledge rehearsal for lifelong learning [34]. We call this strategy Lock and Rehearse (L&R).

Probabilistic Modeling of Fitness Indicators. Fitness Indicators are updated once after every two supernet train-

ing iterations. For each update, the n -th path is individually sampled with $g_n \sim \text{Bernoulli}(p_n)$, where $p_n = \frac{1}{1+e^{-\theta_n}}$. The fitness indicators Θ are updated on the validation set as follows:

$$\Theta^* = \arg \min_{\Theta} \mathbb{E}_{a \sim \mathbb{P}_a(\Theta)} [\mathcal{L}_{\text{val}}(\mathcal{N}(a, W_{\mathcal{A}_s}(a)))] \quad (3)$$

where $\mathbb{P}_a(\Theta)$ denotes the architecture probability distribution parameterized by fitness indicators Θ .

For a layer with K candidate operations, we calculate the layer probability $\hat{p}_{\mathbf{g}}$ for operation configuration \mathbf{g} as the joint probability of K -dimensional Bernoulli random variables based on operation probability p_n :

$$\hat{p}_{\mathbf{g}} = \prod_{n \in \mathbf{K}} (g_n p_n + (1 - g_n)(1 - p_n)), \quad (4)$$

When initialized, $\theta = 0$ for all $\theta \in \Theta$, resulting in $p_n = 0.5$, so that all possible combinations are selected with equal probability at the beginning. When the fitness indicators Θ are optimized with resource constraints as regularization, the formulation in Eq. (3) can be implemented as follows:

$$\Theta^* = \arg \min_{\Theta} \mathbb{E}_{a \sim \mathbb{P}_a(\Theta)} [\mathcal{L}_{CE}(\mathcal{N}(a, W_{\mathcal{A}_s}(a))) + \alpha(\mathcal{R}(a, \tau))^\beta], \quad (5)$$

where $\mathbb{E}[\cdot]$ denotes expectation, \mathcal{R} measures the difference between the resource demand of architecture a and the target demand τ . α and β are application-specific constants. For the latency constraint, we follow [46] to build up a latency lookup table, which records the latency cost of every operation included in the search space.

Simulated Gradients. Since Θ is not directly involved in the computation of \mathcal{L}_{CE} , we are unable to update the first term in Eq. (5) directly through back-propagation, thus a simulated gradient is necessary. Inspired by BinaryConnect [11], we first forward the output of a sampled branch combination $\mathbf{o}_{\mathbf{g}_a}(\mathbf{X})$ to the next layer, where \mathbf{g}_a is a randomly selected configuration. Then we utilize the gradient w.r.t. the output as the simulated gradient. Hence, we have the simulated gradient written as follows:

$$\frac{d\mathcal{L}_{CE}}{d\theta_n} = \sum_{i=1}^{2^K} \frac{d\mathcal{L}_{CE}}{d\hat{p}_{\mathbf{g}_i}} \frac{d\hat{p}_{\mathbf{g}_i}}{d\theta_n} \approx \sum_{i=1}^{2^K} \frac{d\mathcal{L}_{CE}}{d\mathbf{o}_{\mathbf{g}_i}(\mathbf{X})} \mathbf{o}_{\mathbf{g}_i}(\mathbf{X}) \frac{d\hat{p}_{\mathbf{g}_i}}{d\theta_n}, \quad (6)$$

where $\mathbf{o}_{\mathbf{g}_i}(\mathbf{X})$ is defined in Eq. (1). The summation over all \mathbf{g}_i in Eq. (6) is still complex. We use the method in [7] to simplify the computation by selecting only two configurations to reduce the GPU memory and computation required for each iteration. Specifically, we randomly sample another configuration \mathbf{g}_b , and then rescale the layer probabilities $\hat{p}_{\mathbf{g}}$ of configuration \mathbf{g}_a and \mathbf{g}_b to $\tilde{p}_{\mathbf{g}}$ so that $\tilde{p}_{\mathbf{g}_a} + \tilde{p}_{\mathbf{g}_b} = 1$.

Finally, we have the simulated gradient approximated as:

$$\frac{d\mathcal{L}_{CE}}{d\theta_n} \approx \frac{d\mathcal{L}_{CE}}{d\mathbf{o}_{\mathbf{g}_a}(\mathbf{X})} \mathbf{o}_{\mathbf{g}_a}(\mathbf{X}) \frac{d\tilde{p}_{\mathbf{g}_a}}{d\theta_n} + \frac{d\mathcal{L}_{CE}}{d\mathbf{o}_{\mathbf{g}_b}(\mathbf{X})} \mathbf{o}_{\mathbf{g}_b}(\mathbf{X}) \frac{d\tilde{p}_{\mathbf{g}_b}}{d\theta_n}. \quad (7)$$

Likewise, the $\mathcal{R}(a, \tau)$ of second term in Eq. (5) is approximated as:

$$\mathcal{R}(a, \tau) \approx \sum_{l=1}^L (\tilde{p}_{\mathbf{g}_a^l} \mathcal{C}(\mathbf{o}_{\mathbf{g}_a^l}) + \tilde{p}_{\mathbf{g}_b^l} \mathcal{C}(\mathbf{o}_{\mathbf{g}_b^l})) - \tau, \quad (8)$$

where \mathcal{C} maps the combination of operations $\mathbf{o}_{\mathbf{g}}$ to its corresponding resource cost.

3.3.2 Pareto Front Retrieval and Architecture Aggregation

Pareto Front Retrieval. To retrieve the Pareto front needed for architecture aggregation, we evaluate the validation accuracies of sampled architectures by using the weights of the well-trained supernet and the validation dataset. We randomly sample D distinct models based on the probability distribution implied by fitness indicators Θ . The sampled models not fitting the resource constraint are discarded. The Pareto-optimal architectures from the last searching iteration, if exist, will also be evaluated as part of the L&R strategy. In practice, extra D_e samples are needed to overcome the edging effect (see appendix for details). After all samples are evaluated, we can get P Pareto-optimal architectures $\{a_1, \dots, a_p, \dots, a_P\}$. The Pareto-optimal architectures derived from the final round of optimization will be referred to as final results.

Due to the volatility of BN statistics for supernet and the fact that different architectures should adopt different BN statistics, we need to perform recalculation for BN layers. Specifically, before evaluating an architecture on the shared weight supernet, we recalculate the statistics of BN layers by forwarding 20k random training images, which takes about one second.

Aggregation. After Pareto front retrieval, we take the union of operations from all P Pareto-optimal architectures to get the optimized search space $\hat{\mathcal{A}}_s$. Mathematically, we denote $\mathbf{e}_l^p = \{op_n^l | g_n^l = 1, n \in \mathbf{K}_l\}$ as the selected operations of l -th layer for the p -th Pareto-optimal architecture a_p , and denote $\hat{\mathbf{E}}_l^s$ as the optimized search space subset of the layer l in $\hat{\mathcal{A}}_s$. We have $\hat{\mathbf{E}}_l^s = \bigcup_{p=1}^P \mathbf{e}_l^p$.

3.3.3 Inheritance by Search Space Replenishment.

To obtain the new search space subset \mathcal{A}'_s for the next round of evolution, we randomly sample a certain amount of

| Network | Params | FLOPs | Top-1 |
|------------------------------|-------------|-------------|-------------|
| MobileNetV2 1.4 [32] | 6.9M | 585M | 74.7 |
| ShuffleNetV2 2× [24] | - | 591M | 74.9 |
| NASNet-A [59] | 5.3M | 564M | 74.0 |
| DARTS [21] | 4.7M | 574M | 73.3 |
| Proxyless-mobile [7] | 4.1M | 320M | 74.6 |
| FBNet-C [42] | 5.5M | 375M | 74.9 |
| NSENet-27 | 4.6M | 325M | 75.3 |
| NSENet | 4.6M | 330M | 75.5 |
| ShuffleNetV2 2× [24]† | - | 597M | 75.4 |
| MnasNet-A2 [37]† | 4.8M | 340M | 75.6 |
| MixNet-S [39]† | 4.1M | 256M | 75.8 |
| MixNet-M [39]† | 5.0M | 360M | 77.0 |
| MobileNetV3-large/1.25 [15]† | 7.5M | 356M | 76.6 |
| GreedyNAS-B [47]† | 5.2M | 324M | 76.8 |
| EfficientNet-B0 [38]†‡ | 5.3M | 390M | 76.3 |
| NSENet† | 7.6M | 333M | 77.3 |

Table 1. ImageNet results compared with state-of-the-art methods in the *mobile* setting. NSENet-27 denotes the network we discovered in the 27 OPs space. NSENet denotes the network found by exploring second space based on the search space subset we achieved from the 27 OPs space. † denotes model using extra modules such as swish activation [28] and SE module [16]. ‡ denotes model trained with AutoAugment [57].

candidate operations (excluding previously traversed operations) in the overall search space pool \mathcal{A} as the supplemental operations \mathcal{A}_r . Besides, the optimized search space subset $\hat{\mathcal{A}}_s$ obtained by aggregation will be inherited. Specifically, we have $\mathcal{A}'_s = \hat{\mathcal{A}}_s \cup \mathcal{A}_r$, where the new search space subset \mathcal{A}'_s is the union of the supplemental operations \mathcal{A}_r and the optimized search space subset $\hat{\mathcal{A}}_s$. After replenishment, the size of the resulting search space subset \mathcal{A}'_s will be the same as the size of the original \mathcal{A}_s with K candidate operations per layer. Finally we have $\mathcal{A}_s = \mathcal{A}'_s$.

4. Experimental Results

4.1. Search Space and Configurations

Search Space. For FLOPs constrained experiments, our primary search space consists of 27 distinct operations (27 OPs space). By inheriting the final search space subset derived from the 27 OPs space, we continue to search for three extra rounds over a complete new search space (second space) with OPs not covered in 27 OPs space to get the final result. For Latency constrained experiments, we use the 19 operations search space (19 OPs space) [18]. Notably, the number of OPs is quite large compared with [21, 7]. The detailed search space and structure of searched architectures can be found in the appendix.

The order of the candidate operations is randomly shuffled layer-wisely to minimize the potential influence from its ordering, yet all our experiments share the same set of shuffled sequences for a fair comparison. As shown in Figure 2, there could be an imbalanced number of operations being preserved in each layer, which means that some layers

| Network | Params | Latency* | Top-1 |
|-------------------------------|--------------|---------------|-------------|
| MobileNetV2 1.4 [32] | 6.9M | 8.9 ms | 74.7 |
| ShuffleNetV2 2× [24] | - | 6.8 ms | 74.9 |
| NASNet-A [59] | 5.3M | 23 ms | 74.0 |
| PNASNet [20] | 5.1M | 25 ms | 74.2 |
| Proxyless-GPU [7] | 7.1M | 7.9 ms | 75.1 |
| PC-NAS-L [18] | 15.3M | 10.3 ms | 77.5 |
| MixNet-S [39] † | 4.1M | 27 ms | 75.8 |
| MobileNetV3-large/1.25 [15] † | 7.5M | 10.5 ms | 76.6 |
| NSENet-GPU | 15.7M | 8.9 ms | 77.9 |

Table 2. ImageNet results compared with state-of-the-art methods in the *latency constrained* setting. NSENet-GPU is the network we found with the 19 OPs space [18]. † denotes model using extra modules such as swish activation [28] and SE module [16]. * latency for all networks are evaluated under the same setting (batch size 16 on GTX TITAN Xp GPU and TensorRT3 framework).

could be running out of candidates ahead of time. We bypass this issue by taking the moment when such a shortage happened as the end of an NSE search process.

Configurations. The layer-wise size of the search space subset K is set to 5 for FLOPs constraint and 6 for Latency constraint. To get the final model, we randomly sample 5 of the final Pareto optimal points and rescale them to approximate 330M FLOPs, the top-performing model after retraining is selected as the output model (see appendix for detailed setups). All experiments are performed on the ImageNet [31] dataset, where the validation set is constructed by 50K random images sampled from the training set.

4.2. FLOPs Constraint Results

The target τ for FLOPs constraint is 300M FLOPs. As shown in Table 1, our preliminary result (NSENet-27 in Table 1) derived from the 27 OPs space achieves 75.3% Top-1 accuracy with 325M FLOPs, which has already surpassed many of the hand-crafted or automatically designed architectures. When our search space subset continues to accommodate the second space, the result (NSENet in Table 1) further pushes the top performance of the derived model to 75.5% Top-1 accuracy. When auxiliary techniques are considered, our model (NSENet† in Table 1) consistently surpass previous models by a considerable margin.

4.3. Latency Constraint Results

The platform we optimize for is GTX TITAN Xp GPU and TensorRT3 framework. All the latency was evaluated with the batch size set to 16 to fully utilize the GPU resource. The resource constraint is targeted at 8 ms. The result is shown in Table 2, our NSENet-GPU obtains 77.9% top1 accuracy with 8.9 ms latency cost. Notably, while we share the identical search space with PC-NAS-L (see appendix), our searched model performs 0.4% better in terms of Top-1 accuracy with 1.4 ms less latency cost.

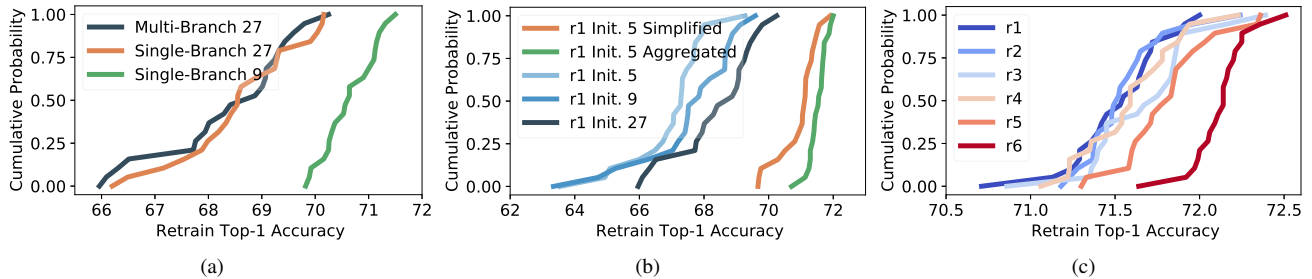


Figure 4. Search space comparisons conditioned on FLOPs. For each experiment, we randomly sample 20 architectures that have FLOPs within the interval of [323M, 327M]. Each model is then trained from scratch for 50 epochs. (b): r1 init n denotes the randomly initialized search space with layer-wise space size $K = n$, “Simplified” and “Aggregated” respectively denote the remained search space after simplification and the search space achieved by Pareto front aggregation. (c): r_n denotes the search space derived from the aggregation of Pareto optimal in the n -th round. The experiment performed in (b) and (c) are both based on the Multi-Branch 27 OPs space.

| Network(s) | Search Cost † | FLOPs | Top-1 |
|----------------------|---------------|---------------|---------------|
| Proxyless-mobile [7] | 200 ‡ | 320M | 74.6 |
| Proxyless [7]-rand6 | 2200* | 327M* | 74.2* |
| Proxyless [7]-27 | 4,000 | 336M | 74.5 |
| One-Shot [2]-27 | 4,000 | 339M | 73.5 |
| NSENet-27 $K = 5$ | 4,000 | 327M** | 75.0** |

Table 3. Comparison with different algorithms under 27 OPs space. † search cost refers to GTX 1080Ti GPU hours, ‡ number cited from [7], which uses GTX V100 GPU. * numbers are the average over 3 runs. ** numbers are the average of final results shown in Figure 5(a) ‘K=5’. The 27 OPs space contains the search space used by Proxyless-mobile and Proxyless-rand6. Proxyless-mobile [7], Proxyless-rand6, and Proxyless-27 are searched using the ProxylessNAS [7] algorithm.

4.4. Necessity of Search Space Evolution

Existing Methods on Large Search Space. In the introduction, we conclude that existing methods cannot deal with large search space effectively, which is our motivation. We assess 4 reasonably fast methods: DARTS [21], One-Shot [2], SPOS [14] and ProxylessNAS [7].

Algorithms like DARTS [21] retain all paths for optimization, making its search cost for our large 27 OPs space prohibitively high (about 27k GPU hours for 100 epochs). Simple one-shot algorithms like SPOS [14] do not retain all paths but fail to converge reasonably under the same setting we used. Specifically, the supernet trained using SPOS [14] on our 27 OPs space does not converge (1% Top-1 accuracy for 100 epochs), even after the attempts to tune the batch size, learning rate, or gradient clips.

With the one-shot strategy described in [2], the supernet trained on our 27 OPs space also converges poorly even after a training schedule that costs 4k GPU hours (aligned by increasing training epochs). The sampled Pareto-optimal architectures with 300M-350M FLOPs have Top-1 validation accuracies less than 30% on supernet and the final result is 73.5% Top-1 accuracy as shown in Table 3.

Algorithms like ProxylessNAS [7] utilizes learnable architecture parameters to gradually narrow the search space,

which helps to get a more reasonable result under a large search space. We run ProxylessNAS on our 27 OPs space and aligned search cost with NSE by increasing the number of training epochs, result denoted by Proxyless-27, is shown in Table 3. We can observe that ProxylessNAS produces a better result than One-Shot-27, but NSE still performs better than ProxylessNAS.

ProxylessNAS with Different Search Spaces. We choose the previous NAS algorithm with the highest 27 OPs space result compared in Table 3, ProxylessNAS, to further demonstrate how search space affects its performance. First, we compare its original result Proxyless-mobile with Proxyless-27. Proxyless-mobile is derived from a manually designed search space in ProxylessNAS [7] with only 6 OPs, a subset of our 27 OPs space. Nevertheless, its accuracy is even slightly higher than Proxyless-27, which means ProxylessNAS is not as effective in 27 OPs space as in 6 OPs space. Then we compare it with the result on random 6 OPs subsets, denoted as Proxyless-rand6. The random subsets are constructed by randomly sampling 6 OPs from our 27 OPs space for each layer of the whole search space. The result shows that the random subset space is not as good as the manually designed 6 OPs space or the full 27 OPs search space. Noticing that the search cost for Proxyless-rand6 is significantly higher than Proxyless-mobile even considering different GPU used, this is caused by less time-efficient OPs within the random subset.

Search Space Simplification Without Search Space Evolution. We also investigate whether search space simplification in Section 3.3.1 alone can handle large search space. We compare the cases of having layer-wise space size $K = 5$ and $K = 9$, which correspond to ‘K=5’ and ‘K=9’ in Figure 5(a) respectively. Note that the initial search space for $K = 5$ is a subset of the initial space used by $K = 9$. The first round of search results for $K = 5$ and $K = 9$ in Figure 5(a) shows that a larger initial search space ($K=9$) leads to an inferior result when compared with a smaller one ($K=5$), although both experiments




Figure 5. Convergence pattern along the NSE optimization trajectory. (a) Point evaluations of intermediate Pareto optimal points. Models are fully trained from scratch. The entire search process does not cover the full search space because of the inequity issue as mentioned in Section 4.1. (b) The progressive improvement of aggregated search space as search proceeds. The plot shares the same set of data with Figure 4(c) and is shown in mean accuracy with 95% confidence intervals.

enable search space simplification. The results show that search space simplification without evolving search space cannot handle large search space well.

5. Ablation Studies

5.1. Comparison of Search Space Quality

To have further insight into whether a search space is competitive for the specific task, we adopt the distribution estimate [27], *i.e.* the distribution of retrain accuracy derived from architectures randomly sampled within a search space, to evaluate search space quality. Figure 4(a) compares three search spaces to discuss the potential influence of enlarging search space in our experiments. The Multi-Branch 27 is the space we used for our key results, Single-Branch 27 stands for the single-branch version of our 27 OPs space, and the Single-Branch 9 is a subset of our 27 OPs space: only DW convolutions with kernel size $\{3, 5, 7\}$ and expand ratio $\{1, 3, 6\}$, a commonly used group of operations for NAS [7, 42, 47], is included. It can be seen that there is no outstanding gap between the single-branch and multi-branch space. Besides, while the Multi-Branch 27 search space theoretically has a higher upper bound for potential architectures, the single-branch 9 space, its subset, has a significantly better architecture distribution in terms of retrain accuracy. This gap, however, is the very issue we want to address through search space evolution.

5.2. Continual Convergence

To demonstrate that our approach can continuously improve the search space subset as well as the searched architectures, we first randomly select 5 Pareto-optimal architectures captured by every round of evolution on the preliminary space, then train them from scratch. As shown in Figure 5(a), a continuous trend of improvement in terms of

the upper bound and the distribution of results is observed. The best model for the entire process can be achieved by inspecting the most recent Pareto-optimal architectures. Furthermore, the quality of the aggregated search space has also been progressively improved as shown in Figure 4(c) and 5(b). After the suspension of an NSE process, we can still restart the process by reusing the final optimized search space subset and replenishing it with a second search space. In this way, our approach (‘K=5 w/ 2nd Space’ in Figure 5(a)) consistently achieves a gain.

5.3. Impact of Components

Layer-wise Size of Search Space Subset. As shown in Figure 4(b), comparing layer-wise space size $K = 5$ (‘r1 Init 5’ in 4(b)) and $K = 9$ (‘r1 Init 9’ in 4(b)), smaller K has a poorer accuracy distribution initially. However, a smaller K can still lead to a better convergence as shown by the comparison between ‘K=5’ and ‘K=9’ in Figure 5(a). This is coherent with our motivation to bypass the difficulty of large search space by progressive search space evolution, as a smaller search space is easier to optimize.

Simplification and Aggregation. To show the effectiveness of our search space simplification and aggregation, we plot the quality of a subset search space, together with its simplified and aggregated space in Figure 4(b). It shows that both search space simplification (‘r1 Init 5 Simplified’ in the figure) and search space aggregation (‘r1 Init 5 Aggregated’ in the figure) improve the search space quality by a considerable margin.

Lock and Rehearse. Another mandatory procedure in our pipeline is the Lock and Rehearse strategy, which prevents the inherited search space from being underestimated. The absence of such a regularization method could lead to a significant performance decrease as shown by the result ‘K=9 w/o L&R’ in Figure 5(a).

6. Conclusions

In this paper, we have introduced a new neural architecture search scheme called NSE. It targets large space architecture search by progressively accommodates new search space while maintaining the previously obtained knowledge. We further extended the flexibility of obtainable architectures by introducing a learnable multi-branch setting. Our proposed NSE scheme provides a consistent performance gain with a stream of incoming search spaces, which has minimized the necessity of search space engineering and leads to a step towards fully automatic neural architecture search.

Acknowledgement This work was supported by the Australian Research Council Grant DP200103223, FT210100228, and Australian Medical Research Future Fund MRFAI000085, Australian Future Fellowship.

References

- [1] Nsganetv2:evolutionary multi-objective surrogate-assisted neural architecture search. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 3
- [2] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *International Conference on Machine Learning*, pages 550–559, 2018. 1, 2, 3, 4, 7
- [3] Gabriel Bender, Hanxiao Liu, Bo Chen, Grace Chu, Shuyang Cheng, Pieter-Jan Kindermans, and Quoc V Le. Can weight sharing outperform random architecture search? an investigation with tunas. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14323–14332, 2020. 2
- [4] James Bergstra, Daniel Yamins, and David Daniel Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. 2013. 2
- [5] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Guttag. What is the state of neural network pruning? *arXiv preprint arXiv:2003.03033*, 2020. 3
- [6] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. In *International Conference on Learning Representations*, 2019. 1, 3
- [7] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv preprint arXiv:1812.00332*, 2018. 1, 2, 3, 4, 5, 6, 7, 8
- [8] Boyu Chen, Peixia Li, Baopu Li, Chen Lin, Chuming Li, Ming Sun, Junjie Yan, and Wanli Ouyang. Bn-nas: Neural architecture search with batch normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 1
- [9] Boyu Chen, Peixia Li, Chuming Li, Baopu Li, Lei Bai, Chen Lin, Ming Sun, Wanli Ouyang, et al. Glit: Neural architecture search for global and local image transformer. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 1
- [10] Xiangxiang Chu, Bo Zhang, Ruijun Xu, and Jixiang Li. Fairnas: Rethinking evaluation fairness of weight sharing neural architecture search. *arXiv preprint arXiv:1907.01845*, 2019. 3, 4
- [11] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pages 3123–3131, 2015. 5
- [12] Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Bichen Wu, Zijian He, Zhen Wei, Kan Chen, Yuandong Tian, Matthew Yu, Peter Vajda, et al. Fbnetv3: Joint architecture-recipe search using neural acquisition function. *arXiv preprint arXiv:2006.02049*, 2020. 3
- [13] Muyuan Fang, Qiang Wang, and Zhao Zhong. Betanas: Balanced training and selective drop for neural architecture search. *arXiv preprint arXiv:1912.11191*, 2019. 2, 3
- [14] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 544–560, 2020. 1, 2, 3, 4, 7
- [15] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1314–1324, 2019. 3, 6
- [16] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 2, 3, 6
- [17] Yiming Hu, Yuding Liang, Zichao Guo, Ruosi Wan, Xiangyu Zhang, Yichen Wei, Qingyi Gu, and Jian Sun. Angle-based search space shrinking for neural architecture search. *arXiv preprint arXiv:2004.13431*, 2020. 2
- [18] Xiang Li, Chen Lin, Chuming Li, Ming Sun, Wei Wu, Junjie Yan, and Wanli Ouyang. Improving one-shot nas by suppressing the posterior fading. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13836–13845, 2020. 6
- [19] Feng Liang, Chen Lin, Ronghao Guo, Ming Sun, Wei Wu, Junjie Yan, and Wanli Ouyang. Computation reallocation for object detection. In *International Conference on Learning Representations*, 2019. 1
- [20] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018. 6
- [21] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. 1, 2, 3, 4, 6, 7
- [22] Jie Liu, Chuming Li, Feng Liang, Chen Lin, Ming Sun, Junjie Yan, Wanli Ouyang, and Dong Xu. Inception convolution with efficient dilation search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11486–11495, 2021. 1
- [23] Zhichao Lu, Ian Whalen, Vishnu Boddeti, Yashesh Dhebar, Kalyanmoy Deb, Erik Goodman, and Wolfgang Banzhaf. Nsga-net: A multi-objective genetic algorithm for neural architecture search. *arXiv preprint arXiv:1810.03522*, 2018. 3
- [24] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 116–131, 2018. 3, 6
- [25] Jieru Mei, Yingwei Li, Xiaochen Lian, Xiaojie Jin, Linjie Yang, Alan Yuille, and Jianchao Yang. Atomnas: Fine-grained end-to-end neural architecture search. In *International Conference on Learning Representations*, 2019. 3
- [26] Hector Mendoza, Aaron Klein, Matthias Feurer, Jost Tobias Springenberg, and Frank Hutter. Towards automatically-tuned neural networks. In *Workshop on Automatic Machine Learning*, pages 58–65, 2016. 2
- [27] Ilija Radosavovic, Justin Johnson, Saining Xie, Wan-Yen Lo, and Piotr Dollár. On network design spaces for visual recog-

- niton. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1882–1890, 2019. 1, 3, 8
- [28] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017. 3, 6
- [29] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. *arXiv preprint arXiv:1802.01548*, 2018. 1, 2, 3
- [30] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc V Le, and Alexey Kurakin. Large-scale evolution of image classifiers. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2902–2911. JMLR. org, 2017. 1, 2
- [31] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 2, 6
- [32] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 6
- [33] Mingzhu Shen, Feng Liang, Ruihao Gong, Yuhang Li, Chuming Li, Chen Lin, Fengwei Yu, Junjie Yan, and Wanli Ouyang. Towards high performance extremely low-bit neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 1
- [34] Daniel L Silver and Robert E Mercer. The task rehearsal method of life-long learning: Overcoming impoverished data. In *Conference of the Canadian Society for Computational Studies of Intelligence*, pages 90–101. Springer, 2002. 4
- [35] Dimitrios Stamoulis, Ruizhou Ding, Di Wang, Dimitrios Lymberopoulos, Bodhi Priyantha, Jie Liu, and Diana Marculescu. Single-path nas: Designing hardware-efficient convnets in less than 4 hours. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 481–497. Springer, 2019. 1
- [36] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017. 3
- [37] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. *arXiv preprint arXiv:1807.11626*, 2018. 3, 6
- [38] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114, 2019. 6
- [39] Mingxing Tan and Quoc V Le. Mixconv: Mixed depthwise convolutional kernels. *CoRR, abs/1907.09595*, 2019. 1, 3, 6
- [40] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuandong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, et al. Fbnetv2: Differentiable neural architecture search for spatial and channel dimensions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12965–12974, 2020. 1, 3
- [41] Martin Wistuba, Ambrish Rawat, and Tejaswini Pedapati. A survey on neural architecture search. *arXiv preprint arXiv:1905.01392*, 2019. 1
- [42] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. *arXiv preprint arXiv:1812.03443*, 2018. 3, 6, 8
- [43] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 3
- [44] Zhicheng Yan, Xiaoliang Dai, Peizhao Zhang, Yuandong Tian, Bichen Wu, and Matt Feiszli. Fp-nas: Fast probabilistic neural architecture search. *arXiv preprint arXiv:2011.10949*, 2020. 3
- [45] Antoine Yang, Pedro M Esperança, and Fabio M Carlucci. Nas evaluation is frustratingly hard. *arXiv preprint arXiv:1912.12522*, 2019. 1
- [46] Tien-Ju Yang, Andrew Howard, Bo Chen, Xiao Zhang, Alec Go, Mark Sandler, Vivienne Sze, and Hartwig Adam. Netadapt: Platform-aware neural network adaptation for mobile applications. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 285–300, 2018. 5
- [47] Shan You, Tao Huang, Mingmin Yang, Fei Wang, Chen Qian, and Changshui Zhang. Greedynas: Towards fast one-shot nas with greedy supernet. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1999–2008, 2020. 6, 8
- [48] Jiahui Yu and Thomas S Huang. Universally slimmable networks and improved training techniques. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1803–1811, 2019. 3
- [49] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas Huang, Xiaodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. *arXiv preprint arXiv:2003.11142*, 2020. 1, 3
- [50] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. *arXiv preprint arXiv:1812.08928*, 2018. 3
- [51] Kaicheng Yu, Christian Sciuto, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. Evaluating the search phase of neural architecture search. *arXiv*, pages arXiv–1902, 2019. 2, 3, 4
- [52] Li Lyna Zhang, Yuqing Yang, Yuhang Jiang, Wenwu Zhu, and Yunxin Liu. Hardware-aware one-shot neural architecture search in coordinate ascent framework. *arXiv preprint arXiv:1910.11609*, 2019. 3
- [53] Xinbang Zhang, Zehao Huang, and Naiyan Wang. You only search once: Single shot neural architecture search via direct sparse optimization. *arXiv preprint arXiv:1811.01567*, 2018. 2, 3

- [54] Yuge Zhang, Quanlu Zhang, and Yaming Yang. How does supernet help in neural architecture search? *arXiv preprint arXiv:2010.08219*, 2020. [2](#)
- [55] Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu. Practical block-wise neural network architecture generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2423–2432, 2018. [1](#), [2](#), [3](#)
- [56] Dongzhan Zhou, Xinchu Zhou, Wenwei Zhang, Chen Change Loy, Shuai Yi, Xuesen Zhang, and Wanli Ouyang. Econas: Finding proxies for economical neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11396–11404, 2020. [1](#)
- [57] Barret Zoph, Ekin D Cubuk, Golnaz Ghiasi, Tsung-Yi Lin, Jonathon Shlens, and Quoc V Le. Learning data augmentation strategies for object detection. *arXiv preprint arXiv:1906.11172*, 2019. [6](#)
- [58] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. [1](#), [2](#), [3](#)
- [59] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8697–8710, 2018. [1](#), [2](#), [3](#), [6](#)