



Evolutionary Algorithm Enhanced Neural Architecture Search for Text-Independent Speaker Verification

Xiaoyang Qu, Jianzong Wang*, Jing Xiao

Ping An Technology (Shenzhen) Co., Ltd.

{quxiaoyang343, wangjianzong347, xiaojing661}@pingan.com.cn

Abstract

State-of-the-art speaker verification models are based on deep learning techniques, which heavily depend on the hand-designed neural architectures from experts or engineers. We borrow the idea of *neural architecture search* (NAS) for the *text-independent speaker verification task*. As NAS can learn deep network structures automatically, we introduce the NAS conception into the well-known x-vector network. Furthermore, this paper proposes an evolutionary algorithm enhanced neural architecture search method called Auto-Vector to automatically discover promising networks for the speaker verification task. The experimental results demonstrate our NAS-based model outperforms state-of-the-art speaker verification models.

Index Terms: speaker verification, deep learning, neural network, neural architecture search.

1. Introduction

Speaker verification is the process of verifying whether an utterance belongs to the same speaker, based on enrolled speaker information. It can be categorized as *text-dependent speaker verification* (TD-SV) and *text-independent speaker verification* (TI-SV). Relatively, TI-SV is more convenient for practical applications, as it poses no constraints, e.g., duration or lexical content, on utterances to verify. However, it is also more difficult to achieve a good performance, due to many potential variabilities of the utterances. In this work, we focus on TI-SV.

In the early years, the i-vector [1] based models with PLDA[2] backend dominated the development of the speaker verification application. In recent years, the deep neural networks (DNN) trained as acoustic models for automatic speech recognition (ASR) are integrated into the i-vector system [3, 4, 5]. Although the ASR DNN can enhance phonetic modeling in the i-vector UBM, it adds a high computational cost to the i-vector system. In the latest years, DL-based techniques can be used as utterance-level speaker feature extractor [6, 7, 8, 9], and enable an end-to-end pipeline to discriminate between speakers [10, 11, 12].

However, these architectures are hand-designed by experts or experienced engineers. It is highly demanding on their knowledge and experiences. As a result, neural architecture search [13][14] techniques are becoming an increasingly popular topic in both academia and industry, because of its great potential to automatically find more effective architectures to outperform hand-crafted ones. The early works on NAS are based on reinforcement learning or evolutionary algorithm, such as [13, 15, 16, 17, 18]. But these approaches are expensive in time. To reduce search time costs, researchers proposed a wide range of optimization paradigms [19, 20, 21, 22], where hyper-network [23, 24, 25, 26] is a typical representative.

*Corresponding author

In this work, we bring the idea of hyper-network based neural architecture search into text-independent speaker verification. We managed to improve its search efficiency by use of a memetic evolutionary algorithm. Our work has several contributions as follows. (1) As NAS can learn deep network structures automatically, we introduce the NAS conception into the x-vector network. (2) To learn more promising structures for speaker verification, we build a large-scale hyper-network with repetitive architecture motifs. (3) To discover more promising candidate networks, we use a memetic evolutionary algorithm. (4) The experiment results demonstrate that our NAS-based x-vector and Auto-Vector outperform state-of-the-art speaker verification methods in two datasets.

2. Proposed Methods

2.1. NAS-based x-vector

First, let us review the well-known x-vector network as shown in Figure 1(a). Suppose the input utterance contains T frames. The first five layers L_1 to L_5 are frame-level information hidden layers. These layers are connected with a time-delay architecture with temporal context windows. The context window over the first layer is set as a range from $T - 2$ to $T + 2$. The second and third layers splice the output of the previous layer at time steps $\{T-2, T, T+2\}$ and $\{T-3, T, T+3\}$, respectively. The statistical pooling layer builds utterance-level feature by calculating mean and standard deviation over frame-level features. Note, the seventh hidden layer L_7 and the final softmax output layer are used for training and discarded in the evaluation process. The sixth layer L_6 is used as the embedding of x-vector.

As NAS can learn deep network structures automatically, we introduce the NAS conception into the x-vector network, as shown in Figure 1(b). For conventional x-vectors, the context windows between frame-level hidden layers are set by experts. Here, we let the number and the size of the context window to be decided by an automatic-decided method. This method is developed from hyper-network-based NAS, which stands out among these efficient NAS approaches because they can significantly reduce the tedious training process by sharing its parameters with all candidate networks. The key point is to specify the search space of hyper-network, which contains all possible candidate networks. As shown in Figure 1(b), we incorporate various choice temporal context windows for the first five layers. Then, we use a memetic evolution search policy shown in Section 2.3.2 to find the optimal candidate network with combination choices for temporal context windows. The static pooling, the sixth, and seventh layers are the same as conventional x-vectors. However, the small search space limits the potentiality of NAS-based x-vector. To enable an ample search space, we designed an Auto-Vector for speaker verification, as shown in Section 2.2.

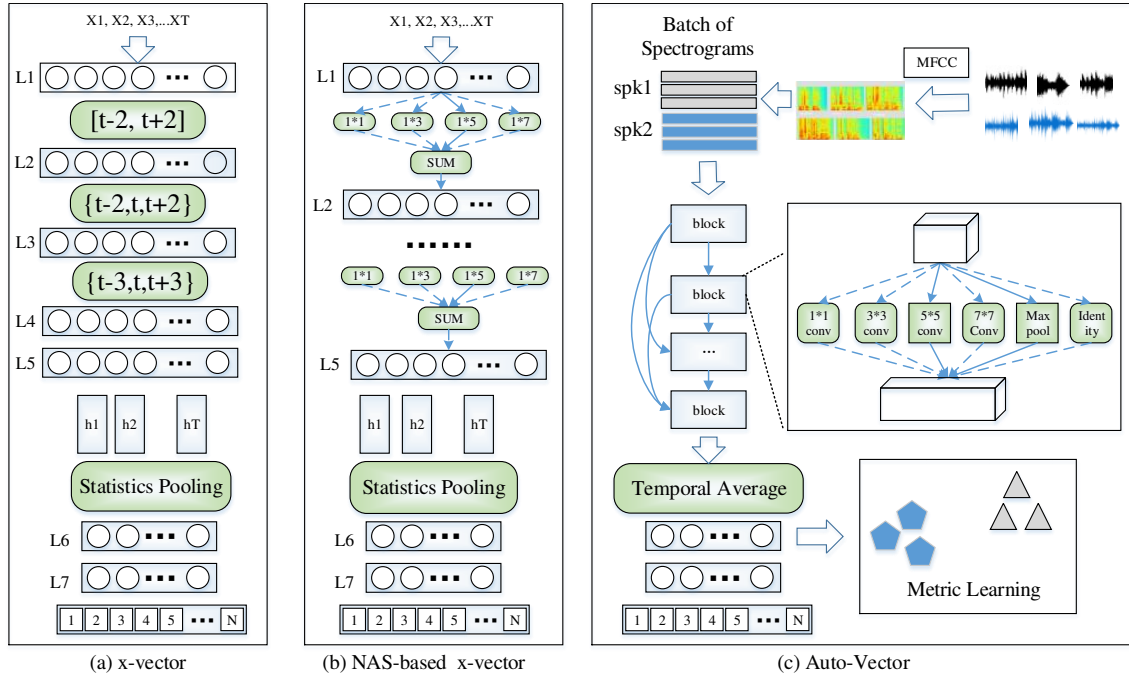


Figure 1: (a)The embedding DNN architecture of x-vector. (b)Our NAS-based x-vector. (c)Our Auto-Vector for speaker verification.

2.2. Auto-Vector

To automatically learn more promising neural architectures for text-independent speaker verification, we build a large-scale hyper-network with repetitive architecture motifs. As shown in Figure 1(c), the framework includes three parts: input features, architecture, and loss.

Input Features. MFCCs (Mel-frequency cepstral coefficients) is used to extract the frame-level acoustic feature vectors from raw waveform signals. Then the frames are converted into input acoustic features of 40-dimensional MFCCs with a frame-length of 25ms. This gives spectrograms of size 40*300 for 3 seconds of speech.

Architecture. As shown in Figure 1(c), we build a hyper-network containing the entire search space of architectures. The architecture of hyper-network is stacked with identical structure but different weights. Assume there are N_B choice blocks, and every block has N_{op} choice operations. Each choice block applies either one or two different operations out of N_{op} possible options. Thus, there are, therefore $N_{op} + \frac{N_{op} * (N_{op} - 1)}{2}$ possible combinations of operations that we can apply in each block. In our experiment, the N_{op} is set as 6, so we have six possible operations: a max-pooling layer, an identity operation, and convolution layers of size 1x1, 3x3, 5x5 and 7x7. The average temporal pooling is implemented by applying a 2D adaptive average pooling several input planes. As the size of the search space grows exponentially with the number of choice blocks N_B , this large-scale search space can enable more possibility of promising networks.

Loss. In addition to softmax pre-training, we also use distance-based loss function, such as triplet loss or generalized end-to-end loss. Among various distance-based loss functions [12, 11, 10], the generalized end-to-end loss function [11] perform best, because it not only learns to rank but also emphasizes the hard examples. The details are shown in Section 2.4.

2.3. Evolutionary Algorithm Enhanced NAS

2.3.1. The Overall Training Procedure

The training procedure of NAS-based x-vector and Auto-Vector consist of four steps: (1) Design a search space. (2) Train the hyper-network. (3) Search the optional sub-networks with their parameters inherited from the hyper-network. (4) Retrain the best-accuracy candidate sub-network as a standalone model. The *first step* have been shown in Section 2.2. The *second step* is to train hyper-network. The training goal is formulated as

$$\theta^*(H) = \underset{\theta}{\operatorname{argmin}} \mathcal{L}_{train}(H, \theta) \quad (1)$$

here θ^* means the weights of the hyper-network H and $\mathcal{L}_{train}(\cdot)$ denotes a loss function on training dataset.

The *third step* is to search for high-quality sub-networks, with their parameters inherited from the hyper-network. The high-quality sub-network search task is a black-box optimization, which aims to find an approximate maximizer of an objective function $f(x)$ using a given budget of N sub-network evaluations. In can be formulated as

$$\tilde{a}^* \approx a^* = \underset{a_n}{\operatorname{argmax}} f(a_n \sim H) \quad (2)$$

note the sub-network a_n is sampling from the search space of hyper-network H . For this optimization goal, we develop a memetic algorithm based evolutionary policy, which is illustrated in Section 2.3.2.

The *last step* is to retrain the obtained optimal sub-network for the best performance. The model parameters learned by minimizing the accumulative loss shown in the equation of

$$\phi^*(\tilde{a}^*) = \underset{\phi}{\operatorname{argmin}} \mathcal{L}_{train}(\tilde{a}^*, \phi) \quad (3)$$

here ϕ^* means the weights of the best-accuracy candidate sub-model \tilde{a}^* . And, $\mathcal{L}_{train}(\cdot)$ denotes a loss function on training dataset. The details of loss function is shown in Section 2.4.

2.3.2. Memetic Evolutionary Search Policy

Our search policy is based on the memetic algorithm. The memetic algorithm is an augmentation of the genetic algorithm. In other words, the memetic algorithm consists of the genetic algorithm and one or more local search components. The memetic algorithm integrates the local search method into the genetic algorithm to reduce the likelihood of premature convergence. Thereby, the promising child individuals are generated by recombination from and adaptation from outstanding individuals.

The search process is shown in Algorithm 1. The inputs include an empty population set Ω with size S , the generation number G , and the well-trained hyper-network H . The key operations of the search process are shown as follows. (1) For *mutation operations*, the selected candidate choose one or two different operations in its every choice block with probability 0.1 to produce a new candidate. Because cross-over operations will result in local operation, we only make mutation operations. (2) The *local search* employs a hill-climbing algorithm to discover high-quality sub-networks by greedily moving in the direction of better-performing sub-networks. (3) The *compete operation* uses an acceptance criterion to pick the better one. (4) The *fitness evaluation* is calculating as $f_i = 1 - \delta_i$, where δ_k means the equal error rate of i individual model. (5) The *selection operation* is based on a tournament selection policy. In the tournament selection policy, a candidate set is randomly selected from the overall population set. Then, the best-fitness individual is chosen from the candidate set rather than the overall population set. This policy can avoid zooming in on good models too early and enable more search space to be explored.

Algorithm 1 Memetic evolutionary algorithm

```

1: Input: The population set  $\Omega$  with size  $S$ 
2: Input: The generation number  $G$ 
3: Input: The well-trained hyper-network  $H$ 
4: Output: A best-accuracy sub-network  $a$ 
5: Initialize population set  $\Omega$ 
6: for  $i \leftarrow 1 \dots S$  do
7:    $a_i \leftarrow \text{uniform-sampling}(H)$ 
8:    $f_i \leftarrow \text{fitness evaluation}(a_i)$ 
9:    $\Omega = \Omega + \{a_i, f_i\}$ 
10: end for
11: for  $j \leftarrow 1 \dots G$  do
12:    $\bar{a}_j \leftarrow \text{mutation}(a_j)$ 
13:    $\bar{a}_j \leftarrow \text{local-search}(\bar{a}_j)$ 
14:    $a_j \leftarrow \text{compete}(\bar{a}_j, a_j)$ 
15:    $f_j \leftarrow \text{fitness-evaluation}(a_j)$ 
16:    $\Omega = \Omega + \{a_j, f_j\}$  # add the promising candidate
17:    $\Omega = \Omega - \text{worst}(\Omega)$  # remove the worst candidate.
18:    $a_{j+1} \leftarrow \text{tournament selection}(\Omega)$ 
19: end for

```

2.4. Backend

The objective of typical cross-entropy loss is to learn to predict directly a label given an input. Metric learning aims to predict relative distance between inputs. In addition to softmax pre-training, we also use distance-based loss function. Assuming N speakers with each M utterance. The loss function $\mathcal{L}(\cdot)$ is shown as follows.

$$\mathcal{L}(e_a, e_p, e_n) = 1 - \sigma(d(e_a, e_p)) + \max_{\substack{1 \leq k \leq N \\ k \neq j}} \sigma(d(e_a, e_n^k)) \quad (4)$$

Table 1: Equal Error Rate Comparison(The Lower, The Better)

	EER (Dataset1)	EER (Dataset2)	Size
LSTM-GE2E[11]	6.2%	8.3%	4.6M
x-vector[27]	4.6%	6.5%	6.14M
NAS-based x-vector	4.3%	5.6%	6.32M
Auto-Vector	1.8%	3.6%	5.17 M

where $d(e_a, e_p)$ means the scaled similarity score between the anchor embedding e_a and the positive embedding e_p . Here, e_a and e_p belongs to the same speaker. The negative embedding e_n^k is the centroid embedding of the k_{th} speakers, which should be evaluated as $e_n^k = \frac{1}{M} \sum_{m=1}^M e_a^k(m)$, using M utterances for the k_{th} speaker. $\sigma(\cdot)$ means the sigmoid function. Here, the scaled similarity score function $d(e_a, e_p)$ is defined as

$$d(e_a, e_p) = w \cdot \cos(e_a, e_p) + b \quad (5)$$

here, w and b are learnable parameters. $\cos(\cdot)$ means the cosine similarity function.

3. EXPERIMENT

3.1. Dataset Collection and Pre-Processing

We use two datasets for Evaluation. **Dataset1** includes 300 speakers with 4527 utterances in total. The duration of which mostly range from 3 to 7 seconds. We split the overall dataset into a training dataset of 270 speakers and a test dataset of 30 speakers. 10 utterances are randomly chosen as enrollment utterances for each speaker, and another 10 randomly chosen utterances are used as evaluation samples.

Dataset2 includes 4000 speakers with 23573 utterances and more than 12,600 hours of speech. This dataset is split into two parts: a training dataset of 3960 speakers and an evaluation dataset of 285 speakers. The evaluation partition consists of 285 speakers that do not overlap with the 3960 speakers for training datasets.

The raw waveform audios with a 16KHz sampling rate are converted into frames using a hamming window of width 25 ms and step 10ms. MFCCs(Mel-frequency cepstral coefficients) is used to extract the frame-level acoustic feature vectors from raw waveform signals. Then the frames are converted into input acoustic features of 40-dimensional MFCCs with a frame-length of 25ms that are mean-normalized over a sliding window of up to 3 seconds. This gives spectrograms of size 40*300 for 3 seconds of speech. An energy-based VAD is employed to filter out non-speech frames from the utterances. There are N speakers each with M utterances.

3.2. Overall Result

Table 1 shows the EER comparison of four models on Dataset1 and Dataset2. Our Auto-Vector performs better than LSTM and x-vector. For two datasets, the equal error rate(EER) of our Auto-Vector is much lower than LSTM and x-vector. This result proves that the neural architecture search network can find a better model than the expert-designed hand-crafted models. We use the same back-end(GE2E) for all evaluated systems to eliminate the impacts of different back-end classifiers.

The configurations of two baseline networks are shown as follows. The first baseline is a 3-layer **LSTM Network**[11] with a projection of size 256. The embedding vector(d-vector)

Table 2: The evaluation results of hyper-network and sub-network on Dataset1. Here, F means the number of filters in the first convolution layer and B means the number of choice blocks. For sub-networks, we retrain top-10 sub-networks and report the mean x and standard deviation y as $x \pm y$.

	size (M)	EER (%)	cost (GPUh)
HyperNet(F=16,B=24)	2.43	3.5	14.6
HyperNet(F=32,B=24)	6.08	2.7	21.7
HyperNet(F=64,B=24)	17.04	1.9	33.9
HyperNet(F=128,B=24)	46.82	1.4	50.7
HyperNet(B=12,F=32)	5.06	3.1	18.9
HyperNet(B=24,F=32)	6.08	2.7	21.7
HyperNet(B=36,F=32)	7.08	2.6	24.3
HyperNet(B=48,F=32)	8.05	2.2	26.2
SubNet(F=16,B=48)	0.8 ± 0.2	3.7 ± 0.3	-
SubNet(F=32,B=48)	2.1 ± 0.5	2.9 ± 0.2	-
SubNet(F=64,B=48)	6.1 ± 1.2	1.9 ± 0.1	-
SubNet(F=128,B=48)	15.2 ± 3.2	1.6 ± 0.1	-

size is the same as the LSTM projection size. There are 768 hidden nodes in the LSTM layer. The expected average moving is used to get the embedding. The second baseline is **x-vector Network**[27]. The first five layers L_1 to L_5 are frame-level information hidden layers. While there are 512 nodes in each of the first four layers L_1 to L_4 , there are 1500 nodes in the fifth layer L_5 . The statistical pooling layer builds utterance-level feature by calculating mean and standard deviation over frame-level features. Two utterance-level layers L_6 and L_7 each have 512 nodes. The sixth layer L_6 is used as embedding.

Our **NAS-based x-vector** is stacked with a repetitive context block. Each block contains 4 choice temporal context windows. As the search space is small, we use random search policy for NAS-based x-vector.


For **Auto-Vector**, the hyper-parameters of hyper-network (number of blocks B and the number of filters F) are analyzed in Section 3.3. The embedding size is set as 512. To decouple the correlation of sub-networks, we set the path dropout rate as 0.1. For the input, the batch size is set as 40 utterances from 8 speakers, each with 5 utterances. For training, we use Adam optimizer and a linear learning rate decay policy with a base learning rate of 0.02. For the memetic evolutionary search, the size of the population set is 100, and the number of generations is 2000.

3.3. The Evaluation Details of Auto-Vector


Hyper-Network Training. First, we parameterize our models based on F , the number of filters in the first convolution layer, as shown in Table 2. When $F = 16$ and $B=24$, we obtain an average EER of 3.5 with about 2.43 M parameters. with a double growth of filter number, the growth of model size is multiplied by nearly three times. Obviously, the equal error rate will decrease with the increase of filters. As we use two reduction blocks in our hyper-network model, the growth model size should be multiplied by four times. However, due to the existing of dense layer in the end of NAS-based model, the growth model size is multiplied by nearly three times. The best model gets 1.4% EER with around 46.82M parameters.

Then, we parameterize our models based on B , the number of choice blocks. When $B= 12$ and $F=32$, we obtain an average

EER of 3.1% with about 5.06M parameters. The best model gets 2.2% EER with around 8.05M parameters. With a double boost of the block number, the model size only increases a little because there is two dense layers in the tail of our model. The weights of the dense layer dominate the model size, so the model size increases at a low rate along with the double increase of the number of filters.



(a) Search hyper-network with $F=32$ and $B=24$



(b) Search hyper-network with $F=64$ and $B=48$

Figure 2: The evaluation histogram of 2000 candidates sub-networks with their parameters inherited from the hyper-network

The Impact of Evolutionary Search. Compared to the random selection algorithm, our hierarchical evolutionary algorithm can generate more high-quality candidate models. The equal error rate (EER) distribution of candidate models is shown in Figure 2. While most of the candidate models searched out by our evolutionary algorithm, have a lower equal error rate (EER) than by random search. Besides, the best-accuracy model is found out by our evolutionary algorithm rather than by a random algorithm. This result further proves that our hierarchical evolutionary algorithm can get more space to be explored to generate more high-quality candidate models.

Sub-Network Re-Training. As shown in Table 2, we retrain top-10 sub-networks and report the mean x and standard deviation y for sub-network training. We aim to find the best-quality model whose model size is smaller than the x-vector network. When $F=64$ and $B=48$, we can discover the optimal model whose model size is 5.17M, and EER is 1.8%.

4. Conclusion

In this paper, we introduce the NAS conception into well-known x-vector network. Enabling more search space to be explored, we use an evolutionary algorithm enhanced neural architecture search framework to search high-quality sub-networks. The experiment shows that our system outperforms two state-of-the-art end-to-end methods in a public dataset. Besides, our NAS method can achieve a reduction of 36%-86% in equal error compared with the state-of-the-art methods.

5. Acknowledgment

This paper is supported by National Key Research and Development Program of China under grant No. 2018YFB1003500, No. 2018YFB0204400 and No. 2017YFB1401202.

6. References

- [1] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.
- [2] T. Stafylakis, P. Kenny, P. Ouellet, J. Perez, M. Kockmann, and P. Dumouchel, "Text-dependent speaker recognition using plda with uncertainty propagation," *matrix*, vol. 500, no. 1, 2013.
- [3] Y. Lei, N. Scheffer, L. Ferrer, and M. McLaren, "A novel scheme for speaker recognition using a phonetically-aware deep neural network," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 1695–1699.
- [4] P. Kenny, T. Stafylakis, P. Ouellet, V. Gupta, and M. J. Alam, "Deep neural networks for extracting baum-welch statistics for speaker recognition," in *Odyssey*, vol. 2014, 2014, pp. 293–298.
- [5] M. McLaren, Y. Lei, and L. Ferrer, "Advances in deep neural network approaches to speaker recognition," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 4814–4818.
- [6] E. Variani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 4052–4056.
- [7] Y. Liu, Y. Qian, N. Chen, T. Fu, Y. Zhang, and K. Yu, "Deep feature for text-depnt speaker verification," *Speech Communication*, vol. 73, pp. 1–13, 2015.
- [8] Z. Shi, M. Wang, L. Liu, H. Lin, and R. Liu, "A double joint bayesian approach for j-vector based text-depnt speaker verification," *arXiv preprint arXiv:1711.06434*, 2017.
- [9] T. Fu, Y. Qian, Y. Liu, and K. Yu, "Tandem deep features for text-dependent speaker verification," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [10] G. Heigold, I. Moreno, S. Bengio, and N. Shazeer, "End-to-end text-dependent speaker verification," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5115–5119.
- [11] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4879–4883.
- [12] C. Li, X. Ma, B. Jiang, X. Li, X. Zhang, X. Liu, Y. Cao, A. Kannan, and Z. Zhu, "Deep speaker: an end-to-end neural speaker embedding system," *arXiv preprint arXiv:1705.02304*, 2017.
- [13] B. Zoph and Q. V. Le, "Neural architecture search with reinforcement learning," *arXiv preprint arXiv:1611.01578*, 2016.
- [14] B. Baker, O. Gupta, N. Naik, and R. Raskar, "Designing neural network architectures using reinforcement learning," *arXiv preprint arXiv:1611.02167*, 2016.
- [15] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [16] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, "Regularized evolution for image classifier architecture search," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 4780–4789.
- [17] R. Miikkulainen, J. Liang, E. Meyerson, A. Rawal, D. Fink, O. Francon, B. Raju, H. Shahrzad, A. Navruzyan, N. Duffy *et al.*, "Evolving deep neural networks," in *Artificial Intelligence in the Age of Neural Networks and Brain Computing*. Elsevier, 2019, pp. 293–312.
- [18] P. R. Lorenzo and J. Nalepa, "Memetic evolution of deep neural networks," in *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 2018, pp. 505–512.
- [19] C. Liu, B. Zoph, M. Neumann, J. Shlens, W. Hua, L.-J. Li, L. Fei-Fei, A. Yuille, J. Huang, and K. Murphy, "Progressive neural architecture search," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 19–34.
- [20] R. Negrinho and G. Gordon, "Deeparchitect: Automatically designing and training deep architectures," *arXiv preprint arXiv:1704.08792*, 2017.
- [21] H. Pham, M. Y. Guan, B. Zoph, Q. V. Le, and J. Dean, "Efficient neural architecture search via parameter sharing," *arXiv preprint arXiv:1802.03268*, 2018.
- [22] J.-H. M. Elsken, Thomas and F. Hutter., "Simple and efficient architecture search for convolutional neural networks," *arXiv preprint arXiv:1711.04528*, 2017.
- [23] A. Brock, T. Lim, J. M. Ritchie, and N. Weston, "Smash: one-shot model architecture search through hypernetworks," *arXiv preprint arXiv:1708.05344*, 2017.
- [24] G. Bender, P.-J. Kindermans, B. Zoph, V. Vasudevan, and Q. Le, "Understanding and simplifying one-shot architecture search," in *International Conference on Machine Learning*, 2018, pp. 549–558.
- [25] C. Zhang, M. Ren, and R. Urtasun, "Graph hypernetworks for neural architecture search," *arXiv preprint arXiv:1810.05749*, 2018.
- [26] Z. Guo, X. Zhang, H. Mu, W. Heng, Z. Liu, Y. Wei, and J. Sun, "Single path one-shot neural architecture search with uniform sampling," *arXiv preprint arXiv:1904.00420*, 2019.
- [27] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.