

Homogeneous Architecture Augmentation for Neural Predictor

Yuqiao Liu^{*1}, Yehui Tang^{*2}, Yanan Sun^{†1}

¹ College of Computer Science, Sichuan University.

² Key Lab of Machine Perception (MOE), Dept. of Machine Intelligence, Peking University.

lyqguitar@gmail.com; yhtang@pku.edu.cn; ysun@scu.edu.cn

Abstract

Neural Architecture Search (NAS) can automatically design well-performed architectures of Deep Neural Networks (DNNs) for the tasks at hand. However, one bottleneck of NAS is the prohibitively computational cost largely due to the expensive performance evaluation. The neural predictors can directly estimate the performance without any training of the DNNs to be evaluated, thus have drawn increasing attention from researchers. Despite their popularity, they also suffer a severe limitation: the shortage of annotated DNN architectures for effectively training the neural predictors. In this paper, we proposed Homogeneous Architecture Augmentation for Neural Predictor (HAAP) of DNN architectures to address the issue aforementioned. Specifically, a homogeneous architecture augmentation algorithm is proposed in HAAP to generate sufficient training data taking the use of homogeneous representation. Furthermore, the one-hot encoding strategy is introduced into HAAP to make the representation of DNN architectures more effective. The experiments have been conducted on both NAS-Benchmark-101 and NAS-Bench-201 dataset. The experimental results demonstrate that the proposed HAAP algorithm outperforms the state of the arts compared, yet with much less training data. In addition, the ablation studies on both benchmark datasets have also shown the universality of the homogeneous architecture augmentation. Our code has been made available at <https://github.com/lyq998/HAAP>.

1. Introduction

Deep Neural Networks (DNNs) have been successfully applied to various challenging real-world problems, including image classification [16], natural language processing [8], speech recognition [40], to name a few. Well-designed architectures of DNNs are generally viewed as

the deciding factor that the DNNs can achieve promising performance. Traditionally, designing architecture is nevertheless time-consuming, and is often exerted by experts with rich knowledge in both the DNNs and the task domain. Neural Architecture Search (NAS) is an automatic way to design the architectures of DNNs without such kind of expertise.

Generally, the NAS algorithm is composed of three different parts: search space, search strategy and performance estimation [10]. Particularly, given a search space predefined, the search strategy should find a well-performed architecture when the search terminates. In literature, there are mainly three techniques used as the search strategies: reinforcement learning [32], gradient-based algorithms [21] and Evolutionary Computation (EC) [1]. However, no matter which technique is used, all NAS algorithms need to estimate the performance of the DNNs searched in order for the effective proceeding of search. In most NAS algorithms, the search strategy evaluates the solutions individually, which inevitably gives rise to heavy computational overheads. For instance, on the commonly used CIFAR-10 benchmark dataset [19], the Large-Scale evolution NAS algorithm [24] used 250 GPU computers for 11 days. Not only that, the NAS algorithm [41] used 800 GPUs for nearly one month. Unfortunately, this is unaffordable for most academic researchers. To address this issue, various algorithms have been designed to speed up the estimation process without numerous computation resource, which can be classified into four different categories: weight inheritance [24], early stopping policy [29, 14], reduced training set [25] and neural predictor [33].

The neural predictor requires a number of well-trained architectures, and then a regression model is trained to map the architectures and the corresponding performance values. When a new DNN is generated, its performance is directly predicted by the regression model. Compared to other speed-up algorithms mentioned above, the neural predictors can provide satisfactory prediction result, thus is popular among the community.

Many previous works focused on designing better re-

*Equal contribution.

†Corresponding author.

gression models to improve the performance of predictors [28, 34]. They try to achieve a better fit from the perspective of the regression model. However, the big issue in neural predictor is actually in the training dataset. In practice, it is prohibitively unaffordable to obtain a large training dataset. Even with the ultra high performance hardware, for example, it will take about 32 minutes to train a neural network on the TPU v2 accelerator [38], which means only nearly 45 annotated training data can be obtained one day. Generally, neural predictor is trained with a small dataset of annotated architectures, e.g., Wen *et al.* [34] only used 119 annotated architectures to build the neural predictor in practice. The main reason for the poor performance of the neural predictor is not the regression model, but the limited training data. As a result, how to make full use of the existing limited data without increasing the computational cost is an important issue in neural predictor.

In this paper, we propose a novel data augmentation strategy in the space of neural architectures by exploring their *Homogeneous* properties. The *Homogeneous* augmentations eliminate the influence of layer order, making the neural predictor pay more attention to the overall layer type. Specifically, the architecture augmentation works by swapping the inner orders to generate a group of homogeneous representations. To effectively represent the intrinsic properties of architectures, one-hot encoding strategy is also developed. The flowchart of the proposed Homogeneous Architecture Augmentation for Neural Predictor (HAAP) is shown in Fig. 1. The experiments on NAS-Bench-101 [38] and NAS-Bench-201 [9] demonstrate the effectiveness of the proposed architecture augmentation method, and show the superiority of HAAP compared with the state of the arts. Not only that, the *Homogeneous* augmentations can also be combined with the state-of-the-art neural predictors to improve their predictions.

The remainder of this paper is organized as follows. The related works are provided in Sec. 2. Sec. 3 gives the implementation details of the proposed approach. Experiments and extensive experimental results demonstrate the effectiveness of HAAP in Sec. 4. Finally, Sec. 5 is for the conclusion and future works.

2. Related Works

In this section, we will first introduce the encoding strategy of NAS algorithms, which is the base of the proposed algorithm. Immediately after, the state-of-the-art algorithms are reviewed and their limitations are summarized to justify the necessity of the proposed algorithm.

2.1. Encoding Strategy of NAS Algorithms

The encoding space (*i.e.*, search space) in NAS can be divided into four different categories based on the basic unit

sought: the layer-based, the block-based, the cell-based and the topology-based.

The basic units in the layer-based encoding space are the primitive layers of DNNs. The layer-based encoding space is often large and building architecture in this space acquires more details and information, resulting in the search strategy putting many efforts. EvoCNN [31] is a typical example of layer-based space. In contrast, the basic units in the block-based encoding space are the blocks which are the combinations of the primitive layers, such as ResNet block and DenseNet block used by AE-CNN [30]. With the help of the integrated blocks, the architectures searched from the block-based encoding space are more likely to be well-designed with fewer parameters. One special case in the block-based encoding space is that all the blocks are all the same, which generates the cell-based encoding space. The architectures found in the cell-based encoding space are made of stacked cells. Since all the cells are the same, the encoding information of one cell can represent a corresponding architecture by stacking the same cell multiple times. The well known AmoebaNet-A [23] is an example of cell-based space. The topology-based encoding space mainly concentrates on the connections of the units rather than their internal structures.

For cell-based encoding space, there are multiple ways to encode the connections of layers in cells. The adjacency matrix is a common way to represent the connections, where “1” means the two layers are connected while “0” means the opposite. This representation method has been widely adopted as the base encoding strategy by existing NAS algorithms. For example, Genetic CNN [36] used a variant of the adjacency matrix where the unimportant elements in the matrix are pruned and the remaining elements are flattened into a binary vector.

2.2. Neural Predictor

There are five representative state-of-the-art neural predictors: Peephole [7], E2EPP [28], Semi-Supervised Assessor of Neural Architectures (SSANA) [33], Neural Predictor for Neural Architecture Search (NPNAS) [34] and ReNAS [37]. All of them have been investigated on image classification tasks owing to various well-designed architectures and available benchmark datasets for classifying images. Recently, we also notice that some public architecture datasets have been made available [38, 9]. This will make the prediction more efficient and the comparison fairer with providing more training data.

As mentioned in Sec. 1, most of the neural predictors mainly aim at the design of the regression models. For example, Peephole and NPNAS used different efficient regression models for better prediction, while ReNAS focused on the redesign of training loss of the model. In this paper, we demonstrate that without a complicated regression model,

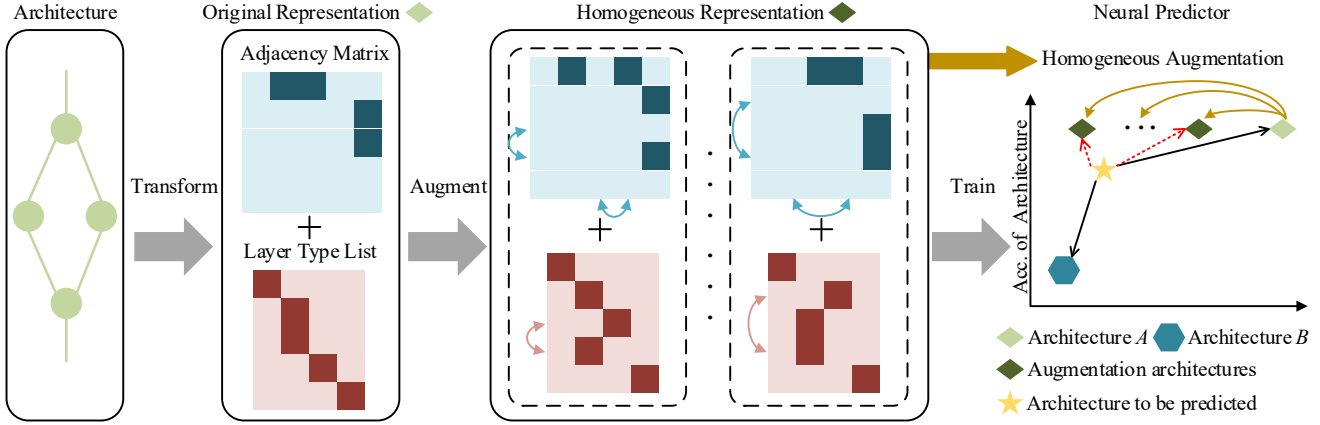


Figure 1. The flowchart of the proposed HAAP. Architectures in different encoding spaces are first transformed into the representation of adjacency matrix and layer type list where one-hot encoding strategy is used. Please note that, we only display one neural architecture in the figure for convenience. Then the original representation is augmented into its homogeneous forms. The original data and the augmented data are all used to train the neural predictor. If the architecture augmentation is not used, the architecture to be predicted will refer more to the accuracy of architecture B . When the architecture augmentation is used, the architecture to be predicted will be more inclined to refer to the accuracy of the augmentation architectures, and then get a more accurate prediction.

the classical regression models (e.g. random forest) can also make an accurate prediction with the utilization of the proposed architecture augmentation method.

3. Approach

Building the neural predictor can be mathematically represented as a regression task. Supposing the predictor is represented by a regression model R . After R has been trained with the training data $[\mathbf{X}, \mathbf{y}]$, where $\mathbf{X} = \{X_1, X_2, \dots, X_N\}$ denotes the encoding of architecture information and $\mathbf{y} = \{y_1, y_2, \dots, y_N\}$ denotes the corresponding performance value of $\{X_1, X_2, \dots, X_N\}$, the prediction $\hat{\mathbf{y}}$ of test data \mathbf{X}^{test} can be predicted by R . Particularly, the training process of R can be described by Equation 1:

$$\min_{T_p} \frac{1}{N} \sum_{n=1}^N \mathcal{L}(R(T_p, X_n), y_n) \quad (1)$$

where T_p is the trainable parameters of the regression model R , and $\mathcal{L}(\cdot)$ denotes the loss function of R . We choose the ℓ_2 norm as the loss function, which is widely used in neural predictors [34, 33].

3.1. Architecture Encoding

In principle, most of the neural architectures can be represented by a standard Directed Acyclic Graph (DAG), where the vertices represent the layers and the edges represent the connections. Therefore, each X_n in \mathbf{X} can be represented by the layer type list $x_n^t = \{x_{n1}^t, x_{n2}^t, \dots, x_{nN_1}^t\}$ and the corresponding adjacency matrix x_n^m , where N_1 denotes the number of layers. Although some architectures in other encoding spaces are not represented by standard DAG,

where the vertices represent the connections while the edges represent the layers, we can use a transformation to change the form into the standard type. For example, to transform the architectures in NAS-Bench-201, we can firstly give a fixed order for edges, and then reduce the useless layers. The detailed transformation method is provided in supplementary materials.

The traditional hard encoding [37] uses an integer vector to encode layer type x_n^t directly. Then, the type vector is broadcasted into the adjacency matrix x_n^m which can be expressed in Equation 2:

$$X_n = x_n^m \times stretch(integer(x_n^t)) \quad (2)$$

where the $integer(\cdot)$ transforms verbal type list to integer vector, and $stretch(\cdot)$ can stretch a vector into a square matrix. As a result, the encoding X_n is a sparse matrix which contains much redundant dimensions. In addition, the relationships and the Euclidean distances between layers are misrepresented by integers. To this end, we come up with a binary encoding strategy which takes the use of the one-hot encoding for type list.

The encoding strategy used in this paper is easy to implement. Equation 3 shows how to convert the original x_n^m and x_n^t into a binary vector.

$$X_n = concat\{flattened(x_n^m), one-hot(x_n^t)\} \quad (3)$$

Specifically, the two-dimensional adjacency matrix x_n^m is flattened (with $flattened(\cdot)$) into an one-dimensional vector by row. Meanwhile, the layer type list is transformed into the one-hot encoding (with $one-hot(\cdot)$). Finally, these two one-dimensional vectors are concatenated into X_n

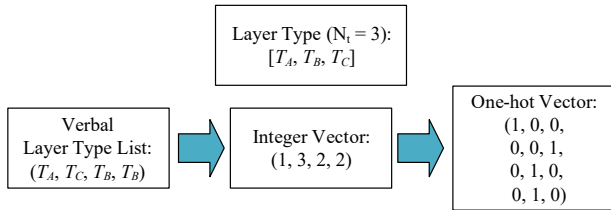


Figure 2. One-hot encoding process for the layer type list.

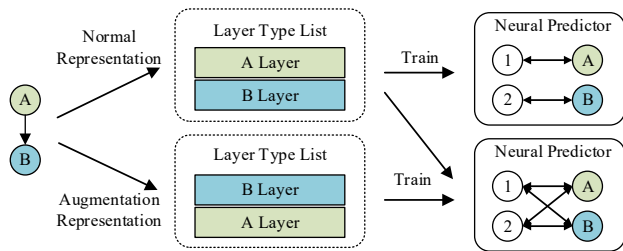


Figure 3. An example to illustrate how the homogeneous augmentation can better train the neural predictor.

(with $\text{concat}(\cdot)$). Fig. 2 provides an example to explain the one-hot encoding strategy operating on the layer type list, where the number of types N_t is 3, and T_A , T_B and T_C denote the three types of layers in the list. To achieve the encoding, the first is to transfer the verbal expression into an integer vector. The second is to employ the one-hot encoding to replace the integers accordingly. The one-hot encoding uses multi dimensional space to represent the layer type to ensure the Euclidean distances between different layer types are the same, which is positive to the regression models.

3.2. Homogeneous Architecture Augmentation

Generally, a common DNN architecture begins with an input layer and ends with an output layer [38, 30]. In this paper, we use *In* and *Out* layers to denote both two layers separately. Please note that the proposed architecture augmentation method is suitable for the case that the neural network has sole input and sole output. This is also the most common case of DNNs. For example, the image classification task only requires one image as input and the neural network outputs the predicted category.

As introduced above, the adjacency matrix and the layer type list make up the neural architecture, where the adjacency matrix describes the topological relationship of the layer type list. Therefore, there should be no order in the layer type list. However, the layers close to the *In* layer are always placed in front of the layer type list, and the layers close to the *Out* layer can only be placed at the back. In addition, null layers are generally placed at the end [34, 37]. In this way, the layers that should have no order and position are sorted in a breadth-first traversal manner. This sort-

ing can mislead the neural predictor and make the predictor pay attention to the relationship between layer type and the specified location which is meaningless. What we have to do is to eliminate the influence of layer position, making the neural predictor pay more attention to the overall layer type and some invariants in the augmentation of the neural architecture, such as the total number of different types of layers, instead of focusing on the type of layer in a specific location. Fig. 3 is an example. Shown on the left is a part of a neural network, including two layers A and B. If the neural predictor is only trained with the normal representation of the architecture, the predictor will bind position 1 to A, and position 2 to B. However, if the predictor is trained with normal and augmented representation together, the influence of the layer order on the predictor will be weakened.

By observing this, the original representation can be augmented with the “unordered” character of the layer type list. However, not all the layers in the layer type list are unordered. The *In* and *Out* layers belong to placeholder set and they are fixed at the first and last positions of the layer type list in all architectures. Thus, it is useless to change their positions. The other layers in the middle of the layer type list belong to operation set which principally determines the performance of an architecture. Actually, only these layers are unordered in layer type list. Assuming that the layer type list is $x^t = \{x_1^t, x_2^t, \dots, x_{N_t}^t\}$, x_1^t is the *In* layer and $x_{N_t}^t$ is the *Out* layer. Since for $\forall x_j^t, j \in \{2, 3, \dots, N_t - 1\}$ could be in any position, therefore there are $(N_t - 2)!$ possible permutations. We use these different permutations to augment the training dataset by a factor of $(N_t - 2)!$.

However, if the layer type list is changed, the adjacency matrix has to make changes accordingly to make sure that the actual architectures of the representation are the same. This is because the order of rows and columns of the adjacency matrix is related to the order in the layer type list. Specifically, the i -th row of the adjacency matrix denotes the output of the i -th layer in the layer type list, and the j -th column denotes the input of the j -th layer. The “Homogeneous Architectures” part of Fig. 1 shows an example to illustrate the architecture augmentation for neural predictors. In this example, we assume the number of layers N_t equals to 5 and the neural architecture contains a null layer. The original representation of the architecture is on the top, and we term these different representation of architecture in dashed box as homogeneous representation, because they represent the same architecture in essence. Specifically, comparing the first homogeneous representation displayed in Fig. 1 with the original representation, the order of 3-rd layer and 4-th layer is swapped from the original one. To maintain the representation of architecture, the corresponding row and column in the adjacency matrix need to swap according to the swapping of the layer type list. Whether

to swap rows or columns first, the matrix ends up with the same effect.

The general process of constructing more training data with the proposed architecture augmentation method is shown in Fig. 1. Specifically, the input can be architectures in different encoding spaces, such as the binary vector used in Genetic CNN [36] or the operations represented by edges in NAS-Bench-201 [9]. The first step is to transform the different encoding strategies into the standard DAG form like NAS-Bench-101 [38]. Because of the character of the regression model, the encoded architecture information must be fixed length. We refer the padding approach in [37] to achieve this goal, by adding zeros into the adjacency matrix and the layer type list. The position of starting to pad zeros is the penultimate element in the layer type list, while the starting positions in the adjacency matrix are the penultimate row and column. We define the zero in the layer type list as null layer type for convenience. The next step is the homogeneous architecture augmentation. Including the original representation, the augmentation ends up with N_c architecture augmentations where $N_c = (N_l - 2)!$ except for *In* and *Out* layers. Next, the unnecessary dimensions in the adjacency matrix and the layer type list are reduced. To be more specific, the first element (*In* layer) and the last element (*Out* layer) in the layer type list, and the first column and the last row in the adjacency matrix will be reduced. Finally, the layer type list is converted to a binary vector and concatenate with flattened adjacency matrix.

The training dataset can be expanded by a factor of N_c in the end. In principle, the generated data are reliable because of the homogeneous representation. It makes sense that all the homogeneous representation of X_n correspond to the same performance value y_n . Similar to the image augmentation (e.g., flip, rotation, etc.), the proposed augmentation can generate a more comprehensive dataset with more possible training points, which reduces the distribution gap between training data and unknown data [27]. As a result, the regression model R , i.e. the neural predictor, can be trained better with more training data augmented.

4. Experiments

We choose the NAS-Bench-101 [38] and the NAS-Bench-201 [9] as the experimental datasets. The NAS-Bench-101 dataset is built for the NAS research to alleviate the need for intensive computing resource. NAS-Bench-101 collects the accuracy of architectures tested on CIFAR-10 [19] with Convolutional Neural Networks (CNNs). The NAS-Bench-201 is another benchmark dataset with different search space. In a nutshell, both benchmark datasets choose the cell-based search space, and the whole architecture is built up by stacking different cells which are shown in Fig. 4. The cells in NAS-Bench-101 use vertices to represent operation layers, and edges to represent the connec-

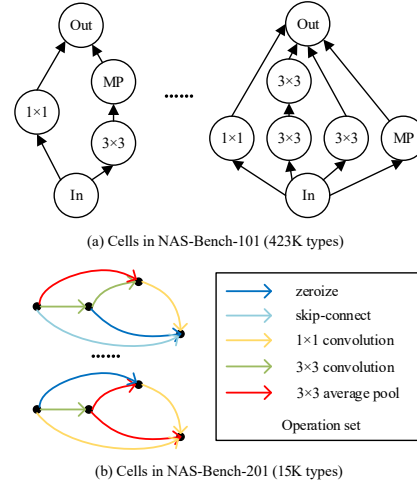


Figure 4. (a) Cells in NAS-Bench-101 where 1×1 denotes 1×1 convolution, 3×3 denotes 3×3 convolution and MP denotes 3×3 max pooling. (b) Cells in NAS-Bench-201.

tions. Whereas in NAS-Bench-201, the operation layers are presented by edges, which could not use the proposed architecture augmentation directly and need additional transformations to get adjacency matrix and layer type list.

In this section, we will first perform the comparison with the state-of-the-art neural predictors to demonstrate the effectiveness of the proposed HAAP algorithm. Second, HAAP is embedded in an Evolutionary Algorithm (EA) [39] based NAS search strategy to search for promising architectures on both two benchmark datasets. Third, we will also perform extensive ablation experiments on different classical regression models and on both datasets to show the effect of proposed architecture augmentation method and the utilization of the one-hot encoding strategy.

4.1. Comparison of Prediction Performance on NAS-Bench-101

The proposed HAAP algorithm is compared with the state-of-the-art algorithms which we have introduced above, i.e., Peephole [7], E2EPP [28], SSANA [33], ReNAS [37] and NPNAS [34]. NAS-Bench-101 is used to providing training and testing architectures. In addition, two indicators are used to measure the performance of the prediction: Kendall’s Tau (KTau) [26] and MSE. Particularly, the KTau reflects the rank correlation of two ordinal variables, and its value varies in $[-1, 1]$. The closer this value gets to 1, the more similar the rankings of the two ordinal variables are. The MSE reflects the error between the predicted accuracy values $\hat{\mathbf{y}}$ and the ground-truth labels \mathbf{y} . Please note that both are widely used by the neural predictors of NAS algorithms [28, 33, 37].

We report the comparison results in Table 1 where the proposed HAAP algorithm is realized by Random Forest

Table 1. Comparison with the state-of-the-art neural predictors. We report the results of HAAP on two different N_{ot} for fair comparison.

Algorithms	N_{ot}	KTau	MSE	Regression Model
Peephole [7]	1K	0.4373 \pm 0.0112	0.0071 \pm 0.0005	LSTM
E2EPP [28]	1K	0.5705 \pm 0.0082	0.0042 \pm 0.0003	RF
SSANA [33]	1K	0.6541 \pm 0.0078	0.0031 \pm 0.0003	AE + GCN
ReNAS [37]	424	0.6574	—	LeNet-5
NPNAS [34]	424	0.6945	—	GCN
NPNAS + HA	424	0.7062	—	GCN
HAAP	424	0.7010 \pm 0.0022	0.0024 \pm 0.0003	RF
	1K	0.7126 \pm 0.0024	0.0023 \pm 0.0003	

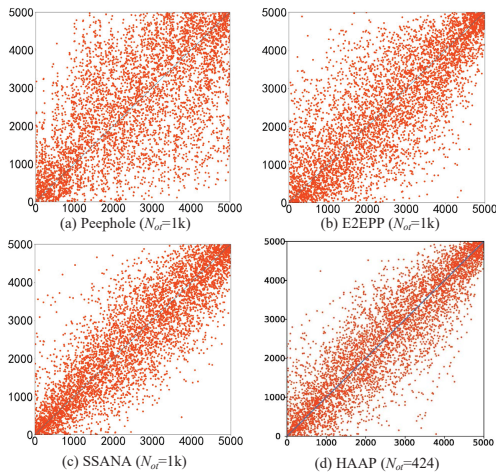


Figure 5. Comparison with state-of-the-art algorithms. The x-axis denotes the true ranking while the y-axis denotes the predicted ranking. Figures (a), (b) and (c) are from [33].

(RF) [4] with fine-tuning. In addition, the symbol “—” in Table 1 implies there is no result publicly reported in the corresponding literature, and N_{ot} denotes the number of the original training data. The best values are shown in bold for the convenience of observation. For a fair comparison, we report the results of HAAP on two different N_{ot} . As can be seen from Table 1, when N_{ot} is set as 424, HAAP can achieve the largest KTau comparing with ReNAS and NPNAS. In addition, if NPNAS is combined with the proposed Homogeneous Augmentation (HA), the prediction performance can be further improved. With N_{ot} increasing to 1K, HAAP also achieves the best values of KTau and MSE comparing with Peephole, E2EPP and SSANA.

Fig. 5 provides the qualitative comparison results, where the x-axis denotes the true ranking while the y-axis denotes the predicted ranking of 5K randomly sampled architectures. If these points are closer to the diagonal, the ranking correlation is stronger, which is equivalent to the values indicated by KTau. Please note that figures (a), (b), (c) are trained with 1K original data whereas N_{ot} in our algorithm is only 424. As can be seen from these figures, the points shown in figure (d) demonstrate the best result,

Table 2. Classification accuracy rate and ranking of the searched architectures of NAS-Bench-101. 1K annotated architectures randomly selected from NAS-Bench-101 are used to train predictors.

Algorithm	Accuracy (%)	Ranking (%)
Peephole [7]	93.41 \pm 0.34	1.64
E2EPP [28]	93.77 \pm 0.13	0.15
SSANA [33]	94.01 \pm 0.12	0.01
HAAP	94.09 \pm 0.11	0.004
Oracle	94.23 \pm 0.0	0.0005

which is achieved by the proposed algorithm. Furthermore, compared with the Neural Networks used by the compared Peephole [7], SSANA [33] and ReNAS [37] algorithms (such as Long Short Term Memory (LSTM) [17], Auto-Encoder (AE), Graph Convolutional Network (GCN) and LeNet-5 [20]), the RF employed by the proposed HAAP algorithm requires few additional tuning of parameters.

4.2. Searching Architectures on NAS-Bench-101 and NAS-Bench-201

We have referred to the NAS experiment in [33], and also take an EA based search strategy to search architectures on NAS-Bench-101. In order to demonstrate the effectiveness of the proposed HA, we do not use the state-of-the-art neural predictor when searching. Specifically, the experiment is repeated for 20 times, and we compare the accuracy and ranking of the top-10 architectures selected by different algorithms which can be seen in Table 2. The second column displays the top-1 classification accuracy on CIFAR-10, and the third column displays the ranking of whole architectures in NAS-Bench-101. The last row of Table 2 shows the oracle where the performance of all architecture is known. The proposed HAAP obtains the highest accuracy and ranking among the neural predictor algorithms. Specifically, the architectures searched by HAAP are significantly better than that searched by Peephole and E2EPP. Furthermore, HAAP is slightly better than SSANA and is very close to the oracle.

We visualize the best architecture searched by different algorithms in Fig. 6 to give an intuitive comparison. The

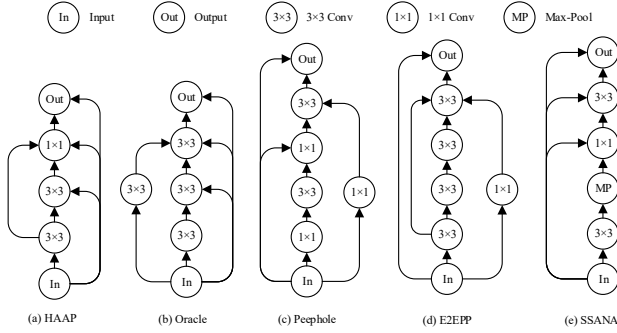


Figure 6. The architectures searched by different neural predictors on NAS-Bench-101. 1K annotated architectures are used to train the neural predictors.

Table 3. Validation and test accuracy rate on CIFAR-10 of the searched architectures on NAS-Bench-201. We report the mean and std of 10 runs for E2EPP and HAAP.

Algorithm	Validation	Test	Time
REA [23]	91.19±0.31	93.92±0.30	12000
RS [2]	90.93±0.36	93.70±0.36	12000
REINFORCE [35]	91.09±0.37	93.85±0.37	12000
BOHB [11]	90.82±0.53	93.61±0.52	12000
E2EPP [28]	90.61±0.89	93.39±0.75	6000
HAAP	91.18±0.25	94.00±0.25	6000
optimal	91.61	94.37	N/A

In node is directly connected to the *Out* node in all architectures displayed. This connection can facilitate gradient propagation [16, 18]. There are five operation layers in architectures searched by Peephole and E2EPP, and four in SSANA and oracle. Whereas there are only three operation layers including a 1×1 convolution in HAAP, which can reduce a lot of computing costs.

In order to compare with other state-of-the-art NAS algorithms, we also conduct an experiment searching architectures on NAS-Bench-201. In Table 3, four algorithms with the top accuracy reported in NAS-Bench-201 and one neural predictor are chosen for comparison. Following the tradition in NAS community, the validation accuracy was used to guide the search during search stage, and the test accuracy was only queried for the searched architecture. To demonstrate the efficiency of HAAP, we stopped HAAP once the time budget of predictor reached the half of the state-of-the-art NAS algorithms (i.e. 6000 seconds). With the half time budget, E2EPP can not exceed any of the competitors. However, HAAP outperforms all the state-of-the-art NAS algorithms, yet with much less search cost.

4.3. Ablation Study

In this subsection, the classical regression models on different cases are compared to show the effectiveness bene-

fitting from the architecture augmentation and the one-hot encoding strategy. They are Decision Tree (DT) [5], Linear Regression (LR), Support Vector Regression (SVR) [6], K-Nearest Neighbors (KNN), RF [4], AdaBoost [12], Gradient Boosted Regression Tree (GBRT) [13], Bagging [3] and ExtraTree [15]. Please note that all the models are implemented by Scikit-learn [22] with the default settings that often give rise to satisfactory performance of the corresponding model in reality.

The results are shown in Table 4, where all experiments use the same training data and test data. Note that Table 4 is split into two parts where the top part uses 424 training data to train the regression model, while the bottom part uses architecture augmented (actually $424 \times (7 - 2)! = 50880$ training data) and the test data are all 5K. We only report the result of a single run in the table. Four cases are designed to show the efficiency of HAAP:

case 1 (baseline): Using only the concatenation of the flattened adjacency matrix and integer vector.

case 2 (one-hot): One-hot encoding for layer type list x^t is used on the basis of case 1.

case 3 (architecture augmentation): Using augmented data while other settings are the same as case 1.

case 4 (architecture augmentation + one-hot): One-hot encoding for x^t is used on the basis of case 3.

The comparison of the four cases is shown in Table 4. As can be seen from the comparison results, almost all the regression models can benefit from these two designed components. When the architecture augmentation is applied, a great promotion is obtained especially using the architecture augmentation and one-hot encoding strategy collectively. Furthermore, RF achieves the best performance in terms of both K τ and MSE indicators which are marked with \star in Table 4. In summary, this group of experiments demonstrates the effectiveness of the proposed architecture augmentation method and utilization of the one-hot encoding strategy.

In addition, this ablation study is also experienced on NAS-Bench-201. As introduced above, we firstly transform the architectures to the standard DAG form. Table 5 shows the K τ and MSE under four cases and different N_{ot} . The values of K τ and MSE in case 1 are much better than that of NAS-Bench-101 as shown in Table 4. This demonstrates that the proposed transformation method is effective, and the architectures in NAS-Bench-201 are more recognizable. The cases with architecture augmentation still get the highest K τ and the lowest MSE which is in bold as shown in Table 5. In addition, under the same setting of N_{ot} , the results are better than that of NAS-Bench-101 which can be seen in Table 1. One important reason is that N_l is 8 in this dataset, and the training data will be expanded by 720 times rather than 120 times in NAS-Bench-101. However, there is no obvious improvement using the one-hot encod-

Table 4. Comparisons on classical regression models on NAS-Bench-101. 424 annotated architectures are used to train the regression models. A and O are short for ‘‘Architecture augmentation’’ and ‘‘One-hot encoding’’ separately. The italic and bold denote the best values (highest K_{Tau} and lowest MSE) of each regression model.

A	O	Regression Models (Neural Predictor)									
		DT	LR	SVR	KNN	RF	AdaBoost	GBRT	Bagging	ExtraTree	
✗	✗	K _{Tau}	0.3380	0.3407	0.2738	0.3041	0.5422	0.3624	0.5559	0.4737	0.3410
		MSE	0.0048	0.0028	0.0051	0.0028	0.0030	0.0036	0.0031	0.0038	0.0039
✗	✓	K _{Tau}	0.4386	0.5074	0.3257	0.4426	0.5909	0.4185	0.5779	0.5405	0.4442
		MSE	0.0048	7.7862	0.0050	0.0027	0.0031	0.0032	0.0030	0.0040	0.0049
✓	✗	K _{Tau}	0.5271	0.3592	0.3256	0.3932	0.6888	0.3158	0.5746	0.6620	0.5154
		MSE	0.0030	0.0029	0.0046	0.0029	0.0023	0.0028	0.0025	0.0024	0.0033
✓	✓	K _{Tau}	0.5424	0.5456	0.3548	0.5068	0.6991 *	0.3309	0.6007	0.6628	0.5189
		MSE	0.0032	0.0026	0.0045	0.0029	0.0022 *	0.0026	0.0025	0.0024	0.0031

Table 5. Results on NAS-Bench-201. The regression model is RF. The bold denotes the best accuracy of each N_{ot} on CIFAR-10.

A	O	N_{ot}		
		424	1K	
✗	✗	K _{Tau}	0.6919 \pm 0.0040	0.7754 \pm 0.0051
		MSE	0.0009 \pm 0.00006	0.0004 \pm 0.00002
✗	✓	K _{Tau}	0.7156 \pm 0.0060	0.7647 \pm 0.0050
		MSE	0.0008 \pm 0.00004	0.0004 \pm 0.00003
✓	✗	K _{Tau}	0.7824 \pm 0.0014	0.8172 \pm 0.0051
		MSE	0.0005 \pm 0.00001	0.0004 \pm 0.00005
✓	✓	K _{Tau}	0.7899 \pm 0.0011	0.8166 \pm 0.0025
		MSE	0.0004 \pm 0.00005	0.0003 \pm 0.00001

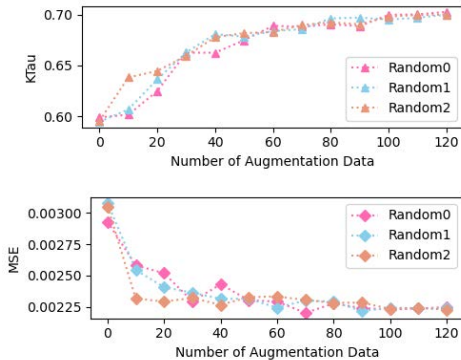


Figure 7. Neural predictor performance results of HAAP w.r.t. different number of architecture augmentation.

ing strategy, and this may be caused by the increasing size of adjacency matrix.

Last but not the least, in order to figure out how many augmentation architectures are needed at least training the predictor to achieve the best performance, we conduct an experiment on NAS-Bench-101, and N_{ot} is set to 424. Fig. 7 shows three separate randomized trials, where architectures are randomly sampled from all the architectures

augmented. The x-axes in Fig. 7 is the number of augmentation data from one original training data, and 0 in x-axis denotes no augmentation is used. As can be observed from Fig. 7, the trials can obtain the highest K_{Tau} and the lowest MSE when all the architectures augmented are used to train the predictor. Because of this, all the architectures augmented are used to train the predictor in all the experiments except this one.

5. Conclusion

The goal of this paper is to develop an efficient neural predictor. The goal has been achieved by two core components. Specifically, we have proposed a homogeneous architecture augmentation method, which could construct more training data from the limited original data. In addition, the one-hot encoding strategy is utilized to transform the DNN architectures, enhance the prediction accuracy made by the neural predictors. The experiments have been conducted on NAS-Bench-101 and NAS-Bench-201, and the significant improvements can be observed in both datasets when architecture augmentation is used. The results compared with five state-of-the-art peer neural predictors show that the proposed HAAP algorithm outperforms all the competitors in terms of both K_{Tau} and MSE indicators. Furthermore, well-performed architectures can be found on NAS-Bench-101 and NAS-Bench-201 with 94.09% accuracy (top 0.004% of the entire search space) and 94.00% accuracy separately using HAAP. Our future work will focus on encoding additional features of the architecture to improve the performance with the help of more comprehensive description of the architectures.

Acknowledgments. This work was supported by National Natural Science Foundation of China under Grant 61803277 and Science and Technology Innovation Talent Project of Sichuan Province under Grant 2021JDRC0001.

References

- [1] Thomas Bäck, David B Fogel, and Zbigniew Michalewicz. Handbook of evolutionary computation. *Release*, 97(1):B1, 1997.
- [2] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *The Journal of Machine Learning Research*, 13(1):281–305, 2012.
- [3] Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- [4] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [5] Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. *Classification and regression trees*. CRC press, 1984.
- [6] Chih-Chung Chang and Chih-Jen Lin. Libsvm: A library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):1–27, 2011.
- [7] Boyang Deng, Junjie Yan, and Dahua Lin. Peephole: Predicting network performance before training. *arXiv preprint arXiv:1712.03351*, 2017.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [9] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *International Conference on Learning Representations (ICLR)*, 2020.
- [10] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *arXiv preprint arXiv:1808.05377*, 2018.
- [11] Stefan Falkner, Aaron Klein, and Frank Hutter. Bohb: Robust and efficient hyperparameter optimization at scale. In *ICML*, 2018.
- [12] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- [13] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- [14] Saya Fujino, Naoki Mori, and Keinosuke Matsumoto. Deep convolutional networks for human sketches by means of the evolutionary deep learning. In *2017 Joint 17th World Congress of International Fuzzy Systems Association and 9th International Conference on Soft Computing and Intelligent Systems (IFSAS-SCIS)*, pages 1–5. IEEE, 2017.
- [15] Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. *Machine learning*, 63(1):3–42, 2006.
- [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] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [18] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [19] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [20] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [21] Yann LeCun, D Touresky, G Hinton, and T Sejnowski. A theoretical framework for back-propagation. In *Proceedings of the 1988 connectionist models summer school*, volume 1, pages 21–28, 1988.
- [22] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- [23] 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.
- [24] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc V Le, and Alexey Kurakin. Large-scale evolution of image classifiers. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2902–2911. JMLR. org, 2017.
- [25] Dolly Sapra and Andy D Pimentel. Constrained evolutionary piecemeal training to design convolutional neural networks. In *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems. Springer*, 2020.
- [26] Pranab Kumar Sen. Estimates of the regression coefficient based on kendall’s tau. *Journal of the American statistical association*, 63(324):1379–1389, 1968.
- [27] Connor Shorten and Taghi M Khoshgofaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):1–48, 2019.
- [28] Yanan Sun, Handing Wang, Bing Xue, Yaochu Jin, Gary G Yen, and Mengjie Zhang. Surrogate-assisted evolutionary deep learning using an end-to-end random forest-based performance predictor. *IEEE Transactions on Evolutionary Computation*, 24(2):350–364, 2020.
- [29] Yanan Sun, Bing Xue, Mengjie Zhang, and Gary G Yen. A particle swarm optimization-based flexible convolutional autoencoder for image classification. *IEEE transactions on neural networks and learning systems*, 30(8):2295–2309, 2018.
- [30] Yanan Sun, Bing Xue, Mengjie Zhang, and Gary G Yen. Completely automated cnn architecture design based on blocks. *IEEE transactions on neural networks and learning systems*, 31(4):1242–1254, 2019.
- [31] Yanan Sun, Bing Xue, Mengjie Zhang, and Gary G Yen. Evolving deep convolutional neural networks for image classification. *IEEE Transactions on Evolutionary Computation*, 2019.
- [32] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

- [33] Yehui Tang, Yunhe Wang, Yixing Xu, Hanting Chen, Boxin Shi, Chao Xu, Chunjing Xu, Qi Tian, and Chang Xu. A semi-supervised assessor of neural architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1810–1819, 2020.
- [34] Wei Wen, Hanxiao Liu, Yiran Chen, Hai Li, Gabriel Bender, and Pieter-Jan Kindermans. Neural predictor for neural architecture search. In *European Conference on Computer Vision*, pages 660–676. Springer, 2020.
- [35] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [36] Lingxi Xie and Alan Yuille. Genetic cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1379–1388, 2017.
- [37] Yixing Xu, Yunhe Wang, Kai Han, Hanting Chen, Yehui Tang, Shangling Jui, Chunjing Xu, Qi Tian, and Chang Xu. Renas: Relativistic evaluation of neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [38] Chris Ying, Aaron Klein, Eric Christiansen, Esteban Real, Kevin Murphy, and Frank Hutter. Nas-bench-101: Towards reproducible neural architecture search. In *International Conference on Machine Learning*, pages 7105–7114, 2019.
- [39] Xinjie Yu and Mitsuo Gen. *Introduction to evolutionary algorithms*. Springer Science & Business Media, 2010.
- [40] Yu Zhang, William Chan, and Navdeep Jaitly. Very deep convolutional networks for end-to-end speech recognition. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4845–4849. IEEE, 2017.
- [41] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.