

A General-Purpose Transferable Predictor for Neural Architecture Search

Fred X. Han^{1*}, Keith G. Mills², Fabian Chudak¹, Parsa Riahi³
Mohammad Salameh¹, Jialin Zhang⁴, Wei Lu¹, Shangling Jui⁴, Di Niu²

¹Huawei Technologies Canada, Edmonton,

² University of Alberta, ³ University of British Columbia,

⁴ Huawei Kirin Solution, Shanghai

Abstract

Understanding and modelling the performance of neural architectures is key to Neural Architecture Search (NAS). Performance predictors have seen widespread use in low-cost NAS and achieve high ranking correlations between predicted and ground truth performance in several NAS benchmarks. However, existing predictors are often designed based on network encodings specific to a predefined search space and are therefore not generalizable to other search spaces or new architecture families. In this paper, we propose a general-purpose neural predictor for NAS that can transfer across search spaces, by representing any given candidate Convolutional Neural Network (CNN) with a Computation Graph (CG) that consists of primitive operators. We further combine our CG network representation with Contrastive Learning (CL) and propose a graph representation learning procedure that leverages the structural information of unlabeled architectures from multiple families to train CG embeddings for our performance predictor. Experimental results on NAS-Bench-101, 201 and 301 demonstrate the efficacy of our scheme as we achieve strong positive Spearman Rank Correlation Coefficient (SRCC) on every search space, outperforming several Zero-Cost Proxies, including Synflow and Jacov, which are also generalizable predictors across search spaces. Moreover, when using our proposed general-purpose predictor in an evolutionary neural architecture search algorithm, we can find high-performance architectures on NAS-Bench-101 and find a MobileNetV3 architecture that attains 79.2% top-1 accuracy on ImageNet.

1 Introduction

Neural Architecture Search (NAS) automates neural network design and has achieved remarkable performance on many computer vision tasks. A NAS strategy typically performs alternated search and evaluation over

candidate networks to maximize a performance metric. While various strategies such as Random Search [14], Differentiable Architecture Search [15], Bayesian Optimization [30], and Reinforcement Learning [21] can be used for search, architecture evaluation is a key bottleneck to identifying better architectures.

To avoid the excessive cost incurred by training and evaluating each candidate network, most current NAS frameworks resort to performance estimation methods to predict accuracy. Popular methods include weight sharing [21, 15, 3], neural predictors [16, 30] and Zero-Cost Proxies (ZCP) [2]. The effectiveness of a performance estimation method is mainly determined by the Spearman Rank Correlation Coefficient (SRCC) between predicted performance and the ground truth. A predictor with higher SRCC can better guide a NAS search algorithm toward finding superior architectures.

While partial training and weight sharing are used extensively in early NAS works [21, 34], thanks to several existing NAS benchmarks that provide an ample amount of labeled networks [7, 33], e.g., NAS-Bench-101 [32] offers 423k networks trained on CIFAR-10, there have been many recent developments in training neural predictors for NAS [29, 25]. As these methods learn to *estimate* performance using labeled architecture representations, they generally enjoy the lowest performance evaluation cost as well as the capacity for continual improvement as more NAS benchmarking data is made available.

However, a major shortcoming of existing neural predictors is that they are not general-purpose. Each predictor is specialized to process networks confined to a specific search space. For example, NAS-Bench-101 limits its search space to a cell, which is a graph of up to 7 internal operators. Each operator can be one of 3 specific operation sequences. In contrast, NAS-Bench-201 [7] and NAS-Bench-301 [33] adopt different candidate operator sets (with more details in Sec. 3) and network topologies. Thus, a neural predictor for NAS-Bench-101 could not predict the performance for NAS-

*Correspondence to: fred.xuefei.han1@huawei.com

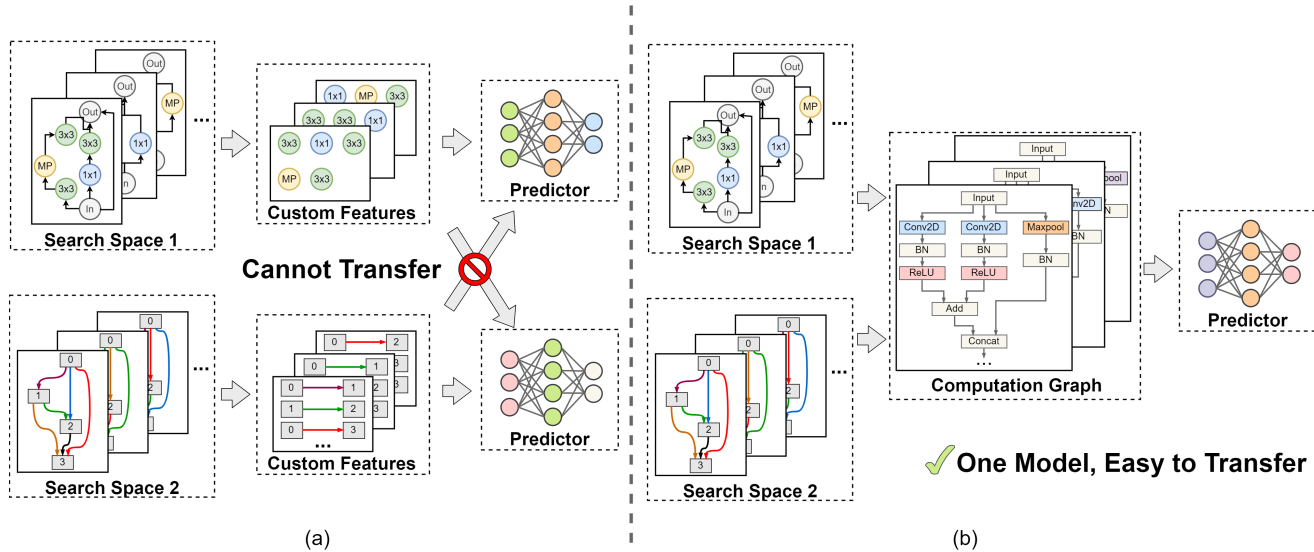


Figure 1: Comparison of architecture representations in neural predictor setups: (a) conventional, space-dependent neural predictors, e.g., BANANAS and SemiNAS; (b) our proposed, general-purpose, space-agnostic predictor using Computational Graphs. Best viewed in color.

Bench-201 or NAS-Bench-301 networks. This search-space-specific design severely limits the practicality and transferability of existing predictors for NAS, since for any new search space that may be adopted in reality, a separate predictor must be re-designed and re-trained based on a large number of labeled networks in the new search space.

Another emerging approach to performance estimation is ZCP methods, which can support any network structure as input. However, the performance of ZCPs may vary significantly depending on the search space. For example, [2] shows that Synflow achieves a high SRCC of 0.74 on the NAS-Bench-201 search space, however the SRCC drops to 0.37 on NAS-Bench-101. Moreover, ZCP methods require the instantiation of a neural network, e.g., performing the forward pass for a batch of training samples, to compute gradient information, and thus incur longer per-network prediction latency. In contrast, neural predictors simply require the network architecture encoding or representation to predict its performance.

In general, neural predictors offer better estimation quality but do not transfer across search spaces, while ZCPs are naturally universal estimation methods, yet are sub-optimal on certain search spaces. To overcome this dilemma, in this paper, we propose a general-purpose predictor for NAS that is transferable across search spaces like ZCPs, while still preserving the benefits of conventional predictors. Our contributions are summarized as follows:

First, we propose the use of *Computation Graphs* to

offer a universal representation of neural architectures from different search spaces. Figure 1 highlights the differences between (a) existing NAS predictors and (b) the proposed framework. The key is to introduce a universal search space representation consisting of graphs of only primitive operators to model any network structure, such that a general-purpose transferable predictor can be learned based on NAS benchmarks available from multiple search spaces.

Second, we propose a framework to learn a *generalizable* neural predictor by combining recent advances in Graph Neural Networks (GNN) [19] and Contrastive Learning (CL) [4]. Specifically, we introduce a graph representation learning process to learn a generalizable neural architecture encoder based on the structural information of vast unlabeled networks in the target search space. The embeddings obtained this way are then fed into a neural predictor, which is trained based on labeled architectures in the source families, achieving transferability to the target search space.

Experimental results on NAS-Bench-101, 201 and 301 show that our predictor can obtain high SRCC on each search space when fine-tuning on no more than 50 labeled architectures in a given target family. Specifically, we outperform several ZCP methods like Synflow and obtain SRCC of 0.917 and 0.892 on NAS-Bench-201 and NAS-Bench-301, respectively. Moreover, we use our predictor for NAS with a simple evolutionary search algorithm and have found a high-performance architecture (with 94.23% accuracy) in NAS-Bench-101 at a low cost of 700 queries to the benchmark,

outperforming other *non-transferable* neural predictors such as SemiNAS [16] and BANANAS [30]. Finally, we further apply our scheme to search for an ImageNet [6] network and have found a Once-for-All MobileNetV3 (OFA-MBv3) architecture that obtains a top-1 accuracy of 79.2% which outperforms the original OFA.

2 Related Work

Neural predictors are a popular choice for performance estimation in low-cost NAS. Existing predictor-based NAS works include SemiNAS [16], which adopts an encoder-decoder setup for architecture encoding/generation, and a simple neural performance predictor that predicts based on the encoder outputs. SemiNAS progressively updates an accuracy predictor during the search. BANANAS [30] relies on an ensemble of accuracy predictors as the inference model in its Bayesian Optimization process. [25] construct a similar auto-encoder-based predictor framework and supplies additional unlabeled networks to the encoder to achieve semi-supervised learning. [29] propose a sample-efficient search procedure with customized predictor designs for NAS-Bench-101 and ImageNet [6] search spaces. NPE-NAS [28], BRP-NAS [8] are also notable predictor-based NAS approaches. By contrast, our approach pre-trains architecture embeddings that are not restricted to a specific underlying search space.

Zero-Cost Proxies (ZCP) are originally proposed as parameter saliency metrics in model pruning techniques. With the recent advances of pruning-at-initialization algorithms, only a single forward/backward propagation pass is needed to assess the saliency. Metrics used in these algorithms are becoming an emerging trend for transferable, low-cost performance estimation in NAS. [2] transfers several ZCP to NAS, such as Synflow [24], Snip [13], Grasp [27] and Fisher [26]. Zero-Cost Proxies could work on any search space and [2] shows that on certain search spaces, they help to achieve results that are comparable to predictor-based NAS algorithms. However, the performance of ZCPs are generally not consistent across different search spaces. Unlike neural predictors such as ours, they require instantiation of candidate networks as well as sample data to compute gradient metrics.

3 A Unified Architecture Representation

In this section, we discuss how to represent neural networks. First, we elaborate on *operator-grouping* within NAS, how it simplifies architecture representation while hindering transferability between search spaces. Then, we introduce our Computational Graph (CG) framework and how it solves the transferability problem.

3.1 Abstract Operation Representations

Without the loss of generality, we consider operations in neural networks and define a *primitive* operator as an atomic node of computation. In other words, a primitive operator is one like Convolution, ReLU, Add, Pooling or Batch Normalization (BN), which are single points of execution that cannot be further reduced.

Operator-grouping is an implicit but widely adopted technique for improving the efficiency of NAS. It is a coarse abstraction where atomic operations are grouped and re-labeled according to pre-defined sequences. For example, convolution operations are typically grouped with BN and ReLU [3]. Rather than explicitly representing all three primitive operations independently, a simpler representation describes all three primitives as one grouped operation pattern that is consistent throughout a search space. We can then use these representations as feature inputs to a search algorithm [22, 16] or neural predictor [29, 8], as prior methods do.

While operator grouping provides a useful method to abstract how we represent neural networks, it hinders transferability as groupings may not be consistent across search spaces. Take the aforementioned convolution example: Although convolutions are typically paired with BN and ReLU operations, the ordering of these primitives can vary. While NAS-Bench-101 use the ‘Conv-BN-ReLU’ ordering, NAS-Bench-201 use ‘ReLU-Conv-BN’, thus forming a discrepancy that current neural predictors do not take into account.

Compounding this issue, the set of operations can differ by search space. While NAS-Bench-101 considers Max Pooling, NAS-Bench-201 only uses Average Pooling, and NAS-Bench-301 uses complex convolutions with larger kernel sizes not found in either 101 nor 201. A neural predictor trained on any one of these search spaces using operator-grouping could not easily transfer to another. Thus, operator-grouping is the main culprit of low transferability in existing predictors. We provide a table enumerating the operator groupings for each search space in the supplementary materials.

3.2 Computational Graphs

To construct a neural predictor that is transferable between search spaces, we consider a representation that can *generalize* across multiple search spaces. We define a Computation Graph (CG) as a detailed representation of a neural network without any customized grouping, i.e., each node in the graph is a primitive operator and the network CG is made up of only such nodes and edges that direct the flow of information. Without operator-grouping, CGs define a search space that could represent any network structure since the number of primitive operators is usually far less than the number of possible groupings.

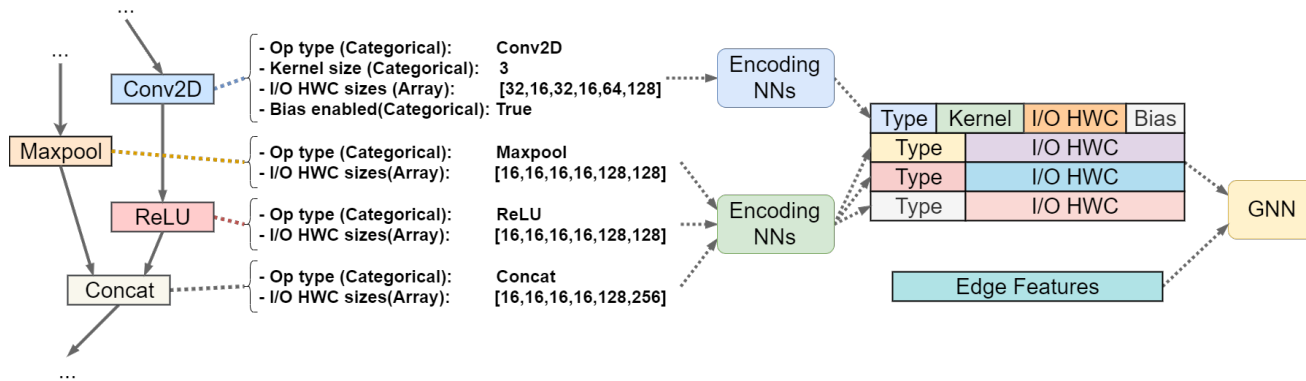


Figure 2: An example illustrating the key graphical features extracted from compute graphs and how we encode them as GNN node features. All nodes contain input and output tensor HWC sizes. Nodes with trainable weights contain additional features on the weight matrix dimensions and bias.

While there are many potential ways to construct a CG, we adopt a simple approach of using the model optimization graph maintained by deep learning frameworks like TensorFlow [1] or PyTorch [20]. As the name suggests, a model optimization graph is originally intended for gradient calculations and weight updates, and it is capable of supporting any network structure. In this work, all the CGs used in our experiments are simplified from TensorFlow model optimization graphs. Specifically, we extract the following nodes from a model optimization graph to form a CG:

- Nodes that refer to trainable neural network weights. For these nodes, we extract the atomic operator type, such as Conv1D, Conv2D, Linear, etc., input/output channel sizes, input/output image height/width sizes, weight matrix dimensions (e.g., convolution kernel tensor) and bias information as node features.
- Nodes that refer to activation functions like ReLU or Sigmoid, pooling layers like Max or Average, as well as Batch Normalization. For these nodes, we extract the operator type, input/output channel and image height/width sizes.
- Key supplementary nodes that indicate how information is processed, e.g., addition, concatenation and element-wise multiplication. For these nodes, we also extract the type, input/output channel and image height/width sizes.

Figure 2 illustrates how we transform a computation graph into learnable feature vectors. Formally, a computation graph G consists of a vertex set $V = \{v_1, v_2, v_3, \dots\}$ and an edge set E , where v refers to a

primitive operator node and E contains pairs of vertices (v_s, v_d) indicating a connection between v_s and v_d . Under this definition, the problem of performance estimation becomes finding a function F , e.g., a Graph Neural Network (GNN), such that for computation graph G_i , which is generated from a candidate neural network, $F(G_i) = Y_i$, where Y_i is the ground truth test score.

Representing networks as CGs enable us to break the barrier imposed by search space definitions and fully utilize all available data for predictor training, regardless of where a labeled network is from, e.g., NAS-Bench-101, 201 or 301. We could also effortlessly transfer a predictor trained on one search space to another, by simply fine-tuning it on additional data from the target space.

4 Neural Predictor via Graph Representation Learning

In this section, we propose a two-stage approach to improve generalization and leverage unlabeled data via Contrastive Learning (CL). We first find a vector representation, i.e., graph embedding, which converts graph features with a variable number of nodes into a fixed-sized latent vector via a graph CL procedure, before feeding the latent vector to an MLP accuracy regressor. Given a target family for performance estimation, a salient advantage of our approach is its ability to leverage unlabeled data, e.g., computation graphs of the target family, which are typically available in abundance. In fact, our approach is able to jointly leverage labeled and unlabeled architectures from multiple search spaces to maximally utilize available information.

For each CG G_i , we would like to produce a vector representation $h_i \in \mathbb{R}^e$ for a fixed hyper-parameter e .

We would like to infer relationships between networks by considering the angles between vector representations. Our CL-based approach learns representations where only *similar* CGs have close vector representations.

4.1 Contrastive Learning Frameworks SimCLR [4, 5] and SupCon [12] apply CL to image classification. The general idea is to learn a base encoder to create vector representations h of images. To train the base encoder, a projection head $Proj(\ast)$ maps the vector representations into a lower-dimensional space $z \in \mathbb{R}^p; p < e$ to optimize a *contrastive loss*.

The contrastive loss forces representations of similar objects to be in agreement. Consider a batch of N images, whose vector representation is $I = \{h_1, h_2, \dots, h_N\} \subset \mathbb{R}^e$. For each G_i , let $z_i = Proj(h_i) \in \{|z| = 1 : z \in \mathbb{R}^p\}$, and let the cosine similarity be $sim(z_i, z_j) = z_i \cdot z_j / \tau$, where the temperature $\tau > 0$ and \cdot is the dot product. The agreement $\chi_{i,j}$ between two arbitrary indices i and j is given by

$$(4.1) \quad \chi_{i,j} = \log \frac{\exp(sim(z_i, z_j))}{\sum_{r \neq i} \exp(sim(z_i, z_r))}.$$

A primary distinction between SimCLR and SupCon is how we determine if two different objects are similar.

SimCLR considers the unsupervised context where we do not have access to label/class information. Data augmentation plays a crucial role. For each *anchor* index i in a batch, we apply a transform or slight perturbation to create an associated *positive* element $j(i)$, while all other samples $r \neq i, j(i)$ are *negative* indices. The SimCLR loss function,

$$(4.2) \quad \mathcal{L}_{SimCLR} = - \sum_{i \in I} \chi_{i,j(i)},$$

serves to maximize agreement between the original and augmented images.

By contrast, if classes are known, we can use the SupCon loss function [12],

$$(4.3) \quad \mathcal{L}_{SupCon} = \sum_{i \in I} \frac{-1}{|P(i)|} \sum_{s \in P(i)} \chi_{i,s},$$

where $P(i)$ is the set of non-anchor indices whose class is the same as i . Thus, not just $j(i)$, but all indices whose class is the same as i contribute to the probability of positive pairs.

Next, we construct a contrastive loss for CGs. We consider the challenges and advantages of using graphs as data as well as the overall problem of regression instead of classification.

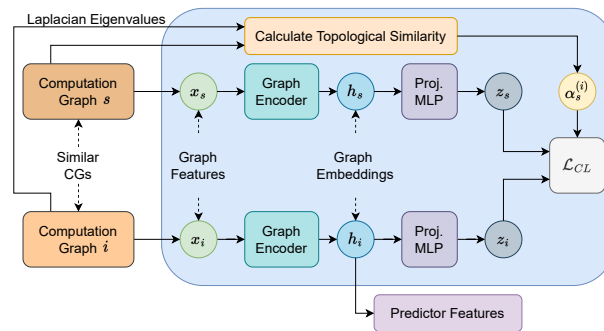


Figure 3: Contrastive Learning framework. We apply a graph encoder to produce embeddings. Next, a projection layer produces representation vectors whose similarity we compare using the CL loss, weighted according to the Laplacian Eigenvalues of each CG.

4.2 Computational Graph Encodings Figure 3 provides a high-level overview of our scheme. When applying CL to CGs, we start by considering the similarity between CGs. We leverage the rich structural information CGs provide by encoding each atomic primitive within a network as a node. Specifically, our approach uses spectral properties of undirected graphs [9]. Given a CG with $|V|$ nodes, we consider its underlying undirected graph G' . Let $A \in \{0, 1\}^{|V| \times |V|}$ be its adjacency matrix and $D \in \mathbb{Z}^{|V| \times |V|}$ be its degree diagonal matrix. The normalized Laplacian matrix is defined as

$$(4.4) \quad \Delta = I - D^{-1/2} A D^{-1/2} = U^T \Lambda U,$$

where $\Lambda \in \mathbb{R}^{|V| \times |V|}$ is the diagonal matrix of eigenvalues and U the matrix of eigenvectors. Eigenvalues Λ encode important connectivity features. For instance, 0 is the smallest eigenvalue and has multiplicity 1 if and only if G' is connected. Smaller eigenvalues focus on general features of the graph, whereas larger eigenvalues focus on features at higher granularity; we refer the reader to [31] for more details.

More generally, we can use the eigenvalues of Δ to measure *pseudo-distance* between graphs. Given two CGs g_1, g_2 , we compute the spectral distance $\sigma_S(g_1, g_2)$ as the Euclidean norm of the corresponding $k = 11$ smallest eigenvalues.

Our contrastive loss incorporates elements from both SimCLR and SupCon in addition to spectral distance. As our task of interest is regression rather than classification, we replace the positive-negative binary relationship between samples in a batch with a probability distribution over all pairs which smoothly favors similar computation graphs.

Table 1: Spearman correlation coefficients (ρ). We compare our CL encoding scheme to a GNN encoder as well as several ZCPs. For our CL and GNN, we report the mean and standard deviation over 5 fine-tuning runs.

Method	NAS-Bench-101	NAS-Bench-201	NAS-Bench-301
Synflow [24]	0.361	0.823	-0.210
Jacov [17]	0.358	0.859	-0.190
Fisher [26]	-0.277	0.687	-0.305
GradNorm [2]	-0.256	0.714	-0.339
Grasp [27]	0.245	0.637	-0.055
Snip [13]	-0.165	0.718	-0.336
GNN-fine-tune	0.542 \pm 0.14	0.884 \pm 0.03	0.872 \pm 0.01
CL-fine-tune	0.553 \pm 0.09	0.917 \pm 0.01	0.892 \pm 0.01

First, if the family of networks each CG belongs to is known, e.g., NAS-Bench-101, 201, etc., we can treat the family affiliation as a class and follow the SupCon approach. Second, rather than using the uniform distribution over $P(i)$ as in Equation 4.3, we use a convex combination over $P(i)$ based on the similarity of the corresponding computation graphs. Overall, our loss function is given as

$$(4.5) \quad \mathcal{L}_{CL} = - \sum_{i \in I} \sum_{s \in P(i)} \alpha_s^{(i)} \chi_{i,s},$$

where $\alpha_s^{(i)} \geq 0$ and $\sum_{s \in P(i)} \alpha_s^{(i)} = 1$. For computation graph i , we simply define $\alpha_*^{(i)}$ to be the softmax of $\sigma_S(i, *)$ with temperature 0.05.

Finally, there are challenges associated with data augmentation for CGs. Slightly perturbing a CG may drastically change its accuracy on a benchmark, e.g., changing an activation function or convolution [18]. In more severe scenarios, arbitrary small changes to a computation graph may make it fall outside of the family of networks of interest or even result in a graph that does not represent a functional neural network at all. To address this, rather than randomly perturbing a CG, we use σ to randomly pick a very similar graph from the training set to form a positive pair. As suggested in [12], we learn the embeddings using large batch sizes. We enumerate the structure of our predictor and other details in the supplementary materials.

5 Experimentation

In this section, we evaluate our proposed transferable predictor by comparing ranking correlations with other transferable Zero-Cost Proxies (ZCP) on popular NAS benchmarks. Then, we compare with other neural predictors in the literature by performing search, showing that methods with higher ranking correlations often produce better results.

5.1 Comparison of ranking correlations We consider the search spaces of NAS-Bench-101 [32], NAS-

Bench-201 [7] and NAS-Bench-301 [33], with 50k, 4096¹ and 10k overall CG samples, respectively. We differentiate between *target* and *source* families depending on our configuration. We treat target families as unseen test domains and assume we only have access to a limited amount of labeled target data, yet a large amount of unlabeled target CGs. We use source families to train our predictors and we assume labels are known for each CG.

We consider the three cases where one of NAS-Bench-101, 201 or 301 is the held-out target family, and use the other two as source families. We use structural information from unlabeled samples in the target family to train our CL encoder in an unsupervised manner. Then, we use labeled data from the source families to train an MLP predictor using supervised regression. Finally, we use a small amount of labeled data from the target family to fine-tune the MLP predictor.

In addition to our CL-based encoder and ZCPs, we consider a simple GNN [19] regressor baseline that we can train and fine-tune in an end-to-end fashion. We provide implementation details for each predictor in the supplementary materials.

We sample 5k instances from NAS-Bench-101, 4096 instances from NAS-Bench-201 and 1k instances from NAS-Bench-301 when they are the target test family, and use all available data when they are a source family. When fine-tuning, we use 50 CGs from NAS-Bench-101 and NAS-Bench-301 and 40 CGs if NAS-Bench-201 is the target family. We execute the ZCPs on the test sets and report the Spearman Rank Correlation Coefficient (SRCC) ($\rho \in [-1, 1]$). SRCC values closer to 1 indicate higher ranking correlations between predictions and the ground truth accuracy.

Table 1 summarizes the results. First, we note that our CL-based encoder achieves the best SRCC in all three target family scenarios. On NAS-Bench-101, only our CG-based schemes can achieve SRCC above 0.5, while some ZCPs fail to even achieve positive

¹CIFAR-10 architectures that do not contain ‘zeroize’.

Table 2: Search results on NAS-Bench-101, 201 and 301 using the same EA search algorithm but with different performance estimation methods. #Q represents the number of unique networks queried during search. Note that the #Q for CL-fine-tune also counts the fine-tuning instances.

Method	NAS-Bench-101			NAS-Bench-201			NAS-Bench-301	
	#Q	Acc. (%)	Rank	#Q	Acc. (%)	Rank	#Q	Acc. (%)
Random	700	94.11 \pm 0.10	26.0	90	93.91 \pm 0.2	104	800	94.75 \pm 0.08
Synflow	700	94.18 \pm 0.05	5.8	90	94.37 \pm 0.0	1.0	800	94.60 \pm 0.11
CL-fine-tune	700	94.23 \pm 0.01	2.2	90	94.37 \pm 0.0	1.0	800	94.83 \pm 0.06

correlation. On NAS-Bench-201, only our CL scheme is able to achieve over 0.9 SRCC while the GNN baseline and the best ZCP methods, Synflow and Jacov, achieve around 0.85 SRCC. Similar to [2], on NAS-Bench-301, all of the ZCP schemes fail to achieve positive ranking correlation. By contrast, both of our predictors achieve over 0.85 SRCC on 301. Moreover, the CL encoder with fine-tuning achieves very low standard deviation on all three benchmarks. This indicates stable, consistent performance. Overall, our findings demonstrate the utility of CGs as generalizable, robust neural network representations and the capacity of our CL scheme to learn rich graph features through unlabeled data.

5.2 Search results We now demonstrate that our predictor is a superior choice for transferable performance estimation in NAS. We vary the performance estimator and report the accuracy and rank of the best architecture found. Specifically, we compare the CL-fine-tune predictor for a given target family against the Synflow ZCP and a random estimation baseline. We pair each predictor with a simple evolutionary approach that we detail in the supplementary materials.

For each method, we conduct 5 search runs on NAS-Bench-101 which has 423,624 labeled candidates, NAS-Bench-201 with 15,625 searchable candidates, and NAS-Bench-301 with over 10^{18} candidates. To establish a fair comparison, we query (#Q) the same number of networks. Intuitively, the number of queries made to the benchmark simulates the real-world cost of performance evaluation in NAS.

Table 2 reports our results. On NAS-Bench-101 and 301, our CL-fine-tune predictors consistently find better architectures than either Synflow or the random estimation baseline. On average, our CL predictor can find the 2nd best NAS-Bench-101 architecture while Synflow finds the 6th best, and does not come close in terms of performance on NAS-Bench-301. On the smallest search space, NAS-Bench-201, both our CL-fine-tune predictor and Synflow easily and consistently find the best CIFAR-10 architecture. We recall from Table 1 how none of the ZCP methods could achieve positive

Table 3: Search performance of our CL-fine-tuning predictor and EA search algorithm against other NAS approaches on NAS-Bench-101. We report the best architecture accuracy and number of queries.

NAS algorithm	#Queries	Best Acc.
Random Search	2000	93.66%
SemiNAS [16]	2000	94.02%
SemiNAS (RE) [16]	2000	94.03%
SemiNAS (RE) [16]	1000	93.97%
BANANAS [30]	800	94.23%
GA-NAS [22]	378	94.23%
Neural-Predictor-NAS [29]	256	94.17%
NPENAS [28]	150	94.14%
BRP-NAS [8]	140	94.22%
EA + CL-fine-tune	700	94.23%

SRCC on NAS-Bench-301 and note how poor estimation performance reflects downstream search results as a random estimation baseline outperforms Synflow on that search space. By contrast, with a small amount of fine-tuning data, our CL-fine-tune predictor achieves exceptional transferability across architecture families and is able to better guide the search process.

Next, we compare our best search results on NAS-Bench-101 to other state-of-art NAS approaches in Table 3. We observe that our EA + CL-fine-tune setup is competitive with other NAS algorithms. For instance, our setup requires fewer queries to find the second-best architecture (94.23%) in NAS-Bench-101 compared to BANANAS. The only other scheme which achieves that level of performance is our previous work, GA-NAS [22]. Moreover, our search result is better than SemiNAS in terms of the number of queries and accuracy. It is also critical to note that, due to the generalizable structure of our CGs, we can transfer our predictor to other search spaces with only a small amount of labeled data for fine-tuning. This is a unique advantage among neural predictors.

Finally, we further test the efficacy of our predictor in NAS by searching on the large-scale classification dataset ImageNet [6]. To reduce carbon footprint, we

Table 4: ImageNet search results. We search in the OFA-MBv3 space and report the top-1 accuracy acquired using the OFA-MBv3 supernet. Most baseline results are provided by [3].

Model	Top-1 Acc. (%)
MobileNetV2 [23]	72.0
MobileNetV2 #1200 [23]	73.5
MobileNetV3-Large [11]	75.2
NASNet-A [34]	74.0
DARTS [15]	73.1
SinglePathNAS [10]	74.7
OFA-Large [3]	79.0
OFA-Base [18]	78.9
OFA-MBv3-CL	79.2

search on the Once-for-All (OFA) [3] design space for MobileNetV3 (MBv3) [11], which allows us to quickly query the accuracy of a found network using a pre-trained supernet. We train a CL encoder and predictor on NAS-Bench-101, 201 and 301, then fine-tune on 50 random networks from the OFA-MBv3 search space.

Table 4 reports our findings. Using the CL predictor, we are able to find an architecture that achieves over 79.0% top-1 accuracy on ImageNet, outperforming works on the same search space such as the original OFA [3] and [18]. These results further reinforce the generalizability of our scheme, as we are able to train a neural predictor on three CIFAR-10 benchmark families, then then transfer it to perform NAS on a MobileNet family designed for ImageNet.

6 Conclusion

In this work, we propose the use of Computational Graphs (CG) as a universal representation for CNN structures. On top of this representation, we design a novel, well-performing, transferable neural predictor that incorporates Contrastive Learning to learn robust embeddings for a downstream performance predictor. Our transferable predictor alleviates the need to manually design and re-train new performance predictors for any new search spaces in the future, which helps to further reduce the computational cost and carbon footprint of NAS. Experiments on the NAS-Bench-101, 201 and 301 search spaces verify our claims. Results show that our predictor is superior to many Zero-Cost Proxy methods on these search spaces. By pairing our predictor with an evolutionary search algorithm we can find a NAS-Bench-101 architecture that obtains 94.23% CIFAR-10 accuracy and a MobileNetV3 architecture that attains 79.2% ImageNet top-1 accuracy.

References

- [1] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. 2015. URL: <https://www.tensorflow.org/>.
- [2] Mohamed S. Abdelfattah et al. “Zero-Cost Proxies for Lightweight NAS”. In: *International Conference on Learning Representations (ICLR)*. 2021.
- [3] Han Cai et al. “Once for All: Train One Network and Specialize it for Efficient Deployment”. In: *International Conference on Learning Representations*. 2020. URL: <https://openreview.net/forum?id=HylxE1HKwS>.
- [4] Ting Chen et al. “A Simple Framework for Contrastive Learning of Visual Representations”. In: *International Conference on Machine Learning*. PMLR. 2020, pp. 1597–1607.
- [5] Ting Chen et al. “Big Self-Supervised Models are Strong Semi-supervised Learners”. In: *Advances in Neural Information Processing Systems 33* (2020), pp. 22243–22255.
- [6] Jia Deng et al. “ImageNet: A Large-Scale Hierarchical Image Database”. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE. 2009, pp. 248–255.
- [7] Xuanyi Dong and Yi Yang. “NAS-Bench-201: Extending the Scope of Reproducible Neural Architecture Search”. In: *International Conference on Learning Representations (ICLR)*. 2020. URL: <https://openreview.net/forum?id=HJxyZkBKDr>.
- [8] Lukasz Dudziak et al. “BRP-NAS: Prediction-based NAS Using GCNs”. In: *Advances in Neural Information Processing Systems 33* (2020), pp. 10480–10490.
- [9] Vijay Prakash Dwivedi and Xavier Bresson. *A Generalization of Transformer Networks to Graphs*. 2021. arXiv: 2012.09699 [cs.LG].
- [10] Zichao Guo et al. “Single Path One-shot Neural Architecture Search with Uniform Sampling”. In: *European Conference on Computer Vision*. Springer. 2020, pp. 544–560.
- [11] Andrew Howard et al. “Searching for MobileNetv3”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, pp. 1314–1324.
- [12] Prannay Khosla et al. “Supervised Contrastive Learning”. In: *Advances in Neural Information Processing Systems 33* (2020), pp. 18661–18673.

- [13] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. “Snip: Single-shot Network Pruning Based on Connection Sensitivity”. In: *arXiv preprint arXiv:1810.02340* (2018).
- [14] Liam Li and Ameet Talwalkar. “Random Search and Reproducibility for Neural Architecture Search”. In: *Uncertainty in Artificial Intelligence*. PMLR. 2020, pp. 367–377.
- [15] Hanxiao Liu, Karen Simonyan, and Yiming Yang. “DARTS: Differentiable Architecture Search”. In: *International Conference on Learning Representations (ICLR)*. 2019. URL: <http://arxiv.org/abs/1806.09055>.
- [16] Renqian Luo et al. “Semi-Supervised Neural Architecture Search”. In: *Advances in Neural Information Processing Systems 33* (2020), pp. 10547–10557.
- [17] Joe Mellor et al. “Neural Architecture Search without Training”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 7588–7598.
- [18] Keith G Mills et al. “Profiling Neural Blocks and Design Spaces for Mobile Neural Architecture Search”. In: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 2021, pp. 4026–4035.
- [19] Christopher Morris et al. “Weisfeiler and Leman Go Neural: Higher-Order Graph Neural Networks”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 2019, pp. 4602–4609.
- [20] Adam Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach et al. Curran Associates, Inc., 2019, pp. 8024–8035.
- [21] Hieu Pham et al. “Efficient Neural Architecture Search via Parameters Sharing”. In: *International Conference on Machine Learning*. PMLR. 2018, pp. 4095–4104.
- [22] Seyed Saeed Changiz Rezaei et al. “Generative Adversarial Neural Architecture Search”. In: *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*. Main Track. International Joint Conferences on Artificial Intelligence Organization, Aug. 2021, pp. 2227–2234. DOI: 10.24963/ijcai.2021/307.
- [23] Mark Sandler et al. “MobileNetV2: Inverted Residuals and Linear Bottlenecks”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018, pp. 4510–4520.
- [24] Hidenori Tanaka et al. “Pruning Neural Networks Without Any Data by Iteratively Conserving Synaptic Flow”. In: *Advances in Neural Information Processing Systems 33* (2020), pp. 6377–6389.
- [25] Yehui Tang et al. “A Semi-Supervised Assessor of Neural Architectures”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 1810–1819.
- [26] Jack Turner et al. “Blockswap: Fisher-Guided Block Substitution for Network Compression on a Budget”. In: *arXiv preprint arXiv:1906.04113* (2019).
- [27] Chaoqi Wang, Guodong Zhang, and Roger Grosse. “Picking Winning Tickets Before Training by Preserving Gradient Flow”. In: *arXiv preprint arXiv:2002.07376* (2020).
- [28] Chen Wei et al. “NPENAS: Neural Predictor Guided Evolution for Neural Architecture Search”. In: *IEEE Transactions on Neural Networks and Learning Systems* (2022).
- [29] Wei Wen et al. “Neural predictor for neural architecture search”. In: *European Conference on Computer Vision*. Springer. 2020, pp. 660–676.
- [30] Colin White, Willie Neiswanger, and Yash Savani. “BANANAS: Bayesian Optimization with Neural Architectures for Neural Architecture Search”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 12. 2021, pp. 10293–10301.
- [31] Peter Wills and François G Meyer. “Metrics for Graph Comparison: A Practitioner’s Guide”. In: *Plos one* 15.2 (2020), e0228728.
- [32] Chris Ying et al. “NAS-Bench-101: Towards Reproducible Neural Architecture Search”. In: *International Conference on Machine Learning*. PMLR. 2019, pp. 7105–7114.
- [33] Arber Zela et al. “Surrogate NAS Benchmarks: Going Beyond the Limited Search Spaces of Tabular NAS Benchmarks”. In: *International Conference on Learning Representations*. 2022. URL: <https://openreview.net/forum?id=OnpFa95RVqs>.
- [34] Barret Zoph et al. “Learning Transferable Architectures for Scalable Image Recognition”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018, pp. 8697–8710.