

CLOSE: Curriculum Learning On the Sharing Extent Towards Better One-shot NAS

Zixuan Zhou^{13*}, Xuefei Ning^{12*}, Yi Cai¹, Jiashu Han¹,
Yiping Deng², Yuhan Dong³, Huazhong Yang¹, and Yu Wang^{1†}

¹ Department of Electronic Engineering, Tsinghua University
*zhouzx21@mails.tsinghua.edu.cn, *foxdoraame@gmail.com,
†yu-wang@tsinghua.edu.cn

² Huawei TCS Lab

³ Tsinghua Shenzhen International Graduate School

Abstract. One-shot Neural Architecture Search (NAS) has been widely used to discover architectures due to its efficiency. However, previous studies reveal that one-shot performance estimations of architectures might not be well correlated with their performances in stand-alone training because of the excessive sharing of operation parameters (i.e., large sharing extent) between architectures. Thus, recent methods construct even more over-parameterized supernet to reduce the sharing extent. But these improved methods introduce a large number of extra parameters and thus cause an undesirable trade-off between the training costs and the ranking quality. To alleviate the above issues, we propose to apply Curriculum Learning On Sharing Extent (CLOSE) to train the supernet both efficiently and effectively. Specifically, we train the supernet with a large sharing extent (an easier curriculum) at the beginning and gradually decrease the sharing extent of the supernet (a harder curriculum). To support this training strategy, we design a novel supernet (CLOSENet) that decouples the parameters from operations to realize a flexible sharing scheme and adjustable sharing extent. Extensive experiments demonstrate that CLOSE can obtain a better ranking quality across different computational budget constraints than other one-shot supernets, and is able to discover superior architectures when combined with various search strategies. Code is available at https://github.com/walkerning/aw_nas.

Keywords: Neural Architecture Search (NAS), One-shot Estimation, Parameter Sharing, Curriculum Learning, Graph-Based Encoding

1 Introduction

Neural Architecture Search (NAS) [39] has achieved great success in automatically designing deep neural networks (DNN) in the past few years. However, traditional NAS methods are extremely time-consuming for discovering the optimal

* Equal contribution.

architectures, since each architecture sampled in the search process needs to be trained from scratch separately. To alleviate the severe problem of search inefficiency, one-shot NAS proposes to share operation parameters among candidate architectures in a “supernet” and train this supernet to evaluate all sampled candidate architectures [4,22,1,18], which reduces the overall search cost from thousands of GPU days to only a few GPU hours.

Despite its efficiency, previous studies reveal that one-shot NAS suffers from the poor ranking correlation between one-shot estimations and stand-alone estimations, which leads to unfair comparisons between the candidate architectures [19,34,35,20]. Ning et al. [20] give some insights on the failure of one-shot estimations. They conclude that one of the main causes of the poor ranking quality is the **large sharing extent of the supernet**. Several recent studies try to improve one-shot NAS by addressing the large sharing extent issue. Zhao et al. [37] confirm the negative impact of co-adaptation of parameters in one supernet. Thus, they split the whole search space into several smaller ones, and train a supernet for each subspace. Su et al. [29] reveal that only one copy of parameters is hard to be maintained for massive architectures. Therefore, they duplicate each parameter of the supernet into several copies, train and then estimate with all the duplicates. However, these improved methods introduce a large number of parameters, which barricades the supernet training. As a result, they have to make a trade-off between the training costs and the ranking quality.

In this paper, we propose to adopt the Curriculum Learning On Sharing Extent (CLOSE) to train the one-shot supernet efficiently and effectively. The underlying intuition behind our method is that training with a large sharing extent can **efficiently bootstrap** the supernet, since the number of parameters to be optimized is much smaller. While in the later training stage, using a supernet with a smaller sharing extent (i.e., a more over-parameterized supernet) can **improve the saturating ranking quality**. Thus, CLOSE uses a relatively large sharing extent in the early training stage of the supernet, then gradually decreases the supernet sharing extent. To support this training strategy, we design a new supernet with an adjustable sharing extent, namely CLOSENet, of which the sharing extent can be flexibly adjusted in the training process. The difference between CLOSENet and the vanilla supernet is that, the construction of vanilla supernets presets the sharing scheme between any architecture pairs, i.e., designates which parameter is shared by which operations in different architectures. In contrast, CLOSENet could flexibly adjust the sharing scheme and extent between architecture pairs, during the training process.

In summary, the contributions of our work are as follows:

1. We propose to apply Curriculum Learning On Sharing Extent (CLOSE) to efficiently and effectively train the one-shot supernet. Specifically, we use a larger sharing extent in the early stages to accelerate the training process, and gradually switch to smaller ones to boost the saturating performances.
2. To fit the CLOSE strategy, we design a novel supernet (CLOSENet) with an adjustable sharing extent. Different from the vanilla supernet with an unad-

justable sharing scheme and sharing extent, CLOSENet can flexibly adapt its sharing scheme and sharing extent during the training process.

3. Extensive experiments on four NAS benchmarks show that CLOSE can achieve a better ranking quality under any computational budgets. When searching for the optimal architectures, CLOSE enables one-shot NAS to find superior architectures compared to existing one-shot NAS methods.

2 Related Work

2.1 One-shot Neural Architecture Search (NAS)

Neural Architecture Search (NAS) is proposed to find optimal architectures automatically. However, the vanilla sample-based NAS methods [39] are extremely time-consuming. To make it more efficient, Pham et al. [22] propose the parameter sharing technique by constructing an over-parameterized network, namely supernet, to share parameters among the candidate architectures. Based on the parameter sharing technique, various one-shot NAS methods are proposed to efficiently search for optimal architectures by only training “one” supernet. Bender et al. [1] propose to directly train the whole supernet with a path-dropout strategy. Liu et al. [18] develop a differentiable search strategy and use it in conjunction with the parameter sharing technique. Guo et al. [10] propose to separate the stages of supernet training and architecture search.

2.2 Weakness and Improvement of One-shot NAS

Despite its high efficiency, one-shot NAS suffers from the poor ranking correlation between the architecture performances using one-shot training and stand-alone training. Sciuto et al. [34] discover that the parameter-sharing rankings do not correlate with stand-alone rankings by conducting a series of experiments in a toy search space. Zela et al. [35] confirm the poor ranking quality in a much larger NAS-Bench-1shot1 search space. Luo et al. [19] make a further investigation of the one-shot NAS, and attribute the poor ranking quality to the insufficient and imbalanced training, and the coupling of training and search phases. Ning et al. [20] provide comprehensive evaluations on multiple NAS benchmarks, and conclude three perspectives to improve the ranking quality of the one-shot NAS, i.e., reducing the temporal variance, sampling bias or parameter sharing extent.

Recent studies adopt the direction of sharing extent reduction to improve the one-shot NAS. Ning et al. [20] prune the search space to reduce the number of candidate architectures, and reveal the improvement of the ranking quality in the pruned search space. But the ranking quality in the overall search space is not improved. Zhao et al. [37] propose *Few-shot NAS* to split the whole search space into several subspaces, and train a single supernet for each subspace. Su et al. [29] propose *K-shot NAS* to duplicate each parameter of the supernet into several copies, and estimate architectures’ performances with all of them. However, these two methods reduce sharing extent with even more over-parameterized supernets, which brings extra computational costs.

2.3 Curriculum Learning

Bengio et al. [2] first propose curriculum learning (CL) strategy based on the learning process of humans and animals in the real world. The basic idea of the CL strategy is to guide models to learn from easier data (tasks) to harder data (tasks). In the past few years, many studies have successfully applied CL strategy in various applications [8,14,23,30,26,7,9], and demonstrated that CL can improve the models’ generalization capacity and convergence speed. Besides common CL methods that adjust the data, there exist CL methods that conduct curriculum learning on the model capacity. Karras et al. [15] propose to progressively increase the model capacity of the GAN to speed up and stabilize the training. Soviany et al. [28] propose a general CL framework at the model level that adjusts the curriculum by gradually increasing the model capacity.

2.4 NAS Benchmarks

NAS benchmarks enable researchers to reproduce the NAS experiments easily and compare different NAS methods fairly. NAS-Bench-201 [6] constructs a cell-based NAS search space containing 15625 architectures and provides their complete training information. NAS-Bench-301 [27] uses a surrogate model to predict the performances of approximately 10^{18} architectures in a more generic search space, with the stand-alone performance of 60k landmark architectures. Different from NAS-Bench-201 and NAS-Bench-301 that focus on topological search spaces, NDS [24] provides benchmarks on two non-topological search spaces (e.g., ResNet [11] and ResNeXt [31]).

3 Method

3.1 Motivation and Preliminary Experiments

In one-shot NAS, many operations in different architectures share the same parameter, while their desired parameters are not necessarily the same. The excessive sharing of parameters, i.e., the large sharing extent, has been widely regarded as the most important factor causing the unsatisfying performance estimation [3,36,20,37,29]. The most recent studies [37,29] improve the ranking quality by reducing the sharing extent. But their methods cause an inevitable trade-off between the training cost and ranking quality at the same time.

Supernets with larger sharing extents (i.e., more parameters) are easier to train in the early training stage. We verify this statement with an experiment on two popular cell-based NAS benchmarks. We construct two supernets with different sharing extents: Supernet-1 (a cell shown in top-left of Fig. 1) is a vanilla supernet adopted by many one-shot NAS methods (e.g., DARTS [18]), in which the compound edges in one cell use different copies of parameters. While Supernet-2 (a cell shown in the bottom-left of Fig. 1) shares only one copy of parameters for all the compound edges in each cell, which leads to a much larger sharing extent than Supernet-1. For example, the parameters of




Fig. 1: Comparison of supernet with small and large sharing extent. The left part shows a cell-architecture that uses Supernet-1 (top) and Supernet-2 (bottom). The right part shows the Kendall’s Tau of Supernet-1 and Supernet-2 throughout the training process on NAS-Bench-201 and NAS-Bench-301.

$conv_{3 \times 3}$ in edge (1,3) and (2,4) are different when using Supernet-1, but the same when using Supernet-2. We train them to convergence and use Kendall’s Tau (see Sec. 4.1 for definition), to evaluate the ranking correlation between the estimated performances by the supernet and the ground-truth performances.

Fig. 1 (right) shows that, on one hand, using a smaller sharing extent (more parameters, larger supernet capacity) can alleviate the undesired coadaptation between architectures, and has the potential to achieve higher saturating performances. The Kendall’s Tau of Supernet-2 is slightly worse than that of Supernet-1 when the supernet are relatively well-trained (epoch 800 to 1000). On the other hand, training the supernet with higher sharing extent than the vanilla one (fewer parameters, smaller supernet capacity) greatly accelerates the training process of parameters, and help the supernet obtain a good ranking quality faster. In the early stage of the training process (epoch 0 to 600), Supernet-2 has a much higher Kendall’s Tau than Supernet-1.

Based on the above results and analysis, **a natural idea to achieve a win-win scenario of supernet training efficiency and high ranking quality is to adapt the sharing extent during the supernet training process.** We can draw parallels between this idea and the CL methods on model capacity [15,28] (see Sec. 2.3), as they all progressively increase the capacity of the model or supernet to achieve training speedup and better performances in the mean time. Based on the above idea, we propose to employ Curriculum Learning On the Sharing Extent (CLOSE) of the supernet. And to enable the adaption of the sharing extent during the training process, we design a novel supernet, CLOSENet, whose sharing extent can be easily adjusted.

In the following, we first demonstrate the construction of CLOSENet in Sec. 3.2. Then, in Sec. 3.3, we describe CLOSE with some necessary training techniques to achieve the best ranking quality.

3.2 CLOSENet: A Supernet with An Adjustable Sharing Extent

The design of CLOSENet is illustrated in Fig. 2. The key idea behind CLOSENet is to decouple the parameters from operations to enable flexible sharing scheme and adjustable sharing extent. Specifically, we design the GLOBal Operation Weight (GLOW) Block to store the parameters, and design a GATE module for assigning the proper GLOW block to each operation.




Fig. 2: CLOSENet contains GLOW blocks to store the parameters of candidate operations, and the GATE module to assign GLOW blocks to operations.

For a better understanding, we take the generic cell-based topological search space as an example. We will show in the appendix that CLOSENet can also adapt to other types of search spaces, such as ResNet-like search spaces.

Generic Cell-based Topological Search Space. In cell-based search space, a complete architecture is stacked by a cell-architecture for multiple times (e.g., 15 times on NAS-Bench-201). A cell-architecture can be represented as a *directed acyclic graph (DAG)*. Each node x^i represents a feature map, while each edge $o^{(i,j)}$ represents an operation that transforms x^i to x^j with the corresponding

parameters. For each node j , the feature x_j is defined as:

$$x^j = \sum_{i < j} o^{(i,j)}(x^i, W^{(i,j)}), \quad (1)$$

where $o(x, W)$ denotes that the operation o transforms the feature x with the parameters from W .

Global Operation Weight (GLOW) Block. The function of GLOW blocks is to store the parameters of candidate operations, as shown in Fig. 2 (middle). In the forward pass (as shown in the bottom of Fig. 2), the GLOW blocks are assigned to each operation via the GATE module (will be introduced in the following), and each operation can use the parameters from its assigned block to process the input feature map.

Specifically, we denote G_i as the i -th GLOW block, and $c^{(i,j)}$ as the index of the assigned block for the operation in edge (i, j) . Then, the computation of the feature x^j in Eq. 1 can be rewritten as:

$$x^j = \sum_{i < j} o^{(i,j)}(x^i, G_{c^{(i,j)}}) \quad (2)$$

GATE Module. We design a GATE module for assigning GLOW blocks to operations. The GATE module consists of an architecture embedder and a MLP module, as shown in Fig. 2 (top). We construct a GCN-based architecture embedder [21], and use it to compute the node embeddings in the architecture. Then, we concatenate the embeddings of the input and output nodes of each operation and feed it into the MLP to get the assignment of the GLOW block.

Specifically, we denote E_i as the embedding of node i , and K as the number of GLOW blocks. For a cell-architecture a with N nodes, we first obtain the node embeddings by the architecture embedder as Eq. 3, and then calculate the probability distribution in edge (i, j) as Eq. 4 and Eq. 5.

$$[E_1, E_2, E_3, \dots, E_N] = \text{ArchEmb}(a) \quad (3)$$

$$[\lambda_1^{(i,j)}, \lambda_2^{(i,j)}, \lambda_3^{(i,j)}, \dots, \lambda_K^{(i,j)}] = \text{MLP}(\text{concat}(E_i, E_j)) \quad (4)$$

$$\Pr(c^{(i,j)} = k) = \frac{\exp(\lambda_k^{(i,j)})}{\sum_{k^0=1}^K \exp(\lambda_{k^0}^{(i,j)})} \quad (5)$$

To allow the back-propagation of gradients, we apply the reparameterization trick on Eq. 2 and Eq. 5, and rewrite the computation of x^j as:

$$x^j = \sum_{i < j} \sum_{k=1}^K h_k^{(i,j)} o^{(i,j)}(x^i, G_k), \quad (6)$$

$$h^{(i,j)} = \arg \max_k (\lambda_k^{(i,j)} + g_k), \quad (7)$$

where $h^{(i,j)}$ is a one-hot vector of dimension K , and g_k are i.i.d samples from Gumbel(0, 1). To make Eq. 7 differentiable, we relax the arg max function to a softmax function as:

$$\hat{h}_k^{(i,j)} = \frac{\exp((\lambda_k^{(i,j)} + g_k)/\tau)}{\sum_{k^0=1}^K \exp((\lambda_{k^0}^{(i,j)} + g_{k^0})/\tau)}, \quad (8)$$

where τ is the Gumbel-Softmax temperature. We use Eq. 7 in the forward pass, and use Eq. 8 in the backward pass to allow gradient propagation.

Adjustment of Sharing Extent Denote $E^{(i,j)}$ as the set of cell-architectures that contain the edge from node i to node j . For edge (i, j) , we define its sharing extent $s^{(i,j)}$ as the average number of architectures that share one GLOW block in CLOSENet. The sharing extent of the supernet s equals the sum of sharing extent of all the edges:

$$s = \sum_{i,j} s^{(i,j)} = \sum_{i,j} \frac{jE^{(i,j)}j}{K} = \frac{\sum_{i,j} jE^{(i,j)}j}{K}. \quad (9)$$

Therefore, we can naturally adjust the sharing extent of CLOSENet by adding or reducing the GLOW blocks. CLOSENet with more GLOW blocks (larger K) has a smaller sharing extent and vice versa.

Strengths Compared to Vanilla Supernets The vanilla supernet and its variants (e.g., *K-shot* and *Few-shot* supernets [29,37]) preset the sharing scheme and extent by attaching a fixed set of parameters to each operation. On the contrary, CLOSENet decouples the parameters from the operations and enables **the dynamic decision of sharing scheme** based on a graph-based encoding of architecture operations. Specifically, the vanilla supernet shares parameters according to the position specified by the node indexes, i.e., the operations in the same ‘‘position’’ share the same parameters across different architectures. This sharing scheme is not flexible and can be suboptimal in some cases. For example, as shown in Fig. 3(upper), the 1 × 1 convolutions on the 0-2 edge share the same parameters between the two architectures, while they should have vastly different optimal parameters. Intuitively, if two operations in two architectures have similar data processing functionality, it might be more reasonable to share their parameters. The design of CLOSENet matches this intuition: The GATE module learns to pick the right GLOW block for each operation based on the graph-based encoding of all operations and topology in the cell architecture. Instead of presetting the sharing scheme according to the position information, CLOSENet takes a more flexible and reasonable way to dynamically determine which block each operation should use, and thereby designates which operations

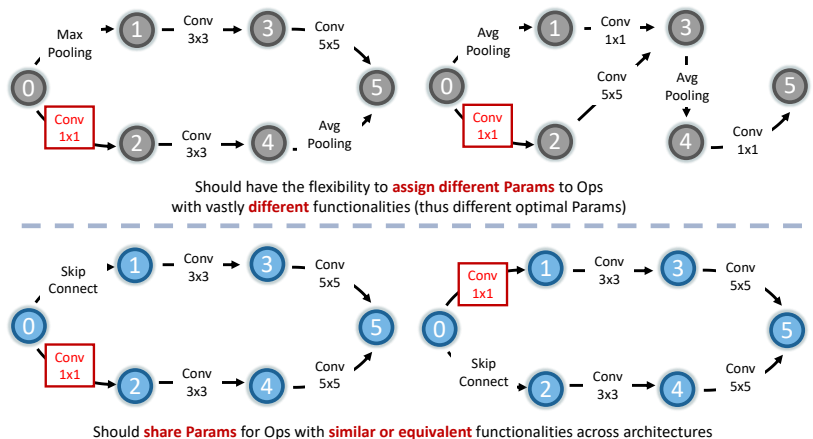


Fig. 3: Two examples that show the strengths of CLOSENet. The sharing scheme of the vanilla supernet, i.e., sharing parameters between operations with the same position indexes, is improper in these two cases. In contrast, CLOSENet designates a more proper sharing pattern between operations according to their graph-based encoding given by GATE.

in different architectures should share their parameters. For example, as shown in Fig. 3(bottom), since the two 1×1 convolutions are equivalent in two isomorphic architectures despite having different position indexes, it is reasonable for them to share parameters. The vanilla supernet uses different parameters for these two convolutions, while CLOSENet assigns the same GLOW block for them.

Moreover, this decoupling enables us to **flexibly adjust the sharing extent** by changing K in Eq. 9. Thus, CLOSENet enables us to apply our curriculum learning-like training strategy. In summary, both the dynamic sharing scheme and the adjustable sharing extent make CLOSENet a more powerful supernet.

3.3 CLOSE: Curriculum Learning On Sharing Extent

We borrow the idea of curriculum learning to design a novel supernet training strategy CLOSE. Specifically, we initialize the CLOSENet with only one GLOW block at the beginning. This large sharing extent helps us to train the supernet much faster. Then, we gradually add GLOW blocks at preset epochs to reduce the sharing extent. In this way, CLOSE not only accelerates the supernet training, but also improves the saturating ranking quality of the supernet.

When switching the curriculum (i.e., increasing the sharing extent), we add a new GLOW block into CLOSENet and a corresponding MLP output unit to the GATE module. How to initialize the newly added parameters is critical to the performance of CLOSENet. Additionally, the regular schedule for the learn-

ing rate does not fit for CLOSE (see below). Correspondingly, we propose two techniques, the Weight Inherited Technique and the Schedule Restart Technique.




Fig. 4: WIT for GLOW blocks (top) and the MLP output units (bottom).

Weight Inherit Technique (WIT). Instead of randomly initializing the new GLOW block and MLP output unit, we make their weights inherit from those of previous GLOW blocks and MLP output units, as shown in Fig. 4. This helps with the more efficient training of the new GLOW block and MLP unit.

Schedule Restart Technique (SRT). In the training process, the learning rate is reduced gradually to approach the optimal solution. That is to say, it will become quite small after many epochs. However, following this schedule, CLOSE might fail to jump out of the local optimal solution of the preceding curriculum. To overcome this problem, we propose to restart the learning rate and schedule at preset epochs. With SRT, CLOSE can quickly reach the new optimal solution after switching to a new curriculum.

4 Experiments

4.1 Evaluation of Ranking Quality

We evaluate our method on four NAS search spaces, including NAS-Bench-201 [6], NAS-Bench-301 [27], NDS ResNet [24] and NDS ResNeXt-A [24]. The training configurations are shown in the appendix. Following previous studies [20,21], we use two evaluation criteria as follows:

- Kendall’s Tau (KD): The relative difference of the number of concordant pairs and discordant pairs, which reflects the overall ranking correlation.
- Precision@topK (P@topK): The proportion of true top-K architectures in the top-K architectures according to the one-shot estimations, which reflects the ability of identifying the top-performing architectures.

Comparison with Vanilla One-Shot Baselines We compare CLOSENet with vanilla supernets on four NAS benchmarks. As shown in Fig. 5, CLOSENet achieves a higher KD and P@top5% on all the NAS benchmarks. Moreover, we can see that throughout the training process, CLOSENet consistently achieves higher ranking quality, which implies CLOSENet’s superiority to the vanilla supernet under any budget for supernet training.





Fig. 5: Comparison of different criteria with the vanilla one-shot supernet on four NAS benchmarks. X-axis: Training epochs. Y-axis: Evaluation criteria.

Comparison with Improved One-Shot Methods Fig. 6 compares CLOSE with previous work on improving NAS evaluation strategy, including *EPEE* [20], *AngleNet* [13], *K-shot NAS* [29] and *Few-shot NAS* [37]. Results show that CLOSE reaches SOTA KDs on all the three datasets of NAS-Bench-201.



Method	Kendall's Tau	
	C-100	IN-16
EPEE	0.5600	0.5400
AngleNet	0.6040	0.5445
K-shot NAS	0.6122	0.5633
CLOSE	0.6693	0.6632

Fig. 6: Comparison with previous improved methods on NAS-Bench-201 for three datasets, i.e. CIFAR-10, CIFAR-100 (C-10) and ImageNet-16 (IN-16).

4.2 Evaluation of Search Performance

We combine CLOSE with various search strategies, including DARTS [18], SNAS [32] and CARS [33], to evaluate whether it improves the search performance.

Results We run DARTS and SNAS search with CLOSE on NAS-Bench-301 and show the derived architecture accuracy in Fig. 7. We can see that CLOSE benefits the search process significantly. In particular, it can alleviate the collapse issue of DARTS caused by the improper preference of parameter-free operations (i.e., *skip_connect*) in early training stages [16,12], as it provides a less biased estimation (see Appendix 1.1).




Fig. 7: Evaluation of CLOSE with two search strategies in the NAS-Bench-301 search space. X-axis: Training epochs. Y-axis: Test accuracy.

We run CARS with CLOSE in the DARTS search space and Tab. 1 shows the performances of the discovered architecture. As can be seen, CLOSE achieves a competitive test error of 2.72% in CIFAR-10. And when transferred to the ImageNet, the found architecture achieves a low test error of 24.7%.

4.3 Ablation Studies

Effect of Number of Curriculums. We conduct an ablation study on the number of curriculums in two different types of search spaces, NAS-Bench-301 (a topological search space) and NDS ResNet (a non-topological search space). Results in Fig. 8 show that, in most cases, using more curriculums can improve the ranking quality of CLOSE.

Effect of WIT and SRT. Tab. 2 demonstrates the effect of WIT and SRT adopted by CLOSE on NAS-Bench-301 and NDS ResNet. Results show that these two techniques are both necessary for the ranking quality of CLOSE.

Effect of GATE Module. We compare the ranking quality of using the GATE module with randomly assigning GLOW blocks to operations. Results in Tab. 3 reveal that our learnable GATE module plays an essential role in CLOSENet.

Method	CIFAR-10			ImageNet	
	Top-1 Error (%)	Param (M)	Search Cost (GPU days)	Top-1 Error (%)	Param (M)
NASNet-A [40]	2.65	3.3	2000	26.0	5.3
AmoebaNet-B [25]	2.55	2.8	3150	26.0	5.3
PNAS [17]	3.41	5.1	225	25.8	5.1
ENAS [22]	2.89	4.6	0.5	-	-
DARTS [18]	2.76	3.3	1.5	26.9	4.9
SNAS [32]	2.85	2.8	1.5	27.3	4.3
BayesNAS [38]	2.81	3.4	0.2	26.5	3.9
GDAS [5]	2.82	2.5	0.17	27.5	4.4
CLOSE (Ours)	2.72 ± 0.04	4.1	0.6	24.7	4.8

Table 1: Comparison of architecture performances on CIFAR-10 and ImageNet.




Fig. 8: The ranking quality of CLOSE with different numbers of curriculums. X-axis: Training epochs. Y-axis: Values of criteria.

WIT	SRT	NAS-Bench-301		NDS ResNet	
		KD	P@top5%	KD	P@top5%
		0.1104	0.1145	0.6339	0.5387
✓		0.1047	0.1122	0.6550	0.5520
	✓	0.2004	0.1610	0.6448	0.5280
✓	✓	0.5168	0.3470	0.6786	0.5667

Table 2: The ranking quality of CLOSE w./w.o. the proposed techniques.

Effect of Gradually Adding GLOW Blocks. To show the benefit of gradually adding GLOW blocks, we conduct two contrast experiments. In the first experiment, we keep a fixed number of GLOW blocks in the training process. Results in Tab. 4 demonstrate CLOSE performs better than fixing the sharing extent. In the second experiment, we gradually add blocks and stop after adding a certain number of blocks. Results in Tab. 5 show that the final ranking quality at 1000 epoch will degrade if GLOW blocks are not sufficiently added.

GATE	NAS-Bench-201		NAS-Bench-301	
	KD	P@top5%	KD	P@top5%
w/o.	0.3627	0.2014	0.2236	0.1924
w.	0.7622	0.5387	0.5168	0.3470

Table 3: The ranking quality of CLOSE w./w.o. the GATE module

Benchmark	Fixed number of blocks				CLOSE
	2	3	4	5	
NB201	0.7320	0.7247	0.7073	-	0.7622
NB301	0.4533	0.3427	0.3301	0.3106	0.5168

Table 4: The ranking quality of supernets that use a *fixed* number of blocks

Benchmark	Number of added blocks in total				
	1	2	3	4	5
NB201	0.7050	0.7072	0.7502	0.7622	-
NB301	0.3990	0.4500	0.4641	0.4879	0.5168

Table 5: The ranking quality of supernets that add *fewer* number of blocks.

5 Conclusions

This work borrows the idea of curriculum learning and proposes a novel training strategy CLOSE to train the NAS supernet both efficiently and effectively. Specifically, CLOSE adopts a curriculum learning-like schedule on the parameter sharing extent of supernets. To support this strategy, we design a novel one-shot supernet, namely CLOSENet, of which the sharing extent can be flexibly adjusted and the sharing scheme is decided based on a graph-based encoding. Extensive experiments demonstrate that equipped with CLOSENet, our proposed method CLOSE reaches a SOTA ranking quality on four NAS benchmarks. When searching in large search spaces, CLOSE can help to discover architectures with superior performances.

Acknowledgments

This work was supported by National Natural Science Foundation of China (No. U19B2019, 61832007), National Key Research and Development Program of China (No. 2019YFF0301500), Tsinghua EE Xilinx AI Research Fund, Beijing National Research Center for Information Science and Technology (BNRist), and Beijing Innovation Center for Future Chips.

References

1. Bender, G., Kindermans, P.J., Zoph, B., Vasudevan, V., Le, Q.: Understanding and simplifying one-shot architecture search. In: International Conference on Machine Learning (ICML). pp. 550–559. PMLR (2018)
2. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: International Conference on Machine Learning (ICML). pp. 41–48 (2009)
3. Benyahia, Y., Yu, K., Smires, K.B., Jaggi, M., Davison, A.C., Salzman, M., Musat, C.: Overcoming multi-model forgetting. In: International Conference on Machine Learning (ICML). pp. 594–603. PMLR (2019)
4. Brock, A., Lim, T., Ritchie, J.M., Weston, N.: Smash: one-shot model architecture search through hypernetworks. In: International Conference on Learning Representations (ICLR) (2018)
5. Dong, X., Yang, Y.: Searching for a robust neural architecture in four gpu hours. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1761–1770 (2019)
6. Dong, X., Yang, Y.: Nas-bench-201: Extending the scope of reproducible neural architecture search. In: International Conference on Learning Representations (ICLR) (2020)
7. Gong, C., Yang, J., Tao, D.: Multi-modal curriculum learning over graphs. *ACM Transactions on Intelligent Systems and Technology (TIST)* **10**(4), 1–25 (2019)
8. Guo, S., Huang, W., Zhang, H., Zhuang, C., Dong, D., Scott, M.R., Huang, D.: Curriculumnet: Weakly supervised learning from large-scale web images. In: European Conference on Computer Vision (ECCV). pp. 135–150 (2018)
9. Guo, Y., Chen, Y., Zheng, Y., Zhao, P., Chen, J., Huang, J., Tan, M.: Breaking the curse of space explosion: Towards efficient nas with curriculum search. In: International Conference on Machine Learning (ICML). pp. 3822–3831. PMLR (2020)
10. Guo, Z., Zhang, X., Mu, H., Heng, W., Liu, Z., Wei, Y., Sun, J.: Single path one-shot neural architecture search with uniform sampling. In: European Conference on Computer Vision (ECCV). pp. 544–560. Springer (2020)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
12. Hong, W., Li, G., Zhang, W., Tang, R., Wang, Y., Li, Z., Yu, Y.: Dropnas: Grouped operation dropout for differentiable architecture search. In: International Joint Conference on Artificial Intelligence (IJCAI). pp. 2326–2332 (2020)
13. Hu, Y., Liang, Y., Guo, Z., Wan, R., Zhang, X., Wei, Y., Gu, Q., Sun, J.: Angle-based search space shrinking for neural architecture search. In: European Conference on Computer Vision (ECCV). pp. 119–134. Springer (2020)
14. Jiang, L., Meng, D., Mitamura, T., Hauptmann, A.G.: Easy samples first: Self-paced reranking for zero-example multimedia search. In: ACM International Multimedia Conference (MM). pp. 547–556 (2014)
15. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. In: International Conference on Learning Representations (ICLR). OpenReview.net (2018)
16. Liang, H., Zhang, S., Sun, J., He, X., Huang, W., Zhuang, K., Li, Z.: Darts+: Improved differentiable architecture search with early stopping. arXiv preprint arXiv:1909.06035 (2019)

17. Liu, C., Zoph, B., Neumann, M., Shlens, J., Hua, W., Li, L.J., Fei-Fei, L., Yuille, A., Huang, J., Murphy, K.: Progressive neural architecture search. In: European Conference on Computer Vision (ECCV). pp. 19–34 (2018)
18. Liu, H., Simonyan, K., Yang, Y.: Darts: Differentiable architecture search. In: International Conference on Learning Representations (ICLR) (2019)
19. Luo, R., Qin, T., Chen, E.: Understanding and improving one-shot neural architecture optimization. CoRR **abs/1909.10815** (2019)
20. Ning, X., Tang, C., Li, W., Zhou, Z., Liang, S., Yang, H., Wang, Y.: Evaluating efficient performance estimators of neural architectures. In: Annual Conference on Neural Information Processing Systems (NIPS) (2021)
21. Ning, X., Zheng, Y., Zhao, T., Wang, Y., Yang, H.: A generic graph-based neural architecture encoding scheme for predictor-based nas. In: European Conference on Computer Vision (ECCV). pp. 189–204. Springer (2020)
22. Pham, H., Guan, M., Zoph, B., Le, Q., Dean, J.: Efficient neural architecture search via parameters sharing. In: International Conference on Machine Learning (ICML). pp. 4095–4104. PMLR (2018)
23. Platanios, E.A., Stretcu, O., Neubig, G., Póczos, B., Mitchell, T.: Competence-based curriculum learning for neural machine translation. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). pp. 1162–1172 (2019)
24. Radosavovic, I., Kosaraju, R.P., Girshick, R., He, K., Dollár, P.: Designing network design spaces. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10428–10436 (2020)
25. Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. In: AAAI Conference on Artificial Intelligence. vol. 33, pp. 4780–4789 (2019)
26. Ren, Z., Dong, D., Li, H., Chen, C.: Self-paced prioritized curriculum learning with coverage penalty in deep reinforcement learning. *IEEE transactions on neural networks and learning systems* **29**(6), 2216–2226 (2018)
27. Siems, J., Zimmer, L., Zela, A., Lukasik, J., Keuper, M., Hutter, F.: Nas-bench-301 and the case for surrogate benchmarks for neural architecture search. arXiv preprint arXiv:2008.09777 (2020)
28. Soviany, P., Ionescu, R.T., Rota, P., Sebe, N.: Curriculum learning: A survey. *International Journal of Computer Vision (IJCV)* pp. 1–40 (2022)
29. Su, X., You, S., Zheng, M., Wang, F., Qian, C., Zhang, C., Xu, C.: K-shot nas: Learnable weight-sharing for nas with k-shot supernets. In: International Conference on Machine Learning (ICML). pp. 9880–9890. PMLR (2021)
30. Tay, Y., Wang, S., Luu, A.T., Fu, J., Phan, M.C., Yuan, X., Rao, J., Hui, S.C., Zhang, A.: Simple and effective curriculum pointer-generator networks for reading comprehension over long narratives. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 4922–4931 (2019)
31. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1492–1500 (2017)
32. Xie, S., Zheng, H., Liu, C., Lin, L.: Snas: stochastic neural architecture search. In: International Conference on Learning Representations (ICLR) (2019)
33. Yang, Z., Wang, Y., Chen, X., Shi, B., Xu, C., Xu, C., Tian, Q., Xu, C.: Cars: Continuous evolution for efficient neural architecture search. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1829–1838 (2020)

34. Yu, K., Sciuto, C., Jaggi, M., Musat, C., Salzmann, M.: Evaluating the search phase of neural architecture search. In: International Conference on Learning Representations (ICLR) (2020)
35. Zela, A., Siems, J., Hutter, F.: Nas-bench-1shot1: Benchmarking and dissecting one-shot neural architecture search. In: International Conference on Learning Representations (ICLR) (2019)
36. Zhang, M., Li, H., Pan, S., Chang, X., Su, S.: Overcoming multi-model forgetting in one-shot nas with diversity maximization. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7809–7818 (2020)
37. Zhao, Y., Wang, L., Tian, Y., Fonseca, R., Guo, T.: Few-shot neural architecture search. In: International Conference on Machine Learning (ICML). pp. 12707–12718. PMLR (2021)
38. Zhou, H., Yang, M., Wang, J., Pan, W.: Bayesnas: A bayesian approach for neural architecture search. In: International Conference on Machine Learning (ICML). pp. 7603–7613. PMLR (2019)
39. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. In: International Conference on Learning Representations (ICLR) (2017)
40. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8697–8710 (2018)