

NN-Former: Rethinking Graph Structure in Neural Architecture Representation

Ruihan Xu¹, Haokui Zhang², Yaowei Wang³, Wei Zeng¹, Shiliang Zhang^{†1}

¹ State Key Laboratory of Multimedia Information Processing,
School of Computer Science, Peking University, China

²Northwestern Polytechnical University, Xi'an, China

³Harbin Institute of Technology (Shenzhen), Shenzhen, China

Abstract

The growing use of deep learning necessitates efficient network design and deployment, making neural predictors vital for estimating attributes such as accuracy and latency. Recently, Graph Neural Networks (GNNs) and transformers have shown promising performance in representing neural architectures. However, each of both methods has its disadvantages. GNNs lack the capabilities to represent complicated features, while transformers face poor generalization when the depth of architecture grows. To mitigate the above issues, we rethink neural architecture topology and show that sibling nodes are pivotal while overlooked in previous research. We thus propose a novel predictor leveraging the strengths of GNNs and transformers to learn the enhanced topology. We introduce a novel token mixer that considers siblings, and a new channel mixer named bidirectional graph isomorphism feed-forward network. Our approach consistently achieves promising performance in both accuracy and latency prediction, providing valuable insights for learning Directed Acyclic Graph (DAG) topology. The code is available at <https://github.com/XuRuihan/NNFormer>.

1. Introduction

Deep neural networks have demonstrated remarkable success across various applications, highlighting the significance of neural architecture design. Designing neural architectures can be quite resource-intensive. Evaluating the performance of a model necessitates training on large datasets. Measuring its inference latency and throughput involves multiple steps such as compilation, deployment, and latency evaluation on various hardware platforms, incurring substantial human effort and resources. One strategy to mitigate these challenges is to predict network attributes with machine learning predictors. By feeding the network structure and hyperparameters into these predictors, valuable characteristics of the network can be estimated with a single feedforward pass, e.g., accuracy on a validation set

or inference times on specific hardware. This predictive approach has been successfully applied in various tasks including neural architecture search [30, 32, 44, 50, 52, 53] and hardware deployment [11, 21, 27, 52, 53, 58], yielding promising outcomes in improving the efficiency and effectiveness of network architecture design.

Previous neural predictors model the neural architecture as a Directed Acyclic Graph (DAG) [9, 11, 25, 27, 34, 39, 44] and utilize Graph Neural Networks (GNNs) or Transformers to extract neural architecture representation. GNNs have emerged as an intuitive solution for learning graph representations [11, 25, 27, 39, 44], which leverage the graph Laplacian and integrate adjacency information to learn the graph topology. GNN-based predictors show strong generalization ability, yet their performance may not be optimal. This is attributed to the structural bias in the message-passing mechanism, which relies solely on adjacency information. As illustrated in Fig. 1(a) and (b), GCNs [22] aggregate the forward and backward adjacent nodes without discrimination, and GATs [43] aggregate them with dynamic weights. Both of them are limited to adjacent information.

With the recent development of transformers, various transformer-based frameworks have been introduced [30, 52, 53]. Transformers have strengths in global modeling and dynamic weight adjustments, hence could extract strong features. Despite the promising performance, they still exhibit several shortcomings. One particular challenge of transformer is related to the long-range receptive field, as depicted in Fig. 1(c) and (d), which can lead to poor generalization performance on deep architectures [52, 53]. The vanilla transformer [30, 42, 52] have a global receptive field, and recent studies proposed transformers on directed transitive closure [9, 34]. Both methods conduct long-range attention that could mix up the information from operations far away, especially when the depth of the input architecture increases to hundreds of layers. For example, NARFormer [52, 53] has illustrated that transformer predictors with global attention struggle in deep network latency prediction, leading to worse performance than GNNs [27, 53].

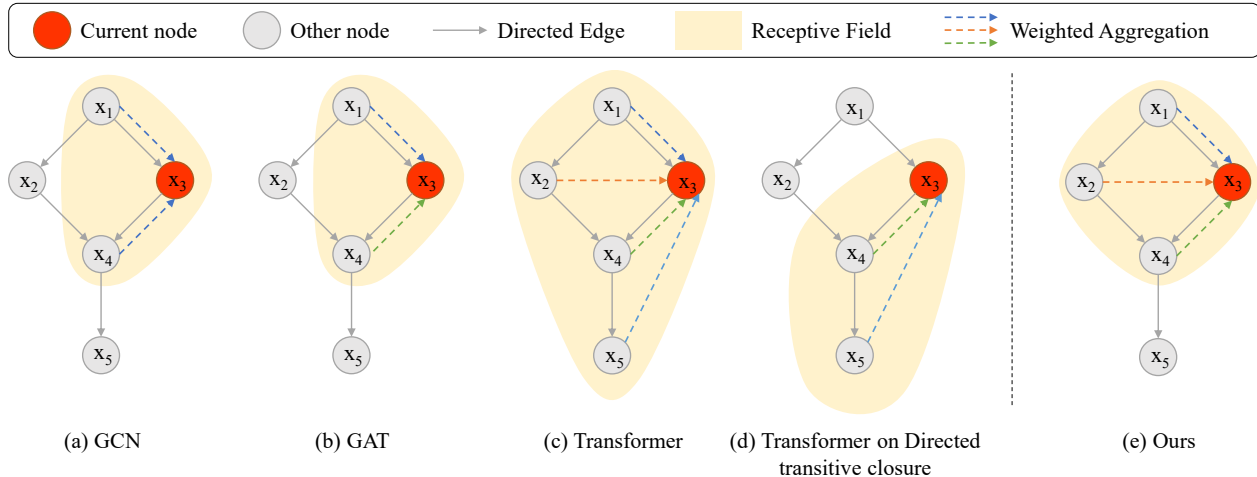


Figure 1. Comparison of different methods on DAG representation of neural architectures. (a) GCNs [22] aggregate adjacent information without discrimination. (b) GATs [43] distinguish adjacent operations, while are still constrained to adjacent nodes. (c) Vanilla transformers [42] aggregate weighted global information, which can result in poor generalization as the network depth increases. (d) Transformers on directed transitive closure [9, 34] aggregate the successor information but still suffer from poor generalization. (e) Our method aggregates sibling information with weighted coefficients. Sibling nodes could extract complementary features in accuracy prediction and allow for concurrent execution in latency prediction.

To study a more effective neural predictor, we rethink the DAG topology and show that the commonly used topological information is not suitable for the neural architecture representation. Most of recent works focus on modeling the relationship of preceding and succeeding operations [9, 34]. However, it is essential to recognize the importance of “sibling nodes”, which share a common parent or child node with the current node as shown in Fig. 1(e). They often exhibit strong connections to the current nodes in neural architecture representation. For example, in the accuracy prediction task, parallel branches may extract complementary features, hence enhancing overall model performance. Furthermore, operations that share the same parent or child node can be executed simultaneously, potentially reducing inference latency. On the contrary, long-range dependency might not be crucial, given that features typically propagate node-by-node within the architecture. Previous methods have not explicitly leveraged sibling cues.

Based on the analysis above, we introduce a new model for neural architecture representation, named Neural Network transFormer (NN-Former). It leverages the strengths of GNNs and transformers, exhibiting good generalization and high performance. For the token mixing module, we utilize a self-attention mechanism of transformers to extract dynamic weights for capturing complex features. We explicitly learn the adjacency and sibling nodes’ features to enhance the topological information. For the channel mixing module, we use a bidirectional graph isomorphism feed-forward network. It learns strong graph topology information such that the position encoding is no longer necessary.

Experiments reveal that 1) our approach surpasses existing methods in both accuracy prediction and latency prediction, and 2) our method has good scalability on both cell-structured architectures and complete neural networks that have hundreds of operations. To the best of our knowledge, this is an original work that leverages sibling cues in neural predictors. Integrating the strengths of GNNs and transformers guarantees its promising performance. The importance of sibling nodes also provides valuable insight into rethinking DAG topology representation in future research.

2. Related Works

Neural Architecture Representation Learning. Neural architecture representation estimates network attributes without actual training or deployment, resulting in significant resource savings. Accuracy predictors forecast the evaluation accuracy, avoiding the resource-intensive process of network training in neural architecture search [2–4, 25, 26, 30, 32, 33, 39, 45, 51–53, 57]. Latency prediction estimates the inference latency without actual deployment, saving time and materials for engineering [11, 21, 27, 58]. Graph-based [3, 25, 39, 51, 57] and transformer-based [30, 52, 53] predictors have been employed to learn the representation of neural architectures. Both methods achieve promising results in neural architecture representation but still face challenges. In this paper, we absorb the strengths of both methods and delve into the topological relationship. **Message-Passing Graph Neural Networks.** Most GNNs can be expressed within the message-passing framework [14, 17, 22, 43, 48, 56]. In this framework, node

representations are computed iteratively by aggregating the embeddings of their neighbors, and a final graph representation can be obtained by aggregating the node embeddings, such as GCN [22], GAT [43], GIN [48], etc. GNN-based models have emerged as a prominent and widely adopted approach for neural network representation learning [3, 25, 39, 51, 57]. GATES [35] and TA-GATES [36] adopt Gated GNNs, aggregating sibling information to some extent. However, they rely on child nodes to sum up sibling features, with no direct interaction between sibling nodes and requiring stronger aggregation. The straightforward structure of GNNs contributes to strong generalization ability, yet also necessitates further improvement in the performance. Enhancing topological information and dynamic bi-directional aggregation is a promising approach.

Transformers on Graphs. Transformer has been introduced into graph representation learning [9, 12, 34, 47], together with network architecture representation learning [30, 52, 53]. TNASP [30] inputs the sum of the operation type embeddings and Laplacian matrix into the standard transformer. NAR-Former [52] encodes each operation and connection into a token and inputs all tokens into a multi-stage fusion transformer. NAR-Former V2 [53] introduced a graph-aided transformer block, which handles both cell-structured networks and entire networks. Flower-Former [18] employs a graph transformer and models the information flow within the networks. ParZC [7] uses a mixer architecture with a Bayesian network to model uncertainty. CAP [19] introduces a subgraph matching-based self-supervised learning method. However, transformers face challenges of poor generalization when the network goes deeper, with global attention mixing up the far away information [53]. To address this limitation, we propose a novel predictor that harnesses the strengths of both GNNs and transformers, allowing it to extract topology features and dynamic weights. This approach enhances the model’s capability and maintains good generalization.

Neural Networks over DAGs. The inductive bias inherent in DAGs has led to specialized neural predictors. GNNs designed for DAGs typically compute graph embeddings using a message-passing framework [41]. On the other hand, transformers applied to DAGs often incorporate the depth of nodes [23, 34] or Laplacian [13] as the position encoding, which may seem non-intuitive for integrating structural information into the model. Additionally, transformer-based models frequently use transitive closure [9, 34] as attention masks, leading to poor generalization [53]. Some hybrid methods with GNNs and Transformers are not tailored to neural architecture representation and also face similar challenges [47, 55]. This paper proposes a novel hybrid model with enhanced topological information from sibling nodes.

3. Methods

3.1. Overview

We adopt a commonly used graph representation of neural architectures [9, 27, 30, 34, 52, 53]. An architecture with n operations is referred to as a Graph $G = (V, E, \mathbf{Z})$, with node set V , edge set $E \subseteq V \times V$, and node features $\mathbf{Z} \in \mathbb{R}^{n \times d}$. Each operation is denoted as a node in V such that $|V| = n$. The edge set E is often given in form of an adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$, where $\mathbf{A}_{ij} = 1$ denotes a directed edge from node i to node j . Each row of \mathbf{Z} represents the feature vector of one node, *i.e.*, operation type and hyperparameters, with the number of nodes n and feature dimension d . Unlike previous methods [30, 52], our predictor is strong and position encoding is unnecessary. For simplicity, \mathbf{Z} is encoded with one-hot encoding for operation type and sinusoidal encoding for operation attributes as [53]. Neural architecture representation [27, 30, 32, 44, 50, 52, 53] uses a predictor $f_{\theta}(\cdot)$ with parameters θ to estimate specific attributes of candidate architectures, *e.g.*, validation accuracy or inference latency:

$$\hat{y} = f_{\theta}(\mathbf{Z}, \mathbf{A}), \quad (1)$$

where \hat{y} denotes the predicted attribute of the architecture.

As illustrated in Figure 2, our approach uses a transformer as the baseline and incorporates discriminative topological features to pursue a strong predictor. Previous transformer-based predictors considered adjacent propagation [30, 52, 53] or transitive closure [9, 34] as the graph structure information. Global attention is effective in shallow network prediction as shown in previous works such as TNASP [30] and NAR-Former [52]. However, as the network depth increases, there exhibits a decrease in the generalization of global attention as shown in [53]. Global attention may be biased towards training data and demonstrate poor generalization performance. To build a general neural predictor for the range of all depths, we propose a non-global neural predictor that outperforms the previous methods on both accuracy and latency predictions.

As discussed in Section 1, sibling nodes have a strong relationship with the current nodes and also provide useful information in accuracy and latency prediction. Thus we introduce an Adjacency-Sibling Multi-head Attention (ASMA) in the self-attention layer to learn the local features. As the sibling relationship can be calculated from adjacency matrix \mathbf{A} , our ASMA is formulated as:

$$\hat{\mathbf{H}}^{l-1} = \text{ASMA} \left(\text{LN} \left(\mathbf{H}^{l-1} \right), \mathbf{A} \right) + \mathbf{H}^{l-1}, \quad (2)$$

where \mathbf{H}^l denotes the feature for the layer l and LN denotes layer normalization. ASMA injects topological information into the transformer, thereby augmenting the

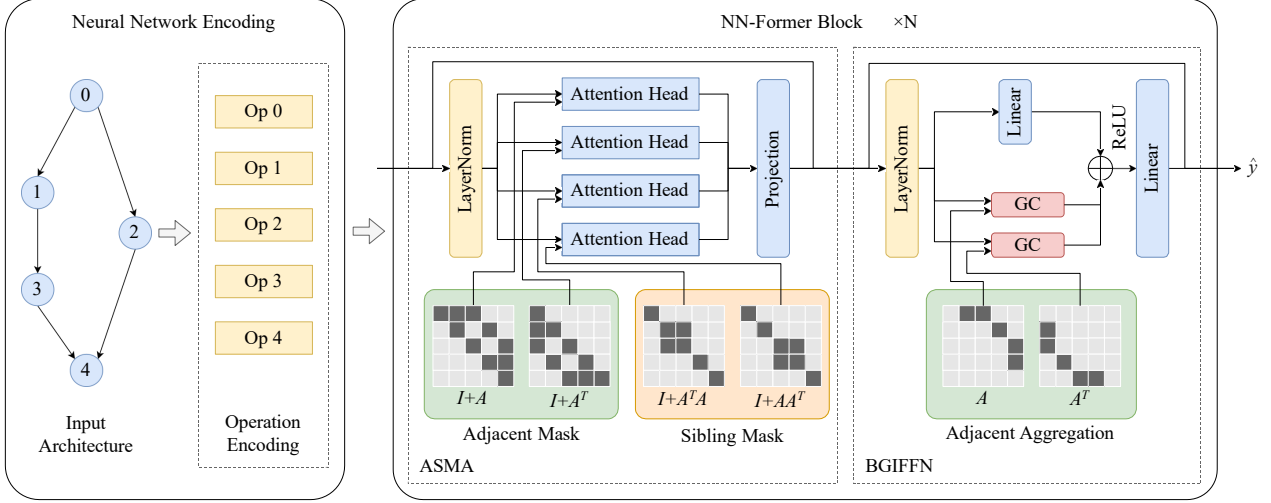


Figure 2. **The proposed NN-Former framework.** We introduce adjacency and sibling attention masks in the Adjacency-Sibling Multihead Attention (ASMA) to learn graph topology information. We also introduce adjacency aggregation in the Bidirectional Graph Isomorphism Feed-Forward Network (BGIFFN) to enhance the topology structure.

capability of Directed Acyclic Graph (DAG) representation learning. In the channel-mixing part, we introduce a Bidirectional Graph Isomorphism Feed-Forward Network (BGIFFN). This module extracts strong topology features and alleviates the necessity of complex position encoding:

$$\mathbf{H}^l = \text{BGIFFN} \left(\text{LN} \left(\hat{\mathbf{H}}^{l-1} \right), \mathbf{A} \right) + \hat{\mathbf{H}}^{l-1}. \quad (3)$$

As for input and output, the first layer feature \mathbf{H}^0 and the last layer feature \mathbf{H}^L are related to the input and output in the following way:

$$\mathbf{H}^0 = \text{LN} \left(\text{FC} \left(\mathbf{Z} \right) \right), \quad (4)$$

$$\hat{y} = \text{FC} \left(\text{ReLU} \left(\text{FC} \left(\mathbf{H}^L \right) \right) \right), \quad (5)$$

where FC denotes fully-connected layer. The following parts proceed to introduce ASMA and BGIFFN in detail.

3.2. Adjacency-Sibling Multihead Attention

Given a node, we define its sibling nodes as those that share the same parents or children. To identify these sibling nodes, we use the adjacency matrix \mathbf{A} and its transpose \mathbf{A}^\top . Specifically, nodes sharing the same parent nodes are indicated by the non-zero positions in the matrix product $\mathbf{A}\mathbf{A}^\top$, reflecting the backward mapping followed by the forward mapping. Similarly, nodes sharing the same children nodes are identified through the matrix product $\mathbf{A}^\top\mathbf{A}$. In this way, we can identify whether there is a sibling relationship between each pair of nodes.

To inject topological information, we introduce a novel multi-head attention module. As shown in Figure 2, we use four-head attention, where each head uses an attention mask

indicating a specific topology. These masks include forward adjacency \mathbf{A} , backward adjacency \mathbf{A}^\top , siblings with the same parents $\mathbf{A}\mathbf{A}^\top$, and siblings with the same children $\mathbf{A}^\top\mathbf{A}$, respectively. The proposed ASMA is denoted as:

$$\text{ASMA}(\mathbf{H}) = \text{Concat}(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{X}_4) \mathbf{W}^O, \quad (6)$$

$$\mathbf{X}_1 = \sigma \left(\left(\mathbf{Q}_1 \mathbf{K}_1^\top \circ (\mathbf{I} + \mathbf{A}) \right) / \sqrt{h} \right) \mathbf{V}_1, \quad (7)$$

$$\mathbf{X}_2 = \sigma \left(\left(\mathbf{Q}_2 \mathbf{K}_2^\top \circ (\mathbf{I} + \mathbf{A}^\top) \right) / \sqrt{h} \right) \mathbf{V}_2, \quad (8)$$

$$\mathbf{X}_3 = \sigma \left(\left(\mathbf{Q}_3 \mathbf{K}_3^\top \circ (\mathbf{I} + \mathbf{A}\mathbf{A}^\top) \right) / \sqrt{h} \right) \mathbf{V}_3, \quad (9)$$

$$\mathbf{X}_4 = \sigma \left(\left(\mathbf{Q}_4 \mathbf{K}_4^\top \circ (\mathbf{I} + \mathbf{A}^\top\mathbf{A}) \right) / \sqrt{h} \right) \mathbf{V}_4, \quad (10)$$

where \mathbf{X}_i denote the i -th head feature and $\mathbf{Q}_i = \mathbf{H}\mathbf{W}_i^Q$, $\mathbf{K}_i = \mathbf{H}\mathbf{W}_i^K$, $\mathbf{V}_i = \mathbf{H}\mathbf{W}_i^V$ denotes the query, key, and value for each head, respectively. σ is the softmax operation, and h denotes the number of head dimension. \mathbf{I} is introduced to contain self-position information. \circ is an elementwise masking operation, which constrains the attention to the non-zero positions of the mask matrix. For example, the first head \mathbf{H}_1 utilizes an attention mask of $\mathbf{I} + \mathbf{A}$, which means it only conducts attention on the self-position and forward adjacency position. ASMA decouples the local topology information into 4 different aspects. This module extracts diverse topological information and enhances feature representation in neural architecture.

3.3. Bidirectional Graph Isomorphism Feed-Forward Network

To further enhance the topology information, we propose a bidirectional graph isomorphism feed-forward network.

We utilize the adjacency matrix \mathbf{A} and its transpose \mathbf{A}^\top to aggregate the forward and backward adjacency positions in the feedforward module. The BGIFFN is formulated as:

$$\text{BGIFFN}(\mathbf{H}, \mathbf{A}) = \text{ReLU}(\mathbf{H}\mathbf{W}_1 + \mathbf{H}_g)\mathbf{W}_2, \quad (11)$$

$$\mathbf{H}_g = \text{Concat}\left(\text{GC}(\mathbf{H}, \mathbf{A}), \text{GC}\left(\mathbf{H}, \mathbf{A}^\top\right)\right), \quad (12)$$

where \mathbf{H}_g denotes the output features of the graph convolution and $\mathbf{W}_1, \mathbf{W}_2$ denote the parameters of linear transformation. GC denotes the graph convolution, which is a simplified form of GCN[22]:

$$\text{GC}(\mathbf{H}, \mathbf{A}) = \mathbf{A}\mathbf{H}\mathbf{W}, \quad (13)$$

where \mathbf{W} denotes the parameters of fully-connected layer. Note that we use the directed adjacency matrix rather than graph Laplacian, which makes it simpler and stronger. With $\text{GC}(\mathbf{H}, \mathbf{A})$ and $\text{GC}(\mathbf{H}, \mathbf{A}^\top)$, we obtain the forward features and backward features. Since we concatenate the two directional features, the BGIFFN will learn forward propagation in one half of the channels, and backward propagation in the other. We will show that BGIFFN demonstrates bidirectional graph isomorphism in Appendix, which enhances topology information.

4. Experiments

We conduct experiments on two tasks, namely accuracy prediction and latency prediction. For accuracy prediction, we evaluate the ranking performance of NN-Former on two benchmarks NAS-Bench-101 [54] and NAS-Bench-201 [8]. For latency performance, we conduct experiments on NNLQ [27]. A series of ablation experiments are conducted to demonstrate the effectiveness of our design. More details will be provided in the appendix.

4.1. Accuracy Prediction

We predict accuracy on NAS-Bench-101 [54] and NAS-Bench-201 [8]. Both datasets use cell-structured architectures. The NAS-Bench-101 [54] dataset contains 423,624 unique architectures, each comprising 9 repeated cells with a maximum of 7 nodes and 9 edges per cell. Like the NAS-Bench-101, the architectures in NAS-Bench-201 [8] are also built using repeated cells. It presents 15,625 distinct cell candidates, each composed of 4 nodes and 6 edges. We report Kendall’s Tau as the previous methods [30, 35, 52].

Experiments on NAS-Bench-101. We implement the configuration in TNASP [30] to train our predictor on subsets of 0.02%, 0.04%, 0.1%, and 1% of the entire dataset. Then we utilized the complete dataset as the test set and computed Kendall’s Tau to evaluate the performance. The results are detailed in Tab. 1. Our predictor consistently outperforms baseline methods, such as CNNs, LSTMs, GNNs,

Table 1. **Accuracy prediction results on NAS-Bench-101 [54].** We use different proportions of data as the training set and report Kendall’s Tau on the whole dataset.

Backbone	Method	Publication	Training Samples			
			0.02% (100)	0.04% (172)	0.1% (424)	1% (4236)
CNN	ReNAS [50]	CVPR 2021	-	-	0.657	0.816
LSTM	NAO [32]	NeurIPS 2018	0.501	0.566	0.666	0.775
GNN	NP [44]	ECCV 2020	0.391	0.545	0.679	0.769
	GATES [35]	ECCV 2020	0.605	0.659	0.691	0.822
	GMAE-NAS [20]	IJCAI 2022	0.666	0.697	0.732	0.775
Transformer	Graphormer [55]	NeurIPS 2021	0.564	0.580	0.611	0.797
	TNASP [30]	NeurIPS 2021	0.600	0.669	0.705	0.820
	NAR-Former [52]	CVPR 2023	0.632	0.653	0.765	0.871
	PINAT [31]	AAAI 2024	0.679	0.715	0.772	0.846
Hybrid	GraphTrans [47]	NeurIPS 2021	0.330	0.472	0.602	0.700
	NAR-Former V2 [53]	NeurIPS 2023	0.663	0.704	0.773	0.861
	NN-Former (Ours)	-	0.709	0.765	0.809	0.877

Table 2. **Accuracy prediction results on NAS-Bench-201 [8].** We use different proportions of data as the training set and report Kendall’s Tau on the whole dataset.

Backbone	Method	Publication	Training Samples			
			1% (156)	3% (469)	5% (781)	10% (1563)
LSTM	NAO [32]	NeurIPS 2018	0.493	0.470	0.522	0.526
GNN	NP [44]	ECCV 2020	0.413	0.584	0.634	0.646
Transformer	Graphormer [55]	NeurIPS 2021	0.630	0.680	0.719	0.776
	TNASP [30]	NeurIPS 2021	0.589	0.640	0.689	0.724
	NAR-Former [52]	CVPR 2023	0.660	0.790	0.849	0.901
	PINAT [31]	AAAI 2024	0.631	0.706	0.761	0.784
Hybrid	GraphTrans [47]	NeurIPS 2021	0.409	0.550	0.588	0.673
	NAR-Former V2 [53]	NeurIPS 2023	0.752	0.846	0.874	0.888
	NN-Former (Ours)	-	0.804	0.860	0.879	0.890

Transformers, and hybrid GNNs and Transformers. This result underscores the superior predictive capability of NN-Former in determining neural architecture performance.

Experiments on NAS-Bench-201. We employ a comparable experimental setup to NAS-Bench-101, *i.e.*, training predictors on different subsets of 1%, 3%, 5%, and 10%, and then evaluating them on the complete dataset. The results are depicted in Tab. 2. NN-Former surpasses other methods in all scenarios except for the 10% subsets. Note that our method aims at unified prediction for both accuracy and latency, it is acceptable that our method achieves comparable results. We outperform NAR-Former V2 [53] for all setups, which has a similar unifying motivation. Additionally, neural architecture search prefers high generalization performance with fewer training samples, resulting in significant resource savings. More details are discussed in the appendix.

4.2. Latency Prediction

We employ NNLQ as the latency prediction task. NNLQ [27] includes 20,000 deep-learning networks and their respective latencies on specific hardware. This dataset encompasses 10 distinct network types, with 2,000 net-

Table 3. **Out of domain latency prediction on>NNLQ [27].** “Test Model = AlexNet” means that only AlexNet models are used for testing, and the other 9 model families are used for training. The best results refer to the lowest MAPE and corresponding ACC (10%) in 10 independent experiments.

Metric	Test Domain	FLOPs	FLOPs+MAC	nn-Meter	TPU	BRP-NAS	NNLP (avg / best)	NAR-Former V2 (avg / best)	Ours (avg / best)
MAPE ↓	AlexNet	44.65	15.45	7.20	10.55	31.68	10.64 / 9.71	24.28 / 18.29	11.47 / 11.17
	EfficientNet	58.36	53.96	18.93	16.74	51.97	21.46 / 18.72	13.20 / 11.37	5.13 / 4.81
	GoogleNet	30.76	32.54	11.71	8.10	25.48	13.28 / 10.90	6.61 / 6.15	6.74 / 6.65
	MnasNet	40.31	35.96	10.69	11.61	17.26	12.07 / 10.86	7.16 / 5.93	2.71 / 2.54
	MobileNetV2	37.42	35.27	6.43	12.68	20.42	8.87 / 7.34	6.73 / 5.65	4.17 / 3.66
	MobileNetV3	64.64	57.13	35.27	9.97	58.13	14.57 / 13.17	9.06 / 8.72	9.07 / 9.03
	NasBench201	80.41	33.52	9.57	58.94	13.28	9.60 / 8.19	9.21 / 7.89	7.93 / 7.71
	ResNet	21.18	18.91	15.58	20.05	15.84	7.54 / 7.12	6.80 / 6.44	7.49 / 7.38
	SqueezeNet	29.89	23.19	18.69	24.60	42.55	9.84 / 9.52	7.08 / 6.56	9.08 / 7.05
	VGG	69.34	66.63	19.47	38.73	30.95	7.60 / 7.17	15.40 / 14.26	20.12 / 19.64
Average		47.70	37.26	15.35	21.20	30.76	11.55 / 10.27	10.55 / 9.13	8.39 / 7.96
Acc(10%) ↑	AlexNet	6.55	40.50	75.45	57.10	15.20	59.07 / 64.40	24.65 / 28.60	56.08 / 57.10
	EfficientNet	0.05	0.05	23.40	17.00	0.10	25.37 / 28.80	44.01 / 50.20	90.85 / 90.90
	GoogleNet	12.75	9.80	47.40	69.00	12.55	36.30 / 48.75	80.10 / 83.35	80.43 / 83.40
	MnasNet	6.20	9.80	60.95	44.65	34.30	55.89 / 61.25	73.46 / 81.60	98.65 / 98.70
	MobileNetV2	6.90	8.05	80.75	33.95	29.05	63.03 / 72.50	78.45 / 83.80	94.90 / 96.85
	MobileNetV3	0.05	0.05	23.45	64.25	13.85	43.26 / 49.65	68.43 / 70.50	74.18 / 74.30
	NasBench201	0.00	10.55	60.65	2.50	43.45	60.70 / 70.60	63.13 / 71.70	69.90 / 71.10
	ResNet	26.50	29.80	39.45	27.30	39.80	72.88 / 76.40	77.24 / 79.70	70.83 / 71.55
	SqueezeNet	16.10	21.35	36.20	25.65	11.85	58.69 / 60.40	75.01 / 79.25	77.85 / 80.95
	VGG	4.80	2.10	26.50	2.60	13.20	71.04 / 73.75	45.21 / 45.30	29.40 / 29.85
Average		7.99	13.20	47.42	34.40	21.34	54.62 / 60.65	62.70 / 67.40	74.31 / 75.47

Table 4. **In domain latency prediction on>NNLQ [27].** Training and testing on the same distribution.

Metric	Test Domain	NNLP avg / best	NAR-Former V2 avg / best	Ours avg / best
MAPE ↓	AlexNet	6.37 / 6.21	6.18 / 5.97	4.69 / 4.61
	EfficientNet	3.04 / 2.82	2.34 / 2.22	2.31 / 2.21
	GoogleNet	4.18 / 4.12	3.63 / 3.46	3.48 / 3.39
	MnasNet	2.60 / 2.46	1.80 / 1.70	1.52 / 1.48
	MobileNetV2	2.47 / 2.37	1.83 / 1.72	1.54 / 1.50
	MobileNetV3	3.50 / 3.43	3.12 / 2.98	3.17 / 2.99
	NasBench201	1.46 / 1.31	1.82 / 1.18	1.11 / 0.96
	SqueezeNet	4.03 / 3.97	3.54 / 3.34	3.09 / 3.08
	VGG	3.73 / 3.63	3.51 / 3.29	2.94 / 2.89
	ResNet	3.34 / 3.25	3.11 / 2.89	2.66 / 2.47
Average		3.47 / 3.44	3.07 / 3.00	2.85 / 2.65
Acc(10%) ↑	AlexNet	81.75 / 84.50	81.90 / 84.00	90.50 / 91.00
	EfficientNet	98.00 / 97.00	98.50 / 100.0	99.00 / 100.0
	GoogleNet	93.70 / 93.50	95.95 / 95.50	97.15 / 97.50
	MnasNet	97.70 / 98.50	99.70 / 100.0	99.50 / 100.0
	MobileNetV2	99.30 / 99.50	99.90 / 100.0	99.60 / 100.0
	MobileNetV3	95.35 / 96.00	96.75 / 98.00	96.50 / 97.00
	NasBench201	100.0 / 100.0	100.0 / 100.0	100.0 / 100.0
	SqueezeNet	93.25 / 93.00	95.95 / 96.50	97.70 / 98.00
	VGG	95.25 / 96.50	95.85 / 96.00	95.80 / 96.50
	ResNet	98.40 / 98.50	98.55 / 99.00	99.45 / 99.50
Average		95.25 / 95.50	96.41 / 96.30	97.45 / 97.85

works for each type. The depth of each architecture varies from tens to hundreds of operations, requiring the scalability of the neural predictor. In line with>NNLQ, the Mean Absolute Percentage Error (MAPE) and Error Bound Accuracy (Acc(δ)) are employed to assess the disparities between latency predictions and actual values.

We conduct the experiments on two scenarios following [53]. In the first in-domain scenario, the training and testing

sets are from the same distribution. The results are shown in Tab. 4. When testing with all test samples, the average MAPE of our methods is 0.62% lower than the>NNLQ [27] and 0.22% lower than the NAR-Former V2 [53]. The average Acc(10%) is 2.20% higher than the>NNLQ and 1.04% higher than the NAR-Former V2. When tested on various types of network data separately, previous methods fail on specific model types, especially on AlexNet, while our method largely mitigates this challenge and obtains a balanced performance on each model type.

The other out-of-domain scenario is more significant, as it involves inferring an unseen network type during the evaluation. The results in Tab. 3 indicate that FLOPs and memory access data is insufficient for predicting latency. Due to the disparity between kernel delay accumulation and actual latency, kernel-based approaches such as nn-Meter [58] and TPU [21] exhibit inferior performance compared to GNNs (>NNLQ [27]) and hybrid models (NAR-Former V2 [53]). Leveraging enriched topological information, our method achieves the highest MAPE and Acc(10%) among the average metrics of the ten experimental sets. In comparison to the runner-up NAR-Former V2 [53], our approach demonstrates a substantial 11.61% increase in average Acc(10%).

4.3. Ablation Studies

In this section, we perform a series of ablation experiments on NAS-Bench-101 [54] and>NNLQ [27] datasets to analyze the effects of the proposed modifications. First, we evaluate ASMA and BGIFN over a vanilla transformer baseline. Second, we verify the design for both modules,

Table 5. Ablation study of ASMA on NNLQ [27]. Results on the NAS-Bench-201 family are reported.

Attention	MAPE↓	Acc(10%)↑
Global	10.83%	58.45%
ASMA	7.93%	69.90%

Table 6. Ablation study of ASMA and BGIFFN on NAS-Bench-101 [54].

Attention	FFN	Kendall’s Tau↑
Global	vanilla	0.4598
ASMA	vanilla	0.6538
Global	BGIFFN	0.7656
ASMA	BGIFFN	0.7654

especially by examining the significance of the topological information. All experiments on the accuracy prediction are conducted on the NAS-Bench-101 0.04% training set.

ASMA on latency prediction. To evaluate the generalization ability of ASMA, we conducted experiments on NNLQ [27] under the out-of-domain setting. We use NAS-Bench-201 as the target model type due to its intricate connections between operations. The results are shown in Tab. 5, where we keep the number of heads unchanged ablate on the attention mask. The model with ASMA shows a large performance enhancement of 11.45% Acc(10%) compared to the one using global attention. It indicates that global attention presents a noticeable disparity, underscoring the advantages of local features when building a unified neural predictor.

ASMA and BGIFFN on accuracy prediction. To evaluate the effectiveness of ASMA and BGIFFN, we conducted accuracy prediction experiments on NAS-Bench-101 [54]. The findings are detailed in Tab. 6. The baseline method utilizes a vanilla transformer and the result is respectable since it does not incorporate any topology information. Introducing ASMA, which integrates adjacency and sibling information, leads to an enhancement of 0.6538. Incorporating BGIFFN leads to further enhancement of 0.7656. It indicates that the ASMA and BGIFFN modules effectively capture structural features within the models. Furthermore, when ASMA and BGIFFN are combined, the model achieves a performance score of 0.7654 which is comparable to the global attention mechanism. However, we have highlighted the limitations of global attention in latency prediction in Tab. 7, and our ASMA achieved the best trade-off across both accuracy and latency prediction tasks.

Topological structure of ASMA. We examine the designs of ASMA in Tab. 7 and Tab. 8. The performance of different attention masks is shown in Tab. 7. Maintaining a consistent number of heads at 4, we modify attention mask for each head. Rows 1 and 2 exclusively utilize forward or

Table 7. Ablation study on the topological structure of ASMA. We explore different attention masks for the 4 heads.

Row	Attention Mask	Kendall’s Tau↑
1	$\mathbf{A}, \mathbf{A}, \mathbf{A}, \mathbf{A}$	0.7522
2	$\mathbf{A}^\top, \mathbf{A}^\top, \mathbf{A}^\top, \mathbf{A}^\top$	0.7545
3	$\mathbf{A}, \mathbf{A}, \mathbf{A}^\top, \mathbf{A}^\top$	0.7566
4	$\mathbf{A}, \mathbf{A}^\top, \mathbf{A}\mathbf{A}, \mathbf{A}^\top\mathbf{A}^\top$	0.7573
5	$\mathbf{A}, \mathbf{A}^\top, \mathbf{A}^\top\mathbf{A}, \mathbf{A}\mathbf{A}^\top$	0.7654

Table 8. Ablation study on the position encoding with ASMA. We explore different ways of position encoding when utilizing ASMA.

ASMA design	Kendall’s Tau↑
Ours	0.7654
+NAR PE [52]	0.7449
+Laplacian [30]	0.7063

Table 9. Ablation study on the topological structure of BGIFFN. We explore different ways of structure aggregation in BGIFFN.

Row	BGIFFN	Kendall’s Tau↑
1	\mathbf{A}	0.7253
2	\mathbf{A}^\top	0.7501
3	$\mathbf{A}, \mathbf{A}^\top, \mathbf{A}^\top\mathbf{A}, \mathbf{A}\mathbf{A}^\top$	0.7470
4	$\mathbf{A}, \mathbf{A}^\top$	0.7654

backward adjacency. Row 3 combines forward and backward adjacency and improves performance. Row 4 investigates the impact of predecessors and successors, indicating marginal enhancement. It shows that empirical topological information in DAG tasks [9, 34] is helpless in neural architecture representation. Row 5 combines the adjacency and sibling nodes and achieves the highest performance, highlighting the significance of sibling nodes. It shows that the design of ASMA is robust and logically sound, showcasing its effectiveness in capturing architectural patterns.

Position Encoding (PE). We explored the impact of PE on our model. Traditionally, transformers heavily rely on PE to capture structural information. However, our approach makes this reliance unnecessary because our method inherently incorporates abundant topological information. As shown in Tab. 8, we experiment with the inclusion of position encoding in our framework, *i.e.*, NAR PE tailored for neural architecture representation in NAR-Former [52], and Laplacian position encoding in TNASP [34]. The results suggest that they have no improvement in the performance of NN-Former. It indicates that our method presents exceptional structural learning capabilities.

Topological structure of BGIFFN. As illustrated in Tab. 9, this experiment maintains the total parameters constant while adjusting the number of splits in the graph convolution branch. In Rows 1 and 2, a single split is retained, and the forward or backward adjacent convolution is con-

Table 10. Ablation study on the BGIFFN design. We explore different ways of adjacency aggregation when utilizing BGIFFN.

BGIFFN design	Kendall’s Tau \uparrow
Ours	0.7654
add \rightarrow multiply	0.7076
\rightarrow GCN (Laplacian)	0.7296
\rightarrow GAT	0.6973

ducted. Moving to Row 3, using 4 splits for adjacency and siblings results in a performance that is even worse than using backward adjacency only. This outcome may be attributed to the considerable strength of topological information provided by ASMA, rendering such a complex graph structure unnecessary. In Row 4, using 2 splits with forward and backward adjacency produces the most favorable result, underscoring the rationale behind our BGIFFN approach. This finding suggests that BGIFFN is well-founded and effective in leveraging topological information.

BGIFFN design. The gating mechanism [35, 49] has demonstrated superior performance to the standard feed-forward layer. However, substituting the elementwise add operation in BGIFFN with the Hadamard product results in a significant performance decrease to 0.7076. This may be attributed to the different features of the two branches, as one represents self-position only, and the other aggregates adjacency features. Directly multiplying the two features yields a decrease in performance. Furthermore, we compare our approach with the conventional GCN [22] and GAT [43]. Both methods lead to a noticeable performance decline, rendering the superiority of our NN-Former.

4.4. Inference Speed

We report the parameters and the latency, memory, and training time on a single RTX 3090 in Tab. 11. Our method has comparable inference latency, memory usage, and training time compared to NAR-Former [52], indicating that the improvement brought by our method is solid.

Compared to the tremendous time spent training candidate architectures, the time of training neural predictors is neglectable. Therefore, the computational resources consumed by predictors do not affect practical applications. In the experiments on NNLQ, our method can make predictions for networks with from 20 to 200 layers, which encompasses the size of practical models. It indicates that our method can be applied to practical use.

4.5. Sibling Nodes Modeling on General DAG Tasks

Our method is dedicated to neural architecture representation learning. However, considering our method is based on DAG modeling, is it possible for the sibling nodes modeling to enhance other DAG tasks? Our experiments show that the DAG tasks are divided into two groups by the importance

Table 11. Computation cost of the proposed modules.

Methods	Params (M)	Latency (ms)	Memory (GB)	Training Time (h)
NAR-Former	4.8	10.31	0.58	0.7
NN-Former w/o ASMA	4.9	11.21	0.67	0.8
NN-Former w/o BGIFFN	3.7	10.17	0.60	0.7
NN-Former	4.9	11.53	0.67	0.8

Table 12. Sibling nodes modeling on general DAG tasks.

Model	Cora($\%$) \uparrow	ogbg-code2($\%$) \uparrow
DAG Transformer	87.39	19.0
DAG Transformer + sibling	88.14	18.9

of sibling nodes. One is that there is a strong relationship between siblings, such as citation prediction. Two papers that cite the same paper might follow similar motivations, methods, or experiments. Similarly, two papers cited by the same paper may also have these in common. The other is that siblings are not as important. For example, Abstract Syntax Tree (AST) uses syntactic construct to aggregate the successors, while siblings do not make practical sense.

We conduct experiments on the sibling effects on general DAG tasks. We use the DAG Transformer [34] as the baseline and add a sibling attention mask without careful calibration. The results in Tab. 12 show that sibling nodes play a crucial role in Cora Citation prediction, while not as important in ogbg-code2 AST prediction. We hope these findings inspire more research on general DAG tasks.

5. Conclusion

This work introduces a novel neural architecture representation model. This model unites the strengths of GCN and transformers, demonstrating strong capability in topology modeling and representation learning. We also conclude that different from the intuition on other DAG tasks, sibling nodes significantly affect the extraction of topological information. Our proposed model performs well on accuracy and latency prediction, showcasing model capability and generalization ability. This work may inspire future efforts in neural architecture representation and neural networks on DAG representation.

Acknowledgement

This work is supported in part by Grant No. 2023-JCJQ-LA-001-088, in part by Natural Science Foundation of China under Grant No. U20B2052, 61936011, in part by the Okawa Foundation Research Award, in part by the Ant Group Research Fund, and in part by the Kunpeng&Ascend Center of Excellence, Peking University.

References

- [1] Mohamed S Abdelfattah, Abhinav Mehrotra, Łukasz Dudziak, and Nicholas D Lane. Zero-cost proxies for lightweight nas. *arXiv preprint arXiv:2101.08134*, 2021. 4
- [2] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791*, 2019. 2
- [3] Yaofu Chen, Yong Guo, Qi Chen, Minli Li, Wei Zeng, Yaowei Wang, and Mingkui Tan. Contrastive neural architecture search with neural architecture comparators. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9502–9511, 2021. 2, 3
- [4] Boyang Deng, Junjie Yan, and Dahua Lin. Peephole: Predicting network performance before training. *arXiv preprint arXiv:1712.03351*, 2017. 2
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 3
- [6] Xiaohan Ding, Xiangyu Zhang, Ningning Ma, Jungong Han, Guiguang Ding, and Jian Sun. Repvgg: Making vgg-style convnets great again. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13733–13742, 2021. 1
- [7] Peijie Dong, Lujun Li, Xinglin Pan, Zimian Wei, Xiang Liu, Qiang Wang, and Xiaowen Chu. Parzc: Parametric zero-cost proxies for efficient nas. *arXiv preprint arXiv:2402.02105*, 2024. 3
- [8] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. *arXiv preprint arXiv:2001.00326*, 2020. 5, 3, 4
- [9] Zehao Dong, Muhan Zhang, Fuhai Li, and Yixin Chen. Pace: A parallelizable computation encoder for directed acyclic graphs. In *International Conference on Machine Learning*, pages 5360–5377. PMLR, 2022. 1, 2, 3, 7
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3, 4
- [11] Łukasz Dudziak, Thomas Chau, Mohamed Abdelfattah, Royson Lee, Hyeji Kim, and Nicholas Lane. Brp-nas: Prediction-based nas using gens. *Advances in Neural Information Processing Systems*, 33:10480–10490, 2020. 1, 2
- [12] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699*, 2020. 3
- [13] Mukul Gagrani, Corrado Rainone, Yang Yang, Harris Teague, Wonseok Jeon, Roberto Bondesan, Herke van Hoof, Christopher Lott, Weiliang Zeng, and Piero Zappi. Neural topological ordering for computation graphs. *Advances in Neural Information Processing Systems*, 35:17327–17339, 2022. 3
- [14] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR, 2017. 2
- [15] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Łukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 3, 4
- [16] Jianyuan Guo, Kai Han, Han Wu, Yehui Tang, Xinghao Chen, Yunhe Wang, and Chang Xu. Cmt: Convolutional neural networks meet vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12175–12185, 2022. 3
- [17] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017. 2
- [18] Dongyeong Hwang, Hyunju Kim, Sunwoo Kim, and Kijung Shin. Flowerformer: Empowering neural architecture encoding using a flow-aware graph transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6128–6137, 2024. 3
- [19] Han Ji, Yuqi Feng, and Yanan Sun. Cap: a context-aware neural predictor for nas. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 4219–4227, 2024. 3
- [20] Kun Jing, Jungang Xu, and Pengfei Li. Graph masked autoencoder enhanced predictor for neural architecture search. In *IJCAI*, pages 3114–3120, 2022. 5
- [21] Sam Kaufman, Phitchaya Phothilimthana, Yanqi Zhou, Charith Mendis, Sudip Roy, Amit Sabne, and Mike Burrows. A learned performance model for tensor processing units. *Proceedings of Machine Learning and Systems*, 3:387–400, 2021. 1, 2, 6
- [22] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2016. 1, 2, 3, 5, 8
- [23] Bhushan Kotnis, Carolin Lawrence, and Mathias Niepert. Answering complex queries in knowledge graphs with bidirectional sequence encoders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4968–4977, 2021. 3
- [24] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images, 2009. 3
- [25] Wei Li, Shaogang Gong, and Xiatian Zhu. Neural graph embedding for neural architecture search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4707–4714, 2020. 1, 2, 3
- [26] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proceedings of the European conference on computer vision (ECCV)*, pages 19–34, 2018. 2
- [27] Liang Liu, Mingzhu Shen, Ruihao Gong, Fengwei Yu, and Hailong Yang. Nlqp: A multi-platform neural network latency query and prediction system with an evolving database. In *Proceedings of the 51st International Conference on Parallel Processing*, pages 1–14, 2022. 1, 2, 3, 5, 6, 7, 4

- [28] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 3
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 3, 4
- [30] Shun Lu, Jixiang Li, Jianchao Tan, Sen Yang, and Ji Liu. Tnasp: A transformer-based nas predictor with a self-evolution framework. *Advances in Neural Information Processing Systems*, 34:15125–15137, 2021. 1, 2, 3, 5, 7
- [31] Shun Lu, Yu Hu, Peihao Wang, Yan Han, Jianchao Tan, Jixiang Li, Sen Yang, and Ji Liu. Pinat: A permutation invariance augmented transformer for nas predictor. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8957–8965, 2023. 5, 3
- [32] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural architecture optimization. *Advances in neural information processing systems*, 31, 2018. 1, 2, 3, 5
- [33] Renqian Luo, Xu Tan, Rui Wang, Tao Qin, Enhong Chen, and Tie-Yan Liu. Semi-supervised neural architecture search. *Advances in Neural Information Processing Systems*, 33:10547–10557, 2020. 2
- [34] Yuankai Luo, Veronika Thost, and Lei Shi. Transformers over directed acyclic graphs. *Advances in Neural Information Processing Systems*, 36, 2023. 1, 2, 3, 7, 8
- [35] Xuefei Ning, Yin Zheng, Tianchen Zhao, Yu Wang, and Huazhong Yang. A generic graph-based neural architecture encoding scheme for predictor-based nas. In *European Conference on Computer Vision*, pages 189–204. Springer, 2020. 3, 5, 8
- [36] Xuefei Ning, Zixuan Zhou, Junbo Zhao, Tianchen Zhao, Yiping Deng, Changcheng Tang, Shuang Liang, Huazhong Yang, and Yu Wang. Ta-gates: An encoding scheme for neural network architectures. *Advances in Neural Information Processing Systems*, 35:32325–32339, 2022. 3
- [37] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992. 3
- [38] Pranab Kumar Sen. Estimates of the regression coefficient based on kendall’s tau. *Journal of the American statistical association*, 63(324):1379–1389, 1968. 3
- [39] Han Shi, Renjie Pi, Hang Xu, Zhenguo Li, James Kwok, and Tong Zhang. Bridging the gap between sample-based and one-shot neural architecture search with bonas. *Advances in Neural Information Processing Systems*, 33:1808–1819, 2020. 1, 2, 3
- [40] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 1
- [41] Veronika Thost and Jie Chen. Directed acyclic graph neural networks. *arXiv preprint arXiv:2101.07965*, 2021. 3
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 1, 2
- [43] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018. 1, 2, 3, 8
- [44] Wei Wen, Hanxiao Liu, Yiran Chen, Hai Li, Gabriel Bender, and Pieter-Jan Kindermans. Neural predictor for neural architecture search. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX*, pages 660–676. Springer, 2020. 1, 3, 5
- [45] 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*, pages 10293–10301, 2021. 2
- [46] Felix Wu, Angela Fan, Alexei Baevski, Yann N Dauphin, and Michael Auli. Pay less attention with lightweight and dynamic convolutions. *arXiv preprint arXiv:1901.10430*, 2019. 3
- [47] Zhanghao Wu, Paras Jain, Matthew Wright, Azalia Mirhoseini, Joseph E Gonzalez, and Ion Stoica. Representing long-range context for graph neural networks with global attention. *Advances in Neural Information Processing Systems*, 34:13266–13279, 2021. 3, 5
- [48] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2018. 2, 3, 1
- [49] Ruihan Xu, Haokui Zhang, Wenzhe Hu, Shiliang Zhang, and Xiaoyu Wang. Parcnv2: Oversized kernel with enhanced attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5762, 2023. 8
- [50] Yixing Xu, Yunhe Wang, Kai Han, Yehui Tang, Shangling Jui, Chunjing Xu, and Chang Xu. Renas: Relativistic evaluation of neural architecture search. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4411–4420, 2021. 1, 3, 5
- [51] Shen Yan, Yu Zheng, Wei Ao, Xiao Zeng, and Mi Zhang. Does unsupervised architecture representation learning help neural architecture search? *Advances in Neural Information Processing Systems*, 33:12486–12498, 2020. 2, 3
- [52] Yun Yi, Haokui Zhang, Wenzhe Hu, Nannan Wang, and Xiaoyu Wang. Nar-former: Neural architecture representation learning towards holistic attributes prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7715–7724, 2023. 1, 2, 3, 5, 7, 8, 4
- [53] Yun Yi, Haokui Zhang, Rong Xiao, Nannan Wang, and Xiaoyu Wang. Nar-former v2: Rethinking transformer for universal neural network representation learning. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 2, 3, 5, 6, 4
- [54] Chris Ying, Aaron Klein, Eric Christiansen, Esteban Real, Kevin Murphy, and Frank Hutter. Nas-bench-101: Towards reproducible neural architecture search. In *International conference on machine learning*, pages 7105–7114. PMLR, 2019. 5, 6, 7, 3, 4
- [55] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do

transformers really perform badly for graph representation?
In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. [3](#), [5](#)

- [56] Jiaxuan You, Jure Leskovec, Kaiming He, and Saining Xie. Graph structure of neural networks. In *International Conference on Machine Learning*, pages 10881–10891. PMLR, 2020. [2](#)
- [57] Chris Zhang, Mengye Ren, and Raquel Urtasun. Graph hypernetworks for neural architecture search. *arXiv preprint arXiv:1810.05749*, 2018. [2](#), [3](#)
- [58] Li Lyna Zhang, Shihao Han, Jianyu Wei, Ningxin Zheng, Ting Cao, Yuqing Yang, and Yunxin Liu. Nn-meter: Towards accurate latency prediction of deep-learning model inference on diverse edge devices. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, pages 81–93, 2021. [1](#), [2](#), [6](#)