This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.

When NAS Meets Trees: An Efficient Algorithm for Neural Architecture Search

Guocheng Qian¹ Xuanyang Zhang² Guohao Li¹ Chen Zhao¹ Yukang Chen³ Xiangyu Zhang² Bernard Ghanem¹ Jian Sun² ¹ King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia ² MEGVII Technology³ The Chinese University of Hong Kong

{guocheng.qian, bernard.ghanem}@kaust.edu.sa

Abstract

The key challenge in neural architecture search (NAS) is designing how to explore wisely in the huge search space. We propose a new NAS method called TNAS (NAS with trees), which improves search efficiency by exploring only a small number of architectures while also achieving a higher search accuracy. TNAS introduces an architecture tree and a binary operation tree, to factorize the search space and substantially reduce the exploration size. TNAS performs a modified bi-level Breadth-First Search in the proposed trees to discover a high-performance architecture. Impressively, TNAS finds the global optimal architecture on CIFAR-10 with test accuracy of 94.37% in four GPU hours in NAS-Bench-201. The average test accuracy is 94.35%, which outperforms the state-of-the-art. Code is available at: https://github.com/guochengqian/TNAS.

1. Introduction

Neural architecture search has spurred increasing interest in both academia and industry for its ability in finding high-performance neural network architectures with minimal human intervention. To achieve the most accurate NAS algorithm, one can explore all candidate architectures, training each one to convergence, and picking the bestperforming architecture. However, this brute-force NAS is infeasible due to the enormous search space. Therefore, one of the key questions towards a successful NAS algorithm is: how to efficiently explore the search space?

One-shot NAS [6, 27, 3, 24] impressively improved the efficiency of NAS. One-shot NAS leverages a weightsharing strategy and approximately trains only one network, called the *supernet*, which subsumes all candidate architectures. Each candidate architecture directly inherits weights from the supernet without training. Despite the efficiency



(a) The entire search space. Each dot represents an architecture.





(b) The pruned search space after the first search stage.



(c) The pruned search space after the second search stage.

(d) The single candidate architecture found after the third stage.

Figure 1: TNAS hierarchically factorizes the search space and gradually prunes the unpromising architectures. The colorbar shows the global rankings of architectures on CIFAR-10 [18] in NAS-Bench-201 [13]. Red stars indicate top-10 architectures.

of one-shot NAS algorithms, they incur architecture evaluation degradation, *i.e.* the architecture performance evaluated using the weight-sharing is not correctthat, which leads to a degraded search accuracy [38, 21].

In this work, we diverge from the paradigms set by early NAS, and instead design a new algorithm to explore the search space in a wiser manner. Consider a search space \mathcal{A} where the number of candidate operations is M and the number of architecture layers to search is L. The size of the entire search space $|\mathcal{A}|$ equals to M^L . If M = 2 or L = 1, $|\mathcal{A}|$ can be drastically reduced to 2^L or M. The intuition

behind our work is to develop a method that factorizes the operation space (size M) and the architecture layers (size L), and thus reduces the exploration size exponentially.

Contributions. (1) We introduce an architecture tree and a binary operation tree to factorize the search space L and M, respectively. By combining the two trees, we iteratively branch a search space into two exclusive subspaces. (2) We propose a novel, flexible, accurate, and efficient NAS algorithm, called TNAS: NAS with trees. TNAS performs a modified bi-level Breadth-First Search (BFS) in the two proposed trees. By adjusting the expansion depths of the BFS, TNAS explicitly controls the exploration size N and is able to exponentially reduce N from M^L to $O(L \log_2 M)$. The essence of TNAS is illustrated in Figure 1. (3) TNAS is is able to find the global optimal architecture on CIFAR-10 [18] (94.37% test acc.) in NAS-Bench-201 [13] within 4 GPU hours on one GTX2080Ti GPU. TNAS outperforms the RL and EA based NAS [45, 28] as well as one-shot NAS [27, 9], with a similar search cost.

2. Related Work

The computational bottleneck of NAS is *exploring* candidate architectures in this huge search space and *exploiting* each one (*i.e.* score the architecture by training to convergence). Through the work, we name the number of architectures to score as the *exploration size*, denoted as N. To alleviate the computational bottleneck, NAS algorithms should consider: (i) how to explore wisely, where time can be saved if the algorithm explores more among the "good" architectures and less on the "bad" ones, and (ii) how to exploit wisely, where training each network to convergence just to know the architecture's performance then throwing weights away is inefficient.

Explore wisely. Early methods adopt Reinforcement Learning [44, 2, 45] or Evolutionary Algorithms [30, 29, 28] to auto-explore the huge search space. Although early NAS methods have been able to discover architectures that outperform manually designed networks, they consume significant computational resources. This is primarily because these algorithms require a large exploration size to achieve a decent search accuracy. Progressive NAS is a method that factorizes the search space into a product of smaller search spaces and can greatly reduce the exploration size. PNAS [23] and P-DARTS [9] start searching with shallow models and gradually progress to deeper ones. Li et al. propose block-wise progressive NAS [19, 20] that consider the architectures is built by sequential blocks and search the architecture block by block. SGAS [21], GreedyNAS [36], and [17, 41, 34] progressively shrink the search space by dropping unpromising candidates. These progressive NAS methods require a much smaller exploration size, but their greedy nature hampers their search accuracy. Our TNAS

designs a new paradigm for exploring wisely by introducing two trees to factorize the search space.

Exploit wisely. A straightforward idea of reducing exploitation is to train fewer epochs as done in Block-QNN [42]. A more advanced solution is to share weights among child networks, apart from training them from scratch. This weight-sharing strategy was first proposed by ENAS [6, 27] and has inspired many following works, including one-shot NAS [3, 24, 15, 32, 26, 37, 10, 14]. To alleviate the evaluation degradation [38, 4, 21] issues of oneshot NAS caused by weight-sharing, few-shot NAS [40, 31] were proposed by training k supernets instead of training only one. Another line of work to exploit wisely is accuracy prediction [23, 11, 43], where an accuracy predictor is learned to directly estimate an architecture's accuracy without training it completely. Recently, metric-based NAS methods [25, 8, 1, 39, 17, 7] have emerged, using welldesigned metrics to score the sampled architectures quickly with significantly less training or even no training. Since our paper focuses on how to wisely explore the NAS search space, wise exploitation is an orthogonal direction. In fact, we highlight here that our TNAS can be applied with nearly all the aforementioned exploit-wise NAS methods.

3. Methodology

We present TNAS (NAS with trees) to efficiently find a high-performance architecture by performing a modified bilevel Breadth-First Search in the proposed architecture tree T_A and binary operation tree T_O .

3.1. Architecture Tree T_A

Given a search space with L layers and M operations per layer, we propose an architecture tree \mathcal{T}_A to factorize the one-shot architecture and to exponentially reduce the exploration size. The architecture tree T_A is illustrated in Figure 2(a). Each node in the tree represents an architecture. The root node is the *M*-path *L*-layer one-shot architecture. Each path in a layer denotes a distinct operation from M candidate operations. T_A has a maximum depth level equal to L. For each node (architecture) at depth i $(i \in [0, 1, \dots, L-1))$, the tree separates the M operations in layer i into M branches each with a single operation. Such branching is repeated for each node, until the leaf nodes are reached. Each leaf node represents a distinct single-path architecture. The union of the leaf nodes is the set of all candidate architectures. Note that if layer *i* contains multiple operations, the output of this layer will be the summation of the outputs of all operations at this layer, as inspired by the one-shot NAS [3] and is formulated as:

$$\bar{o}^i(x) = \sum_{o_j \in \mathcal{O}} o^i_j(x), \tag{1}$$



Figure 2: Illustration of the architecture tree T_A and the proposed Breadth-First Search (BFS).

where $\mathcal{O} = \{o_j \mid j = 1, 2, ..., M\}$ denotes M different operations and x denotes the input feature map.

Breadth-First Search (BFS) in \mathcal{T}_A . Here, we show that the architecture search can be done by performing our modified BFS in the architecture tree \mathcal{T}_A . Our BFS requires a hyperparameter, the expansion depth denoted as d_a , where the subscript a denotes "architecture". BFS starts at the root node (the one-shot model) at depth 0, expands all its successors until depth d_a , and obtains up to M^{d_a} leaf nodes after expansion. BFS scores the subnets defined by these leaf nodes, and picks the node with the highest score as the root node for the next step. The above procedure is defined as a decision step, and is repeated until a single-path architecture is determined. The score function can be chosen to be the validation performance after training, or a metric function proposed by any metric-based NAS method such as the number of linear regions [25]. For simplicity, we choose the scoring function to be validation performance in our experiments. The expansion depth d_a of the BFS denotes how many layers to branch at each decision step. As illustrated in Figure 2(b), the BFS with $d_a = 1$ is a sequential, greedy NAS algorithm that decides the operation for the architecture layer by layer, similar to the progressive NAS method SGAS [21]. The BFS with $d_a = L$ as shown in Figure 2(c) works as the brute-force NAS, where only 1 decision step is required. The BFS explores all M^L subnets and decides the operation for all of the layers at the same decision step. When $d_a = k \in \{2, \dots, L-1\}, k$ layers are branched in each decision step, M^k subnets need to be scored, and $\left\lceil \frac{L}{k} \right\rceil$ decision steps are required. This case works similar to the block-wise NAS [19], while our BFS does not require any block-level supervision.

3.2. Binary Operation Tree T_O

We propose a binary operation tree T_O that hierarchically factorizes the operation space to further reduce the exploration size. Each node in T_O is an *operation group* con-



Figure 3: The binary operation tree T_O .

sisting of one or more distinct operations. The root node represents \mathcal{O} , the entire operation space containing all Moperations. \mathcal{T}_O starts from the root node and branches it into two child nodes that represent two exclusive operation groups. Such branching is repeated for each node until a leaf node that represents a single operation is reached. \mathcal{T}_O has M leaf nodes. The union of leaf nodes is \mathcal{O} . Taking NAS-Bench-201 [13] operation space as an example, we illustrate the \mathcal{T}_O in Figure 3.

Breadth-First Search (BFS) in T_O . The expansion depth of our modified BFS in \mathcal{T}_O is denoted as d_o . BFS starts at the root node (the entire operation space) at depth 0, expands all its successors until depth d_o , and obtains up to 2^{d_o} leaf nodes after expansion. These leaf nodes represent the current candidate operation groups. BFS scores the architectures equipped with these different operation groups, and picks the node defined by the operation group with the highest score as the root node for the next stage. The above procedure is defined as a decision stage, and is repeated until a single operation is picked. Note that each architecture layer can choose different operation groups at each decision stage. If $d_o = 1$, the BFS decides the operation groups per depth following T_O . In this case, BFS consists of $\left\lceil \frac{\log_2(M-1)+1}{d} \right\rceil = 3$ decision stages. At the 1st stage, BFS decides among None or Not None for each architecture layer. At the 2^{nd} stage, for those layers that chose Not None, the algorithm decides among the Convolution group



Figure 4: Illustration of TNAS ($d_a = 2, d_o = 1$).

or *Topology* group. At the final stage, the algorithm will pick a single operation for each layer. If $d_o = 3$, BFS only needs one decision stage to decide which single operation to choose for each layer.

3.3. TNAS

We present a new NAS algorithm: Neural Architecture Search with Trees (TNAS). Given a search space with Mcandidate operations and L layers, TNAS constructs a binary operation tree \mathcal{T}_O and an architecture tree \mathcal{T}_A . TNAS starts from the M-path L-layer one-shot model, and performs bi-level Breadth-First Search on \mathcal{T}_O and \mathcal{T}_A . At the outer loop, TNAS performs BFS with the expansion depth $d_o = 1$ on \mathcal{T}_O by default, to make a large d_a feasible. The outer loop requires $\lceil \log_2 (M-1) + 1 \rceil$ decision stages. Each stage branches each operation group of the chosen layers into two child operation groups, which define the operation search space for the inner loop. The outer loop repeats the decision stage until every architecture layer reaches a leaf node of \mathcal{T}_O , *i.e.* all the layers pick a single operation. In the inner loop, TNAS performs BFS with an expansion depth d_a on \mathcal{T}_A . The inner loop takes $\left| \frac{L}{d_a} \right|$ decision steps. Each step chooses d_a undecided layers to branch, obtains 2^{d_a} subnets, scores each subnet, and then chooses the highest scoring one. The chosen subnet will be used to replace the one-shot model and become the starting point for the next step. The inner loop repeats the above decision step until it chooses a leaf node of \mathcal{T}_A , *i.e.* all layers of the architecture have decided their operation group at the current decision stage. We illustrate the TNAS algorithm $(d_o = 1, d_a = 2)$ in Figure 4. The NAS-Bench-201 [13]

Table 1: **State-of-the-art comparison on NAS-Bench-201**. Top-1 test accuracy (mean and standard deviation over 5 runs) are reported. For each dataset, **optimum** indicates the best test accuracy achievable in the NAS-Bench-201 search space.

Architecture	CIFAR-10	CIFAR-100	ImageNet-16-120	Search Cost (hours)	Search Method
optimum	94.37	73.51	47.31	-	-
ResNet [16]	93.97	70.86	43.63	-	-
REA [28]	93.92 ± 0.30	71.84 ± 0.99	45.54 ± 1.03	3.3	EA
REINFORCE [35]	93.85 ± 0.37	71.71 ± 1.09	45.24 ± 1.18	3.3	RL
RS [5]	93.70 ± 0.36	71.04 ± 1.07	44.57 ± 1.25	3.3	random
NAS w.o. Training [25]	91.78 ± 1.45	67.05 ± 2.89	37.07 ± 6.39	-	training-free
TE-NAS [8]	93.90 ± 0.47	71.24 ± 0.56	42.38 ± 0.46	-	training-free
RSPS [22]	87.66 ± 1.69	58.33 ± 4.34	31.14 ± 3.88	2.2	random
ENAS [27]	54.30 ± 0.00	15.61 ± 0.00	16.32 ± 0.00	3.7	EA
DARTS (2nd) [24]	54.30 ± 0.00	15.61 ± 0.00	16.32 ± 0.00	8.3	gradient
GDAS [12]	93.61 ± 0.09	70.70 ± 0.30	41.84 ± 0.90	8.0	gradient
DARTS- [10]	93.80 ± 0.40	71.53 ± 1.51	45.12 ± 0.82	3.2	gradient
VIM-NAS [33]	94.31 ± 0.11	$\textbf{73.07} \pm 0.58$	46.27 ± 0.17	-	gradient
TNAS (ours)	94.35±0.03	73.02±0.34	46.31±0.24	3.6	tree
TNAS (best)	94.37	73.09	46.33	3.6	tree

search space (*i.e.* M = 5 and L = 6) is used as an example. **Exploration size analysis.** Given a search space with M operations and L layers, TNAS reduces the exploration size exponentially from M^L to:

$$N = O\left(2^{d_o d_a} \times \left\lceil \frac{L}{d_a} \right\rceil \times \left\lceil \frac{\log_2\left(M-1\right)+1}{d_o} \right\rceil\right)$$
(2)

4. Experiments

Setup. We evaluate TNAS on NAS-Bench-201 [13] with $(d_o = 1, d_a = 6)$. We train each architecture over 2 epochs and use the top-1 accuracy on validation set as the score for the architecture. If the architecture consists of a layer with multiple operations, the output of this layer is the sum of all outputs as Equation 1. Note that other scoring methods aforementioned in Section 2 can also be applied.

Results. Table 1 compares TNAS with SOTA. *TNAS finds the global optimal architecture in CIFAR-10 [18] within 4 GPU hours.* TNAS achieves 94.35% average test accuracy, outperforming all other NAS methods. We highlight that TNAS outperforms the REA [28], REIN-FORCE [35] and random search (RS [5]) with a similar search cost, which clearly demonstrates the benefit of our NAS paradigm. TNAS also performs significantly better than the one-shot based methods, such as ENAS [27],GDAS [12] and DARTS- [10], while being more efficient.

5. Conclusion

We present a novel NAS algorithm, TNAS, that performs bi-level BFS on the proposed binary operation tree and the architecture tree. By adjusting the search depths on the trees, TNAS can explicitly control the exploration size. TNAS finds the global optimal architecture in NAS-Bench-201 [13] with a search cost of less than 4 GPU hours.

Acknowledgments This work was done when Guocheng was remotely interned at Megvii technology. This work was also supported by the KAUST Office of Sponsored Research (OSR) through VCC funding.

References

- Mohamed S. Abdelfattah, Abhinav Mehrotra, Lukasz Dudziak, and Nicholas Donald Lane. Zero-cost proxies for lightweight NAS. In *International Conference on Learning Representations (ICLR)*, 2021.
- [2] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architectures using reinforcement learning. In *ICLR (Poster)*. OpenReview.net, 2017.
- [3] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *International Conference on Machine Learning*, pages 549–558, 2018.
- [4] Gabriel Bender, Hanxiao Liu, Bo Chen, Grace Chu, Shuyang Cheng, Pieter-Jan Kindermans, and Quoc V. Le. Can weight sharing outperform random architecture search? an investigation with tunas. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 14311– 14320, 2020.
- [5] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. J. Mach. Learn. Res., 13:281– 305, 2012.
- [6] Andrew Brock, Theodore Lim, James M. Ritchie, and Nick Weston. SMASH: one-shot model architecture search through hypernetworks. In *ICLR (Poster)*. OpenReview.net, 2018.
- [7] Boyu Chen, Peixia Li, Baopu Li, Chen Lin, Chuming Li, Ming Sun, Junjie Yan, and Wanli Ouyang. Bn-nas: Neural architecture search with batch normalization. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision (ICCV), pages 307–316, October 2021.
- [8] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. In *International Conference* on Learning Representations (ICLR), 2021.
- [9] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [10] Xiangxiang Chu, Xiaoxing Wang, Bo Zhang, Shun Lu, Xiaolin Wei, and Junchi Yan. DARTS-: robustly stepping out of performance collapse without indicators. In *ICLR*. Open-Review.net, 2021.
- [11] Xiaoliang Dai, Peizhao Zhang, Bichen Wu, Hongxu Yin, Fei Sun, Yanghan Wang, Marat Dukhan, Yunqing Hu, Yiming Wu, Yangqing Jia, Peter Vajda, Matt Uyttendaele, and Niraj K. Jha. Chamnet: Towards efficient network design through platform-aware model adaptation. In *CVPR*, pages 11398–11407, 2019.
- [12] Xuanyi Dong and Yi Yang. Searching for a robust neural architecture in four gpu hours. In *Proceedings of the IEEE Conference on computer vision and pattern recognition*, pages 1761–1770, 2019.
- [13] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *ICLR*. OpenReview.net, 2020.

- [14] Yu-Chao Gu, Li-Juan Wang, Yun Liu, Yi Yang, Yu-Huan Wu, Shao-Ping Lu, and Ming-Ming Cheng. Dots: Decoupling operation and topology in differentiable architecture search. In *CVPR*, 2021.
- [15] Zichao Guo, X. Zhang, Haoyuan Mu, Wen Heng, Z. Liu, Y. Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In *ECCV*, 2020.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.
- [17] Yiming Hu, Yuding Liang, Zichao Guo, Ruosi Wan, X. Zhang, Yichen Wei, Qingyi Gu, and Jian Sun. Angle-based search space shrinking for neural architecture search. *ArXiv*, abs/2004.13431, 2020.
- [18] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- [19] Changlin Li, Jiefeng Peng, Liuchun Yuan, Guangrun Wang, Xiaodan Liang, Liang Lin, and Xiaojun Chang. Blockwisely supervised neural architecture search with knowledge distillation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020.
- [20] Changlin Li, Tao Tang, Guangrun Wang, Jiefeng Peng, Bing Wang, Xiaodan Liang, and Xiaojun Chang. BossNAS: Exploring hybrid CNN-transformers with block-wisely selfsupervised neural architecture search. In *ICCV*, 2021.
- [21] Guohao Li, Guocheng Qian, Itzel C. Delgadillo, Matthias Müller, Ali K. Thabet, and Bernard Ghanem. SGAS: sequential greedy architecture search. In *CVPR*, pages 1617–1627. IEEE, 2020.
- [22] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *Uncertainty in Artificial Intelligence*, pages 367–377. PMLR, 2020.
- [23] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In Proceedings of the European Conference on Computer Vision (ECCV), pages 19–34, 2018.
- [24] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *ICLR (Poster)*. Open-Review.net, 2019.
- [25] Joe Mellor, Jack Turner, Amos J. Storkey, and Elliot J. Crowley. Neural architecture search without training. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 7588–7598. PMLR, 2021.
- [26] Houwen Peng, Hao Du, Hongyuan Yu, Qi Li, Jing Liao, and Jianlong Fu. Cream of the crop: Distilling prioritized paths for one-shot neural architecture search. In *NeurIPS*, 2020.
- [27] Hieu Pham, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 4092–4101. PMLR, 2018.
- [28] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4780–4789, 2019.

- [29] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc V. Le, and Alexey Kurakin. Large-scale evolution of image classifiers. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 2902–2911. PMLR, 2017.
- [30] K. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.
- [31] Xiu Su, Shan You, Mingkai Zheng, Fei Wang, Chen Qian, Changshui Zhang, and Chang Xu. K-shot NAS: learnable weight-sharing for NAS with k-shot supernets. In Marina Meila and Tong Zhang, editors, *ICML*, volume 139, pages 9880–9890, 2021.
- [32] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuandong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, Peter Vajda, and Joseph E. Gonzalez. Fbnetv2: Differentiable neural architecture search for spatial and channel dimensions. In *CVPR*, pages 12962–12971. Computer Vision Foundation / IEEE, 2020.
- [33] Yaoming Wang, Yuchen Liu, Wenrui Dai, Chenglin Li, Junni Zou, and Hongkai Xiong. Learning latent architectural distribution in differentiable neural architecture search via variational information maximization. In *ICCV*, pages 12292– 12301. IEEE, 2021.
- [34] Junru Wu, Xiyang Dai, Dongdong Chen, Yinpeng Chen, Mengchen Liu, Ye Yu, Zhangyang Wang, Zicheng Liu, Mei Chen, and Lu Yuan. Stronger nas with weaker predictors. arXiv preprint arXiv:2102.10490, 2021.
- [35] Chris Ying, Aaron Klein, Esteban Real, Eric Christiansen, Kevin P. Murphy, and Frank Hutter. Nas-bench-101: Towards reproducible neural architecture search. In *ICML*, 2019.
- [36] Shan You, Tao Huang, Mingmin Yang, Fei Wang, Chen Qian, and Changshui Zhang. Greedynas: Towards fast one-shot nas with greedy supernet. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1999–2008, 2020.
- [37] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas S. Huang, Xiaodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. In ECCV (7), volume 12352 of Lecture Notes in Computer Science, pages 702–717. Springer, 2020.
- [38] Kaicheng Yu, Christian Sciuto, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. Evaluating the search phase of neural architecture search. In *ICLR*. OpenReview.net, 2020.
- [39] Xuanyang Zhang, Pengfei Hou, Xiangyu Zhang, and Jian Sun. Neural architecture search with random labels. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10907–10916, 2021.
- [40] Yiyang Zhao, Linnan Wang, Yuandong Tian, Rodrigo Fonseca, and Tian Guo. Few-shot neural architecture search. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 12707–12718. PMLR, 2021.
- [41] Xiawu Zheng, Rongrong Ji, Qiang Wang, Qixiang Ye, Zhenguo Li, Yonghong Tian, and Qi Tian. Rethinking performance estimation in neural architecture search. 2020

IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 11353–11362, 2020.

- [42] Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu. Practical block-wise neural network architecture generation. In *The IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), pages 2423–2432, 2018.
- [43] Zhao Zhong, Zichen Yang, Boyang Deng, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu. Blockqnn: Efficient block-wise neural network architecture generation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(7):2314–2328, 2021.
- [44] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. In *ICLR*. OpenReview.net, 2017.
- [45] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 8697– 8710, 2018.