

This ICCV workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Graph-based Neural Architecture Search with Operation Embeddings

Michail Chatzianastasis^{*†} National Technical University of Athens mixalisx97@gmail.com

Georgios Siolas National Technical University of Athens gsiolas@islab.ntua.gr

Abstract

Neural Architecture Search (NAS) has recently gained increased attention, as a class of approaches that automatically searches in an input space of network architectures. A crucial part of the NAS pipeline is the encoding of the architecture that consists of the applied computational blocks, namely the operations and the links between them. Most of the existing approaches either fail to capture the structural properties of the architectures or use hand-engineered vector to encode the operator information. In this paper, we propose the replacement of fixed operator encoding with learnable representations in the optimization process. This approach, which effectively captures the relations of different operations, leads to smoother and more accurate representations of the architectures and consequently to improved performance of the end task. Our extensive evaluation in ENAS benchmark demonstrates the effectiveness of the proposed operation embeddings to the generation of highly accurate models, achieving state-of-the-art performance. Finally, our method produces top-performing architectures that share similar operation and graph patterns, highlighting a strong correlation between the structural properties of the architecture and its performance.

1. Introduction

Deep learning has been in the middle of a research outburst during the last years and the constant need for highly accurate models requires extensive architecture engineering. Neural Architecture Search (NAS) has emerged as the most promising field for the efficient automated search and generation of state-of-the-art models. Its contribution has been studied for a variety of tasks, ranging from medical imGeorge Dasoulas[†] École Polytechnique george.dasoulas1@gmail.com

Michalis Vazirgiannis Ecole Polytechnique mvazirg@lix.polytechnique.fr

age segmentation [11] to objection detection [9] and speech recognition [4].

The research interest has focused on latent space optimization techniques [2], due to their efficiency with respect to the search space and the optimization [32]. Specifically, a generative model learns continuous representations of neural network architectures, and then the objective function (i.e., the performance of the architecture) is optimized on the latent space. Recent work has shown that the representation of the architecture is crucial for the overall performance of the NAS method [28, 29]. The most promising approaches represent the architecture as a graph, in which every node is associated with a layer operation. However they assume a fixed encoding of the operations, such as onehot vectors [36]. This assumption puts a limitation on both the expressivity of the operation information and the possible relations between the operations, as it employs orthogonal representations with equal distances between them.

In this work, we suggest the replacement of the fixed representation of the operations with learnable embeddings that are integrated as parameters into the optimization. Our goal is to produce more accurate and smooth architecture representations, that can take into account how the different operators interact with each other. Our contributions can be summarized as follows:

- We propose the operation embeddings as a continuous representation of the applied operators and we integrate them as parameters into the NAS pipeline.
- We experimentally show that the parameterized representations of the operations lead to the generation of state-of-the-art architectures.
- We observe that the top-performing generated architectures share similar structural patterns, with the clustering coefficient and the average path length being strong indicators of the model performance.

^{*}Work done as intern at École Polytechnique in Data Science and Mining (DaSciM) team group.

[†]Equal Contribution

The rest of the paper is structured as follows: In Section 2 we highlight previously developed work in the area of NAS. In particular, we emphasize on the application of graph learning algorithms for the encoding of the neural network architecture. In Section 3 we present our main contribution, which is the introduction of *operation embeddings* into Graph VAEs. Finally, in Section 4 we evaluate the contribution of the operation embeddings to VAE models of the neural network architectures through an experimental study in the ENAS search space. Moreover, we investigate how several structural characteristics of the network encoding affect the performance of the generated architectures.

2. Related Work

Neural Architecture Search. In the last years, significant progress has been made in automating architecture engineering. Neural Architecture Search has proved its ability to construct architectures that achieve state-of-theart results in various tasks, with little human intervention [37, 38, 24, 22, 8]. The NAS task can be formulated as an optimization problem in an input space of network architectures. Common techniques for solving this optimization problem as reinforcement learning [1, 37] and evolutionary methods [23, 22, 26] operate in a discrete search space. Directly searching an architecture within this space is inefficient given its exponential growth as the number of operations and layers increases [17, 7]. To tackle this challenge, recent research works have introduced differentiable search methods, that operate on a continuous relaxation of the search space [17, 16]. In particular, Neural Architecture Optimization (NAO) has been proposed as a framework that trains an auto-encoder and a performance predictor using gradient descent [17]. However, the model is trained in a supervised manner, limiting the ability to transfer the latent space in other datasets. In this work, we are showing how our method can be applied to unsupervised NAS models.

Graph Representation Learning for Neural Architecture Search. Unsupervised graph representation learning methods have shown promising results in neural architecture search and specifically in the accurate and expressive architecture encoding [15, 32, 36]. The basic idea is to represent a neural network architecture as a graph and to learn a smooth continuous latent space, such that high-performance architectures be mapped close to each other. Given a continuous and smooth architecture representation, various strategies can be efficiently applied, such as the bayesian optimization [10, 29].

Modeling of network architecture as a graph can be achieved in various ways. String-based methods have been proposed before in order to provide representations of the architecture as a graph [3, 34]. These methods represent the graph as a sequence of strings and apply Recurrent Neural Network (RNN) models to process the sequence. The disadvantage of these approaches is that they do not preserve the permutation invariance property [35, 31], imposing restrictions to the expressiveness of the representations.

In contrast to the string-based methods, recent research works leverage the structure of the graph and operate directly on it, using message-passing operations [36, 15]. D-VAE proposed as a graph-based autoencoder for Directed Acyclic Graphs (DAGs) [36]. It applies a graph neural network model with an asynchronous message passing process to encode the architectures. The main limitation of this approach is the utilization of a fixed one-hot encoding of the operation blocks, not being able to capture possible operation relations. Variational Graph Isomorphism auto-encoder is proposed to obtain unsupervised representations of neural network architectures [32]. It leverages Graph Isomorphism Networks (GIN) [31] to encode the graph architectures into the latent space. However, it decodes the whole graph in one shot and also the operations are represented with fixed uninformative vectors. A recent work proposes the utilization of operation information into a generic graph-based framework for encoding neural network architectures, without applying it, though, to an unsupervised setting of architecture generation [20].

Finally, the authors in [33] utilize relation graphs for the representation of neural network architectures. They apply network generators in order to construct relation graphs with specific structural characteristics. Our work as well, studies structural properties that affect the performance of neural network architecture, but instead of relation graphs, we use DAGs with operation embeddings for the representation of the network architectures.

3. Operation Embeddings in Variational Graph Auto-Encoders

In this section we first introduce the necessary notation for: a) modeling neural network architectures as graph structures and b) building various graph generative models. Then, we describe our proposed method that replaces the fixed vector encodings of the operations with learnable representations, called as *operation embeddings*. The proposed embeddings can be easily incorporated into a variety of graph-based models and enhance their performance in neural architecture search tasks.

3.1. Neural Network Architecture as a Directed Acyclic Graph

A neural network architecture represents a computation, that is applied to an input signal using a fixed set of operations. We can define the computational graph of an architecture \mathcal{A} as $G_{\mathcal{A}} = (V, E)$, where V is the set of nodes or the applied operations and E is the set of edges or the links that



Figure 1. Example of architecture encoding as a computational graph. The used operations are the 3×3 convolution, the 5×5 convolution and the max pooling operator.

define the signal flow among the applied operations. We assume that K is the set of the available architecture operations (e.g an example of K can be {max, min, $conv_{3\times3}, ...$ }). $G_{\mathcal{A}}$ is a directed, acyclic (i.e., a finite number of performed operations) and labeled graph, where each node $u \in V$ is associated with label x_u that corresponds to the operation of node u. The most common representation scheme of a labeled graph is the adjacency matrix $\mathbf{A} \in \{0, 1\}^{|V| \times |V|}$ and the label matrix $\mathbf{X} \in \mathbb{Z}^{|V| \times |K|}$, which corresponds to the one-hot encoding of the operations. We note that assuming a DAG structure, \mathbf{A} is not symmetric, imposing an ordering of the nodes and the sequence of processing them. An example of a computational graph of a neural network architecture is visualized in Figure 1.

3.2. Variational Graph Auto-Encoders

Let $G = (\mathbf{A}, \mathbf{X})$ be an input graph that represents a neural network architecture. According to the standard Variational Graph Auto-Encoder (VGAE) definition [12], our goal is to learn a probabilistic encoder model $q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X})$ which provides a distribution over latent representations, and a probabilistic decoder model $p_{\theta}(\mathbf{A}, \mathbf{X}|\mathbf{Z})$ from which we can generate new graphs. We also assume a prior normal distribution over the latent space $\mathbf{Z} \sim N(\mathbf{0}, \mathbf{1})$. We train the whole system by minimizing the *evidence lower bound*:

$$L(\phi, \theta; \mathbf{A}, \mathbf{X}) = \mathbb{E}_{q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X})}[\log p_{\theta}(\mathbf{A}, \mathbf{X}|\mathbf{Z})] - \mathrm{KL}[q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X})||p(\mathbf{Z})], \quad (1)$$

where KL denotes the Kullback–Leibler divergence. Equation 1 indicates that the model does not take into account the performance of an neural architecture and is trained in an unsupervised manner. We make the assumption that architectures with structural similarities and similar operators have similar performance.

3.2.1 Encoder

In the VGAE framework, the encoder uses a graph representation learning model to project G into a representation space with lower dimensionality. More specifically, we use a Graph Neural Network (GNN) model to obtain the representation of the nodes and then we apply a second neural network model to produce the mean and the variance of the posterior approximation. Let a GNN model $\phi : \mathbb{Z}^{|V| \times |V|} \times \mathbb{Z}^{|V| \times |K|} \to \mathbb{R}^{|V| \times d}$ denote a graph neural network that takes as input the connections and the operations of the nodes, and outputs an representation of every node. Also, let $\psi_1, \psi_2 : \mathbb{R}^{|V| \times d} \to \mathbb{R}^l$ denote two differentiable pooling functions that take as input the node representations and output a single representation vector for the whole graph. The encoder can be described via the following equations:

$$\mu_G = \psi_1(\phi(\mathbf{A}, \mathbf{X}))), \qquad (2a)$$

$$\sigma_G = \psi_2(\phi(\mathbf{A}, \mathbf{X}))), \tag{2b}$$

where μ_G and σ_G denote the mean and the variance of the approximation of the posterior distribution respectively. Note that this formulation expresses multiple variational graph auto-encoder models, that have been used before and utilize either synchronous or asynchronous messagepassing processes [36, 25, 32]. Moreover, standard choices of ψ_1, ψ_2 functions are pooling operators followed by Multi-Layer Perceptrons (MLP).

3.2.2 Decoder

The decoder is responsible for translating the latent representation into graph structures. For this work, we use the autoregressive decoder defined in [36]. We, now, briefly describe the decoder. Given a time step t, when node u_t is generated we have the following iterative procedure:

- 1. We apply an MLP model, which uses as input the current state of the graph, to determine the type of node *u*.
- 2. We update the hidden state of node u using a Gated Recurrent Unit (GRU) model [5]: $h_{u_t} = gru(\mathbf{x}_{u_t}, h_{\text{pred}})$, with h_{pred} denoting the aggregated representation from the predecessors of node u_t .
- 3. For all time steps k = t 1, t 2, ..., 1 we apply an MLP model that, given as input the states h_{u_t} and h_{u_k} , computes the probability p_{edge} of existence of edge (u_k, u_t) . In case that $p_{edge} > 0.5$ we add the edge (u_k, u_t) into the DAG and we perform the second step to update the representation h_{u_t} .

The iteration stops when the examined node is ending type, and then we output the generated graph structure.

3.3. Operation Embeddings

In the formulation of the models, described in Section 3.2, the operation matrix $\mathbf{X} \in \mathbb{Z}^{V \times |K|}$ is a fixed vector representation. Usually, this representation corresponds to the one-hot encoding of the operation set, so that there is an unordered representation of the operators. Given that, the auto-encoder treats all the operations equally.

The limitations of the one-hot encoding are twofold:

- a) It does not take into account the computational relationships and structural dependencies of the different operations. For example, a 5×5 convolutional layer is more similar to a 3×3 convolutional layer rather than to a max-pool layer in terms of the computational level.
- b) It cannot exploit information from the data, as the onehot vectors are fixed. This means that the optimization cannot affect the way that the model chooses operations.

Inspired by the success of word embeddings [18], we propose the incorporation of the embedding $\mathcal{O} : K \to \mathbb{R}^{|K| \times d_{op}}$ into the encoder model, to tackle the aforementioned limitations. The mapping $\mathcal{O}(\cdot)$ projects the set of available operations into a d_{op} -dimensional continuous space in a differentiable manner. We call the mapping \mathcal{O} operation embedding. Equations 2a and 2b are transformed as follows:

$$\mu_G = \psi_1(\phi(\mathbf{A}, \mathcal{O}(\mathbf{X}))), \tag{3a}$$

$$\sigma_G = \psi_2(\phi(\mathbf{A}, \mathcal{O}(\mathbf{X}))). \tag{3b}$$

In order to learn the operation embeddings used in equations 3a and 3b, we treat them as parameters of the autoencoder and optimize them with gradient descent along with the other weights. The incorporation of the operation embeddings into the architecture generation pipeline is visualized in Figure 2. We note that the same embedding $O(\cdot)$ is shared among the encoder and the decoder.

Latent Space Our ultimate goal is to produce smooth and accurate latent representations of neural architectures. Essentially, we want architectures with similar performance to be mapped in latent representations that are close to each other. This can help the downstream search algorithm to efficiently discover a distribution of high-performing architectures. Since the parameters of the embeddings matrix are changing throughout the training, a variable representation, based on the end task, of the operations is possible. The gradients of $\mathcal{O}(\mathbf{X})$'s weights affect the model training procedure as well. Using the equations 3a and 3b, the model

is able to map computationally similar operations close to each other. Consequently, architectures with similar structures and operation choices can have similar representations, leading to a smooth latent space.

Implementation For this study, we choose lowdimensionality for the produced operation embeddings with $d_{op} = 3$, as the number of different operations is small. We fully train the autoencoder model for N epochs and we repeat the process for T iterations. Let $\mathcal{O}_{n,t}(\mathbf{X})$ denote the operation embeddings matrix in n-th epoch of the t-th iteration. In the first iteration, we initialize the weights of $\mathcal{O}_{1,1}(\mathbf{X})$ from $\mathcal{N}(0, 1)$. In the iteration $t = T_i$, we initialize the operation embeddings using the output of the last epoch in the previous iteration $\mathcal{O}_{N,t-1}(\mathbf{X})$. Using this pre-training schema, we manage to achieve faster convergence of the model among the iterations, as the operation embeddings include more prior knowledge, based on the examined task.

4. Experiments

In this section, we empirically evaluate our proposed operation embeddings method. The experimentation details and the code are provided in the supplementary material.

Baselines. To demonstrate the effectiveness of our approach, we incorporate operation embeddings into two variants of a well known variational graph auto-encoder model for DAGs, that use as encoders either asynchronous message-passing operations (D-VAE) [36] or simultaneous graph convolutions (GCN) [13]. We refer to the models with operation embeddings as DVAE-EMB and GCN-EMB respectively. In DVAE-EMB we repeat the model training for T = 4 iterations, and in GCN-EMB for T = 1 as described in Section 3. We also include S-VAE [3] and GraphRNN [34] as baselines, which represent the architecture as a sequence of strings, and do not operate directly on the graph structure.

Tasks. In order to have a fair comparison with other approaches, we follow the experimental setup of [14, 36]. First, we compute basic effectiveness metrics of the variational auto-encoder models and we measure the *predictive performance* of the latent representations. Next, we present the best-performing architectures obtained with Bayesian optimization on the latent space, and note the observed similarities of several graph characteristics of them. Finally, we visualize the learned latent representations to show their smoothness.

Dataset. We train the variational graph auto-encoder models in ENAS search space [21] for 300 epochs. The



Figure 2. Workflow of Variational Graph Auto-Encoder for NAS with integration of Operation Embeddings. First, the corresponding directed acyclic graph $G_A = (A, \mathbf{X})$ of the input architecture is constructed. Then, the operations X are given as input in the operation embeddings layer, to obtain the embeddings $O(\mathbf{X})$. Finally, the adjacency matrix and the operation embeddings are passed to the auto-encoder.



Figure 3. The reconstruction loss during the training process.

dataset contains 19,020 neural architectures. Each architecture has 6 layers besides one input and one output layer. Each layer is associated with one operation. There are six available operations: 3×3 and 5×5 convolutions, 3×3 and 5×5 depthwise-separable convolutions [6], 3×3 max pooling and 3×3 average pooling. We use 90% of the dataset as training data, and the remaining 10% for evaluation. For the evaluation of their true performance, we fully train the architectures on CIFAR-10, using the same experimental setup with [21].

4.1. Basic abilities of Variational Graph Auto-Encoders

In this experiment, we evaluate the reconstructive abilities and the generative properties of the auto-encoders. We use the following metrics proposed by [36] :

- 1. *Accuracy*. The percentage of perfectly reconstructed architectures.
- 2. *Validity*. The percentage of valid architectures generated from the prior distribution.



Figure 4. The KL divergence during the training process.

3. *Uniqueness*. The proportion of unique architectures out of the valid generations.

We present the results in Table 4.2. DVAE-EMB and GCN-EMB outperform their counterparts in terms of reconstruction accuracy and validity, demonstrating the effectiveness of operation embeddings. D-VAE and GCN have smaller reconstruction accuracy, because their one-hot vector representation fails to capture the operation information accurately.

We, also, visualize in Figures 3 and 4 the reconstruction loss and the KL divergence during the training of D-VAE and our proposed model DVAE-EMB. We observe that the convergence of the reconstruction loss of DVAE-EMB is faster than D-VAE. Intuitively, the incorporation of operation embeddings helps the model to acquire extra information as it captures the relations between the operations. These relations can not be discovered in the D-VAE model, which uses one hot-vectors for representing the operations. Therefore, our model can converge in fewer epochs achieving lower training loss.

Model	Accuracy	Validity	Uniqueness
D-VAE	99.96	100.00	37.26
GCN	98.70	99.53	34.00
S-VAE	99.98	100.00	37.03
GraphRNN	99.85	99.84	29.77
DVAE-EMB	99.99	100.00	39.15
GCN-EMB	98.87	99.95	32.63

Table 1. Reconstruction accuracy, prior validity and uniqueness results (%) for the baselines and our method.

Moreover, in Figure 4, we observe a common pattern in the KL-divergence between the two models. In the first epochs, the encoder is quite simple therefore the posterior approximation $q_{\phi}(z|x)$ is close to the prior p(z). Consequently, the KL divergence has small values. During the optimization of the auto-encoder, the training of the encoder proceeds and the posterior approximation diverges from the prior. As a result, the KL divergence grows. After 100 epochs, when the reconstruction loss is close to zero for each model, the KL divergence starts decreasing because it is the only factor that affect the loss function.

4.2. Predictive performance of encoded latent representations

Next, we evaluate the representation power of the learned latent representations with respect to the performance of the generated neural network architectures. If we can accurately predict the performance based on the latent representations, then we can easily discover the best architectures from the latent space using a downstream strategy.

Following the experimentation setup in [36], we train a Sparse Gaussian Process (SGP) with 500 inducing points on the latent representations of the training data, in order to predict the accuracy of the test architectures. We use two evaluation metrics, the Root Mean Square Error (RMSE) between the Gaussian process predictions and the true performances, and the Pearson correlation coefficient (Pearson's r). Pearson correlation coefficient measures the linear correlation between the predictions and the true performances. Therefore, a model with a small RMSE and a high Pearson's r has strong predictive abilities. The experiments are repeated 10 times and we report the mean and the standard deviation.

We show the results in Table 4.2. The models incorporated with operation embeddings (DVAE-EMB and GCN-EMB) outperform the rest of the models in both metrics. This indicates that the latent spaces of DVAE-EMB and GCN-EMB are more suitable for searching highperformance neural architectures. Comparing DVAE-EMB with GCN-EMB, we observe that DVAE-EMB has the best performance due to its asynchronous message-passing scheme. However, GCN-EMB is significantly better than

Model	RMSE	Pearson's r
D-VAE	0.384 ± 0.002	0.920 ±0.001
GCN	0.485 ± 0.006	0.870 ± 0.001
S-VAE	0.478 ± 0.002	0.873 ± 0.001
GraphRNN	0.726 ± 0.002	0.669 ± 0.001
DVAE-EMB	0.371 ±0.003	0.925 ±0.001
GCN-EMB	0.441 ±0.002	0.892 ± 0.001

Table 2. Predictive performance of encoded latent representations(%)



Figure 5. Illustration of the top-5 performing architectures found by **DVAE-EMB**. The reported test accuracies (from left to right) are: 95.35%, 95.33%, 95.17%, 95.11%, 95.10%.



Figure 6. Illustration of the top 5 performing architectures found by **D-VAE**. The reported test accuracies (from left to right) are: 94.80%, 94.74%, 94.70%, 94.63%, 94.63%.

GCN, an outcome that highlights the contribution of the operation embeddings in learning predictive latent representations. S-VAE and GraphRNN, which leverage neither the graph structure nor the operation embeddings, present low predictive performance.

4.3. Best performing architectures obtained from Bayesian optimization (BO)

In this experiment, we perform Bayesian optimization in order to generate high-performing architectures, using DVAE-EMB and D-VAE. Following the methodology adopted by [36], we perform 10 iterations of batch Bayesian Optimization and we report the average accuracy results across 10 trials. We use expected improvement (EI) [19] as the acquisition function. Moreover, we use the SGP from the previous experiment to model the distribution of the objective function. In every iteration, we select a 50-sample batch by maximizing the acquisition function. For each batch, the selected latent representations are decoded into network architectures and are evaluated using their weightsharing accuracy on CIFAR-10. The decoded architectures are added to the training set and the SGP is retrained, initiating the next BO iteration. Finally, we select the 5 topperforming generated architectures and we fully train them on CIFAR-10 to evaluate their true performance, following the training procedure of [21].

In Figures 5, 6, we visualize the top-5 architectures discovered by DVAE-EMB and D-VAE and make the following observations:

- DVAE-EMB generate better architectures than D-VAE. Our best architecture achieves accuracy equal to 95.33%, while D-VAE's highest accuracy is 94.80%. This indicates that our proposed method leads to the construction of a very efficient latent space for searching neural network architectures. Moreover, our top-5 architectures achieve accuracy higher than 95%, therefore our model is able to learn not only a single highperforming architecture, but a distribution of such architectures.
- 2. The top architectures generated from DVAE-EMB present a smoother operation transition than those generated from D-VAE. Specifically, we observe that DVAE's architectures present a diversity of operations, in contrast with DVAE-EMB in which two operations (convolution 5x5 and separable convolution 5x5) are mostly used. Intuitively, DVAE-EMB learned the computational similarity of those two operations, and encoded the top-performing architectures in similar points on latent space.
- **3.** The top-performing architectures share the same structural patterns. This observation is supported by previous works that highlight the strong effect of the wiring patterns of the layers on the performance of the architecture [30, 33].

4.4. Architecture Performance and Graph Properties

The common graph structures that we observed in Figures 5 and 6 lead us to investigate which graph characteristics are the most informative about the performance of the architecture. For this reason, we monitored the structural patterns of the 19,020 architectures generated in the ENAS



Figure 7. Average path length with respect to model performance. Pearson's r = 0.32



Figure 8. Clustering coefficient with respect to model performance. Pearson's r = -0.39

benchmark [21] and we computed various graph metrics. Two of these metrics, a) the *clustering coefficient* and b) the *average path length* reveal a correlation with the architecture performance. Specifically, we clustered the architectures into six groups according to their performance and we measured their distributions with respect to the examined properties.

The results are visualized in Figures 7 and 8. We can observe that the mean average path length is increasing, whereas the clustering coefficient is decreasing, as we move to groups of architectures with higher performance. These findings are aligned with the corresponding Pearson correlation coefficients of the two metrics with the model performance. In particular, Pearson's r between the average path length and the performance is 0.32 indicating a positive correlation, while Pearson's r between the clustering coefficient and the performance is -0.39 indicating a negative correlation.

4.5. Latent Space Visualization

In this experiment, we compare the produced latent space of D-VAE with that of DVAE-EMB, projected in 2D space using t-SNE [27]. The visualizations are presented



Figure 9. 2-D Visualization of the latent space learned by **D-VAE** during training.



Figure 10. 2-D Visualization of the latent space learned by our model DVAE-EMB during training.

in Figures 9 and 10, where the weight sharing accuracy is color encoded.

In the first epoch (Figures 9(b) and 10(a)) the autoencoder is not yet trained and the representations do not form a continuous latent space. Therefore, multiple architectures are mapped to the same representation making the continuous optimization method not feasible. As the training proceeds, we observe that the architecture representations of both models span the whole latent space. This indicates that the auto-encoders are able to produce a 1 - 1 correspondence between the architectures and the latent representations.

Regarding the smoothness of the latent space, DVAE-EMB can accurately cluster together the high-performing architectures, as shown in Figure 10(c). This is the most important property in our application, as a smooth latent space can significantly enhance the performance of the search strategy. Note that the latent space was constructed in a fully-unsupervised manner, without having an accuracy signal of the architectures. Therefore, the smoothness is achieved by leveraging only the graph structure and the operation information of the architectures. In contrast, in the D-VAE's latent space the transition of accuracy is not smooth. The high performing architectures are located all over the latent space, without forming clusters. This indicates that our operation embeddings method benefits the process of mapping similar operations together and hence mapping similar performance architectures together.

5. Conclusion

Graph-based NAS methods have focused so far on encoding the structural properties of architectures, assuming a fixed representation of the performed operations. In this work, we introduce operation embeddings as a way to replace one-hot encodings of operations with learnable continuous representations that are incorporated into the optimization process. Our method enables the NAS framework to learn the computational and structural relationships of different operations, leading to a more accurate architecture latent space. The introduced approach has been evaluated on an exhaustive experimental study in ENAS benchmark, highlighting the effectiveness of operation embeddings. Our findings indicate that operation embeddings lead to shorter training time, smoother architecture representations and enhanced performance of various NAS models. We hope that the effectiveness and the flexibility of operation embeddings can motivate future studies to explore the representation power of the operation encoding.

References

- Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architectures using reinforcement learning. *CoRR*, abs/1611.02167, 2016. 2
- [2] Piotr Bojanowski, Armand Joulin, David Lopez-Pas, and Arthur Szlam. Optimizing the latent space of generative networks. In *Proceedings of the 35th International Conference* on Machine Learning, 2018. 1
- [3] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, 2016. 2, 4
- [4] Yi-Chen Chen, Jui-Yang Hsu, Cheng-Kuang Lee, and Hung yi Lee. Darts-asr: Differentiable architecture search for multilingual speech recognition and adaptation, 2020. 1
- [5] Kyunghyun Cho, Bart van Merrienboer, Çaglar Gülçehre, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoderdecoder for statistical machine translation. *CoRR*, abs/1406.1078, 2014. 3
- [6] F. Chollet. Xception: Deep learning with depthwise separable convolutions. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1800–1807, 2017. 5
- [7] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *Journal of Machine Learning Research*, 20(55):1–21, 2019. 2
- [8] Golnaz Ghiasi, Tsung-Yi Lin, Ruoming Pang, and Quoc V. Le. Nas-fpn: Learning scalable feature pyramid architecture for object detection. *CoRR*, 2019. 2
- [9] C. Jiang, H. Xu, W. Zhang, X. Liang, and Z. Li. Sp-nas: Serial-to-parallel backbone search for object detection. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 11860–11869, 2020. 1
- [10] Kirthevasan Kandasamy, Willie Neiswanger, Jeff Schneider, Barnabas Poczos, and Eric Xing. Neural architecture search with bayesian optimisation and optimal transport. *CoRR*, 2018. 2
- [11] Sungwoong Kim, Ildoo Kim, Sungbin Lim, Woonhyuk Baek, Chiheon Kim, Hyungjoo Cho, Boogeon Yoon, and Taesup Kim. Scalable neural architecture search for 3d medical image segmentation. *CoRR*, abs/1906.05956, 2019. 1
- [12] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In 2nd International Conference on Learning Representations, 2014. 3
- [13] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *CoRR*, 2017.
 4
- [14] Matt J. Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational autoencoder. *CoRR*, 2017. 4
- [15] Jian Li, Yong Liu, Jiankun Liu, and Weiping Wang. Neural architecture optimization with graph vae. CoRR, 2020. 2
- [16] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *CoRR*, 2018. 2

- [17] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural architecture optimization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. 2
- [18] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26, pages 3111–3119. Curran Associates, Inc., 2013. 4
- [19] J. Močkus. On bayesian methods for seeking the extremum. In G. I. Marchuk, editor, *Optimization Techniques IFIP Technical Conference Novosibirsk, July 1–7, 1974*, pages 400–404, Berlin, Heidelberg, 1975. Springer Berlin Heidelberg.
 7
- [20] Xuefei Ning, Yin Zheng, Tianchen Zhao, Yu Wang, and Huazhong Yang. A generic graph-based neural architecture encoding scheme for predictor-based nas. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 189–204, Cham, 2020. Springer International Publishing. 2
- [21] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *Proceedings of the 35th International Conference on Machine Learning*, 2018. 4, 5, 7
- [22] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. *CoRR*, 2019. 2
- [23] 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*, 2017. 2
- [24] Frederic Runge, Danny Stoll, Stefan Falkner, and Frank Hutter. Learning to design RNA. In *International Conference on Learning Representations*, 2019. 2
- [25] Martin Simonovsky and Nikos Komodakis. Graphvae: Towards generation of small graphs using variational autoencoders. *CoRR*, 2018. 3
- [26] Masanori Suganuma, Shinichi Shirakawa, and Tomoharu Nagao. A genetic programming approach to designing convolutional neural network architectures, 2017. 2
- [27] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008. 7
- [28] Colin White, Willie Neiswanger, Sam Nolen, and Yash Savani. A study on encodings for neural architecture search, 2020. 1
- [29] Colin White, Willie Neiswanger, and Yash Savani. BA-NANAS: bayesian optimization with neural architectures for neural architecture search. *CoRR*, abs/1910.11858, 2019. 1, 2
- [30] Mitchell Wortsman, Ali Farhadi, and Mohammad Rastegari. Discovering neural wirings. *CoRR*, 2019. 7

- [31] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. 2
- [32] Shen Yan, Yu Zheng, Wei Ao, Xiao Zeng, and Mi Zhang. Does unsupervised architecture representation learning help neural architecture search? *CoRR*, 2020. 1, 2, 3
- [33] Jiaxuan You, Jure Leskovec, Kaiming He, and Saining Xie. Graph structure of neural networks. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 10881–10891. PMLR, 13–18 Jul 2020. 2, 7
- [34] Jiaxuan You, Rex Ying, Xiang Ren, William L. Hamilton, and Jure Leskovec. Graphrnn: Generating realistic graphs with deep auto-regressive models. *CoRR*, 2018. 2, 4
- [35] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan Salakhutdinov, and Alexander J. Smola. Deep sets. *CoRR*, abs/1703.06114, 2017. 2
- [36] Muhan Zhang, Shali Jiang, Zhicheng Cui, Roman Garnett, and Yixin Chen. D-vae: A variational autoencoder for directed acyclic graphs. volume 32, 2019. 1, 2, 3, 4, 5, 6, 7
- [37] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. *CoRR*, 2017. 2
- [38] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. *CoRR*, 2018. 2