

Weight-Sharing Neural Architecture Search: A Battle to Shrink the Optimization Gap

Lingxi Xie, Xin Chen, Kaifeng Bi, Longhui Wei, Yuhui Xu, Zhengsu Chen, Lanfei Wang, An Xiao, Jianlong Chang, Xiaopeng Zhang, and Qi Tian, *Fellow, IEEE*

Abstract—Neural architecture search (NAS) has attracted increasing attentions in both academia and industry. In the early age, researchers mostly applied **individual** search methods which sample and evaluate the candidate architectures separately and thus incur heavy computational overheads. To alleviate the burden, **weight-sharing** methods were proposed in which exponentially many architectures share weights in the same **super-network**, and the costly training procedure is performed only once. These methods, though being much faster, often suffer the issue of instability. This paper provides a literature review on NAS, in particular the weight-sharing methods, and points out that the major challenge comes from the **optimization gap** between the super-network and the sub-architectures. From this perspective, we summarize existing approaches into several categories according to their efforts in bridging the gap, and analyze both advantages and disadvantages of these methodologies. Finally, we share our opinions on the future directions of NAS and AutoML. Due to the expertise of the authors, this paper mainly focuses on the application of NAS to computer vision problems and may bias towards the work in our group.

Index Terms—AutoML, Neural Architecture Search, Weight-Sharing, Super-Network, Optimization Gap, Computer Vision.

1 INTRODUCTION

THE rapid development of deep learning [1] has claimed its domination in the area of artificial intelligence. In particular, in the computer vision community, deep neural networks have been successfully applied to a wide range of challenging problems including image classification [2], [3], [4], [5], [6], object detection [7], [8], [9], [10], semantic segmentation [11], [12], [13], boundary detection [14], pose estimation [15], *etc.*, and surpassed conventional approaches based on hand-designed features (*e.g.*, SIFT [16]) by significant margins.

Despite the success of these efforts, researchers quickly encountered the difficulty that the network design space has become much more complex than that in the early deep learning era. Several troubles emerged. For example, the manually designed networks do not generalize to various scenarios, *e.g.*, a network designed for holistic image classification may produce inferior performance in dense image prediction tasks (semantic segmentation, image super-resolution, *etc.*). In addition, it is sometimes required to design networks under specific hardware constraints (*e.g.*, FLOPs, memory, latency, *etc.*), and manual design is often insufficient to explore a large number of possibilities. Motivated by these factors, researchers started to seek for

the solution that can flexibly adjusted to different scenarios in minimal efforts.

Automated machine learning (AutoML [17]) appeared recently as a practical tool which enables researchers and even engineers without expertise in machine learning to use off-the-shelf algorithms conveniently for different purposes. As a subarea of AutoML, **neural architecture search** (NAS) focuses on designing effective neural networks in an automatic manner [18], [19], [20]. By defining a large set of possible network architectures (the set is often referred to as the search space), NAS can explore and evaluate a large number of networks which have never been studied before. Nowadays, NAS has achieved remarkable performance gain beyond manually designed architectures, in particular, in the scenarios that computational costs are constrained, *e.g.*, in the mobile setting [21] of ImageNet classification.

In the early age, NAS algorithms often sampled a large number of architectures from the search space and then trained each of them from scratch to validate its performance. Despite their effectiveness in finding high-quality architectures, these approaches often require heavy computational overhead (*e.g.*, tens of thousands of GPU days [20]), which hinders them from being transplanted to various applications. To alleviate this issue, researchers proposed to re-use the weights of previously optimized architectures [22] or share computation among different but relevant architectures [23] sampled from the search space. This direction eventually led to the idea that the search space is formulated into an over-parameterized **super-network** so that the sampled **sub-architectures** get evaluated without additional optimization [24], [25]. Consequently, the search costs have been reduced by several orders of magnitudes, *e.g.*, recent advances allow an effective architecture to be found within 0.1 GPU-days on the CIFAR10 dataset [26], or within 2.0 GPU-days on the ImageNet dataset [27].

- Lingxi Xie, Xin Chen, Kaifeng Bi, Longhui Wei, Zhengsu Chen, An Xiao, Jianlong Chang, Jianlong Chang, Xiaopeng Zhang, and Qi Tian are with Huawei Inc., China.
E-mail: 198808xc@gmail.com, 1410452@tongji.edu.cn, bikafeng1@huawei.com, weilonghui1@huawei.com, chen-zhengsu1@huawei.com, xiaoan1@huawei.com, jianlong.chang@huawei.com, zxphistory@gmail.com, tian.qi1@huawei.com
- Yuhui Xu is with Shanghai Jiao Tong University, China.
E-mail: yuhuiXu@sjtu.edu.cn
- Lanfei Wang is with Beijing University of Posts and Telecommunications, China.
E-mail: wanglanfei@bupt.edu.cn

Manuscript received Month Date, 2020.

This paper focuses on the weight-sharing methodology for NAS, and investigate the critical weakness of it, **instability**, which refers to the phenomenon that individual runs of search can lead to very different performance. We owe the instability to the **optimization gap** between the super-network and its sub-architectures – in other words, the search procedure tries to optimize the super-network, but a well-optimized super-network does not necessarily produce high-quality sub-architectures. This is somewhat similar to the well-known over-fitting problem in machine learning, and thus can be alleviated from different aspects. We try to summarize the existing efforts into a unified framework, and thus we can answer the questions of where the community is, what the main difficulties are, what to do in the near future.

We noticed previous surveys [28], [29], [30], [31], [32], [33] which used three factors (*i.e.*, search space, search strategy, evaluation method) to categorize NAS approaches. We inherit this framework but merges the latter two factors since they are highly coupled in recent works. Differently, this paper focuses on the weight-sharing methodology and investigate it from the perspective of shrinking the optimization gap. As a disclaimer, this paper discusses the modern neural architecture search methods, *i.e.*, in the context of deep learning, so we do not cover the early work on finding efficient neural networks [34], [35], [36], [37].

The remainder of this paper is organized as follows. Section 2 first formulates NAS into the problem of finding the optimal sub-architecture(s) in a super-network, and introduces both individual and weight-sharing approaches to explore the search space efficiently. Then, Section 3 is the main part of this paper, in which we formalize the optimization gap to be the main challenge of weight-sharing NAS, based on which we review a few popular but preliminary solutions to shrink the gap. Next, in Section 4, we discuss a few other directions that are related to NAS. Finally, in Section 5, we look into the future of NAS and AutoML, and put forward a few open problems in the research field.

2 NEURAL ARCHITECTURE SEARCH

In this section, we provide a generalized formulation for neural architecture search (NAS) and build up a notation system which will be used throughout this paper. We first introduce the fundamental concepts of NAS in Section 2.1. Then, two basic elements, the search space and the search strategy (including the evaluation method), are introduced in Sections 2.2 and 2.3. Lastly, there is a brief summary in Section 2.4.

2.1 Overall Framework

The background of NAS is identical to regular machine learning tasks, in which a dataset \mathcal{D} with N elements is given, and each element in it is denoted by an input-output pair, $(\mathbf{x}_n, \mathbf{y}_n)$. The goal is to find a model that works best on the dataset and, more importantly, transfers well to the testing scenario. For later convenience, we use $\mathcal{D}_{\text{train}}$, \mathcal{D}_{val} and $\mathcal{D}_{\text{test}}$ to indicate the training, validation and testing subsets of \mathcal{D} , respectively. There are some standard datasets for NAS. For image classification, the most popular datasets

Algorithm 1: A Generalized Pipeline of NAS

Input : The search space \mathcal{S} , the training and validation datasets $\mathcal{D}_{\text{train}}$ and \mathcal{D}_{val} , the evaluation method $\text{Eval}(\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle; \mathcal{D}_{\text{val}})$;
Output: The optimal model $\langle \mathbb{S}^*, \omega^*(\mathbb{S}^*) \rangle$;
1 Initialize a probabilistic distribution \mathfrak{P} on \mathcal{S} ;
2 **repeat**
3 Sample an architecture \mathbb{S} from \mathfrak{P} ;
4 Train \mathbb{S} on $\mathcal{D}_{\text{train}}$ to obtain $\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle$;
5 Compute $\text{Acc} = \text{Eval}(\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle; \mathcal{D}_{\text{val}})$;
6 **if** Acc surpasses the best seen value **then**
7 Update the best model as $\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle$;
8 **end**
9 Using the evaluation result to update \mathfrak{P} ;
10 **until** convergence **or** the time limit is reached;
Return: The best seen model $\langle \mathbb{S}^*, \omega^*(\mathbb{S}^*) \rangle$.

are CIFAR10 [38] and ImageNet-1K [39], [40]. For natural language modeling, the most popular ones are Penn Treebank and WikiText-2 [41].

NAS starts with defining a **search space**, \mathcal{S} . Each element $\mathbb{S} \in \mathcal{S}$ is a network architecture that receives an input \mathbf{x} and delivers an output \mathbf{y} , which we denote as $\mathbf{y} = \mathbf{f}^{\mathbb{S}}(\mathbf{x}; \omega)$ where $\omega \doteq \omega(\mathbb{S})$ indicates the network weights (*e.g.*, in the convolutional layers). Note that the form and dimensionality of ω depends on \mathbb{S} . We use $\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle$ to denote an architecture equipped with the optimal network weights. The goal of NAS is to find the optimal architecture \mathbb{S}^* which produces the best performance (*e.g.*, classification accuracy) on the testing set:

$$\mathbb{S}^* = \arg \max_{\mathbb{S} \in \mathcal{S}} \text{Eval}(\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle; \mathcal{D}_{\text{test}}), \quad (1)$$

where $\text{Eval}(\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle; \mathcal{D}_{\text{test}})$ is the function that measures how $\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle$ behaves on the testing set, $\mathcal{D}_{\text{test}}$. But, in real practice, $\mathcal{D}_{\text{test}}$ is hidden in the search procedure, so the training and validation sets are often used instead, and thus the objective becomes:

$$\begin{aligned} \mathbb{S}^* &= \arg \max_{\mathbb{S} \in \mathcal{S}} \text{Eval}(\langle \mathbb{S}, \omega^*(\mathbb{S}) \rangle; \mathcal{D}_{\text{val}}) \\ \text{s.t.} \quad \omega^*(\mathbb{S}) &= \arg \min_{\omega} \mathcal{L}(\omega(\mathbb{S}); \mathcal{D}_{\text{train}}), \end{aligned} \quad (2)$$

where $\mathcal{L}(\omega(\mathbb{S}); \mathcal{D}_{\text{train}})$ is the loss function (*e.g.*, cross-entropy) computed on the training set.

Most often, the search space \mathcal{S} is sufficiently large to enable powerful architectures to be found. Therefore, it is unlikely that an exhaustive search method can explore the large space with reasonable computational costs. Therefore, a practical solution is to use a heuristic search method which defines \mathfrak{P} , a probabilistic distribution on \mathcal{S} , and adjusts the probability, $p(\mathbb{S})$, that an element \mathbb{S} is sampled according to the rewards from those already sampled architectures. This is to say, there is a more important factor, the **search strategy**, which determines the way of adjusting the probabilistic distribution. Note that each search strategy is equipped with an **evaluation method** which judges the quality of a sampled architecture.

A generalized framework of NAS is described in Algorithm 1. In what follows, we briefly review the popular choices of the search space and search strategy.

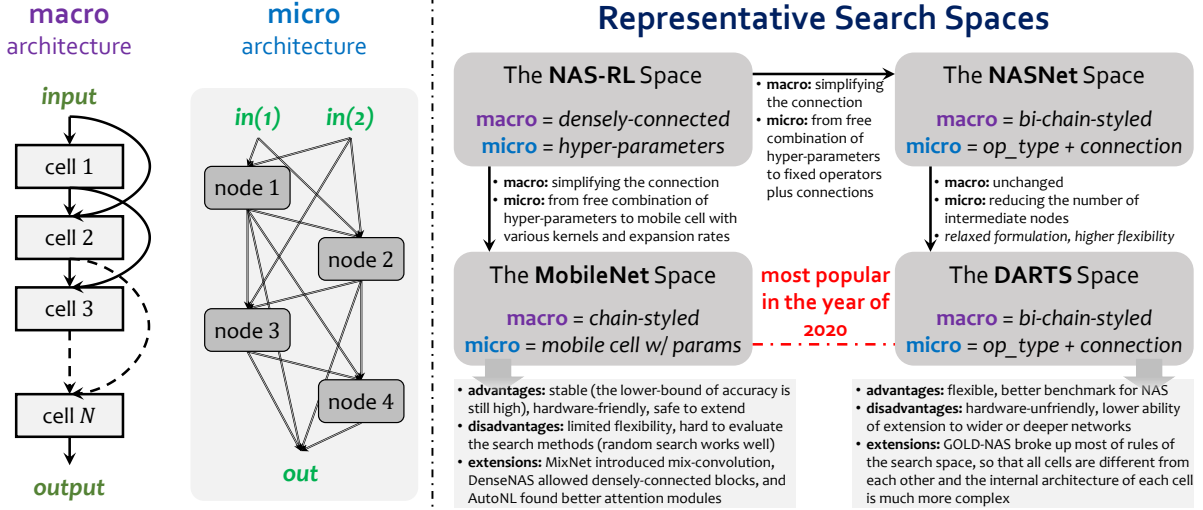


Fig. 1. **Left:** an example of the search space. For simplicity, we show a macro architecture where each cell is allowed to receive input from two precursors, and a micro architecture (cell) where the type and existence of the inter-node connections are searchable. **Right:** the relationship between four representative search spaces introduced in Sections 2.2.2–2.2.5, respectively. For the two most popular search spaces to date, MobileNet and DARTS, we summarize the advantages and disadvantages below, and point out some extensions.

2.2 Search Space

The search space, \mathcal{S} , determines the range of architecture search. Increasing the scale of the search space, $|\mathcal{S}|$, is both intriguing and challenging. On the one hand, a larger space enables the search algorithm to cover more architectures so that the chance of finding a powerful architecture is increased. On the other hand, a larger search space makes the search algorithm more difficult to converge, because $\text{Eval}(\langle \mathcal{S}, \omega \rangle; \mathcal{D}_{\text{val}})$ can be a very complicated function with respect to \mathcal{S} – more importantly, there is no guarantee of any of its properties in optimization. Hence, the number of samples required for depicting the function is often positively correlated to $|\mathcal{S}|$.

There exist many types of search spaces. There are often two steps of constructing a search space, namely, the macro and micro architectures of the space. As shown in Figure 1, the **macro** architecture determines the overall backbone of the network, and the **micro** architecture determines the details of each network unit. Sometimes, the micro architectures are referred to as **cells** and the number of cells can vary between the search and re-training phases or to adjust different tasks. Each NAS algorithm searches for either the macro or micro architecture but not necessarily both of them. Today, most NAS methods fix the macro architecture and investigate better micro architectures, which largely constrains the flexibility of the search space.

In what follows, we instantiate the macro and micro architectures using some popular examples. Note that some methods considered different search spaces for vision and language modeling and we will mainly focus on the vision part for which the search space is often more complicated. The relationship among these popular search spaces is illustrated in Figure 1.

2.2.1 The *EvoNet* Search Space

EvoNet [18] is one of the first works that investigated an open search space, *i.e.*, the architectures are not limited

within the combination of a fixed number of components, and instead, the architecture is allowed to grow arbitrarily with requirement. Thus, the search space is latently defined by a set of **actions** on the architecture, *e.g.*, inserting a convolutional layer to a specific position, altering the kernel size of a convolutional layer, adding a skip connection between two layers, *etc.*, and some actions can be reversely performed, *e.g.*, removing the convolutional layer from a specific position.

In an open search space, there is not a fixed number of architectures because the network shape (*e.g.*, depth) is only constrained by the hardware platform – even provided the length of the search procedure and the hardware limit, the accurate space size is difficult to estimate. A very loose estimation is obtained by assuming that the search procedure starts with a fixed architecture and has T steps in total. Thus, if each step can take up to C actions, the search space size satisfies $|\mathcal{S}| \leq 1 + C + C^2 + \dots + C^T$.

Other open search spaces. There are other efforts that tried to explore an open space. For example, [42] formulated the search procedure in a hierarchical way and allow the connection between some layers to be replaced by small modules named motifs; and [43] learned an action space to generate complex architectures. As we shall explain in Section 5, open-space search is likely to be the future trend of NAS [44], but the community needs a robust search strategy that can efficiently explore the challenging spaces.

2.2.2 The *NAS-RL* Search Space

NAS-RL [20] is the work that defined the modern formulation of NAS. It also sculptured the framework that uses the macro and micro architectures to determine the network. The **densely-connected** macro architecture has N layers, each of which can be connected to any of its precursors. Besides the direct precursor, other connections can be present or absent, bringing a total of $2^{(N-1)(N-2)/2}$ possibilities. Each layer has a micro architecture involving the hyper-parameters of the filter height, filter width, stride height,

stride width, and the number of filters. For simplicity, each hyper-parameter is chosen from a pre-defined set, *e.g.*, the number of filters is one among $\{24, 36, 48, 64\}$. Let C be the number of combinations for each layer, then the total number of architectures is $|\mathcal{S}| = C^N \times 2^{N(N-1)/2}$. In the original paper, C varies with different search options and the largest C is 576.

The search space for language modeling. The goal is to construct a cell that connects the neighboring states of a recurrent network, in which there are three input variables and two output variables. Similarly, the recurrent controller is used to determine the architecture with a fixed number of intermediate nodes. For each node that is related to the input data and the hidden state, there is an operator (*e.g.*, add, elem-mult, *etc.*) and an activation function (*e.g.*, tanh, ReLU, sigmoid, *etc.*), and the controller also determines how to connect the memory states to the temporary variables inside the tree.

Related search spaces. Genetic-CNN [19] shared the similar ideology of building the architecture in a macro-then-micro manner. The architecture is composed of a fixed number of stages and each stage has a fixed number of layers. Pooling layers are inserted between neighboring stages. Only the regular 3x3-convolution is used, but the inter-layer connectivity in the same stage can be searched. The number of architectures in each stage is $2^{N(N-1)/2}$ where N is the number of layers in the stage.

2.2.3 The NASNet Search Space

NASNet [45] introduced the idea that the macro architecture contains repeatable cells, so that the search procedure is performed within a relatively shallow network but the final architecture can be much deeper. In addition, the densely-connected macro architecture is simplified into **bi-chain-styled**, *i.e.*, each layer steadily receives input from two precursors. Regarding the micro architecture, each cell contains N hidden nodes, indexed from 0 to $N - 1$, and the two input nodes are indexed as -2 and -1 . Each hidden node is connected to two nodes with smaller indices, with the operator of each connection chosen from a pre-defined set of C_1 candidates (*e.g.*, dil-conv-3x3, sep-conv-5x5, identity, max-pool-3x3, *etc.*) and the summarization function chosen from a pre-defined set of C_2 candidates (*e.g.*, sum, concat, product, *etc.*). The outputs of all intermediate nodes are concatenated into the output of the cell. C_1 is 13 in NASNet, but is often reduced to a smaller number in the follow-up methods because researchers noticed that there exist some weak operators that are rarely used (*e.g.*, max-pool-7x7 and 1x7-then-7x1-conv), plus, shrinking the search space can reduce the cost of search. We denote C_3 as the number of different topologies in a cell, which is computed via $\binom{2}{2} \times \binom{3}{2} \times \dots \times \binom{N+1}{2}$. The number of architectures in a normal cell or a reduction cell is thus $(C_1^2 \times C_2)^N \times C_3$.

Related search spaces. Many follow-up methods have used the same search space but simplified either the candidate set or the cell topology. For example, PNASNet [46] and AmoebaNet [47] reduced C_1 to 8, and DARTS [25] reduced N from 5 to 4. In addition, DARTS did not use individual summarization functions for different nodes and used sum as the default choice.

2.2.4 The DARTS Search Space

DARTS [25] has a very similar search space to NASNet and follows a weight-sharing search pipeline (see Section 2.3) so that a more flexible super-network is constructed. In the standard DARTS space, the set of candidate operators have been defined as sep-conv-3x3, sep-conv-5x5, dil-conv-3x3, dil-conv-5x5, max-pool-3x3, avg-pool-3x3, and skip-connect. A dummy zero operator is added to the search procedure but is not allowed to appear in the final architecture. To ease the formulation of differentiable architecture search, each layer receives input data from an arbitrary number of previous layers in the same cell though only two of them are allowed to survive. Provided that 4 intermediate cells are present, the number of architectures for both the normal and reduction cells is $\binom{2}{2} \binom{3}{2} \binom{4}{2} \binom{5}{2} \times 7^8 \approx 1.0 \times 10^9$, and the entire space has around 1.1×10^{18} architectures.

The search space for language modeling. Similar to NAS-RL, the DARTS space is also used for finding an improved recurrent cell in the language models. The input sources, the number of intermediate nodes, and the connectivity between these nodes are adjusted accordingly, but still, the outputs of all intermediate nodes are concatenated.

Related search spaces. The super-network formulation of DARTS has offered opportunities to create some new search spaces. For example, one can shrink the search space by reducing the number of operators (*e.g.*, only sep-conv-3x3 and skip-connect used [48]), or expanding the search space by enabling multiple operators to be preserved in each edge [49], [50], adding new operators [51], allowing the cells of the same type to be individual from each other [52], and relax the constraints that each layer receives input from two previous layers [27].

2.2.5 The MobileNet Search Space

An alternative way [137], [138] of simplifying the NAS-RL search space is to borrow the design of MobileNet. The macro architecture is set to be **chain-styled**, *i.e.*, each layer receives input from its direct precursor. The micro architecture is also constrained to be one of the MobileNet (MB) cells [21]. Changeable options include the type of separable-convolution [342] (the regular convolution can be replaced by de-convolution), the type of skip-layer connection (can be either identity or some kinds of pooling), the number of channels, the kernel size, the expansion ratio (how the intermediate layer of the MB cell magnitudes the number of basic number of channels). Beyond these, the number of layers within each cell (or referred to as a block) can also be searched [139], [151]. So, the number of architectures in each cell is basically $C^1 + C^2 + \dots + C^L$ where C is the number of combinations of each layer and L is the maximal number of layers in a cell. The above number is further powered up by N , the number of cells in the final architecture.

Related search spaces. MNASNet [139] was the predecessor of the MobileNet search space, where there are more searchable options including the number of cells in each stage, *i.e.*, L is relatively large. After that, researchers soon noticed that (i) there are duplicates in using different numbers of layers in neighboring cells and (ii) preserving a small number of MB cells is sufficient to achieve good performance. So, a few follow-up methods [143], [144], [343] have

Search Space / Task	Important Work	Other Work
EvoNet [18] & open	[18], [42]	[44]
NAS-RL [20]	[20]	[53], [54]
NASNet [45]	[23], [45], [46], [47], [55], [56], [57]	[58], [59], [60], [61], [62], [63], [64], [65], [66], [67], [68], [69], [70], [71], [72], [73], [74], [75], [76], [77]
DARTS [25]	[25], [26], [48], [49], [78], [79], [80], [81], [82], [83], [84], [85]	[27], [43], [50], [51], [52], [86], [87], [88], [89], [90], [91], [92], [93], [94], [95], [96], [97], [98], [99], [100], [101], [102], [103], [104], [105], [106], [107], [108], [109], [110], [111], [112], [113], [114], [115], [116], [117], [118], [119], [120], [121], [122], [123], [124], [125], [126], [127], [128], [129], [130], [131], [132], [133], [134], [135], [136]
MobileNet [137], [138]	[137], [138], [139], [140], [141], [142], [143], [144], [145], [146]	[147], [148], [149], [150], [151], [152], [153], [154], [155], [156], [157], [158], [159], [160], [161], [162], [163], [164], [165], [166], [167], [168], [169], [170], [171], [172], [173], [174], [175], [176], [177], [178], [179], [180], [181], [182], [183], [184], [185], [186]
Other Spaces	[19], [22], [24], [187], [188], [189], [190], [191], [192], [193], [194], [195]	[196], [197], [198], [199], [200], [201], [202], [203], [204], [205], [206], [207], [208], [209], [210], [211], [212], [213], [214], [215], [216], [217], [218], [219], [220], [221], [222], [223], [224], [225], [226], [227], [228], [229], [230], [231], [232], [233], [234], [235], [236], [237], [238], [239], [240], [241], [242], [243], [244], [245], [246], [247], [248], [249]
Object Detection	[250], [251], [252]	[253], [254], [255], [256], [257], [258], [259], [260]
Semantic Segmentation	[261], [262]	[255], [263], [264], [265], [266], [267], [268]
Generative Models	[269]	[270], [271], [272], [273], [274], [275]
Medical Image Analysis	[276]	[277], [278], [279], [280], [281], [282], [283], [284], [285], [286]
Adversarial Learning	–	[283], [287], [288]
Low-level Vision	–	[289], [290], [291], [292], [293], [294], [295], [296], [297], [298]
Video Processing	–	[299], [300], [301], [302]
Multi-task Learning	–	[303], [304], [305], [306]
Language Modeling	[307]	[308], [309], [310], [311], [312], [313], [314], [315], [316]
Speech Recognition	–	[317], [318], [319]
Graph Networks	–	[320], [321], [322]
Recommendation	–	[323], [324], [325], [326]
Miscellaneous	[327] (captioning)	[328] (attribute), [329], [330] (person re-id), [331] (inpainting), [332] (pose), [333] (reconstruction), [334] (knowledge graph), [335] (style transfer), [336] (saliency), [337] (privacy), [338] (deep image prior), [339] (counting), [340] (text recognition), [341] (discussion)

TABLE 1

Summary of search spaces. Note that (i) there are two categorization methods, *i.e.*, by space configuration (NAS-RL, NASNet, MobileNet, DARTS, open, others) and by non-classification task (detection, segmentation, *etc.*); (ii) one work may appear in multiple rows because it deals with more than one space or tasks. We use both the total number of Google scholar citations and the number per time unit, according to the statistics on the July 31st, 2020, to judge if this work is an ‘important work’ (around 10% top-cited papers are considered important).

reduced L into 1 and reduced C into 6 (*i.e.*, MB cells with a 3×3 , 5×5 , or 7×7 convolution, and an expansion ratio of 3 or 6). The MobileNet space is also friendly to augmentation, including introducing mixed convolution [145], [161], [162], allowing dense connections between layers [151], [173], or adding attention blocks [170].

2.2.6 Other Search Spaces

Besides the previous search spaces that were widely explored, there exist other interesting methods. For example, some early efforts have either applied the technique of network transformation [344] or morphism [345] to generate complex architectures from the basic one [22], [58], [327], or made use of shared computation to combine exponentially many architectures with cheap computation [346], [347], [348]. Researchers have also designed other methods for architecture design such as repeating one-of-many choices [187], [188] or defining a language for encoding [190], [349]. However, in the current era, increasing the search space arbitrarily seems not the optimal choice, since the search strategy is not sufficiently robust and may produce weak results that are on par with random search [195].

There are efforts in applying NAS methods to other tasks, for which the search spaces need to adjust accordingly. For **object detection**, there are efforts of transferring powerful backbones in classification [252], searching for

a new backbone [250], [260], [350], considering the way of fusing multi-stage features [251], [253], and integrating multiple factors into the search space [254], [256]. Regarding **semantic segmentation** in which spatial resolution is more sensitive [351], the most interesting part seems to design of the down-sampling and up-sampling path [262], [352], though searching for efficient blocks is also verified useful [152], [266], [276]. NAS is also applied for **image generation**, in particular, designing network architectures for the generative adversarial networks (GAN) [269], [270], [353]. Recently, this method has been extended to conditional GAN where the generator varies among classes and thus the search space becomes exponentially larger [274]. Beyond the search spaces that are extendable (see Sections 2.2.2 and 2.2.4), there were efforts that directly search for powerful architectures for **RNN design** [354], [355], [356], [357] or **language modeling** [307], [308], [309], [314].

Table 1 provides a summary of search spaces for classification and the spaces transferred to different tasks.

2.3 Search Strategy and Evaluation Method

We combine the review of search strategies and evaluation methods into one part, because these two factors are closely correlated, *e.g.*, the individual heuristic search strategies often lead to an separate evaluation on each sampled ar-

chitecture, while the weight-sharing search strategies often integrate architecture optimization into architecture search.

2.3.1 Individual Heuristic Search Strategies

The most straightforward idea of architecture search is to enumerate all possible architectures from the search space and evaluate them individually. However, since the search space is often very large so that exhaustive search is intractable, researchers turned to heuristic search. The fundamental idea is to make use of the relationship between architectures, so that if an architecture reports good/bad performance, it is natural to tune up/down the probability of sampling its close neighbors. The algorithm starts with formulating a function, $p: \mathcal{S} \mapsto \mathbb{R}$, so that each architecture \mathbb{S} in the search space corresponds to a probability density that it is sampled, and the function gets updated after the algorithm gets rewards from the evaluation result of the sampled architectures. This basic framework is illustrated in Figure 2.

There are two popular choices of the heuristic algorithm, both of which maintain a latent form of the probabilistic distribution $p(\cdot)$. The first type, named the **genetic algorithm** or **evolutionary algorithm**, constructs a family of architectures during the search procedure [18], [19], and allows each architecture (or named an individual) to generate other architectures sharing similar properties. Typical genetic operations include crossover, mutation, selection or elimination, *etc.* The other type follows a different path, known as **reinforcement learning** [20], [45], [188], which constructs an architecture using a learnable policy, and updates the parameters in the policy according to the final performance. There have been efforts in combining the evolutionary algorithm and reinforcement learning for NAS [59], or comparing the efficiency between these two options [311] and/or against the random search baseline [47].

Regardless of the heuristic algorithm being used, such search methods suffer the problem of heavy computational overhead. This is mainly due to the need of re-training each sampled architecture from scratch. There are a few standard ways of alleviating the computational burden in the search procedure, including using a smaller number of channels, stacking fewer cells, or exploring on a smaller dataset. However, it can still take long for the algorithm to thoroughly traverse the search space and find a satisfying architecture, in particular when the training dataset is large. **The random search baseline.** When the probability function does not change with the evaluation results, the individual heuristic algorithm degenerates to random search in which the probability of each architecture being sampled is equal and does not change with time. Note that random search can find any high-quality architecture given sufficient computational costs, but it is verified less efficient than heuristic search [47] or weight-sharing search [179] methods. Random search is widely used as a baseline test of search efficiency, *e.g.*, evaluating the difficulty of the search space and thus validating the high cost-performance ratio of the proposed approach [25], [27], [48], [79].

2.3.2 Weight-Sharing Heuristic Search Strategies

A smart solution of fast search lies in sharing computation among different architectures, as many of them have similar




Fig. 2. The heuristic search method that maintains a probabilistic distribution on the search space. In each iteration, the algorithm samples an architecture following the distribution, performs evaluation individually, and updates the distribution accordingly. The design of this figure borrows the idea from [28].

properties, *e.g.*, beyond a well-trained architecture, only one layer is added or one layer is made wider. In these scenarios, it is possible to copy the weights of the trained network with a small part being transformed [22], so that the number of epochs required for optimizing the new network is significantly reduced.

Moving one step forward, an elegant solution of weight-sharing search is to construct the search space into a **super-network** [23], [24], [347], from which each architectures can be sampled as a **sub-architecture**, and different sub-architectures share corresponding modules. Consider a network that has a fixed topology with L edges and each edge has K options to be considered. The super-network contains LK units to be optimized, but after the optimization, there are up to K^L sub-architectures that can be sampled from the super-network. The search procedure is therefore largely accelerated because a large amount of computation has been shared among the sub-architectures.

However, it is worth emphasizing that sampling a sub-architecture with the weights directly copied from the super-network does not reflect its accurate performance on the validation set, because the weight-sharing training procedure was focused on optimizing the super-network as an entire but not the individual sub-architectures. This difficulty is named the optimization gap which we will formalize in Section 3.2. To alleviate the gap, typical efforts include fine-tuning the sub-architectures for more accurate performance prediction [23], facilitating the fairness during super-network training [144], and modeling the search space using graph-based methods [171].

Relationship to individual heuristic search. Note that, despite the super-network that organizes the search space in a more efficient manner, there still needs a standalone module to sample sub-architectures from the super-network. That is to say, either the evolutionary algorithm [142], [144] or the reinforcement learning algorithm [23], [56] can be used. Due to the effectiveness of sub-architecture sampling, it is also possible to split the search space into subspaces and enumerate all architectures in each of them [171].

2.3.3 Weight-Sharing Differentiable Search Strategies

An alternative way of weight-sharing search is to allow a relaxed formulation of the super-network [25], [57], [204] so that it is differentiable to both the network weights (mainly the convolutional kernels and batch normalization [381] coefficients) and architectural parameters (modeling the likelihood that the candidate edges and/or operators are preserved in the final architecture). The most intriguing

Search Strategy	Important Work	Other Work
Individual – Reinforcement	[20], [45], [46], [55], [139], [143], [188], [190], [251], [252]	[53], [54], [60], [61], [63], [203], [207], [209], [217], [223], [225], [243], [263], [264], [265], [269], [277], [280], [281], [303], [308], [309], [317], [320], [321], [331], [337], [354]
Individual – Evolutionary	[18], [19], [42], [47], [189], [194], [307], [327]	[44], [58], [71], [147], [200], [213], [214], [215], [216], [219], [240], [241], [242], [246], [278], [289], [290], [291], [292], [293], [299], [300], [304], [311], [312], [326], [333], [335], [336], [338], [355], [356]
Individual – Others	[187], [192], [193], [195], [206], [261]	[59], [65], [76], [149], [196], [197], [198], [201], [205], [208], [212], [220], [221], [222], [229], [230], [232], [239], [247], [328], [357], [358]
Weight-Sharing Heuristic	[22], [23], [24], [56], [142], [144], [146], [199]	[43], [62], [64], [66], [67], [69], [72], [77], [110], [123], [145], [148], [151], [152], [153], [156], [157], [158], [159], [160], [161], [162], [163], [164], [165], [166], [167], [168], [170], [171], [172], [173], [174], [175], [176], [177], [179], [182], [183], [184], [185], [186], [202], [211], [218], [226], [227], [228], [233], [235], [237], [238], [244], [245], [248], [249], [250], [254], [258], [259], [271], [274], [275], [284], [285], [287], [288], [295], [298], [305], [310], [314], [318], [325], [340], [359], [360], [361], [362], [363]
Weight-Sharing Differentiable	[25], [26], [48], [57], [78], [79], [80], [81], [82], [83], [137], [138], [262]	[27], [49], [50], [51], [52], [68], [70], [74], [84], [85], [86], [87], [88], [89], [90], [91], [92], [93], [94], [95], [96], [97], [98], [99], [100], [101], [102], [103], [104], [105], [106], [107], [108], [109], [111], [112], [113], [114], [115], [116], [117], [118], [119], [120], [121], [122], [124], [125], [126], [127], [128], [129], [130], [131], [132], [133], [134], [135], [136], [141], [150], [154], [155], [169], [178], [180], [181], [204], [224], [231], [234], [253], [255], [256], [257], [266], [267], [268], [270], [272], [273], [276], [279], [282], [283], [286], [294], [296], [297], [301], [302], [306], [313], [315], [316], [319], [322], [323], [324], [329], [330], [332], [334], [339], [341], [364], [365]
Predictor-based Search	[140], [191], [366]	[73], [75], [210], [236], [367], [368], [369], [370], [371], [372], [373], [374], [375], [376], [377], [378], [379], [380]

TABLE 2

Summary of search strategies. We categorize the individual heuristic search strategies into three parts, *i.e.*, using reinforcement learning, using evolutionary algorithms, and using other methods (*e.g.*, Bayesian optimization, *etc.*). Note that the categorization is not absolute, since the boundary between different methods is not always clear, *e.g.*, an individual search method may apply weight-sharing in a minor part, and ‘other’ individual methods may apply the ideology of reinforcement learning or evolutionary algorithms. For the judgment of ‘important work’, please see the caption of Table 1.

property of this formulation is that the architectural parameters can be optimized together with the network weights, so that at the end of the search procedure, the optimal sub-architecture is directly produced without the need of performing an additional sampling procedure, therefore, the computational costs are largely reduced.

The most popular example of differentiable search is DARTS [25], where the goal is to find the optimal sub-architecture for the normal and reduction cells as defined in Section 2.2.4. At the start, each edge is individually initialized so that all candidate operators have the same weight (*e.g.*, $1/C$ where C is the number of candidate operators). Then, during the search procedure, these architectural parameters are updated in a continuous space, so that the stronger (more effective) operators gradually take the lead. At the end of the search procedure, each edge undergoes a so-called discretization process, during which the strongest operator is preserved (and assigned with the weight of 1.0) while others are eliminated. The preserved sub-architecture (the architectural parameters are fixed) is fed into the re-training procedure for the optimal network weights.

Relationship to weight-sharing heuristic search. First, we emphasize that weight-sharing differentiable search also works on the MobileNet search space [137], [138], though it is more commonly used to explore the DARTS space. Essentially, differentiable search serves as an alternative way of approximately evaluating exponentially many architectures at the same time. The evaluation procedure is embedded into the search procedure and thus not explicitly performed. Therefore, the issue of optimization gap also exists in this framework. In particular, due to the inevitable approximation error of gradient computation, the differentiable search

procedure can sometimes run into a dramatic failure [52], [83], and researchers developed various methods to alleviate this problem including performing early termination [48], [83] or using progressive optimization [81], [102] to reduce the discretization error [133]. We will summarize these methods from another perspective in Section 3.3.

2.3.4 Predictor-Based Search Methods

Besides the aforementioned search methods, there are also efforts in exhaustively enumerating all architectures in the search space (often relatively smaller) and then using them as the ground-truth to test the efficiency of different search strategies. The most popular benchmarks include NAS-Bench-101 [382] that contains $\sim 423K$ architectures evaluated on CIFAR10 and NAS-Bench-201 [383] that contains 15,526 architectures evaluated on CIFAR10, CIFAR100, and ImageNet-16-120. We also noticed an extension of the similar idea to the NLP scope [384].

These benchmarks offer the opportunities of training predictors, *i.e.*, the input is an encoded neural architecture and the output is the network accuracy [366]. There have been efforts of exploring the internal relationship between architectures (an effective way is to design a good encoding method for the architectures [369], [371], [380]) for better prediction performance, in particular under a limited number of sampled architectures [75], [385]. The trained predictors can inspire algorithms beyond the toy search spaces [171], [371]. We will discuss more about this issue in Section 3.3.5.

Table 2 provides a summary of the frequently used search strategies, where we find that weight-sharing search is the current mainstream and future direction.

2.4 Summary

There is a close relationship between the search space and the search strategy together with evaluation method. When the search space is small and ‘tidy’ (*i.e.*, similar architectures are closely correlated in performance), individual heuristic search methods often work very well. However, when the search space becomes larger and more flexible, individual search methods often require heavy computational burden that is not affordable to most researchers and developers. Thus, since researchers believe that larger search spaces are the future trend of NAS, it has become an important topic to design efficient and stabilized weight-sharing search strategies. In the next part, we will owe the instability of weight-sharing architecture search to the optimization gap and discuss the solution to this problem.

As a side note, there have been research [386], [387] on designing the search space for architecture search or design. From the perspective of NAS, this is similar to the idea that gradually removes the candidate connections/operators during the search procedure [27], [81], [102], or undergoes a two-stage NAS algorithm that the network topology and operators are determined separately [52], [91], [284]. We carefully put forward the opinion that these multi-stage algorithms reflect the instability and immaturity of NAS algorithms, and we advocate for an algorithm that can explore a very large search space in an end-to-end manner.

In evaluating the quality of the search space and/or algorithms, choosing a good benchmark as well as a set of proper hyper-parameters is very important for NAS research. We suggest the readers to refer to [388] for more practices of research on NAS. Though, it was believed that evaluating the performance of NAS methods is often hard [389], [390], [391]. Different settings beyond supervised learning have been investigated in NAS, including like semi-supervised learning [392], self-supervised learning [131], unsupervised learning [115], [377], incremental learning [361], [393], federated learning [394], [395], *etc.*, showing the promising transferability of NAS methods. Last but not least, there are several toolkits for AutoML [17], [396], [397], [398], [399] that can facilitate the reproducibility of NAS methods.

3 TOWARDS STABILIZED WEIGHT-SHARING NAS

In this section, we delve deep into the weight-sharing NAS algorithms and try to answer the key question: why do these methods suffer search instability? Here, by **instability** we mean that the same algorithm can produce quite different results, in terms of either the searched architecture or the test accuracy, when it is executed for several times [79]. For this purpose, we start with formulating the weight-sharing NAS methods in Section 3.1 and formalizing the optimization gap in Section 3.2. Then, in Section 3.3, we summarize the existing solutions to alleviate the optimization gap, followed by a summary of unsolved issues in Section 3.4.

3.1 An Alternative Formulation of Weight-Sharing NAS

Weight-sharing NAS methods start with defining a super-network that contains all learnable parameters in the search space. This involves using a vectorized form, α , to encode each architecture $\mathbb{S} \in \mathcal{S}$, so that the super-network is

parameterized using both α and ω . We follow the convention [25] to rewrite the super-network as $\mathbf{y} = \mathbf{F}(\mathbf{x}; \alpha, \omega)$. Throughout the remaining part of this paper, we refer to α and ω as the architectural parameters and network weights, respectively. Though α takes discrete values, a common practice is to slack the search space and correspond each element of α to the probability that an operator appears in the final sub-architecture. Hence, finding the optimal sub-architecture \mathbb{S}^* comes down to finding the optimal architectural parameters α and performing discretization (mapping the possibly continuous α to a discrete form that corresponds to a valid \mathbb{S}), and thus Eqn (2) becomes:

$$\begin{aligned} \mathbb{S}^* &= \text{disc}(\alpha^*), \quad \alpha^* = \arg \min_{\alpha} \mathcal{L}(\omega^*(\alpha), \alpha; \mathcal{D}_{\text{val}}) \quad (3) \\ \text{s.t.} \quad \omega^*(\alpha) &= \arg \min_{\omega} \mathcal{L}(\omega, \alpha; \mathcal{D}_{\text{train}}), \end{aligned}$$

where $\mathcal{L}(\omega, \alpha; \mathcal{D}) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}^*) \in \mathcal{D}} [\text{CE}(\mathbf{F}(\mathbf{x}; \alpha; \omega), \mathbf{y}^*)]$ are the expected cross-entropy loss (computed in any training subset). We expect $\text{Eval}(\langle \mathbb{S}, \omega(\mathbb{S}) \rangle | \mathcal{D}) = 1 - \mathcal{L}(\omega, \alpha; \mathcal{D})$ when α corresponds to \mathbb{S} , but $\mathcal{L}(\omega, \alpha; \mathcal{D})$ is more flexible in dealing with the continuous form of α . Both $\mathcal{D}_{\text{train}}$ and \mathcal{D}_{val} are subsets of the training data but they are not necessarily different, *e.g.*, in one-level optimization methods [27], [97], both $\mathcal{D}_{\text{train}}$ and \mathcal{D}_{val} are the entire training set. The discretization function, $\text{disc}(\cdot)$, does not take an essential operation when the search space is not slacked, *e.g.*, in the weight-sharing heuristic search, α does not take continuous values and \mathbb{S}^* is directly derived from α^* .

Note that Eqn (3) can be solved using non-weight-sharing methods where $\omega(\alpha)$ have individual weights for different α . For weight-sharing methods, the key to optimize Eqn (3) is to share $\omega(\alpha)$ over different α and perform the optimization with respect to ω and α over all sub-architectures. In other words, this is to achieve an amortized optimum over all sub-architectures and α plays the role of the amortizing distribution. With this formulation, we provide an alternative perspective in discriminating weight-sharing heuristic and differentiable search. The **heuristic search** methods often keep all elements of α identical, *i.e.*, during training the super-network, the distribution \mathfrak{P} remains uniform over the entire search space. A standalone heuristic search procedure is required to find the best sub-architecture. On the contrary, the **differentiable search** methods slack the search space and update α and ω simultaneously, assuming that α gradually approaches the one-hot distribution corresponding to the best sub-architecture. So, after the super-network is well-trained, the best sub-architecture is often determined using a greedy algorithm on α while ω is often directly discarded.

3.2 The Devil Lies in the Optimization Gap

Both types of weight-sharing search algorithms suffer the instability issue, in particular, the search results are sensitive to random initialization and/or hyper-parameters. We owe this problem to the **optimization gap**, *i.e.*, a well-optimized super-network does not always produce high-quality sub-architectures. Mathematically, it implies that optimizing Eqn (3) yields an α^* and the corresponding $\mathbb{S}^* = \text{disc}(\alpha^*)$, but there exists another architecture \mathbb{S}' that achieves far higher accuracy. Below, we show examples of how the

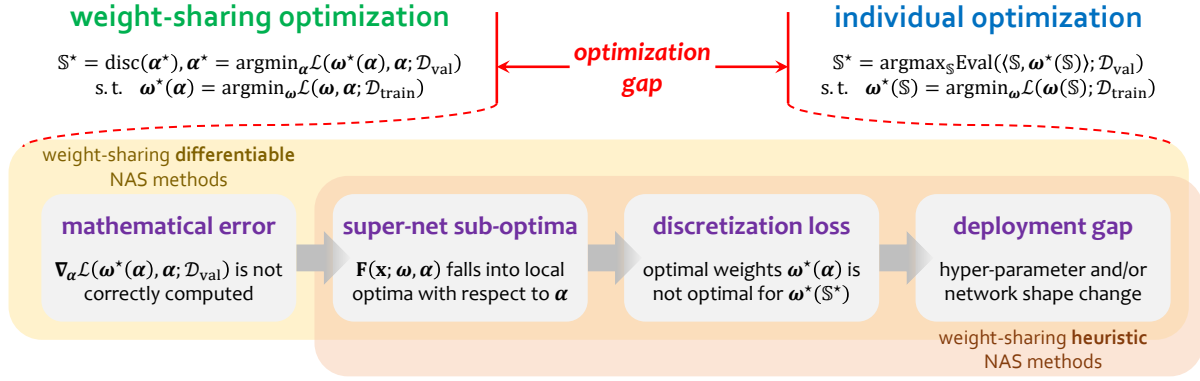


Fig. 3. The optimization gap between the goals of weight-sharing optimization and individual optimization, where the latter is accurate while the former is approximated. We summarize four major reasons to compose of the optimization gap, sorted by their appearance in the NAS flowchart and corresponding to the contents in Sections 3.3.1–3.3.4, respectively. Note that weight-sharing heuristic methods may not involve the mathematical error because $\nabla_{\alpha} \mathcal{L}(\omega^*(\alpha), \alpha; \mathcal{D}_{\text{val}})$ is often not computed.

optimization gap appears as different phenomena according to the type of weight-sharing NAS algorithms.

For **weight-sharing heuristic search**, the optimization gap can be quantified by sampling a set of sub-architectures, $\{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_M\}$, and observing the true accuracy (by training each of them from scratch) and estimated accuracy (predicted by the super-network). Unfortunately, the relationship is often quite weak, *e.g.*, the Kendall’s τ -coefficient between the rank lists according to the true and estimated accuracy is close to zero [171], [400] when the candidates have similar complexity. This is partly due to the strategy that optimizes part of the super-network in each iteration [137], [142], [144], so that whether a sub-architecture performs well in the final evaluation depends on the frequency that it appears, in particular the final iterations of the training procedure. This is mostly random and uncontrollable.

On the other hand, for **weight-sharing differentiable search**, the optimization gap can be quantified by observing the accuracy of the optimal sub-architecture without re-training the network weights, ω . It is shown in [133] that even when the super-network achieves around 90% accuracy (in the CIFAR10 training set), the pruned sub-architecture, without parameter re-training, often reports less than 20% accuracy in the same dataset, which is dramatically low. This implies that the super-network is trained without ‘knowing’ that a discretization stage will take effect to produce the final sub-architecture.

Summarizing the above, the optimization gap likely adds a non-observable noise, $\epsilon(\alpha)$, to the true accuracy of $\text{Eval}(\alpha) \doteq 1 - \mathcal{L}(\omega^*(\alpha), \alpha; \mathcal{D}_{\text{val}})$. Sometimes, for two candidates, say α_1 and α_2 , the noise impacts even heavier, *i.e.*, $|\epsilon(\alpha_1) - \epsilon(\alpha_2)| > |\text{Eval}(\alpha_1) - \text{Eval}(\alpha_2)|$. Due to the randomness of $\epsilon(\alpha)$, the judgment becomes unreliable. This can largely downgrade the search results, or cause the results sensitive to initialization or hyper-parameters.

3.3 Shrinking the Optimization Gap

Essentially, the optimization gap is the consequence of improving search speed with the price of lower evaluation accuracy. Shrinking the optimization gap forms the core challenge of weight-sharing NAS, for which researchers

have proposed various solutions. Before going into details, we point out that the optimization gap is inevitable because the network weights are shared by exponentially many sub-architectures and the optimization procedure focuses on minimizing the amortized loss over the entire search space, not a specific sub-architecture.

We first summarize the typical reasons that cause the optimization gap into four categories, shown in Figure 3. First, when the differentiable methods are used, the end-to-end training procedure may suffer dramatic errors in estimating the gradients with respect to α . Second, the search procedure may cause the super-network fall into a local optimum, in particular the joint optimization of α and ω increases the complexity of the high-dimensional space. Third, the discretization process that derives the optimal sub-architecture may a dramatic accuracy drop on the network training accuracy. Fourth, there are some differences in the network shape or hyper-parameters between the search and evaluation phases, making the searched architecture difficult in deployment. Note that most weight-sharing NAS approaches may involve more than one of the above issues. In what follows, we review the solutions to shrink the optimization gap by considering these factors individually or applying a learning-based approach that deals with the entire system as a black-box.

3.3.1 Fixing the Mathematical Error

The optimization error mostly exists in the differentiable search method, in particular, DARTS [25], in which a bi-level optimization method is used for updating ω and α separately. The reason for bi-level optimization lies in the imbalance between α and ω . Most often, the number of learnable parameters of α is typically hundreds or thousands, but the that of ω is often millions. So, minimizing Eqn (3) can easily bias towards optimizing ω due to the higher dimension of the parameter space¹. To avoid such bias, DARTS [25] proposed to evaluate the network accuracy

1. A direct consequence of such bias is that α is not well optimized which introduces random noise to NAS. In an extreme situation, the super-network can achieve satisfying accuracy by only optimizing ω (*i.e.*, α remains the randomly initialized status), but this delivers zero information to NAS.

with respect to α and ω separately (*i.e.*, $\mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{val}} = \emptyset$), which leads to the iterative optimization:

$$\omega_{t+1} \leftarrow \omega_t - \eta_\omega \cdot \nabla_\omega \mathcal{L}(\omega_t, \alpha_t; \mathcal{D}_{\text{train}}), \quad (4)$$

$$\alpha_{t+1} \leftarrow \alpha_t - \eta_\alpha \cdot \nabla_\alpha \mathcal{L}(\omega_{t+1}, \alpha_t; \mathcal{D}_{\text{val}}), \quad (5)$$

where t is the iteration index starting with 0, and η_α and η_ω are learning rates. Bi-level optimization strategy avoids the bias, but the correctness of bi-level optimization depends on (i) ω_{t+1} has arrived at the optimum, *i.e.*, $\omega_{t+1} = \omega^*(\alpha_t)$, which is very difficult to satisfy especially when the dimensionality of ω is high; (ii) the gradient with respect to α is accurately computed, but this is computationally intractable due to the requirement of computing \mathbf{H}_ω^{-1} , the inverse Hessian matrix of ω . The solution that DARTS [25] used is to ignore the optimality of ω_{t+1} and use the identity matrix to approximate \mathbf{H}_ω^{-1} , but, as pointed out in [48], [52], this leads to dramatic errors in gradient computation² and forms the major reason for mode collapse, *e.g.*, the searched sub-architecture are mostly occupied by the skip-connect operator and perform quite bad [48], [52], [83], [128].

A straightforward solution to alleviate the search instability is to perform early termination [25], [48], [51], [83] or further constrain the property of the sub-architecture, *e.g.*, containing a fixed number of skip-connect operators [51], [81], [83]. But, not allowing the search procedure to converge brings the problem that the search result is sensitive to initialization, which reflects in the unsatisfying stability, *e.g.*, most differentiable search methods can run into bad sub-architectures with a probability of 20%–80%.

There are some work that tried to stabilize bi-level optimization from the mathematical perspective. Typical efforts include using a better approximation (\mathbf{H}_ω to replace \mathbf{H}_ω^{-1}) to achieve an estimation error bound [52], applying adversarial perturbation as regularization to smooth the search space [113], or applying mixed-level optimization to bypass the dramatic approximation [114]. While these methods do not fix the error completely, an alternative path of research focuses on replacing bi-level optimization with one-level optimization which is free of the mathematical burden but the optimization bias towards ω needs to be addressed. For this purpose, researchers proposed to partition the operator set into parameterized and non-parameterized groups to facilitates the fair comparison between similar operators [97]. A more elegant solution is provided by [27] which shows that simply adding regularization (*e.g.*, data augmentation [401] or Dropout [402]) or using a larger dataset (*e.g.*, ImageNet [39]) in the search procedure is sufficient to balance α and ω and improve one-level optimization. Nevertheless, since the validation set no longer exists, one-level optimization may suffers the difficulty of judging the extent that the super-network fits the training set, which may incur a higher risk of over-fitting.

3.3.2 Avoiding the Sub-Optima of Super-Network

Training a super-network is not easy, because the space determined by α is complicated and its property changes significantly with ω . This can be understood as the competition

2. It was proved [52] that, even in optimizing a toy super-network, the angle between the true and estimated gradient vectors of α can be larger than 90°, *i.e.*, the optimization is going towards an uncontrolled, wrong direction.

among different operators, and thus whether an operator is competitive largely depends on how its weights (part of ω) has been optimized. That being said, the operators that are easy to optimize (*e.g.*, non-parameterized ones like skip-connect and pooling and the ones with fewer learnable weights like convolution with small kernels) may dominate in the early stage, which makes it even more difficult for the powerful operators (*e.g.*, convolution of large kernels) to catch up with them [81]. In other words, the search algorithm runs into a local minimum in the space of α .

One of the effective solutions is to warm up the search procedure [26], [81], [82], [109], *i.e.*, freezing the update of α in the starting stage and allowing ω (mostly the weights of all convolution operators) to be well initialized. In essence, this is to postpone the optimization of α . In the heuristic search methods, it behaves as performing fair sampling [142], [144] for super-network training followed by a standalone sampling procedure starting with a uniform distribution on \mathcal{S} . Besides, there are also other kinds of regularization including partitioning the operators into parameterized and non-parameterized groups to avoid unfair competition [97], [235], using operator-level Dropout to weaken the superiority of non-parameterized operators [81], and dynamically adjusting the candidate operator set to allow some perturbation in \mathcal{S} [82].

Note that bi-level optimization also belongs to the larger category of regularization, which decomposes the joint optimization into the subspaces of α and ω . This explains why regularization plays an important role in one-level optimization [27]. Finally, note that regularization often incurs slower convergence speed, but fortunately, the search cost of weight-sharing methods is relatively low and the increased overhead is often acceptable.

3.3.3 Reducing the Discretization Loss

After the super-network has been well optimized, discretization plays an important role in deriving the optimal sub-architecture from the super-network, *i.e.*, $\mathcal{S}^* = \text{disc}(\alpha^*)$ as in Eqn (3). A straightforward solution is to directly eliminate the weak operators from the super-network, but it has been shown that discretization can dramatically downgrade the training accuracy of the super-network (without re-training), especially for the differentiable search methods [101], [133]. There are mainly three solutions. First, converting α from continuous to discrete so that the search procedure is always optimizing the discretized sub-architectures. This strategy is used in both heuristic [142], [144] and differentiable [78], [86], [93] search methods, and it can be interpreted as a kind of super-network regularization. Second, pushing the values of α towards either ends so that the eliminated elements have very low weights [88], [101], [133]. Third, weakening one-time discretization as gradual network pruning [27], [88], [94], [102] so that the super-network is not pulled too far from the normal status (ω is close to the optimum) and recovers after some training epochs (before the next pruning operation).

Note that, when the computational overhead is not considered, the safest way to reduce the discretization loss is to re-train all sub-architectures individually. Obviously, this strategy makes the algorithm degenerate to the heuristic search method with the distribution \mathfrak{P} provided by the

weight-sharing search procedure and the advantage in efficiency no longer exists. To make better use of the search procedure, a common practice is to fine-tune the top-ranked sub-architecture for a short period of time [22], [367]. The length of fine-tuning is an adjustable parameter and plays the role of a tradeoff between purely weight-sharing methods (when the length is zero) and purely individual methods (when the length goes to infinity).

3.3.4 Bridging the Deployment Gap

The last step is to deploy the sub-architecture to the target scenario. In weight-sharing NAS, re-training is often required to achieve better performance, but there may be significant differences between the search and re-training settings. The major reasons are: (i) the search stage may need additional memory to save all network components and thus the network shape is shrunk; and (ii) the re-training stage may need to fit different constraints (*e.g.*, the mobile setting [21] that requires the network FLOPs is just below 600M). A practical way is to adjust the network shape (*e.g.*, the depth, width, and input resolution) and change the hyper-parameters (*e.g.*, the number of epochs, the learning rate, whether some regularization methods have been applied, *etc.*) accordingly. This causes a deployment gap. Conceptually, the algorithm finds the optimal sub-architecture under the setting used in the search procedure, but this does not imply its optimality under the hyper-parameters used in the re-training stage³. Note that even with the same network shape, the difference in hyper-parameters can lead to unstable search results, while this factor is often concealed by other deployment gaps [52].

The first revealed deployment gap is named the **depth gap** [81], *i.e.*, in the search space that contains repeatable cells, the searched sub-architecture can be extended in depth to achieve higher accuracy. However, it was shown that the optimal sub-architecture varies under different depths, and it proposed a progressive method to close the depth gap. Despite the progressively growing super-network, a more essential solution to the depth gap is to directly search in the target network shape. This is natural for most heuristic search methods [139], [142], [144], [343] while the differentiable search methods can encounter the difficulty of training very deep super-networks – this is mainly due to the mathematical error [52] elaborated in Section 3.3.1. Going one step further, researchers have been investigating the relationship between the shallow super-network and the deep sub-architectures [126], [166], [248], and a likely conclusion is that they can benefit from each other.

Another important deployment gap comes from transferring the searched architecture to other tasks, *e.g.*, from image classification to object detection or semantic segmentation. Though many NAS methods have shown the ability of direct architecture transfer [26], [45], [144], we argue that this is not the optimal solution since the optimal architecture may be task-specific. An example is that EfficientDet [252]

is difficult to be trained yet does not show dominating detection accuracy, although it borrows the architecture and pre-trained weights from EfficientNet [143], the state-of-the-art image classification architecture. To improve the transferability of the searched architectures, researchers have applied meta learning [403], [404], [405] (see Section 4.3) or performed network adaptation [255], but we point out that many aspects remain unexplored in this new direction.

3.3.5 Learning-based Approaches

To handle the complicated and coupled factors of the optimization gap, researchers have also investigated learning-based approaches that judge the quality of each sub-architecture based on both the super-network and other kinds of auxiliary information. We point out that this is closely related to the heuristic search methods that keep the distribution \mathfrak{P} of the search space uniform. The difference lies in that the heuristic search methods often update a controller to latently assign a probability to each sub-architecture, but the learning-based methods often encode each sub-architecture into a vector and then use a model to predict the accuracy explicitly.

The performance of learning-based approaches largely depends on the way of representing the sub-architectures, as the encoding method determines the relationship between different sub-architectures [369], [371], [380]. Sometimes, the auxiliary information does not come from a standalone predictor, but from a super-network, *e.g.*, via knowledge distillation [146], [185]. Unlike formulating the entire space for the toy-sized NAS benchmarks [382], [383], in a large search space, the predictor is often embedded into the search procedure to guide the super-network optimization [166], [171], [367]. Besides improving accuracy, the predictor can also filter out weak candidates earlier and thus accelerate the search procedure [74], [121].

We emphasize that the learning-based approaches tend to consider NAS as a black-box predictor and use a small number of training samples to approximate the accuracy distribution on the search space. This is somewhat dangerous as the actual property of the distribution remains unclear. From a higher perspective, these approaches alleviate the optimization gap globally, but do not guarantee to capture the local property very well (*i.e.*, the gap still exists).

3.4 Unsolved Problems

Despite the aforementioned efforts in shrinking the optimization gap, we notice that several important problems of weight-sharing NAS remain unsolved.

First, weight-sharing NAS methods are built upon the assumption that the search space is relatively smooth. However, the actual property of the search space remains unclear. For example, does the search space contain exponentially many sub-architectures that are equally good? Given one of the best sub-architecture, \mathbb{S}^* , is it possible that the neighborhood of \mathbb{S}^* is mostly filled with bad sub-architectures? These issues become increasingly important as the search space gets more complicated. But, due to the huge computational costs to accurately evaluate all sub-architectures (except for the toy search spaces [382], [383], [406]), this problem remains mostly uncovered.

3. For example, if the search stage contains much fewer cells, the searched cell architecture may have very deep connections which increases the difficulty of convergence in the re-training stage. Similarly, if the search stage is equipped with lighter regularization (*e.g.*, Dropout), the searched architecture may contain fewer learnable parameters which downgrades its upper-bound in the re-training stage.

Second, weight-sharing NAS methods facilitate different operators and connections to compete with each other, but it is unknown if the comparison criterion is made fair. For example, almost all search spaces involve the competition among convolutions with various kernel sizes (*e.g.*, 3×3 vs. 5×5 or 7×7), and sometimes, network compression forces competition among different channel and/or bit widths. Although the ‘heavy’ operators are intuitively stronger than the ‘light’ operators, they often suffer slower convergence during the training procedure. That is to say, the search result is sensitive to the time point that comparison is made between these operators. This can cause considerable uncertainty especially in very large search spaces.

Third, weight-sharing NAS methods are able to prune a large super-network into small sub-architectures, but may encounter difficulties in the reverse direction, *i.e.*, deriving a large and/or deep sub-architecture from a relatively small super-network⁴. This is necessary when higher recognition accuracy is desired, and existing approaches are mostly rule-based, including copying the repeatable cells (see Section 3.3.4) and applying simple magnification rules upon the base network. We expect weight-sharing NAS methods to have the ability of extending the search space when needed, which may require the assistance of dynamic routing networks (see Section 4.2).

4 RELATED RESEARCH DIRECTIONS

We briefly review some related research directions to NAS, including (i) compressing the network towards a higher cost-performance ratio, (ii) designing dynamic routing network architectures so as to fit different input data, (iii) learning efficient learning strategies (meta learning), and (iv) looking for the optimal hyper-parameters for network training.

4.1 Model Compression

As the deep networks become more and more complicated, deploying them to edge devices (*e.g.*, mobile phones) requires a new direction named model compression [407], [408]. Two of the most popular techniques are network pruning [409], [410], [411], [412] that eliminates less useful channels or operators from a trained network, and network quantization [413], [414], [415] that uses low-precision (*e.g.*, 1-bit [416]) computation to replace the frequently used FP32 (*i.e.*, 32-bit floating point number) operations. Both methodologies have been verified to achieve a better tradeoff between accuracy and complexity.

Conceptually, model compression shares a similar formulation to NAS, *i.e.*, the generalized formulation in Section 2.1 directly applies with either a regularization term for model complexity or a hard constraint for the maximal resource. Therefore, NAS approaches are often easily transferred for model compression [358], [417], including pruning [244], [359], [360], [364], [418], [419], quantization [224], [420], [421], [422], and joint optimization [423], [424], [425]. Sometimes, the searched configuration or connectivity mask

can improve recognition accuracy [426]. Yet, the optimization gap also exists in these scenarios, and the methods introduced in Section 3.3 are still useful, *e.g.*, super-network regularization [419], [427], [428], network fine-tuning [420], re-training [429], [430], *etc.* In addition, the unfairness issue is significant when the operators of different channel and/or bit widths exist [431]. Nevertheless, it is believed that NAS will form an important module in the future direction named software-hardware co-design [432], [433], [434], [435], [436], [437], [438], [439].

4.2 Dynamic Routing Networks

Most neural architectures are static, *i.e.*, not changing with input data. As a counterpart to save computation, researchers have investigated the option of adjusting the network architecture according to test data [440]. This strategy, named dynamic routing, brings various benefits, including higher amortized efficiency by assigning heavy computation to easy samples and light computation to hard samples [440], [441], [442] information for higher recognition accuracy [443]. There are also discussion on the training strategy for the dynamic routing networks [444], [445].

From a generalized perspective, a dynamic network can be considered as a variant of a static network. Note that even for a static network, the input data can cause part of neurons (after the activation function, *e.g.*, ReLU [446]) to output a value of zero, that is, the connections to these neurons are eliminated. Therefore, the difference lies in that dynamic networks are switching off larger units (*e.g.*, network layers or blocks [447]). This implies that NAS methods, with simple modification, can be applied for dynamic network design [264], [448]. More importantly, dynamic routing networks may pave the way of weight-sharing NAS in an enlargeable search space [187], [218], which we believe is a new research direction.

4.3 Meta Learning

Meta learning [449], [450] is sometimes referred to as ‘learning to learn’, which aims to formulate the common practices in a learning procedure, so that the learned model can quickly adapt to new tasks and/or environments. It originates from the idea of mimicking the learning behavior of human beings [451], yet eases the machine learning algorithms to deploy to various scenarios, *e.g.*, the few-shot learning setting [452], [453], [454] that allows the model to learn new concepts with a small number of training samples. There are mainly three types of meta learners, namely, model-based [455], [456], [457], metric-based [452], [458], and optimization-based [459], [460], [461] methods, differing from each other in the way of formulating the discriminative function⁵. For reviews of meta learning, please refer to [462], [463].

Meta learning has a longer history than NAS, and modern NAS was considered a special case of meta learning [24], [191], [199], namely, learning to design a computational model that assists learning. In particular, meta learning plays an important role in the early age of NAS, when the

4. Non-weight-sharing methods can achieve this goal [18], [44], but the price is to re-train the enlarged network every time, bringing heavy computational overheads.

5. Borrowed from <http://metalearning-symposium.ml/files/vinyals.pdf>.

NAS methods need to learn from a small number of sub-architectures that are sampled and individually evaluated. As weight-sharing NAS methods become more and more popular, meta learning can be used for various purposes, including (i) enhancing the few-shot learning ability of NAS methods [125], [464], [465], (ii) allowing NAS methods to transfer to different tasks [403], [404], [405], and (iii) providing complementary information to shrink the optimization gap [237], [466].

4.4 Hyper-Parameter Optimization

Hyper-parameter optimization (HPO) is another important direction of AutoML, with the aim of finding proper hyper-parameters to improve the quality of model training. This is a classic problem in machine learning [467], but in the deep learning era, the optimization goal has been focused on the important settings that impact network training including the learning rate [468], [469], [470], [471], mini-batch size [472], [473], weight decay [474], [475], momentum [476], [477], or a combination of them [398], [478], [479]. For more thorough reviews of HPO, please refer to [480], [481], [482].

We point out that NAS is essentially a kind of HPO, *i.e.*, the network architecture is encoded as a set of hyper-parameters. Indeed, the NAS approaches in the early years have mostly inherited the ideology of HPO methods including random search [480], [483], heuristic search [484], [485], [486], [487], [488], [489], and meta learning [490]. On the other hand, NAS approaches have been transplanted to HPO [491], and sometimes NAS and HPO are performed together [492], [493]. But, one of the important challenges that HPO algorithms have encountered is the much slower optimization speed compared to the weight-sharing NAS algorithms. This is because each step of HPO requires a complete training procedure, unlike NAS that updates the architectural parameters, α , in each step of the search procedure. Essentially, this is due to the missing of the **greedy optimization assumption**: during NAS training, it is always a good choice to minimize the loss function in a short period of time, but this does not hold for HPO training⁶. A direct consequence is the limited degree of freedom of the HPO search space. We expect the future research of HPO can alleviate this difficulty and thus achieve higher flexibility of optimizing hyper-parameters. The recent efforts to combine NAS and HPO [494] provided a promising direction.

In the field of AutoML, we notice an important subarea of HPO named automated data augmentation, which is currently one of the most effective and widely-applied HPO techniques. We introduce it separately in the next part.

4.5 Automated Data Augmentation

Data augmentation is a technique that enlarges the training set using transformed copies of the existing training images. The assumption is that slightly perturbed (*e.g.*, rotated, flipped, distorted, *etc.*) samples share the same label as the original one. AutoAugment [401] first formulated data augmentation into an HPO problem in a large search

space and applied reinforcement learning for finding the optimal strategy. It achieved good practice in many image classification datasets, but the search cost is very high (*e.g.*, hundreds of GPU-days even for CIFAR10). There are a few follow-ups that tried to accelerate AutoAugment using approximated [495], [496], [497] or weight-sharing [498] optimization, or simply used randomly generated strategies which were verified equally good [499]. Without careful control, the aggressively augmented training samples may contaminate the training procedure. To weaken the impact of these samples and alleviate the empirical risk, researchers proposed to increase the training batch size [500], compute individual batch statistics [501], or use knowledge distillation [502] for case-specific training label refinement.

There is a survey [503] that reviews a wide range of data augmentation techniques for deep learning. Applying AutoML techniques for data augmentation was believed to have a high potential for future research.

5 CONCLUSIONS AND FUTURE PROSPECTS

Neural architecture search (NAS) is an important topic of automated machine learning (AutoML) that aims to find replacement for the manually designed architectures. With three-year rapid development, NAS shows the potential of becoming a standard tool of deep learning algorithms applied to different AI problems. Currently, NAS has contributed significantly to some state-of-the-art computer vision algorithms, but NAS also suffers significant drawbacks that hinders its broader applications.

However, in the past months, the research of NAS seems to have slowed down, arguably due to both the limited search space and unstable search algorithms. These factors are highly correlated. The conservative design of the search spaces makes it difficult to evaluate and compare the performance of different search algorithms, on the other hand, the instability of the search algorithms obstacles the exploration in larger search spaces. Today, two most popular search spaces, the DARTS search space (Section 2.2.4) and the MobileNet search space (Section 2.2.5) have been so thoroughly explored that some tricks (*e.g.*, preserving a fixed number of skip-connect operators in each cell [81], [83]) can produce highly competitive search results, or some inaccurate conclusions can be made [195], [363], yet these results do not show significant advantages over the random search methods with a reasonable amount of trials.

As the last but not least part of this survey, we share our opinions on the future directions of NAS. We believe that the future prospect of NAS lies in whether researchers can propose a new framework to systematically solve the aforementioned problems, in particular, designing more flexible search spaces and developing reliable search algorithms. These two factors facilitate each other, but in the current time point (July 31st, 2020), the former seems more urgent.

A good **search space** for NAS research is expected to satisfy four conditions. First, the space is sufficiently large to include both chain-styled [3] and multi-path [4] networks, popular hand-designed modules (*e.g.*, the residual block [5], group convolution [504], depthwise separable convolution [342], shuffle modules [505], attention units [506],

6. A clear example lies in the learning rate. If the goal is to cause the loss function drop rapidly, the algorithm may choose to use a small learning rate at the start of training, but this strategy is not good for the global optimality.

[507], activation functions [51], [508], etc.), and the combination of them [21]. According to our rough estimation, the cardinality of such a search space is expected to exceed 10^{1000} , but the largest search space to date (GOLD-NAS [27]) only has a cardinality of 10^{110} – 10^{170} . Second, the space should be adjustable, allowing the search algorithm to increase or decrease the model size without restarting the search procedure from scratch. Third, the space can be applied to a wide range of vision problems without heavy modifications – preferably, the searched architecture transfers to different tasks with merely light-weighted heads to be replaced (or searched), like HRNet [509]. Fourth, the best architecture in the space should achieve significantly better accuracy-complexity tradeoff than that found by random search (with reasonable computational costs) or manual design (including pure manual design and adding simple rules to the search procedure).

Based on the search space, a good **search algorithm** should, first of all, be able to explore the search space without getting lost, e.g., running into a set of sub-optimal architectures that are even worse than random search or manual design. We believe that various network training methods will be useful, including knowledge distillation [185], [510], [511], [512], [513], automated data augmentation (see Section 4.5), dynamic routing networks (see Section 4.2), and transfer learning [174], [514]. In addition, the algorithm should allow different constraints (e.g., hardware factors like latency [109], [138], [184] and other factors [141], [515]) to be embedded. We believe that such NAS algorithms, once developed, can largely facilitate the development of deep learning and even the next generation of AI computing.

ACKNOWLEDGMENTS

The authors would like to thank the colleagues in Huawei Noah’s Ark Lab, in particular, Prof. Tong Zhang, Dr. Wei Zhang, Dr. Chunjing Xu, Jianzhong He, and Zewei Du, for instructive discussions. We would also like to thank the organizers of AutoML.org for providing an awesome list of AutoML papers⁷.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*, 2012.
- [3] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *International Conference on Learning Representations*, 2015.
- [4] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *Computer Vision and Pattern Recognition*, 2015.
- [5] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Computer Vision and Pattern Recognition*, 2016.
- [6] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *Computer Vision and Pattern Recognition*, 2017.
- [7] R. Girshick, “Fast r-cnn,” in *International Conference on Computer Vision*, 2015.
- [8] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in *Advances in Neural Information Processing Systems*, 2015.
- [9] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *Computer Vision and Pattern Recognition*, 2016.
- [10] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, “Ssd: Single shot multibox detector,” in *European Conference on Computer Vision*, 2016.
- [11] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Computer Vision and Pattern Recognition*, 2015.
- [12] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, 2017.
- [13] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask r-cnn,” in *International Conference on Computer Vision*, 2017.
- [14] S. Xie and Z. Tu, “Holistically-nested edge detection,” in *International Conference on Computer Vision*, 2015.
- [15] A. Toshev and C. Szegedy, “DeepPose: Human pose estimation via deep neural networks,” in *Computer Vision and Pattern Recognition*, 2014.
- [16] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [17] C. Thornton, F. Hutter, H. H. Hoos, and K. Leyton-Brown, “Auto-weka: Combined selection and hyperparameter optimization of classification algorithms,” in *SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2013.
- [18] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, J. Tan, Q. V. Le, and A. Kurakin, “Large-scale evolution of image classifiers,” in *International Conference on Machine Learning*, 2017.
- [19] L. Xie and A. Yuille, “Genetic cnn,” in *International Conference on Computer Vision*, 2017.
- [20] B. Zoph and Q. V. Le, “Neural architecture search with reinforcement learning,” in *International Conference on Learning Representations*, 2017.
- [21] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, “Mobilenets: Efficient convolutional neural networks for mobile vision applications,” *arXiv preprint arXiv:1704.04861*, 2017.
- [22] H. Cai, T. Chen, W. Zhang, Y. Yu, and J. Wang, “Efficient architecture search by network transformation,” in *AAAI Conference on Artificial Intelligence*, 2018.
- [23] H. Pham, M. Y. Guan, B. Zoph, Q. V. Le, and J. Dean, “Efficient neural architecture search via parameter sharing,” in *International Conference on Machine Learning*, 2018.
- [24] A. Brock, T. Lim, J. M. Ritchie, and N. Weston, “Smash: one-shot model architecture search through hypernetworks,” *arXiv preprint arXiv:1708.05344*, 2017.
- [25] H. Liu, K. Simonyan, and Y. Yang, “Darts: Differentiable architecture search,” in *International Conference on Learning Representations*, 2019.
- [26] Y. Xu, L. Xie, X. Zhang, X. Chen, G.-J. Qi, Q. Tian, and H. Xiong, “Pc-darts: Partial channel connections for memory-efficient architecture search,” in *International Conference on Learning Representations*, 2020.
- [27] K. Bi, L. Xie, X. Chen, L. Wei, and Q. Tian, “Gold-nas: Gradual, one-level, differentiable,” *arXiv preprint arXiv:2007.03331*, 2020.
- [28] T. Elsken, J. H. Metzen, and F. Hutter, “Neural architecture search: A survey,” *Journal of Machine Learning Research*, vol. 20, no. 55, pp. 1–21, 2019.
- [29] X. He, K. Zhao, and X. Chu, “Automl: A survey of the state-of-the-art,” *arXiv preprint arXiv:1908.00709*, 2019.
- [30] Y. Jaafra, J. L. Laurent, A. Deruyver, and M. S. Naceur, “Reinforcement learning for neural architecture search: A review,” *Image and Vision Computing*, vol. 89, pp. 57–66, 2019.
- [31] M. Wistuba, A. Rawat, and T. Pedapati, “A survey on neural architecture search,” *arXiv preprint arXiv:1905.01392*, 2019.
- [32] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, X. Chen, and X. Wang, “A comprehensive survey of neural architecture search: Challenges and solutions,” *arXiv preprint arXiv:2006.02903*, 2020.
- [33] E.-G. Talbi, “Optimization of deep neural networks: a survey and unified taxonomy,” 2020.

7. This survey is gratefully helped by the list of literature on <https://www.automl.org/automl/literature-on-neural-architecture-search/>.

- [34] X. Yao, "Evolving artificial neural networks," *Proceedings of the IEEE*, vol. 87, no. 9, pp. 1423–1447, 1999.
- [35] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [36] J. Bayer, D. Wierstra, J. Togelius, and J. Schmidhuber, "Evolving memory cell structures for sequence learning," in *International Conference on Artificial Neural Networks*, 2009.
- [37] S. Ding, H. Li, C. Su, J. Yu, and F. Jin, "Evolutionary artificial neural networks: a review," *Artificial Intelligence Review*, vol. 39, no. 3, pp. 251–260, 2013.
- [38] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Tech. Rep., 2009.
- [39] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition*, 2009.
- [40] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein et al., "Imagenet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [41] S. Merity, C. Xiong, J. Bradbury, and R. Socher, "Pointer sentinel mixture models," *arXiv preprint arXiv:1609.07843*, 2016.
- [42] H. Liu, K. Simonyan, O. Vinyals, C. Fernando, and K. Kavukcuoglu, "Hierarchical representations for efficient architecture search," in *International Conference on Learning Representations*, 2018.
- [43] L. Wang, S. Xie, T. Li, R. Fonseca, and Y. Tian, "Sample-efficient neural architecture search by learning action space," *arXiv preprint arXiv:1906.06832*, 2019.
- [44] E. Real, C. Liang, D. R. So, and Q. V. Le, "Automl-zero: Evolving machine learning algorithms from scratch," *arXiv preprint arXiv:2003.03384*, 2020.
- [45] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," in *Computer Vision and Pattern Recognition*, 2018.
- [46] C. Liu, B. Zoph, M. Neumann, J. Shlens, W. Hua, L.-J. Li, L. Fei-Fei, A. Yuille, J. Huang, and K. Murphy, "Progressive neural architecture search," in *European Conference on Computer Vision*, 2018.
- [47] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, "Regularized evolution for image classifier architecture search," in *AAAI Conference on Artificial Intelligence*, 2019.
- [48] A. Zela, T. Elsken, T. Saikia, Y. Marrakchi, T. Brox, and F. Hutter, "Understanding and robustifying differentiable architecture search," in *International Conference on Learning Representations*, 2020.
- [49] H. Zhou, M. Yang, J. Wang, and W. Pan, "Bayesnas: A bayesian approach for neural architecture search," in *International Conference on Machine Learning*, 2019.
- [50] M. Cho, M. Soltani, and C. Hegde, "One-shot neural architecture search via compressive sensing," *arXiv preprint arXiv:1906.02869*, 2019.
- [51] L. Wang, L. Xie, T. Zhang, J. Guo, and Q. Tian, "Scalable nas with factorizable architectural parameters," *arXiv preprint arXiv:1912.13256*, 2019.
- [52] K. Bi, C. Hu, L. Xie, X. Chen, L. Wei, and Q. Tian, "Stabilizing darts with amended gradient estimation on architectural parameters," *arXiv preprint arXiv:1910.11831*, 2019.
- [53] Y. Zhou, S. Ebrahimi, S. Ö. Arik, H. Yu, H. Liu, and G. Diamos, "Resource-efficient neural architect," *arXiv preprint arXiv:1806.07912*, 2018.
- [54] P. Bashivan, M. Tensen, and J. J. DiCarlo, "Teacher guided architecture search," in *International Conference on Computer Vision*, 2019.
- [55] Z. Zhong, J. Yan, W. Wu, J. Shao, and C.-L. Liu, "Practical block-wise neural network architecture generation," in *Computer Vision and Pattern Recognition*, 2018.
- [56] G. Bender, P.-J. Kindermans, B. Zoph, V. Vasudevan, and Q. Le, "Understanding and simplifying one-shot architecture search," in *International Conference on Machine Learning*, 2018.
- [57] R. Luo, F. Tian, T. Qin, E. Chen, and T.-Y. Liu, "Neural architecture optimization," in *Advances in Neural Information Processing Systems*, 2018.
- [58] J. Liang, E. Meyerson, and R. Miikkulainen, "Evolutionary architecture search for deep multitask networks," in *Genetic and Evolutionary Computation Conference*, 2018.
- [59] Y. Chen, G. Meng, Q. Zhang, S. Xiang, C. Huang, L. Mu, and X. Wang, "Renas: Reinforced evolutionary neural architecture search," in *Computer Vision and Pattern Recognition*, 2019.
- [60] J.-M. Pérez-Rúa, M. Baccouche, and S. Pateux, "Efficient progressive neural architecture search," *arXiv preprint arXiv:1808.00391*, 2018.
- [61] Z. Zhong, Z. Yang, B. Deng, J. Yan, W. Wu, J. Shao, and C.-L. Liu, "Blockqnn: Efficient block-wise neural network architecture generation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [62] X. Zhang, Z. Huang, and N. Wang, "You only search once: Single shot neural architecture search via direct sparse optimization," *arXiv preprint arXiv:1811.01567*, 2018.
- [63] M. Guo, Z. Zhong, W. Wu, D. Lin, and J. Yan, "Irlas: Inverse reinforcement learning for architecture search," in *Computer Vision and Pattern Recognition*, 2019.
- [64] K. A. Laube and A. Zell, "Shufflenasnets: Efficient cnn models through modified efficient neural architecture search," in *International Joint Conference on Neural Networks*, 2019.
- [65] L. Wang, Y. Zhao, Y. Jinnai, Y. Tian, and R. Fonseca, "Alphax: exploring neural architectures with deep neural networks and monte carlo tree search," *arXiv preprint arXiv:1903.11059*, 2019.
- [66] V. Macko, C. Weill, H. Mazzawi, and J. Gonzalvo, "Improving neural architecture search image classifiers via ensemble learning," *arXiv preprint arXiv:1903.06236*, 2019.
- [67] M. Wistuba and T. Pedapati, "Inductive transfer for neural architecture optimization," *arXiv preprint arXiv:1903.03536*, 2019.
- [68] Y. Chen, K. Zhu, L. Zhu, X. He, P. Ghamisi, and J. A. Benediktsson, "Automatic design of convolutional neural network for hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 9, pp. 7048–7066, 2019.
- [69] G. Adam and J. Lorraine, "Understanding neural architecture search techniques," *arXiv preprint arXiv:1904.00438*, 2019.
- [70] Y. Akimoto, S. Shirakawa, N. Yoshinari, K. Uchida, S. Saito, and K. Nishida, "Adaptive stochastic natural gradient method for one-shot neural architecture search," *arXiv preprint arXiv:1905.08537*, 2019.
- [71] C. Saltori, S. Roy, N. Sebe, and G. Iacca, "Regularized evolutionary algorithm for dynamic neural topology search," in *International Conference on Image Analysis and Processing*, 2019.
- [72] R. Pasunuru and M. Bansal, "Continual and multi-task architecture search," in *Annual Meeting of the Association for Computational Linguistics*, 2019.
- [73] P. Liu, B. Wu, H. Ma, P. K. Chundi, and M. Seok, "Memnet: Memory-efficiency guided neural architecture search with augment-trim learning," *arXiv preprint arXiv:1907.09569*, 2019.
- [74] Y. Zhou, P. Wang, S. Arik, H. Yu, S. Zawad, F. Yan, and G. Diamos, "Epnas: Efficient progressive neural architecture search," in *British Machine Vision Conference*, 2019.
- [75] H. Shi, R. Pi, H. Xu, Z. Li, J. T. Kwok, and T. Zhang, "Multi-objective neural architecture search via predictive network performance optimization," *arXiv preprint arXiv:1911.09336*, 2019.
- [76] T. Zhang, H.-P. Cheng, Z. Li, F. Yan, C. Huang, H. H. Li, and Y. Chen, "Autoshrink: A topology-aware nas for discovering efficient neural architecture," in *AAAI Conference on Artificial Intelligence*, 2020.
- [77] R. Ardywibowo, S. Boluki, X. Gong, Z. Wang, and X. Qian, "Nads: Neural architecture distribution search for uncertainty awareness," *arXiv preprint arXiv:2006.06646*, 2020.
- [78] S. Xie, H. Zheng, C. Liu, and L. Lin, "Snas: stochastic neural architecture search," in *International Conference on Learning Representations*, 2019.
- [79] L. Li and A. Talwalkar, "Random search and reproducibility for neural architecture search," *arXiv preprint arXiv:1902.07638*, 2019.
- [80] X. Dong and Y. Yang, "Searching for a robust neural architecture in four gpu hours," in *Computer Vision and Pattern Recognition*, 2019.
- [81] X. Chen, L. Xie, J. Wu, and Q. Tian, "Progressive differentiable architecture search: Bridging the depth gap between search and evaluation," in *International Conference on Computer Vision*, 2019.
- [82] N. Nayman, A. Noy, T. Ridnik, I. Friedman, R. Jin, and L. Zelnik, "Xnas: Neural architecture search with expert advice," in *Advances in Neural Information Processing Systems*, 2019.
- [83] H. Liang, S. Zhang, J. Sun, X. He, W. Huang, K. Zhuang, and Z. Li, "Darts+: Improved differentiable architecture search with early stopping," *arXiv preprint arXiv:1909.06035*, 2019.

- [84] Y. Guo, Y. Zheng, M. Tan, Q. Chen, J. Chen, P. Zhao, and J. Huang, "Nat: Neural architecture transformer for accurate and compact architectures," in *Advances in Neural Information Processing Systems*, 2019.
- [85] X. Dong and Y. Yang, "One-shot neural architecture search via self-evaluated template network," in *International Conference on Computer Vision*, 2019.
- [86] F. P. Casale, J. Gordon, and N. Fusi, "Probabilistic neural architecture search," *arXiv preprint arXiv:1902.05116*, 2019.
- [87] A. Hundt, V. Jain, and G. D. Hager, "sharpdarts: Faster and more accurate differentiable architecture search," *arXiv preprint arXiv:1903.09900*, 2019.
- [88] A. Noy, N. Nayman, T. Ridnik, N. Zamir, S. Doveh, I. Friedman, R. Giryes, and L. Zelnik, "Asap: Architecture search, anneal and prune," in *International Conference on Artificial Intelligence and Statistics*, 2020.
- [89] Y. Weng, T. Zhou, L. Liu, and C. Xia, "Automatic convolutional neural architecture search for image classification under different scenes," *IEEE Access*, vol. 7, pp. 38 495–38 506, 2019.
- [90] X. Zheng, R. Ji, L. Tang, B. Zhang, J. Liu, and Q. Tian, "Multinomial distribution learning for effective neural architecture search," in *International Conference on Computer Vision*, 2019.
- [91] H. Hu, J. Langford, R. Caruana, S. Mukherjee, E. J. Horvitz, and D. Dey, "Efficient forward architecture search," in *Advances in Neural Information Processing Systems*, 2019.
- [92] Q. Yao, J. Xu, W.-W. Tu, and Z. Zhu, "Efficient neural architecture search via proximal iterations," in *AAAI Conference on Artificial Intelligence*, 2020.
- [93] J. Chang, Y. Guo, G. MENG, S. XIANG, C. Pan *et al.*, "Data: Differentiable architecture approximation," in *Advances in Neural Information Processing Systems*, 2019.
- [94] X. Zheng, R. Ji, L. Tang, Y. Wan, B. Zhang, Y. Wu, Y. Wu, and L. Shao, "Dynamic distribution pruning for efficient network architecture search," *arXiv preprint arXiv:1905.13543*, 2019.
- [95] M. Zhang, H. Li, S. Pan, T. Liu, and S. Su, "Efficient novelty-driven neural architecture search," *arXiv preprint arXiv:1907.09109*, 2019.
- [96] Z. Yang, Y. Wang, X. Chen, B. Shi, C. Xu, C. Xu, Q. Tian, and C. Xu, "Cars: Continuous evolution for efficient neural architecture search," in *Computer Vision and Pattern Recognition*, 2020.
- [97] G. Li, X. Zhang, Z. Wang, Z. Li, and T. Zhang, "Stacnas: Towards stable and consistent optimization for differentiable neural architecture search," *arXiv preprint arXiv:1909.11926*, 2019.
- [98] F. M. Carlucci, P. Esperanca, R. Tutunov, M. Singh, V. Gabillon, A. Yang, H. Xu, Z. Chen, and J. Wang, "Manas: multi-agent neural architecture search," *arXiv preprint arXiv:1909.01051*, 2019.
- [99] R. Luo, T. Qin, and E. Chen, "Understanding and improving one-shot neural architecture optimization," *arXiv preprint arXiv:1909.10815*, 2019.
- [100] H.-C. Lee, D.-G. Kim, and B. Han, "Efficient decoupled neural architecture search by structure and operation sampling," in *International Conference on Acoustics, Speech and Signal Processing*, 2020.
- [101] X. Chu, T. Zhou, B. Zhang, and J. Li, "Fair darts: Eliminating unfair advantages in differentiable architecture search," in *European Conference on Computer Vision*, 2020.
- [102] G. Li, G. Qian, I. C. Delgadillo, M. Muller, A. Thabet, and B. Ghanem, "Sgas: Sequential greedy architecture search," in *Computer Vision and Pattern Recognition*, 2020.
- [103] H. Dong, B. Zou, L. Zhang, and S. Zhang, "Automatic design of cnns via differentiable neural architecture search for polsar image classification," *IEEE Transactions on Geoscience and Remote Sensing*, 2020.
- [104] S. Green, C. M. Vineyard, R. Helinski, and C. K. Koç, "Rapdarts: Resource-aware progressive differentiable architecture search," *arXiv preprint arXiv:1911.05704*, 2019.
- [105] H. G. Hong, P. Ahn, and J. Kim, "Edas: Efficient and differentiable architecture search," *arXiv preprint arXiv:1912.01237*, 2019.
- [106] X. Xie, Y. Zhou, and S.-Y. Kung, "Exploiting operation importance for differentiable neural architecture search," *arXiv preprint arXiv:1911.10511*, 2019.
- [107] X. Jin, J. Wang, J. Slocum, M.-H. Yang, S. Dai, S. Yan, and J. Feng, "Rc-darts: Resource constrained differentiable architecture search," *arXiv preprint arXiv:1912.12814*, 2019.
- [108] A. Vahdat, A. Mallya, M.-Y. Liu, and J. Kautz, "Unas: Differentiable architecture search meets reinforcement learning," in *Computer Vision and Pattern Recognition*, 2020.
- [109] Y. Xu, L. Xie, X. Zhang, X. Chen, B. Shi, Q. Tian, and H. Xiong, "Latency-aware differentiable neural architecture search," *arXiv preprint arXiv:2001.06392*, 2020.
- [110] D. Zhou, X. Zhou, W. Zhang, C. C. Loy, S. Yi, X. Zhang, and W. Ouyang, "Econas: Finding proxies for economical neural architecture search," in *Computer Vision and Pattern Recognition*, 2020.
- [111] S. Hu, S. Xie, H. Zheng, C. Liu, J. Shi, X. Liu, and D. Lin, "Dsnas: Direct neural architecture search without parameter retraining," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 12 084–12 092.
- [112] K. P. Singh, D. Kim, and J. Choi, "Learning architectures for binary networks," in *European Conference on Computer Vision*, 2020.
- [113] X. Chen and C.-J. Hsieh, "Stabilizing differentiable architecture search via perturbation-based regularization," *arXiv preprint arXiv:2002.05283*, 2020.
- [114] C. He, H. Ye, L. Shen, and T. Zhang, "Milenas: Efficient neural architecture search via mixed-level reformulation," in *Computer Vision and Pattern Recognition*, 2020.
- [115] C. Liu, P. Dollár, K. He, R. Girshick, A. Yuille, and S. Xie, "Are labels necessary for neural architecture search?" in *European Conference on Computer Vision*, 2020.
- [116] S. Niu, J. Wu, Y. Zhang, Y. Guo, P. Zhao, J. Huang, and M. Tan, "Disturbance-immune weight sharing for neural architecture search," *arXiv preprint arXiv:2003.13089*, 2020.
- [117] A. Bulat, B. Martinez, and G. Tzimiropoulos, "Bats: Binary architecture search," in *European Conference on Computer Vision*, 2020.
- [118] X. Zhang, J. Chen, and B. Gu, "Adwpnas: Architecture-driven weight prediction for neural architecture search," *arXiv preprint arXiv:2003.01335*, 2020.
- [119] L. Li, M. Khodak, M.-F. Balcan, and A. Talwalkar, "Geometry-aware gradient algorithms for neural architecture search," *arXiv preprint arXiv:2004.07802*, 2020.
- [120] X. Chu, B. Zhang, and X. Li, "Noisy differentiable architecture search," *arXiv preprint arXiv:2005.03566*, 2020.
- [121] X. Zheng, R. Ji, Q. Wang, Q. Ye, Z. Li, Y. Tian, and Q. Tian, "Re-thinking performance estimation in neural architecture search," in *Computer Vision and Pattern Recognition*, 2020.
- [122] L. Zhuo, B. Zhang, H. Chen, L. Yang, C. Chen, Y. Zhu, and D. Doremann, "Cp-nas: Child-parent neural architecture search for 1-bit cnns," in *International Joint Conference on Artificial Intelligence*, 2020.
- [123] S. Zaidi, A. Zela, T. Elsken, C. Holmes, F. Hutter, and Y. W. Teh, "Neural ensemble search for performant and calibrated predictions," *arXiv preprint arXiv:2006.08573*, 2020.
- [124] X. Chen, R. Wang, M. Cheng, X. Tang, and C.-J. Hsieh, "Drnas: Dirichlet neural architecture search," *arXiv preprint arXiv:2006.10355*, 2020.
- [125] J. Wang, J. Wu, H. Bai, and J. Cheng, "M-nas: Meta neural architecture search," in *AAAI Conference on Artificial Intelligence*, 2020.
- [126] H. Yu and H. Peng, "Cyclic differentiable architecture search," *arXiv preprint arXiv:2006.10724*, 2020.
- [127] Y. Zhao, L. Wang, Y. Tian, R. Fonseca, and T. Guo, "Few-shot neural architecture search," *arXiv preprint arXiv:2006.06863*, 2020.
- [128] P. Zhou, C. Xiong, R. Socher, and S. C. Hoi, "Theory-inspired path-regularized differential network architecture search," *arXiv preprint arXiv:2006.16537*, 2020.
- [129] Y. Guo, Y. Chen, Y. Zheng, P. Zhao, J. Chen, J. Huang, and M. Tan, "Breaking the curse of space explosion: Towards efficient nas with curriculum search," *arXiv preprint arXiv:2007.07197*, 2020.
- [130] W. Hong, G. Li, W. Zhang, R. Tang, Y. Wang, Z. Li, and Y. Yu, "Dropnas: Grouped operation dropout for differentiable architecture search," in *International Joint Conference on Artificial Intelligence*, 2020.
- [131] S. Kaplan and R. Giryes, "Self-supervised neural architecture search," *arXiv preprint arXiv:2007.01500*, 2020.
- [132] Y. Li, M. Dong, Y. Wang, and C. Xu, "Neural architecture search in a proxy validation loss landscape," in *International Conference on Machine Learning*, 2020.
- [133] Y. Tian, C. Liu, L. Xie, J. Jiao, and Q. Ye, "Discretization-aware architecture search," *arXiv preprint arXiv:2007.03154*, 2020.
- [134] X. Wang, C. Xue, J. Yan, X. Yang, Y. Hu, and K. Sun, "Mergenase: Merge operations into one for differentiable architecture search," in *International Joint Conference on Artificial Intelligence*, 2020.

- [135] Y. Wang, W. Dai, C. Li, J. Zou, and H. Xiong, "Si-vdnas: Semi-implicit variational dropout for hierarchical one-shot neural architecture search," in *International Joint Conference on Artificial Intelligence*, 2020.
- [136] M. Zhang, H. Li, S. Pan, T. Liu, and S. Su, "One-shot neural architecture search via novelty driven sampling," in *International Joint Conference on Artificial Intelligence*, 2020.
- [137] H. Cai, L. Zhu, and S. Han, "Proxylessnas: Direct neural architecture search on target task and hardware," in *International Conference on Learning Representations*, 2019.
- [138] B. Wu, X. Dai, P. Zhang, Y. Wang, F. Sun, Y. Wu, Y. Tian, P. Vajda, Y. Jia, and K. Keutzer, "Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search," in *Computer Vision and Pattern Recognition*, 2019.
- [139] M. Tan, B. Chen, R. Pang, V. Vasudevan, M. Sandler, A. Howard, and Q. V. Le, "Mnasnet: Platform-aware neural architecture search for mobile," in *Computer Vision and Pattern Recognition*, 2019.
- [140] X. Dai, P. Zhang, B. Wu, H. Yin, F. Sun, Y. Wang, M. Dukhan, Y. Hu, Y. Wu, Y. Jia *et al.*, "Chamnet: Towards efficient network design through platform-aware model adaptation," in *Computer Vision and Pattern Recognition*, 2019.
- [141] D. Stamoulis, R. Ding, D. Wang, D. Lymberopoulos, B. Priyantha, J. Liu, and D. 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*, 2019.
- [142] Z. Guo, X. Zhang, H. Mu, W. Heng, Z. Liu, Y. Wei, and J. Sun, "Single path one-shot neural architecture search with uniform sampling," *arXiv preprint arXiv:1904.00420*, 2019.
- [143] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International Conference on Machine Learning*, 2019.
- [144] X. Chu, B. Zhang, R. Xu, and J. Li, "Fairnas: Rethinking evaluation fairness of weight sharing neural architecture search," *arXiv preprint arXiv:1907.01845*, 2019.
- [145] M. Tan and Q. V. Le, "Mixconv: Mixed depthwise convolutional kernels," in *British Machine Vision Conference*, 2019.
- [146] H. Cai, C. Gan, and S. Han, "Once for all: Train one network and specialize it for efficient deployment," in *International Conference on Learning Representations*, 2020.
- [147] J. Fang, Y. Chen, X. Zhang, Q. Zhang, C. Huang, G. Meng, W. Liu, and X. Wang, "Eat-nas: elastic architecture transfer for accelerating large-scale neural architecture search," *arXiv preprint arXiv:1901.05884*, 2019.
- [148] A. Wong, Z. Qiu Lin, and B. Chwyl, "Attonets: Compact and efficient deep neural networks for the edge via human-machine collaborative design," in *Computer Vision and Pattern Recognition Workshops*, 2019.
- [149] Y. Xiong, R. Mehta, and V. Singh, "Resource constrained neural network architecture search: Will a submodularity assumption help?" in *International Conference on Computer Vision*, 2019.
- [150] S. Han, H. Cai, L. Zhu, J. Lin, K. Wang, Z. Liu, and Y. Lin, "Design automation for efficient deep learning computing," *arXiv preprint arXiv:1904.10616*, 2019.
- [151] J. Fang, Y. Sun, Q. Zhang, Y. Li, W. Liu, and X. Wang, "Densely connected search space for more flexible neural architecture search," in *Computer Vision and Pattern Recognition*, 2020.
- [152] A. Shaw, D. Hunter, F. Landola, and S. Sidhu, "Squeezenas: Fast neural architecture search for faster semantic segmentation," in *International Conference on Computer Vision workshops*, 2019.
- [153] X. Chu, B. Zhang, and R. Xu, "Moga: Searching beyond mobilenetv3," in *International Conference on Acoustics, Speech and Signal Processing*, 2020.
- [154] X. Chu, B. Zhang, J. Li, Q. Li, and R. Xu, "Scarletnas: Bridging the gap between scalability and fairness in neural architecture search," *arXiv preprint arXiv:1908.06022*, 2019.
- [155] S. Yan, B. Fang, F. Zhang, Y. Zheng, X. Zeng, M. Zhang, and H. Xu, "Hm-nas: Efficient neural architecture search via hierarchical masking," in *International Conference on Computer Vision Workshops*, 2019.
- [156] L. Sun, X. Yu, L. Wang, J. Sun, H. Inakoshi, K. Kobayashi, and H. Kobashi, "Automatic neural network search method for open set recognition," in *International Conference on Image Processing*, 2019.
- [157] J. Yu, P. Jin, H. Liu, G. Bender, P.-J. Kindermans, M. Tan, T. Huang, X. Song, R. Pang, and Q. Le, "Bignas: Scaling up neural architecture search with big single-stage models," in *European Conference on Computer Vision*, 2020.
- [158] X. Li, C. Lin, C. Li, M. Sun, W. Wu, J. Yan, and W. Ouyang, "Improving one-shot nas by suppressing the posterior fading," in *Computer Vision and Pattern Recognition*, 2020.
- [159] C. Li, J. Peng, L. Yuan, G. Wang, X. Liang, L. Lin, and X. Chang, "Block-wisely supervised neural architecture search with knowledge distillation," in *Computer Vision and Pattern Recognition*, 2020.
- [160] M. Kang, J. Mun, and B. Han, "Towards oracle knowledge distillation with neural architecture search," in *AAAI Conference on Artificial Intelligence*, 2020.
- [161] J. Mei, Y. Li, X. Lian, X. Jin, L. Yang, A. Yuille, and J. Yang, "Atomnas: Fine-grained end-to-end neural architecture search," in *International Conference on Learning Representations*, 2020.
- [162] S. Chen, Y. Chen, S. Yan, and J. Peng, "Efficient differentiable neural architecture search with meta kernels," *arXiv preprint arXiv:1912.04749*, 2019.
- [163] X. Chu, X. Li, Y. Lu, B. Zhang, and J. Li, "Mixpath: A unified approach for one-shot neural architecture search," *arXiv preprint arXiv:2001.05887*, 2020.
- [164] Z. Ding, Y. Chen, N. Li, D. Zhao, and C. Chen, "Bnas: An efficient neural architecture search approach using broad scalable architecture," *arXiv preprint arXiv:2001.06679*, 2020.
- [165] X. Dai, D. Chen, M. Liu, Y. Chen, and L. Yuan, "Da-nas: Data adapted pruning for efficient neural architecture search," in *European Conference on Computer Vision*, 2020.
- [166] S.-Y. Huang and W.-T. Chu, "Ponas: Progressive one-shot neural architecture search for very efficient deployment," *arXiv preprint arXiv:2003.05112*, 2020.
- [167] Z. Shen, J. Qian, B. Zhuang, S. Wang, and J. Xiao, "Bs-nas: Broadening-and-shrinking one-shot nas with searchable numbers of channels," *arXiv preprint arXiv:2003.09821*, 2020.
- [168] S. You, T. Huang, M. Yang, F. Wang, C. Qian, and C. Zhang, "Greedydynam: Towards fast one-shot nas with greedy supernet," in *Computer Vision and Pattern Recognition*, 2020.
- [169] A. Wan, X. Dai, P. Zhang, Z. He, Y. Tian, S. Xie, B. Wu, M. Yu, T. Xu, K. 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*, 2020, pp. 12 965–12 974.
- [170] Y. Li, X. Jin, J. Mei, X. Lian, L. Yang, C. Xie, Q. Yu, Y. Zhou, S. Bai, and A. L. Yuille, "Neural architecture search for lightweight non-local networks," in *Computer Vision and Pattern Recognition*, 2020.
- [171] X. Chen, L. Xie, J. Wu, L. Wei, Y. Xu, and Q. Tian, "Fitting the search space of weight-sharing nas with graph convolutional networks," *arXiv preprint arXiv:2004.08423*, 2020.
- [172] Y. Hu, Y. Liang, Z. Guo, R. Wan, X. Zhang, Y. Wei, Q. Gu, and J. Sun, "Angle-based search space shrinking for neural architecture search," in *European Conference on Computer Vision*, 2020.
- [173] R. Guo, C. Lin, C. Li, K. Tian, M. Sun, L. Sheng, and J. Yan, "Powering one-shot topological nas with stabilized share-parameter proxy," in *European Conference on Computer Vision*, 2020.
- [174] Z. Lu, G. Sreeksumar, E. Goodman, W. Banzhaf, K. Deb, and V. N. Boddeti, "Neural architecture transfer," *arXiv preprint arXiv:2005.05859*, 2020.
- [175] H. Phan, Z. Liu, D. Huynh, M. Savvides, K.-T. Cheng, and Z. Shen, "Binarizing mobilenet via evolution-based searching," in *Computer Vision and Pattern Recognition*, 2020.
- [176] Y. Wang, Y. Xu, and D. Tao, "Dc-nas: Divide-and-conquer neural architecture search," *arXiv preprint arXiv:2005.14456*, 2020.
- [177] X. Xia and W. Ding, "Hnas: Hierarchical neural architecture search on mobile devices," *arXiv preprint arXiv:2005.07564*, 2020.
- [178] Z. Yang, Y. Wang, D. Tao, X. Chen, J. Guo, C. Xu, C. Xu, and C. Xu, "Hourglass: Extremely fast neural architecture search through an hourglass lens," *arXiv preprint arXiv:2005.14446*, 2020.
- [179] G. Bender, H. Liu, B. Chen, G. Chu, S. Cheng, P.-J. Kindermans, and Q. V. Le, "Can weight sharing outperform random architecture search? an investigation with tunas," in *Computer Vision and Pattern Recognition*, 2020.
- [180] S. R. Chaudhuri, E. Eban, H. Li, M. Moroz, and Y. Movshovitz-Attias, "Fine-grained stochastic architecture search," *arXiv preprint arXiv:2006.09581*, 2020.
- [181] X. Dai, A. Wan, P. Zhang, B. Wu, Z. He, Z. Wei, K. Chen, Y. Tian, M. Yu, P. Vajda *et al.*, "Fbnetv3: Joint architecture-recipe search using neural acquisition function," *arXiv preprint arXiv:2006.02049*, 2020.

- [182] Z. Li, T. Xi, J. Deng, G. Zhang, S. Wen, and R. He, "Gp-nas: Gaussian process based neural architecture search," in *Computer Vision and Pattern Recognition*, 2020.
- [183] P. Liu, B. Wu, H. Ma, and M. Seok, "Memnas: Memory-efficient neural architecture search with grow-trim learning," in *Computer Vision and Pattern Recognition*, 2020.
- [184] L. Lyna Zhang, Y. Yang, Y. Jiang, W. Zhu, and Y. Liu, "Fast hardware-aware neural architecture search," in *Computer Vision and Pattern Recognition Workshops*, 2020.
- [185] I. Trofimov, N. Klyuchnikov, M. Salnikov, A. Filippov, and E. Burnaev, "Multi-fidelity neural architecture search with knowledge distillation," *arXiv preprint arXiv:2006.08341*, 2020.
- [186] Z. Lu, K. Deb, E. Goodman, W. Banzhaf, and V. N. Boddeti, "Nsganetv2: Evolutionary multi-objective surrogate-assisted neural architecture search," *arXiv preprint arXiv:2007.10396*, 2020.
- [187] C. Cortes, X. Gonzalvo, V. Kuznetsov, M. Mohri, and S. Yang, "Adanet: Adaptive structural learning of artificial neural networks," in *International Conference on Machine Learning*, 2017.
- [188] B. Baker, O. Gupta, N. Naik, and R. Raskar, "Designing neural network architectures using reinforcement learning," in *International Conference on Learning Representations*, 2017.
- [189] M. Suganuma, S. Shirakawa, and T. Nagao, "A genetic programming approach to designing convolutional neural network architectures," in *Genetic and Evolutionary Computation Conference*, 2017.
- [190] R. Negrinho and G. Gordon, "Deeparchitect: Automatically designing and training deep architectures," *arXiv preprint arXiv:1704.08792*, 2017.
- [191] B. Baker, O. Gupta, R. Raskar, and N. Naik, "Accelerating neural architecture search using performance prediction," *arXiv preprint arXiv:1705.10823*, 2017.
- [192] K. Kandasamy, W. Neiswanger, J. Schneider, B. Póczos, and E. P. Xing, "Neural architecture search with bayesian optimisation and optimal transport," in *Advances in Neural Information Processing Systems*, 2018.
- [193] T. Elsken, J. H. Metzen, and F. Hutter, "Efficient multi-objective neural architecture search via lamarckian evolution," in *International Conference on Learning Representations*, 2019.
- [194] K. O. Stanley, J. Clune, J. Lehman, and R. Miikkulainen, "Designing neural networks through neuroevolution," *Nature Machine Intelligence*, vol. 1, no. 1, pp. 24–35, 2019.
- [195] S. Xie, A. Kirillov, R. Girshick, and K. He, "Exploring randomly wired neural networks for image recognition," in *International Conference on Computer Vision*, 2019.
- [196] S. C. Smithson, G. Yang, W. J. Gross, and B. H. Meyer, "Neural networks designing neural networks: multi-objective hyperparameter optimization," in *International Conference on Computer-Aided Design*, 2016.
- [197] T. Wei, C. Wang, and C. W. Chen, "Modularized morphing of neural networks," *arXiv preprint arXiv:1701.03281*, 2017.
- [198] B. Wang, Y. Sun, B. Xue, and M. Zhang, "Evolving deep convolutional neural networks by variable-length particle swarm optimization for image classification," in *IEEE Congress on Evolutionary Computation*, 2018.
- [199] T. Elsken, J.-H. Metzen, and F. Hutter, "Simple and efficient architecture search for convolutional neural networks," *arXiv preprint arXiv:1711.04528*, 2017.
- [200] Y. Sun, G. G. Yen, and Z. Yi, "Evolving unsupervised deep neural networks for learning meaningful representations," *IEEE Transactions on Evolutionary Computation*, vol. 23, no. 1, pp. 89–103, 2018.
- [201] Y. Sun, B. Xue, M. Zhang, and G. G. Yen, "A particle swarm optimization-based flexible convolutional autoencoder for image classification," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 8, pp. 2295–2309, 2018.
- [202] M. Wistuba, "Finding competitive network architectures within a day using uct," *arXiv preprint arXiv:1712.07420*, 2017.
- [203] J. K. Dutta, J. Liu, U. Kurup, and M. Shah, "Effective building block design for deep convolutional neural networks using search," *arXiv preprint arXiv:1801.08577*, 2018.
- [204] R. Shin, C. Packer, and D. Song, "Differentiable neural network architecture search," in *International Conference on Learning Representations: Workshops*, 2018.
- [205] P. Kamath, A. Singh, and D. Dutta, "Neural architecture construction using envelopes," *arXiv preprint arXiv:1803.06744*, 2018.
- [206] H. Cai, J. Yang, W. Zhang, S. Han, and Y. Yu, "Path-level network transformation for efficient architecture search," in *International Conference on Machine Learning*, 2018.
- [207] J.-D. Dong, A.-C. Cheng, D.-C. Juan, W. Wei, and M. Sun, "Dpp-net: Device-aware progressive search for pareto-optimal neural architectures," in *European Conference on Computer Vision*, 2018.
- [208] H. Jin, Q. Song, and X. Hu, "Efficient neural architecture search with network morphism," *arXiv preprint arXiv:1806.10282*, 2018.
- [209] C.-H. Hsu, S.-H. Chang, J.-H. Liang, H.-P. Chou, C.-H. Liu, S.-C. Chang, J.-Y. Pan, Y.-T. Chen, W. Wei, and D.-C. Juan, "Monas: Multi-objective neural architecture search using reinforcement learning," *arXiv preprint arXiv:1806.10332*, 2018.
- [210] R. Istrate, F. Scheidegger, G. Mariani, D. Nikolopoulos, C. Bekas, and A. C. I. Malossi, "Tapas: Train-less accuracy predictor for architecture search," in *AAAI Conference on Artificial Intelligence*, 2019.
- [211] M. Wistuba, "Deep learning architecture search by neuro-cell-based evolution with function-preserving mutations," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2018.
- [212] R. Y. Rohkar, S. Nisimov, Y. Gurwicz, G. Koren, and G. Novik, "Constructing deep neural networks by bayesian network structure learning," in *Advances in Neural Information Processing Systems*, 2018.
- [213] J. Prellberg and O. Kramer, "Lamarckian evolution of convolutional neural networks," in *International Conference on Parallel Problem Solving from Nature*, 2018.
- [214] P. R. Lorenzo and J. Nalepa, "Memetic evolution of deep neural networks," in *Genetic and Evolutionary Computation Conference*, 2018.
- [215] Y. Sun, B. Xue, M. Zhang, and G. G. Yen, "Evolving deep convolutional neural networks for image classification," *IEEE Transactions on Evolutionary Computation*, 2019.
- [216] B. Wang, Y. Sun, B. Xue, and M. Zhang, "A hybrid differential evolution approach to designing deep convolutional neural networks for image classification," in *Australasian Joint Conference on Artificial Intelligence*, 2018.
- [217] A.-C. Cheng, J.-D. Dong, C.-H. Hsu, S.-H. Chang, M. Sun, S.-C. Chang, J.-Y. Pan, Y.-T. Chen, W. Wei, and D.-C. Juan, "Searching toward pareto-optimal device-aware neural architectures," in *International Conference on Computer-Aided Design*, 2018.
- [218] C. Zhang, M. Ren, and R. Urtasun, "Graph hypernetworks for neural architecture search," in *International Conference on Learning Representations*, 2019.
- [219] Z. Lu, I. Whalen, V. Boddeti, Y. Dhebar, K. Deb, E. Goodman, and W. Banzhaf, "Nsga-net: neural architecture search using multi-objective genetic algorithm," in *Genetic and Evolutionary Computation Conference*, 2019.
- [220] Y. Sun, B. Xue, M. Zhang, and G. G. Yen, "Automatically evolving cnn architectures based on blocks," *arXiv preprint arXiv:1810.11875*, 2018.
- [221] B. Fielding and L. Zhang, "Evolving image classification architectures with enhanced particle swarm optimisation," *IEEE Access*, vol. 6, pp. 68 560–68 575, 2018.
- [222] Y. Geifman and R. El-Yaniv, "Deep active learning with a neural architecture search," in *Advances in Neural Information Processing Systems*, 2019.
- [223] A.-C. Cheng, C. H. Lin, D.-C. Juan, W. Wei, and M. Sun, "Instanas: Instance-aware neural architecture search," in *AAAI Conference on Artificial Intelligence*, 2020.
- [224] B. Wu, Y. Wang, P. Zhang, Y. Tian, P. Vajda, and K. Keutzer, "Mixed precision quantization of convnets via differentiable neural architecture search," *arXiv preprint arXiv:1812.00090*, 2018.
- [225] W. Jiang, X. Zhang, E. H.-M. Sha, L. Yang, Q. Zhuge, Y. Shi, and J. Hu, "Accuracy vs. efficiency: Achieving both through fpga-implementation aware neural architecture search," in *Annual Design Automation Conference*, 2019.
- [226] G. Dikov, P. van der Smagt, and J. Bayer, "Bayesian learning of neural network architectures," *arXiv preprint arXiv:1901.04436*, 2019.
- [227] P. Savarese and M. Maire, "Learning implicitly recurrent cnns through parameter sharing," *arXiv preprint arXiv:1902.09701*, 2019.
- [228] B. Wang, Y. Sun, B. Xue, and M. Zhang, "A hybrid ga-pso method for evolving architecture and short connections of deep convolutional neural networks," in *Pacific Rim International Conference on Artificial Intelligence*, 2019.

- [229] F. E. F. Junior and G. G. Yen, "Particle swarm optimization of deep neural networks architectures for image classification," *Swarm and Evolutionary Computation*, vol. 49, pp. 62–74, 2019.
- [230] I. Radosavovic, J. Johnson, S. Xie, W.-Y. Lo, and P. Dollár, "On network design spaces for visual recognition," in *International Conference on Computer Vision*, 2019.
- [231] T. Saikia, Y. Marrakchi, A. Zela, F. Hutter, and T. Brox, "Autodispnet: Improving disparity estimation with automl," in *International Conference on Computer Vision*, 2019.
- [232] L. Ma, J. Cui, and B. Yang, "Deep neural architecture search with deep graph bayesian optimization," in *International Conference on Web Intelligence*, 2019.
- [233] H. Zhu, Z. An, C. Yang, K. Xu, E. Zhao, and Y. Xu, "Eena: efficient evolution of neural architecture," in *International Conference on Computer Vision Workshops*, 2019.
- [234] K. Nguyen, C. Fookes, and S. Sridharan, "Constrained design of deep iris networks," *IEEE Transactions on Image Processing*, 2020.
- [235] Y. Jiang, C. Zhao, Z. Dou, and L. Pang, "Neural architecture refinement: A practical way for avoiding overfitting in nas," *arXiv preprint arXiv:1905.02341*, 2019.
- [236] Y. Sun, H. Wang, B. Xue, Y. Jin, G. G. Yen, and M. Zhang, "Surrogate-assisted evolutionary deep learning using an end-to-end random forest-based performance predictor," *IEEE Transactions on Evolutionary Computation*, vol. 24, no. 2, pp. 350–364, 2019.
- [237] H.-P. Cheng, T. Zhang, Y. Yang, F. Yan, S. Li, H. Teague, H. Li, and Y. Chen, "Swiftnet: Using graph propagation as meta-knowledge to search highly representative neural architectures," *arXiv preprint arXiv:1906.08305*, 2019.
- [238] Y. Zhou, X. Sun, C. Luo, Z.-J. Zha, and W. Zeng, "Posterior-guided neural architecture search," in *AAAI Conference on Artificial Intelligence*, 2020.
- [239] B. Wang, B. Xue, and M. Zhang, "Particle swarm optimisation for evolving deep neural networks for image classification by evolving and stacking transferable blocks," *arXiv preprint arXiv:1907.12659*, 2019.
- [240] M. Mörtens and D. Izzo, "Neural network architecture search with differentiable cartesian genetic programming for regression," in *Genetic and Evolutionary Computation Conference Companion*, 2019.
- [241] W. Irwin-Harris, Y. Sun, B. Xue, and M. Zhang, "A graph-based encoding for evolutionary convolutional neural network architecture design," in *IEEE Congress on Evolutionary Computation*, 2019.
- [242] M. Shen, K. Han, C. Xu, and Y. Wang, "Searching for accurate binary neural architectures," in *International Conference on Computer Vision Workshops*, 2019.
- [243] P. Balaprakash, R. Egele, M. Salim, S. Wild, V. Vishwanath, F. Xia, T. Brettin, and R. Stevens, "Scalable reinforcement-learning-based neural architecture search for cancer deep learning research," in *International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019.
- [244] R. Zhao and W. Luk, "Efficient structured pruning and architecture searching for group convolution," in *International Conference on Computer Vision Workshops*, 2019.
- [245] H.-P. Cheng, T. Zhang, Y. Yang, F. Yan, H. Teague, Y. Chen, and H. Li, "Msnnet: Structural wired neural architecture search for internet of things," in *International Conference on Computer Vision Workshops*, 2019.
- [246] S. Hassantabar, X. Dai, and N. K. Jha, "Steerage: Synthesis of neural networks using architecture search and grow-and-prune methods," *arXiv preprint arXiv:1912.05831*, 2019.
- [247] J. Jiang, F. Han, Q. Ling, J. Wang, T. Li, and H. Han, "Efficient network architecture search via multiobjective particle swarm optimization based on decomposition," *Neural Networks*, vol. 123, pp. 305–316, 2020.
- [248] R. Meng, W. Chen, D. Xie, Y. Zhang, and S. Pu, "Neural inheritance relation guided one-shot layer assignment search," *arXiv preprint arXiv:2002.12580*, 2020.
- [249] H.-P. Cheng, T. Zhang, S. Li, F. Yan, M. Li, V. Chandra, H. Li, and Y. Chen, "Nasgem: Neural architecture search via graph embedding method," *arXiv preprint arXiv:2007.04452*, 2020.
- [250] Y. Chen, T. Yang, X. Zhang, G. Meng, X. Xiao, and J. Sun, "Detnas: Backbone search for object detection," in *Advances in Neural Information Processing Systems*, 2019.
- [251] G. Ghiasi, T.-Y. Lin, and Q. V. Le, "Nas-fpn: Learning scalable feature pyramid architecture for object detection," in *Computer Vision and Pattern Recognition*, 2019.
- [252] M. Tan, R. Pang, and Q. V. Le, "Efficientdet: Scalable and efficient object detection," in *Computer Vision and Pattern Recognition*, 2020.
- [253] H. Xu, L. Yao, W. Zhang, X. Liang, and Z. Li, "Auto-fpn: Automatic network architecture adaptation for object detection beyond classification," in *International Conference on Computer Vision*, 2019.
- [254] L. Yao, H. Xu, W. Zhang, X. Liang, and Z. Li, "Sm-nas: Structural-to-modular neural architecture search for object detection," in *AAAI Conference on Artificial Intelligence*, 2020.
- [255] J. Fang, Y. Sun, K. Peng, Q. Zhang, Y. Li, W. Liu, and X. Wang, "Fast neural network adaptation via parameter remapping and architecture search," in *International Conference on Learning Representations*, 2020.
- [256] J. Guo, K. Han, Y. Wang, C. Zhang, Z. Yang, H. Wu, X. Chen, and C. Xu, "Hit-detector: Hierarchical trinity architecture search for object detection," in *Computer Vision and Pattern Recognition*, 2020.
- [257] C. Jiang, S. Wang, X. Liang, H. Xu, and N. Xiao, "Elixirnet: Relation-aware network architecture adaptation for medical lesion detection," in *AAAI*, 2020, pp. 11 093–11 100.
- [258] Y. Xiong, H. Liu, S. Gupta, B. Akin, G. Bender, P.-J. Kindermans, M. Tan, V. Singh, and B. Chen, "Mobiledets: Searching for object detection architectures for mobile accelerators," *arXiv preprint arXiv:2004.14525*, 2020.
- [259] C. Jiang, H. Xu, W. Zhang, X. Liang, and Z. Li, "Sp-nas: Serial-to-parallel backbone search for object detection," in *Computer Vision and Pattern Recognition*, 2020.
- [260] B. Chen, G. Ghiasi, H. Liu, T.-Y. Lin, D. Kalenichenko, H. Adam, and Q. V. Le, "Mnasfpn: Learning latency-aware pyramid architecture for object detection on mobile devices," in *Computer Vision and Pattern Recognition*, 2020.
- [261] L.-C. Chen, M. Collins, Y. Zhu, G. Papandreou, B. Zoph, F. Schroff, H. Adam, and J. Shlens, "Searching for efficient multi-scale architectures for dense image prediction," in *Advances in Neural Information Processing Systems*, 2018.
- [262] C. Liu, L.-C. Chen, F. Schroff, H. Adam, W. Hua, A. L. Yuille, and L. Fei-Fei, "Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation," in *Computer Vision and Pattern Recognition*, 2019.
- [263] V. Nekrasov, H. Chen, C. Shen, and I. Reid, "Fast neural architecture search of compact semantic segmentation models via auxiliary cells," in *Computer Vision and Pattern Recognition*, 2019.
- [264] —, "Architecture search of dynamic cells for semantic video segmentation," in *Winter Conference on Applications of Computer Vision*, 2020.
- [265] V. Nekrasov, C. Shen, and I. Reid, "Template-based automatic search of compact semantic segmentation architectures," in *Winter Conference on Applications of Computer Vision*, 2020.
- [266] P. Lin, P. Sun, G. Cheng, S. Xie, X. Li, and J. Shi, "Graph-guided architecture search for real-time semantic segmentation," in *Computer Vision and Pattern Recognition*, 2020.
- [267] K. C. Wong and M. Moradi, "Segnas3d: Network architecture search with derivative-free global optimization for 3d image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*, 2019.
- [268] X. Zhang, H. Xu, H. Mo, J. Tan, C. Yang, and W. Ren, "Dcnas: Densely connected neural architecture search for semantic image segmentation," *arXiv preprint arXiv:2003.11883*, 2020.
- [269] X. Gong, S. Chang, Y. Jiang, and Z. Wang, "Autogan: Neural architecture search for generative adversarial networks," in *International Conference on Computer Vision*, 2019.
- [270] C. Gao, Y. Chen, S. Liu, Z. Tan, and S. Yan, "Adversarialnas: Adversarial neural architecture search for gans," in *Computer Vision and Pattern Recognition*, 2020.
- [271] M. Li, J. Lin, Y. Ding, Z. Liu, J.-Y. Zhu, and S. Han, "Gan compression: Efficient architectures for interactive conditional gans," in *Computer Vision and Pattern Recognition*, 2020.
- [272] Y. Fu, W. Chen, H. Wang, H. Li, Y. Lin, and Z. Wang, "Autogan-distiller: Searching to compress generative adversarial networks," *arXiv preprint arXiv:2006.08198*, 2020.
- [273] Y. Tian, L. Shen, G. Su, Z. Li, and W. Liu, "Alphagan: Fully differentiable architecture search for generative adversarial networks," *arXiv preprint arXiv:2006.09134*, 2020.
- [274] P. Zhou, L. Xie, X. Zhang, B. Ni, and Q. Tian, "Searching towards class-aware generators for conditional generative adversarial networks," *arXiv preprint arXiv:2006.14208*, 2020.
- [275] Y. Tian, Q. Wang, Z. Huang, W. Li, D. Dai, M. Yang, J. Wang, and O. Fink, "Off-policy reinforcement learning for efficient and

- effective gan architecture search," *arXiv preprint arXiv:2007.09180*, 2020.
- [276] Y. Weng, T. Zhou, Y. Li, and X. Qiu, "Nas-unet: Neural architecture search for medical image segmentation," *IEEE Access*, vol. 7, pp. 44 247–44 257, 2019.
- [277] A. Mortazi and U. Bagci, "Automatically designing cnn architectures for medical image segmentation," in *International Workshop on Machine Learning in Medical Imaging*, 2018.
- [278] M. Baldeon-Calisto and S. K. Lai-Yuen, "Adaresu-net: Multiobjective adaptive convolutional neural network for medical image segmentation," *Neurocomputing*, vol. 392, pp. 325–340, 2020.
- [279] Z. Zhu, C. Liu, D. Yang, A. Yuille, and D. Xu, "V-nas: Neural architecture search for volumetric medical image segmentation," in *International Conference on 3D Vision*, 2019.
- [280] W. Bae, S. Lee, Y. Lee, B. Park, M. Chung, and K.-H. Jung, "Resource optimized neural architecture search for 3d medical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*, 2019.
- [281] Z. Zhang, L. Zhou, L. Gou, and Y. N. Wu, "Neural architecture search for joint optimization of predictive power and biological knowledge," *arXiv preprint arXiv:1909.00337*, 2019.
- [282] S. Kim, I. Kim, S. Lim, W. Baek, C. Kim, H. Cho, B. Yoon, and T. Kim, "Scalable neural architecture search for 3d medical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*, 2019.
- [283] N. Dong, M. Xu, X. Liang, Y. Jiang, W. Dai, and E. Xing, "Neural architecture search for adversarial medical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 828–836.
- [284] Q. Yu, D. Yang, H. Roth, Y. Bai, Y. Zhang, A. L. Yuille, and D. Xu, "C2fnas: Coarse-to-fine neural architecture search for 3d medical image segmentation," in *Computer Vision and Pattern Recognition*, 2020.
- [285] D. Guo, D. Jin, Z. Zhu, T.-Y. Ho, A. P. Harrison, C.-H. Chao, J. Xiao, and L. Lu, "Organ at risk segmentation for head and neck cancer using stratified learning and neural architecture search," in *Computer Vision and Pattern Recognition*, 2020.
- [286] X. Yan, W. Jiang, Y. Shi, and C. Zhuo, "Ms-nas: Multi-scale neural architecture search for medical image segmentation," *arXiv preprint arXiv:2007.06151*, 2020.
- [287] D. V. Vargas and S. Kotyan, "Evolving robust neural architectures to defend from adversarial attacks," *arXiv preprint arXiv:1906.11667*, 2019.
- [288] M. Guo, Y. Yang, R. Xu, Z. Liu, and D. Lin, "When nas meets robustness: In search of robust architectures against adversarial attacks," in *Computer Vision and Pattern Recognition*, 2020.
- [289] G. J. van Wyk and A. S. Bosman, "Evolutionary neural architecture search for image restoration," in *International Joint Conference on Neural Networks*, 2019.
- [290] X. Chu, B. Zhang, H. Ma, R. Xu, J. Li, and Q. Li, "Fast, accurate and lightweight super-resolution with neural architecture search," *arXiv preprint arXiv:1901.07261*, 2019.
- [291] X. Chu, B. Zhang, R. Xu, and H. Ma, "Multi-objective reinforced evolution in mobile neural architecture search," *arXiv preprint arXiv:1901.01074*, 2019.
- [292] P. Liu, M. D. El Basha, Y. Li, Y. Xiao, P. C. Sanelli, and R. Fang, "Deep evolutionary networks with expedited genetic algorithms for medical image denoising," *Medical Image Analysis*, vol. 54, pp. 306–315, 2019.
- [293] D. Song, C. Xu, X. Jia, Y. Chen, C. Xu, and Y. Wang, "Efficient residual dense block search for image super-resolution," in *AAAI Conference on Artificial Intelligence*, 2020.
- [294] Y. Guo, Y. Luo, Z. He, J. Huang, and J. Chen, "Hierarchical neural architecture search for single image super-resolution," *arXiv preprint arXiv:2003.04619*, 2020.
- [295] M. Mozejko, T. Latkowski, L. Treszczotko, M. Szafraniuk, and K. Trojanowski, "Superkernel neural architecture search for image denoising," in *Computer Vision and Pattern Recognition Workshops*, 2020.
- [296] R. Li, R. T. Tan, and L.-F. Cheong, "All in one bad weather removal using architectural search," in *Computer Vision and Pattern Recognition*, 2020.
- [297] H. Zhang, Y. Li, H. Chen, and C. Shen, "Memory-efficient hierarchical neural architecture search for image denoising," in *Computer Vision and Pattern Recognition*, 2020.
- [298] R. Lee, Ł. Dudziak, M. Abdelfattah, S. I. Venieris, H. Kim, H. Wen, and N. D. Lane, "Journey towards tiny perceptual super-resolution," in *European Conference on Computer Vision*, 2020.
- [299] A. Piergiovanni, A. Angelova, A. Toshev, and M. S. Ryoo, "Evolving space-time neural architectures for videos," in *International Conference on Computer Vision*, 2019.
- [300] M. S. Ryoo, A. Piergiovanni, M. Tan, and A. Angelova, "AssembleNet: Searching for multi-stream neural connectivity in video architectures," in *International Conference on Learning Representations*, 2020.
- [301] W. Peng, X. Hong, and G. Zhao, "Video action recognition via neural architecture searching," in *International Conference on Image Processing*, 2019.
- [302] Z. Qiu, T. Yao, Y. Zhang, Y. Zhang, and T. Mei, "Scheduled differentiable architecture search for visual recognition," *arXiv preprint arXiv:1909.10236*, 2019.
- [303] J.-M. Pérez-Rúa, V. Vielzeuf, S. Pateux, M. Baccouche, and F. Jurie, "Mfas: Multimodal fusion architecture search," in *Computer Vision and Pattern Recognition*, 2019.
- [304] Y. Yu, Y. Li, S. Che, N. K. Jha, and W. Zhang, "Software-defined design space exploration for an efficient dnn accelerator architecture," *IEEE Transactions on Computers*, 2020.
- [305] Y. Liu, B. Zhuang, C. Shen, H. Chen, and W. Yin, "Training compact neural networks via auxiliary overparameterization," *arXiv preprint arXiv:1909.02214*, 2019.
- [306] Y. Gao, H. Bai, Z. Jie, J. Ma, K. Jia, and W. Liu, "Mtl-nas: Task-agnostic neural architecture search towards general-purpose multi-task learning," in *Computer Vision and Pattern Recognition*, 2020.
- [307] D. So, Q. Le, and C. Liang, "The evolved transformer," in *International Conference on Machine Learning*, 2019.
- [308] J. Chen, K. Chen, X. Chen, X. Qiu, and X. Huang, "Exploring shared structures and hierarchies for multiple nlp tasks," *arXiv preprint arXiv:1808.07658*, 2018.
- [309] C. Wong, N. Houlsby, Y. Lu, and A. Gesmundo, "Transfer learning with neural automl," in *Advances in Neural Information Processing Systems*, 2018.
- [310] T. Véniat, O. Schwander, and L. Denoyer, "Stochastic adaptive neural architecture search for keyword spotting," in *International Conference on Acoustics, Speech and Signal Processing*, 2019.
- [311] K. Maziarsz, M. Tan, A. Khorlin, K.-Y. S. Chang, and A. Gesmundo, "Evolutionary-neural hybrid agents for architecture search," *arXiv preprint arXiv:1811.09828*, 2018.
- [312] H. Mazzawi, X. Gonzalvo, A. Kracun, P. Sridhar, N. Subrahmanya, I. Lopez-Moreno, H.-J. Park, and P. Violette, "Improving keyword spotting and language identification via neural architecture search at scale," in *INTERSPEECH*, 2019.
- [313] Y. Jiang, C. Hu, T. Xiao, C. Zhang, and J. Zhu, "Improved differentiable architecture search for language modeling and named entity recognition," in *Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing*, 2019.
- [314] Y. Wang, Y. Yang, Y. Chen, J. Bai, C. Zhang, G. Su, X. Kou, Y. Tong, M. Yang, and L. Zhou, "Textnas: A neural architecture search space tailored for text representation," in *AAAI Conference on Artificial Intelligence*, 2020.
- [315] D. Chen, Y. Li, M. Qiu, Z. Wang, B. Li, B. Ding, H. Deng, J. Huang, W. Lin, and J. Zhou, "Adabert: Task-adaptive bert compression with differentiable neural architecture search," *arXiv preprint arXiv:2001.04246*, 2020.
- [316] Y. Li, C. Hu, Y. Zhang, N. Xu, Y. Jiang, T. Xiao, J. Zhu, T. Liu, and C. Li, "Learning architectures from an extended search space for language modeling," *arXiv preprint arXiv:2005.02593*, 2020.
- [317] A. Baruwu, M. Abisiga, I. Gbadegesin, and A. Fakunle, "Leveraging end-to-end speech recognition with neural architecture search," *arXiv preprint arXiv:1912.05946*, 2019.
- [318] J. Li, C. Liang, B. Zhang, Z. Wang, F. Xiang, and X. Chu, "Neural architecture search on acoustic scene classification," *arXiv preprint arXiv:1912.12825*, 2019.
- [319] S. Ding, T. Chen, X. Gong, W. Zha, and Z. Wang, "Autospeech: Neural architecture search for speaker recognition," *arXiv preprint arXiv:2005.03215*, 2020.
- [320] Y. Gao, H. Yang, P. Zhang, C. Zhou, and Y. Hu, "Graphnas: Graph neural architecture search with reinforcement learning," *arXiv preprint arXiv:1904.09981*, 2019.

- [321] K. Zhou, Q. Song, X. Huang, and X. Hu, "Auto-gnn: Neural architecture search of graph neural networks," *arXiv preprint arXiv:1909.03184*, 2019.
- [322] Y. Zhao, D. Wang, X. Gao, R. Mullins, P. Lio, and M. Jamnik, "Probabilistic dual network architecture search on graphs," *arXiv preprint arXiv:2003.09676*, 2020.
- [323] T. Li, J. Zhang, K. Bao, Y. Liang, Y. Li, and Y. Zheng, "Autost: Efficient neural architecture search for spatio-temporal prediction," in *Knowledge Discovery & Data Mining*, 2020.
- [324] W. Cheng, Y. Shen, and L. Huang, "Differentiable neural input search for recommender systems," *arXiv preprint arXiv:2006.04466*, 2020.
- [325] P. Zhao, K. Xiao, Y. Zhang, K. Bian, and W. Yan, "Amer: Automatic behavior modeling and interaction exploration in recommender system," *arXiv preprint arXiv:2006.05933*, 2020.
- [326] Q. Song, D. Cheng, H. Zhou, J. Yang, Y. Tian, and X. Hu, "Towards automated neural interaction discovery for click-through rate prediction," *arXiv preprint arXiv:2007.06434*, 2020.
- [327] R. Miikkulainen, J. Liang, E. Meyerson, A. Rawal, D. Fink, O. Francon, B. Raju, H. Shahrzad, A. Navruzyan, N. Duffy *et al.*, "Evolving deep neural networks," in *Artificial Intelligence in the Age of Neural Networks and Brain Computing*, 2019, pp. 293–312.
- [328] S. Huang, X. Li, Z.-Q. Cheng, Z. Zhang, and A. Hauptmann, "Gnas: A greedy neural architecture search method for multi-attribute learning," in *ACM International Conference on Multimedia*, 2018.
- [329] R. Quan, X. Dong, Y. Wu, L. Zhu, and Y. Yang, "Auto-reid: Searching for a part-aware convnet for person re-identification," in *International Conference on Computer Vision*, 2019.
- [330] S. Zhang, R. Cao, X. Wei, P. Wang, and Y. Zhang, "Person re-identification with neural architecture search," in *Pattern Recognition and Computer Vision*, 2019.
- [331] Y. Li and I. King, "Architecture search for image inpainting," in *International Symposium on Neural Networks*, 2019.
- [332] S. Yang, W. Yang, and Z. Cui, "Pose neural fabrics search," *arXiv preprint arXiv:1909.07068*, 2019.
- [333] W. Zhang, L. Zhao, Q. Li, S. Zhao, Q. Dong, X. Jiang, T. Zhang, and T. Liu, "Identify hierarchical structures from task-based fmri data via hybrid spatiotemporal neural architecture search net," in *Medical Image Computing and Computer-Assisted Intervention*, 2019.
- [334] Y. Zhang, Q. Yao, and L. Chen, "Neural recurrent structure search for knowledge graph embedding," *arXiv preprint arXiv:1911.07132*, 2019.
- [335] J. An, H. Xiong, J. Huan, and J. Luo, "Ultrafast photorealistic style transfer via neural architecture search," in *AAAI Conference on Artificial Intelligence*, 2020.
- [336] S. Bianco, M. Buzzelli, G. Ciocca, and R. Schettini, "Neural architecture search for image saliency fusion," *Information Fusion*, vol. 57, pp. 89–101, 2020.
- [337] S. Zhang, L. Xiang, C. Li, Y. Wang, Z. Liu, Q. Zhang, and B. Li, "Preventing information leakage with neural architecture search," *arXiv preprint arXiv:1912.08421*, 2019.
- [338] K. Ho, A. Gilbert, H. Jin, and J. Collomosse, "Neural architecture search for deep image prior," *arXiv preprint arXiv:2001.04776*, 2020.
- [339] Y. Hu, X. Jiang, X. Liu, B. Zhang, J. Han, X. Cao, and D. Doermann, "Nas-count: Counting-by-density with neural architecture search," *arXiv preprint arXiv:2003.00217*, 2020.
- [340] H. Zhang, Q. Yao, M. Yang, Y. Xu, and X. Bai, "Efficient backbone search for scene text recognition," *arXiv preprint arXiv:2003.06567*, 2020.
- [341] E. Kokiopoulou, A. Hauth, L. Sbaiz, A. Gesmundo, G. Bartok, and J. Berent, "Fast task-aware architecture inference," *arXiv preprint arXiv:1902.05781*, 2019.
- [342] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Computer Vision and Pattern Recognition*, 2017.
- [343] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan *et al.*, "Searching for mobilenetv3," in *Computer Vision and Pattern Recognition*, 2019.
- [344] T. Chen, I. Goodfellow, and J. Shlens, "Net2net: Accelerating learning via knowledge transfer," in *International Conference on Learning Representations*, 2016.
- [345] T. Wei, C. Wang, Y. Rui, and C. W. Chen, "Network morphism," in *International Conference on Machine Learning*, 2016.
- [346] J. Wang, Z. Wei, T. Zhang, and W. Zeng, "Deeply-fused nets," *arXiv preprint arXiv:1605.07716*, 2016.
- [347] S. Saxena and J. Verbeek, "Convolutional neural fabrics," in *Advances in Neural Information Processing Systems*, 2016.
- [348] L. Zhao, M. Li, D. Meng, X. Li, Z. Zhang, Y. Zhuang, Z. Tu, and J. Wang, "Deep convolutional neural networks with merge-and-run mappings," in *International Joint Conference on Artificial Intelligence*, 2018.
- [349] R. Negrinho, M. Gormley, G. J. Gordon, D. Patil, N. Le, and D. Ferreira, "Towards modular and programmable architecture search," in *Advances in Neural Information Processing Systems*, 2019.
- [350] X. Du, T.-Y. Lin, P. Jin, G. Ghiasi, M. Tan, Y. Cui, Q. V. Le, and X. Song, "Spinenet: Learning scale-permuted backbone for recognition and localization," in *Computer Vision and Pattern Recognition*, 2020.
- [351] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in *Computer Vision and Pattern Recognition*, 2019.
- [352] H. Wu, J. Zhang, and K. Huang, "Sparsemask: Differentiable connectivity learning for dense image prediction," in *International Conference on Computer Vision*, 2019.
- [353] H. Wang and J. Huan, "Agan: Towards automated design of generative adversarial networks," *arXiv preprint arXiv:1906.11080*, 2019.
- [354] M. Schrimpf, S. Merity, J. Bradbury, and R. Socher, "A flexible approach to automated rnn architecture generation," *arXiv preprint arXiv:1712.07316*, 2017.
- [355] A. Rawal and R. Miikkulainen, "From nodes to networks: Evolving recurrent neural networks," *arXiv preprint arXiv:1803.04439*, 2018.
- [356] A. Ororbia, A. ElSaid, and T. Desell, "Investigating recurrent neural network memory structures using neuro-evolution," in *Genetic and Evolutionary Computation Conference*, 2019.
- [357] Z. Huang and B. Xiang, "Wenet: Weighted networks for recurrent network architecture search," *arXiv preprint arXiv:1904.03819*, 2019.
- [358] I. Fedorov, R. P. Adams, M. Mattina, and P. Whatmough, "Sparse: Sparse architecture search for cnns on resource-constrained microcontrollers," in *Advances in Neural Information Processing Systems*, 2019.
- [359] X. Li, Y. Zhou, Z. Pan, and J. Feng, "Partial order pruning: for best speed/accuracy trade-off in neural architecture search," in *Computer Vision and Pattern Recognition*, 2019.
- [360] J. Yu and T. Huang, "Network slimming by slimmable networks: Towards one-shot architecture search for channel numbers," *arXiv preprint arXiv:1903.11728*, 2019.
- [361] S. Huang, V. François-Lavet, and G. Rabusseau, "Neural architecture search for class-incremental learning," *arXiv preprint arXiv:1909.06686*, 2019.
- [362] K. Yu, R. Ranftl, and M. Salzmann, "How to train your super-net: An analysis of training heuristics in weight-sharing nas," *arXiv preprint arXiv:2003.04276*, 2020.
- [363] T. D. Ottelander, A. Dushatskiy, M. Virgolin, and P. A. Bosman, "Local search is a remarkably strong baseline for neural architecture search," *arXiv preprint arXiv:2004.08996*, 2020.
- [364] X. Dong and Y. Yang, "Network pruning via transformable architecture search," in *Advances in Neural Information Processing Systems*, 2019.
- [365] H. Chen, B. Z. Li'an Zhuo, X. Zheng, J. Liu, D. Doermann, and R. Ji, "Binarized neural architecture search," in *AAAI Conference on Artificial Intelligence*, 2020.
- [366] B. Deng, J. Yan, and D. Lin, "Peephole: Predicting network performance before training," *arXiv preprint arXiv:1712.03351*, 2017.
- [367] C. White, W. Neiswanger, and Y. Savani, "Bananas: Bayesian optimization with neural architectures for neural architecture search," *arXiv preprint arXiv:1910.11858*, 2019.
- [368] Y. Xu, Y. Wang, K. Han, H. Chen, Y. Tang, S. Jui, C. Xu, Q. Tian, and C. Xu, "Renas: Architecture ranking for powerful networks," *arXiv preprint arXiv:1910.01523*, 2019.
- [369] D. Friede, J. Lukasik, H. Stuckenschmidt, and M. Keuper, "A variational-sequential graph autoencoder for neural architecture performance prediction," *arXiv preprint arXiv:1912.05317*, 2019.
- [370] C. Wei, C. Niu, Y. Tang, and J. Liang, "Npenas: Neural predictor guided evolution for neural architecture search," *arXiv preprint arXiv:2003.12857*, 2020.
- [371] X. Ning, Y. Zheng, T. Zhao, Y. Wang, and H. Yang, "A generic graph-based neural architecture encoding scheme for predictor-based nas," in *European Conference on Computer Vision*, 2020.

- [372] Y. Tang, Y. Wang, Y. Xu, H. Chen, B. Shi, C. Xu, C. Xu, Q. Tian, and C. Xu, "A semi-supervised assessor of neural architectures," in *Computer Vision and Pattern Recognition*, 2020.
- [373] J. Li, Y. Liu, J. Liu, and W. Wang, "Neural architecture optimization with graph vae," *arXiv preprint arXiv:2006.10310*, 2020.
- [374] J. Mellor, J. Turner, A. Storkey, and E. J. Crowley, "Neural architecture search without training," *arXiv preprint arXiv:2006.04647*, 2020.
- [375] V. Nguyen, T. Le, M. Yamada, and M. A. Osborne, "Optimal transport kernels for sequential and parallel neural architecture search," *arXiv preprint arXiv:2006.07593*, 2020.
- [376] B. Ru, C. Lyle, L. Schut, M. van der Wilk, and Y. Gal, "Revisiting the train loss: an efficient performance estimator for neural architecture search," *arXiv preprint arXiv:2006.04492*, 2020.
- [377] S. Yan, Y. Zheng, W. Ao, X. Zeng, and M. Zhang, "Does unsupervised architecture representation learning help neural architecture search?" *arXiv preprint arXiv:2006.06936*, 2020.
- [378] T. Chau, Ł. Dudziak, M. S. Abdelfattah, R. Lee, H. Kim, and N. D. Lane, "Brp-nas: Prediction-based nas using gcns," *arXiv preprint arXiv:2007.08668*, 2020.
- [379] R. Luo, X. Tan, R. Wang, T. Qin, E. Chen, and T.-Y. Liu, "Neural architecture search with gbdtd," *arXiv preprint arXiv:2007.04785*, 2020.
- [380] C. White, W. Neiswanger, S. Nolen, and Y. Savani, "A study on encodings for neural architecture search," *arXiv preprint arXiv:2007.04965*, 2020.
- [381] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Machine Learning*, 2015.
- [382] C. Ying, A. Klein, E. Christiansen, E. Real, K. Murphy, and F. Hutter, "Nas-bench-101: Towards reproducible neural architecture search," in *International Conference on Machine Learning*, 2019.
- [383] X. Dong and Y. Yang, "Nas-bench-201: Extending the scope of reproducible neural architecture search," in *International Conference on Learning Representations*, 2020.
- [384] N. Klyuchnikov, I. Trofimov, E. Artemova, M. Salnikov, M. Fedorov, and E. Burnaev, "Nas-bench-nlp: Neural architecture search benchmark for natural language processing," *arXiv preprint arXiv:2006.07116*, 2020.
- [385] Y. Xu, Y. Wang, K. Han, S. Jui, C. Xu, Q. Tian, and C. Xu, "Renas: Relativistic evaluation of neural architecture search," in *AAAI Conference on Artificial Intelligence*, 2020.
- [386] X. Zhou, D. Dou, and B. Li, "Searching for stage-wise neural graphs in the limit," *arXiv preprint arXiv:1912.12860*, 2019.
- [387] I. Radosavovic, R. P. Kosaraju, R. Girshick, K. He, and P. Dollár, "Designing network design spaces," in *Computer Vision and Pattern Recognition*, 2020.
- [388] M. Lindauer and F. Hutter, "Best practices for scientific research on neural architecture search," *arXiv preprint arXiv:1909.02453*, 2019.
- [389] C. Sciuto, K. Yu, M. Jaggi, C. Musat, and M. Salzmann, "Evaluating the search phase of neural architecture search," *arXiv preprint arXiv:1902.08142*, 2019.
- [390] A. Yang, P. M. Esperança, and F. M. Carlucci, "Nas evaluation is frustratingly hard," in *International Conference on Learning Representations*, 2020.
- [391] K. Yu, C. Sciuto, M. Jaggi, C. Musat, and M. Salzmann, "Evaluating the search phase of neural architecture search," in *International Conference on Learning Representations*, 2020.
- [392] R. Luo, X. Tan, R. Wang, T. Qin, E. Chen, and T.-Y. Liu, "Semi-supervised neural architecture search," *arXiv preprint arXiv:2002.10389*, 2020.
- [393] Q. Gao, Z. Luo, and D. Klabjan, "Efficient architecture search for continual learning," *arXiv preprint arXiv:2006.04027*, 2020.
- [394] C. He, M. Annamaram, and S. Avestimehr, "Fednas: Federated deep learning via neural architecture search," *arXiv preprint arXiv:2004.08546*, 2020.
- [395] M. Xu, Y. Zhao, K. Bian, G. Huang, Q. Mei, and X. Liu, "Federated neural architecture search," *arXiv preprint arXiv:2002.06352*, 2020.
- [396] H. Jin, Q. Song, and X. Hu, "Auto-keras: An efficient neural architecture search system," in *SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019.
- [397] N. Erickson, J. Mueller, A. Shirikov, H. Zhang, P. Larroy, M. Li, and A. Smola, "Autogluon-tabular: Robust and accurate automl for structured data," *arXiv preprint arXiv:2003.06505*, 2020.
- [398] K. Kandasamy, K. R. Vysyaraju, W. Neiswanger, B. Paria, C. R. Collins, J. Schneider, B. Póczos, and E. P. Xing, "Tuning hyperparameters without grad students: Scalable and robust bayesian optimisation with dragonfly," *Journal of Machine Learning Research*, vol. 21, no. 81, pp. 1–27, 2020.
- [399] L. Zimmer, M. Lindauer, and F. Hutter, "Auto-pytorch tabular: Multi-fidelity metalearning for efficient and robust autodl," *arXiv preprint arXiv:2006.13799*, 2020.
- [400] Y. Zhang, Z. Lin, J. Jiang, Q. Zhang, Y. Wang, H. Xue, C. Zhang, and Y. Yang, "Deeper insights into weight sharing in neural architecture search," *arXiv preprint arXiv:2001.01431*, 2020.
- [401] E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le, "Autoaugment: Learning augmentation policies from data," in *Computer Vision and Pattern Recognition*, 2019.
- [402] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [403] J. G. Robles and J. Vanschoren, "Learning to reinforcement learn for neural architecture search," *arXiv preprint arXiv:1911.03769*, 2019.
- [404] A. Shaw, W. Wei, W. Liu, L. Song, and B. Dai, "Meta architecture search," in *Advances in Neural Information Processing Systems*, 2019.
- [405] X. Chen, Y. Duan, Z. Chen, H. Xu, Z. Chen, X. Liang, T. Zhang, and Z. Li, "Catch: Context-based meta reinforcement learning for transferrable architecture search," in *European Conference on Computer Vision*, 2020.
- [406] A. Zela, J. Siems, and F. Hutter, "Nas-bench-1shot1: Benchmarking and dissecting one-shot neural architecture search," in *International Conference on Learning Representations*, 2020.
- [407] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," in *International Conference on Learning Representations*, 2016.
- [408] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size," *arXiv preprint arXiv:1602.07360*, 2016.
- [409] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in *Advances in Neural Information Processing Systems*, 2015.
- [410] T.-J. Yang, Y.-H. Chen, and V. Sze, "Designing energy-efficient convolutional neural networks using energy-aware pruning," in *Computer Vision and Pattern Recognition*, 2017.
- [411] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
- [412] Z. Huang and N. Wang, "Data-driven sparse structure selection for deep neural networks," in *European Conference on Computer Vision*, 2018.
- [413] Y. Gong, L. Liu, M. Yang, and L. Bourdev, "Compressing deep convolutional networks using vector quantization," *arXiv preprint arXiv:1412.6115*, 2014.
- [414] W. Chen, J. Wilson, S. Tyree, K. Weinberger, and Y. Chen, "Compressing neural networks with the hashing trick," in *International Conference on Machine Learning*, 2015.
- [415] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Quantized neural networks: Training neural networks with low precision weights and activations," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6869–6898, 2017.
- [416] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "Xnor-net: Imagenet classification using binary convolutional neural networks," in *European Conference on Computer Vision*, 2016.
- [417] A. Ashok, N. Rhinehart, F. Beainy, and K. M. Kitani, "N2n learning: Network to network compression via policy gradient reinforcement learning," in *International Conference on Learning Representations*, 2018.
- [418] Y. He, J. Lin, Z. Liu, H. Wang, L.-J. Li, and S. Han, "Amc: Automl for model compression and acceleration on mobile devices," in *European Conference on Computer Vision*, 2018.
- [419] X. Xiao, Z. Wang, and S. Rajasekaran, "Autoprune: Automatic network pruning by regularizing auxiliary parameters," in *Advances in Neural Information Processing Systems*, 2019.
- [420] H. Cai, T. Wang, Z. Wu, K. Wang, J. Lin, and S. Han, "On-device image classification with proxyless neural architecture search and quantization-aware fine-tuning," in *International Conference on Computer Vision Workshops*, 2019.

- [421] M. G. d. Nascimento, T. W. Costain, and V. A. Prisacariu, "Finding non-uniform quantization schemes using multi-task gaussian processes," in *European Conference on Machine Learning*, 2020.
- [422] H. Yu, Q. Han, J. Li, J. Shi, G. Cheng, and B. Fan, "Search what you want: Barrier panelty nas for mixed precision quantization," in *European Conference on Computer Vision*, 2020.
- [423] Y. Chen, G. Meng, Q. Zhang, X. Zhang, L. Song, S. Xiang, and C. Pan, "Joint neural architecture search and quantization," *arXiv preprint arXiv:1811.09426*, 2018.
- [424] X. Lu, H. Huang, W. Dong, X. Li, and G. Shi, "Beyond network pruning: a joint search-and-training approach," in *International Joint Conference on Artificial Intelligence*, 2020.
- [425] T. Wang, K. Wang, H. Cai, J. Lin, Z. Liu, H. Wang, Y. Lin, and S. Han, "Apq: Joint search for network architecture, pruning and quantization policy," in *Computer Vision and Pattern Recognition*, 2020.
- [426] K. Ahmed and L. Torresani, "Maskconnect: Connectivity learning by gradient descent," in *European Conference on Computer Vision*, 2018.
- [427] Y. Wang, X. Zhang, L. Xie, J. Zhou, H. Su, B. Zhang, and X. Hu, "Pruning from scratch," in *AAAI Conference on Artificial Intelligence*, 2020.
- [428] Y. Li, W. Li, M. Danelljan, K. Zhang, S. Gu, L. Van Gool, and R. Timofte, "The heterogeneity hypothesis: Finding layer-wise dissimilated network architecture," *arXiv preprint arXiv:2006.16242*, 2020.
- [429] A. Gordon, E. Eban, O. Nachum, B. Chen, H. Wu, T.-J. Yang, and E. Choi, "Morphnet: Fast & simple resource-constrained structure learning of deep networks," in *Computer Vision and Pattern Recognition*, 2018.
- [430] Z. Chen, J. Niu, L. Xie, X. Liu, L. Wei, and Q. Tian, "Network adjustment: Channel search guided by flops utilization ratio," in *Computer Vision and Pattern Recognition*, 2020.
- [431] Z. Cai and N. Vasconcelos, "Rethinking differentiable search for mixed-precision neural networks," in *Computer Vision and Pattern Recognition*, 2020.
- [432] H. Cai, J. Lin, Y. Lin, Z. Liu, K. Wang, T. Wang, L. Zhu, and S. Han, "Automl for architecting efficient and specialized neural networks," *IEEE Micro*, vol. 40, no. 1, pp. 75–82, 2019.
- [433] Q. Lu, W. Jiang, X. Xu, Y. Shi, and J. Hu, "On neural architecture search for resource-constrained hardware platforms," *arXiv preprint arXiv:1911.00105*, 2019.
- [434] M. S. Abdelfattah, Ł. Dudziak, T. Chau, R. Lee, H. Kim, and N. D. Lane, "Best of both worlds: Automl codesign of a cnn and its hardware accelerator," *arXiv preprint arXiv:2002.05022*, 2020.
- [435] S. Gupta and B. Akin, "Accelerator-aware neural network design using automl," *arXiv preprint arXiv:2003.02838*, 2020.
- [436] W. Jiang, Q. Lou, Z. Yan, L. Yang, J. Hu, X. S. Hu, and Y. Shi, "Device-circuit-architecture co-exploration for computing-in-memory neural accelerators," *IEEE Transactions on Computers*, 2020.
- [437] W. Jiang, L. Yang, E. H.-M. Sha, Q. Zhuge, S. Gu, S. Dasgupta, Y. Shi, and J. Hu, "Hardware/software co-exploration of neural architectures," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2020.
- [438] W. Jiang, L. Yang, S. Dasgupta, J. Hu, and Y. Shi, "Standing on the shoulders of giants: Hardware and neural architecture co-search with hot start," *arXiv preprint arXiv:2007.09087*, 2020.
- [439] J. Lin, W.-M. Chen, Y. Lin, J. Cohn, C. Gan, and S. Han, "Mcunet: Tiny deep learning on iot devices," *arXiv preprint arXiv:2007.10319*, 2020.
- [440] L. Liu and J. Deng, "Dynamic deep neural networks: Optimizing accuracy-efficiency trade-offs by selective execution," in *AAAI Conference on Artificial Intelligence*, 2018.
- [441] G. Huang, D. Chen, T. Li, F. Wu, L. van der Maaten, and K. Q. Weinberger, "Multi-scale dense networks for resource efficient image classification," in *International Conference on Learning Representations*, 2018.
- [442] X. Wang, F. Yu, Z.-Y. Dou, T. Darrell, and J. E. Gonzalez, "Skipnet: Learning dynamic routing in convolutional networks," in *European Conference on Computer Vision*, 2018.
- [443] Y. Li, L. Song, Y. Chen, Z. Li, X. Zhang, X. Wang, and J. Sun, "Learning dynamic routing for semantic segmentation," in *Computer Vision and Pattern Recognition*, 2020.
- [444] M. McGill and P. Perona, "Deciding how to decide: Dynamic routing in artificial neural networks," *arXiv preprint arXiv:1703.06217*, 2017.
- [445] H. Li, H. Zhang, X. Qi, R. Yang, and G. Huang, "Improved techniques for training adaptive deep networks," in *International Conference on Computer Vision*, 2019.
- [446] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *International Conference on Machine Learning*, 2010.
- [447] Z. Wu, T. Nagarajan, A. Kumar, S. Rennie, L. S. Davis, K. Grauman, and R. Feris, "Blockdrop: Dynamic inference paths in residual networks," in *Computer Vision and Pattern Recognition*, 2018.
- [448] Z. Yuan, B. Wu, Z. Liang, S. Zhao, W. Bi, and G. Sun, "S2dnas: Transforming static cnn model for dynamic inference via neural architecture search," *arXiv preprint arXiv:1911.07033*, 2019.
- [449] J. Schmidhuber, "Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook," Ph.D. dissertation, Technische Universität München, 1987.
- [450] T. Schaul and J. Schmidhuber, "Metalearning," *Scholarpedia*, vol. 5, no. 6, p. 4650, 2010.
- [451] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people," *Behavioral and Brain Sciences*, vol. 40, 2017.
- [452] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," in *Advances in Neural Information Processing Systems*, 2016.
- [453] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine, "One-shot visual imitation learning via meta-learning," *arXiv preprint arXiv:1709.04905*, 2017.
- [454] M. Ren, E. Triantafillou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel, "Meta-learning for semi-supervised few-shot classification," *arXiv preprint arXiv:1803.00676*, 2018.
- [455] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap, "Meta-learning with memory-augmented neural networks," in *International Conference on Machine Learning*, 2016.
- [456] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel, "RL²: Fast reinforcement learning via slow reinforcement learning," *arXiv preprint arXiv:1611.02779*, 2016.
- [457] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel, "A simple neural attentive meta-learner," in *International Conference on Learning Representations*, 2018.
- [458] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, 2017.
- [459] K. Li and J. Malik, "Learning to optimize," *arXiv preprint arXiv:1606.01885*, 2016.
- [460] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," in *International Conference on Learning Representations*, 2017.
- [461] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International Conference on Machine Learning*, 2017.
- [462] C. Lemke, M. Budka, and B. Gabrys, "Metalearning: a survey of trends and technologies," *Artificial Intelligence Review*, vol. 44, no. 1, pp. 117–130, 2015.
- [463] J. Vanschoren, "Meta-learning: A survey," *arXiv preprint arXiv:1810.03548*, 2018.
- [464] S. Doveh, E. Schwartz, C. Xue, R. Feris, A. Bronstein, R. Giryes, and L. Karlinsky, "Metadapt: Meta-learned task-adaptive architecture for few-shot classification," *arXiv preprint arXiv:1912.00412*, 2019.
- [465] T. Elsken, B. Staffler, J. H. Metzen, and F. Hutter, "Meta-learning of neural architectures for few-shot learning," in *Computer Vision and Pattern Recognition*, 2020.
- [466] A. Dubatovka, E. Kokiopoulou, L. Sbaiz, A. Gesmundo, G. Bartok, and J. Berent, "Ranking architectures using meta-learning," *arXiv preprint arXiv:1911.11481*, 2019.
- [467] M. Claesen and B. De Moor, "Hyperparameter search in machine learning," *arXiv preprint arXiv:1502.02127*, 2015.
- [468] M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, and N. De Freitas, "Learning to learn by gradient descent by gradient descent," in *Advances in Neural Information Processing Systems*, 2016.
- [469] I. Loshchilov and F. Hutter, "Sgdr: Stochastic gradient descent with warm restarts," in *International Conference on Learning Representations*, 2017.

- [470] Z. Li and S. Arora, "An exponential learning rate schedule for deep learning," in *International Conference on Learning Representations*, 2020.
- [471] N. Loizou, S. Vaswani, I. Laradji, and S. Lacoste-Julien, "Stochastic polyak step-size for sgd: An adaptive learning rate for fast convergence," *arXiv preprint arXiv:2002.10542*, 2020.
- [472] N. Mu, Z. Yao, A. Gholami, K. Keutzer, and M. Mahoney, "Parameter re-initialization through cyclical batch size schedules," *arXiv preprint arXiv:1812.01216*, 2018.
- [473] Z. Yao, A. Gholami, Q. Lei, K. Keutzer, and M. W. Mahoney, "Hessian-based analysis of large batch training and robustness to adversaries," in *Advances in Neural Information Processing Systems*, 2018.
- [474] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2019.
- [475] G. Zhang, C. Wang, B. Xu, and R. Grosse, "Three mechanisms of weight decay regularization," in *International Conference on Learning Representations*, 2019.
- [476] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *International Conference on Machine Learning*, 2013.
- [477] T. Lancewicki and S. Kopru, "Automatic and simultaneous adjustment of learning rate and momentum for stochastic gradient-based optimization methods," in *International Conference on Acoustics, Speech and Signal Processing*, 2020.
- [478] H. Mendoza, A. Klein, M. Feurer, J. T. Springenberg, and F. Hutter, "Towards automatically-tuned neural networks," in *Workshop on Automatic Machine Learning*, 2016.
- [479] I. Loshchilov and F. Hutter, "Cma-es for hyperparameter optimization of deep neural networks," *arXiv preprint arXiv:1604.07269*, 2016.
- [480] J. S. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for hyper-parameter optimization," in *Advances in Neural Information Processing Systems*, 2011.
- [481] M. Feurer and F. Hutter, "Hyperparameter optimization," in *Automated Machine Learning*. Springer, Cham, 2019, pp. 3–33.
- [482] T. Yu and H. Zhu, "Hyper-parameter optimization: A review of algorithms and applications," *arXiv preprint arXiv:2003.05689*, 2020.
- [483] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *Journal of Machine Learning Research*, vol. 13, no. 1, pp. 281–305, 2012.
- [484] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Advances in neural information processing systems*, 2012, pp. 2951–2959.
- [485] J. Bergstra, D. Yamins, and D. Cox, "Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures," in *International Conference on Machine Learning*, 2013.
- [486] D. Maclaurin, D. Duvenaud, and R. Adams, "Gradient-based hyperparameter optimization through reversible learning," in *International Conference on Machine Learning*, 2015.
- [487] S. R. Young, D. C. Rose, T. P. Karnowski, S.-H. Lim, and R. M. Patton, "Optimizing deep learning hyper-parameters through an evolutionary algorithm," in *Workshop on Machine Learning in High-Performance Computing Environments*, 2015.
- [488] S. Falkner, A. Klein, and F. Hutter, "Bohb: Robust and efficient hyperparameter optimization at scale," *arXiv preprint arXiv:1807.01774*, 2018.
- [489] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A novel bandit-based approach to hyperparameter optimization," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6765–6816, 2017.
- [490] M. Feurer, J. T. Springenberg, and F. Hutter, "Initializing bayesian hyperparameter optimization via meta-learning," in *AAAI Conference on Artificial Intelligence*, 2015.
- [491] I. Bello, B. Zoph, V. Vasudevan, and Q. V. Le, "Neural optimizer search with reinforcement learning," in *International Conference on Machine Learning*, 2017.
- [492] A. Zela, A. Klein, S. Falkner, and F. Hutter, "Towards automated deep learning: Efficient joint neural architecture and hyperparameter search," *arXiv preprint arXiv:1807.06906*, 2018.
- [493] A. Klein and F. Hutter, "Tabular benchmarks for joint architecture and hyperparameter optimization," *arXiv preprint arXiv:1905.04970*, 2019.
- [494] X. Dong, M. Tan, A. W. Yu, D. Peng, B. Gabrys, and Q. V. Le, "Autohas: Differentiable hyper-parameter and architecture search," *arXiv preprint arXiv:2006.03656*, 2020.
- [495] S. Lim, I. Kim, T. Kim, C. Kim, and S. Kim, "Fast autoaugment," in *Advances in Neural Information Processing Systems*, 2019.
- [496] C. Lin, M. Guo, C. Li, X. Yuan, W. Wu, J. Yan, D. Lin, and W. Ouyang, "Online hyper-parameter learning for auto-augmentation strategy," in *International Conference on Computer Vision*, 2019.
- [497] D. Ho, E. Liang, X. Chen, I. Stoica, and P. Abbeel, "Population based augmentation: Efficient learning of augmentation policy schedules," in *International Conference on Machine Learning*, 2019.
- [498] R. Hataya, J. Zdenek, K. Yoshizoe, and H. Nakayama, "Faster autoaugment: Learning augmentation strategies using backpropagation," *arXiv preprint arXiv:1911.06987*, 2019.
- [499] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, "Randaugment: Practical automated data augmentation with a reduced search space," in *Computer Vision and Pattern Recognition Workshops*, 2020.
- [500] X. Zhang, Q. Wang, J. Zhang, and Z. Zhong, "Adversarial autoaugment," in *International Conference on Learning Representations*, 2020.
- [501] C. Xie, M. Tan, B. Gong, J. Wang, A. L. Yuille, and Q. V. Le, "Adversarial examples improve image recognition," in *Computer Vision and Pattern Recognition*, 2020.
- [502] L. Wei, A. Xiao, L. Xie, X. Chen, X. Zhang, and Q. Tian, "Circumventing outliers of autoaugment with knowledge distillation," in *European Conference on Computer Vision*, 2020.
- [503] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, p. 60, 2019.
- [504] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Computer Vision and Pattern Recognition*, 2017.
- [505] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "Shufflenet v2: Practical guidelines for efficient cnn architecture design," in *European Conference on Computer Vision*, 2018.
- [506] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Computer Vision and Pattern Recognition*, 2018.
- [507] H. Zhang, C. Wu, Z. Zhang, Y. Zhu, Z. Zhang, H. Lin, Y. Sun, T. He, J. Mueller, R. Manmatha *et al.*, "Resnest: Split-attention networks," *arXiv preprint arXiv:2004.08955*, 2020.
- [508] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," *arXiv preprint arXiv:1710.05941*, 2017.
- [509] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang *et al.*, "Deep high-resolution representation learning for visual recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [510] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [511] F. P. Such, A. Rawal, J. Lehman, K. O. Stanley, and J. Clune, "Generative teaching networks: Accelerating neural architecture search by learning to generate synthetic training data," *arXiv preprint arXiv:1912.07768*, 2019.
- [512] J. Gu and V. Tresp, "Search for better students to learn distilled knowledge," *arXiv preprint arXiv:2001.11612*, 2020.
- [513] Y. Liu, X. Jia, M. Tan, R. Vemulapalli, Y. Zhu, B. Green, and X. Wang, "Search to distill: Pearls are everywhere but not the eyes," in *Computer Vision and Pattern Recognition*, 2020.
- [514] R. Panda, M. Merler, M. Jaiswal, H. Wu, K. Ramakrishnan, U. Finkler, C.-F. Chen, M. Cho, D. Kung, R. Feris *et al.*, "Nastransfer: Analyzing architecture transferability in large scale neural architecture search," *arXiv preprint arXiv:2006.13314*, 2020.
- [515] J. Fernandez-Marques, P. N. Whatmough, A. Mundy, and M. Mattina, "Searching for winograd-aware quantized networks," *arXiv preprint arXiv:2002.10711*, 2020.