

---

# Do Not Train It: A Linear Neural Architecture Search of Graph Neural Networks

---

Peng Xu<sup>1\*</sup> Lin Zhang<sup>2\*</sup> Xuanzhou Liu<sup>3</sup> Jiaqi Sun<sup>3</sup> Yue Zhao<sup>4</sup> Haiqin Yang<sup>2</sup> Bei Yu<sup>1</sup>

## Abstract

Neural architecture search (NAS) for Graph neural networks (GNNs), called NAS-GNNs, has achieved significant performance over manually designed GNN architectures. However, these methods inherit issues from the conventional NAS methods, such as high computational cost and optimization difficulty. More importantly, previous NAS methods have ignored the uniqueness of GNNs, where GNNs possess expressive power without training. With the randomly-initialized weights, we can then seek the optimal architecture parameters via the sparse coding objective and derive a novel NAS-GNNs method, namely neural architecture coding (NAC). Consequently, our NAC holds a no-update scheme on GNNs and can efficiently compute in linear time. Empirical evaluations on multiple GNN benchmark datasets demonstrate that our approach leads to state-of-the-art performance, which is up to  $200\times$  faster and 18.8% more accurate than the strong baselines.

## 1. Introduction

Remarkable progress of graph neural networks (GNNs) has boosted research in various domains, such as traffic prediction, recommender systems, etc., as summarized in (Wu et al., 2021). The central paradigm of GNNs is to generate node embeddings through the message-passing mechanism (Hamilton, 2020), including passing, transforming, and aggregating node features across the input graph. Despite its effectiveness, designing GNNs requires laborious efforts to choose and tune neural architectures for different tasks

<sup>\*</sup>Equal contribution <sup>1</sup>The Chinese University of Hong Kong, China <sup>2</sup>International Digital Economy Academy, China <sup>3</sup>Shenzhen International Graduate School, Tsinghua University, China <sup>4</sup>Carnegie Mellon University, USA. Correspondence to: Bei Yu <byu@cse.cuhk.edu.hk>, Lin Zhang <linzhang0529@gmail.com>.

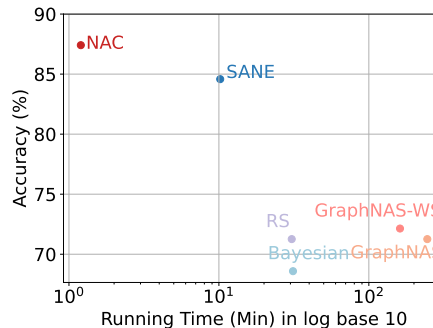


Figure 1: Accuracy vs. running time on Cora. NAC (ours) outperforms the leading methods significantly in both accuracy and speed (in minutes).

and datasets (You et al., 2020), which limits the usability of GNNs. To automate the process, researchers have made efforts to leverage neural architecture search (NAS) (Liu et al., 2019a; Zhang et al., 2021b) for GNNs, including GraphNAS (Gao et al., 2020), Auto-GNN (Zhou et al., 2019), PDNAS (Zhao et al., 2020) and SANE (Zhao et al., 2021b). In this work, we refer the problem of NAS for GNNs as *NAS-GNNs*.

While NAS-GNNs have shown promising results, they inherit issues from general NAS methods and fail to account for the unique properties of GNN operators. It is important to understand the difficulty in general NAS training (e.g., architecture searching and weight evaluation). Based on the searching strategy, NAS methods can be categorized into three types: reinforcement learning-based methods (Zoph & Le, 2017), evolutionary algorithms-based methods (Jozefowicz et al., 2015), and differential-based methods (Liu et al., 2019a; Wu et al., 2019a). Both reinforcement learning-based and evolutionary algorithm-based methods suffer from high computational costs due to the need to re-train sampled architectures from scratch. On the contrary, the weight-sharing differential-based paradigm reuses the neural weights to reduce the search effort and produces the optimal sub-architecture directly without excessive processes, such as sampling, leading to significant computational cost reduction and becoming the new frontier of NAS.

However, the weight sharing paradigm requires the neural weights to reach optimality so as to obtain the optimal sub-architecture based on its bi-level optimization (BLO)

strategy (Liu et al., 2019a), which alternately optimizes the network weights (outputs of operators) and architecture parameters (importance of operators). First, it is hard to achieve the optimal neural weights in general due to the curse of dimensionality in deep learning, leading to unstable searching results, also called the optimization gap (Xie et al., 2022). Second, this paradigm often shows a sloppy gradient estimation (Bi et al., 2020a;b; Guo et al., 2020b) due to the alternating optimization, softmax-based estimation, and unfairness in architecture updating. This type of work suffers from slow convergence during training and is sensitive to initialization due to the wide use of early termination. If not worse, it is unclear why inheriting weights for a specific architecture is still efficient—the weight updating and sharing lack interpretability.

*Is updating the GNN weights necessary? Or, does updating weights contribute to optimal GNN architecture searching?* Existing NAS-GNN methods rely on updating the weights, and in fact, all these issues raised are due to the need to update weights to the optimum. Unlike other deep learning structures, graph neural networks behave almost linearly, so they can be simplified as linear networks while maintaining superior performance (Wu et al., 2019b). Inspired by this, we find that the untrained GNN model is nearly optimal in theory. **Note that the first paper on modern GNNs, i.e., GCN (Kipf & Welling, 2017a), already spotted this striking phenomenon in the experiment, but gained little attention.** To the best of our knowledge, our work is the first to unveil this and provide theoretical proof. The issues mentioned before may not be as much of a concern given no weight update is needed, making NAS-GNN much simpler.

In this paper, we formulate the NAS-GNN problem as a sparse coding problem by leveraging the untrained GNNs, called neural architecture coding (NAC). We prove that untrained GNNs have built-in orthogonality, making the output dependent on the linear output layer. With no-update scheme, we only need to optimize the architecture parameters, resulting in a single-level optimization strategy as opposed to the bi-level optimization in the weight-sharing paradigm, which reduces the computational cost significantly and improves the optimization stability. Much like the sparse coding problem (Zhang et al., 2015), our goal is also to learn a set of sparse coefficients for selecting operators when treating these weights collectively as a dictionary, making sharing weights straightforward and understandable. Through extensive experiments on multiple challenging benchmarks, we demonstrate that our approach is competitive with the state-of-the-art baselines, while decreasing the computational cost significantly, as shown in Figure 1.

In summary, our main contributions are:

- **Problem Formulation:** We present (to our best knowledge) the first linear complexity NAS algorithm for GNNs,

namely NAC, which is solved by sparse coding.

- **Theoretical Analysis:** Our NAC holds a no-update scheme, which is theoretically justified by the built-in model linearity in GNNs and orthogonality in the model weights.
- **Effectiveness and Efficiency:** We compare NAC with state-of-the-art baselines and show superior performance in both accuracy and speed. Especially, NAC brings up to 18.8% improvement in terms of accuracy and is 200× faster than baselines.

## 2. Related Work and Preliminaries

**Graph Neural Networks** (GNNs) are powerful representation learning techniques (Xu et al., 2019) with many key applications (Hamilton et al., 2017). Early GNNs are motivated from the spectral perspective, such as Spectral GNN (Bruna et al., 2014) that applies the Laplacian operators directly. ChebNet (Defferrard et al., 2016) approximates these operators using summation instead to avoid a high computational cost. GCN (Kipf & Welling, 2017b) further simplifies ChebNet by using its first order, and reaches the balance between efficiency and effectiveness, revealing the message-passing mechanism of modern GNNs. Concretely, recent GNNs aggregate node features from neighbors and stack multiple layers to capture long-range dependencies. For instance, GraphSAGE (Hamilton et al., 2017) concatenates nodes features with mean/max/LSTM pooled neighbouring information. GAT (Velickovic et al., 2018) aggregates neighbor information using learnable attention weights. GIN (Xu et al., 2019) converts the aggregation as a learnable function based on the Weisfeiler-Lehman test instead of prefixed ones as other GNNs, aiming to maximize the power of GNNs.

GNNs consist of two major components, where the *aggregation* step aggregates node features of target nodes’ neighbors and the *combination* step passes previous aggregated features to networks to generate node embeddings. Mathematically, we can update node  $v$ ’s embedding at the  $l$ -th layer by

$$\mathbf{h}_v^l = \phi(\mathbf{W}^l \cdot o(\{\mathbf{h}_u^{l-1}, \forall u \in N(v)\})), \quad (1)$$

where  $N(v)$  denotes the neighbours of  $v$ .  $\mathbf{W}^l$  denotes the trainable weight shared by all nodes at the  $l$ -layer.  $\phi$  is an activation function (with ReLU as default). A main difference of different GNNs lies in the design of the aggregation functions,  $o(\cdot)$ . As GNNs have a large design space (You et al., 2020), this makes it challenging for the architecture search.

**Neural Architecture Search** (Xie & Yuille, 2017) generally includes three major parts, i.e., search spaces, search strategies, and evaluation methods. Recent advances adopt weight-sharing (Xie et al., 2022) to avoid (re-)training from scratch as in previous methods, where any architecture can share the weights directly for evaluation without re-training.

In particular, making the weights and architectural parameters differentiable (Liu et al., 2019a; Luo et al., 2018) further improves the training efficiency. Although efficient, weight-sharing NAS has genetic issues that are caused by its bi-level optimization, including over-fitting and inaccurate gradient estimation (Xie et al., 2022), making these methods collapse unexpectedly. A variety of methods are proposed to address this from different perspectives, while mainly adding regularization towards weights, such as using  $k$ -shot (Su et al., 2021), paths dropout (Bender et al., 2018), sampling paths (Guo et al., 2020a), sparsity (Zhang et al., 2021a). However, these cures are still within the bi-level optimization framework, which can at best alleviate the problem but not directly tackle it. In this work, we propose NAC, a linear (i.e., single-level) optimization paradigm, to address the issues in weight-sharing NAS.

Concretely, consider WS-NAS defined by  $\{\mathcal{A}, \mathcal{W}\}$ , where  $\mathcal{A}$  and  $\mathcal{W}$  denote the search space and weights of neural networks, respectively. WS-NAS addresses two subproblems of these two space alternatively using a bi-level optimization framework (Liu et al., 2019a), including weight optimization and architecture optimization as shown in the following.

$$\begin{aligned} \alpha^* &= \operatorname{argmax}_{\alpha} \mathcal{M}_{val}(w^*(\alpha), \alpha), \\ \text{s.t. } w^*(\alpha) &= \operatorname{argmin}_w \mathcal{L}_{train}(\alpha, w), \end{aligned} \quad (2)$$

where  $\mathcal{L}$  denotes a loss function in the training set.  $\mathcal{M}$  denotes an evaluation metric, such as accuracy, in the validation set. Optimizing separately, WS-NAS avoids the potential biased caused by imbalance dimensionality between  $w$  and  $\alpha$ , i.e., the neural weights and the architecture parameters. DARTS (Liu et al., 2019a) further converts the discrete search space into a differentiable one with the help of the softmax function that is writing as  $\frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in \mathcal{O}} \exp(\alpha_{o'}^{(i,j)})}$ , where  $\mathcal{O}$  represents the search space consisting of all operators. Thus, we can get the importance of each operator in the format of probability. However, it is difficult to reach the optimal state for the objectives defined in Equation (2) simultaneously, because achieving optimality in deep neural networks is still an open question, and the gradient w.r.t.  $\alpha$  is intractable (Xie et al., 2022). In practice, researchers often apply early stopping to tackle it, which leads to unexpected model collapse (Zela et al., 2020) and initialization sensitivity (Xie et al., 2022). More related work can be found at Appendix A.13.

**NAS for Graph Neural Networks:** Many NAS methods have been proposed and adapted for GNNs, such as GraphNAS (Gao et al., 2020), Auto-GNN (Zhou et al., 2019), and SANE (Zhao et al., 2021b), and thus enable learning an adaptive GNN architecture with less human effort. Methods like Auto-GNN and GraphNAS adopt reinforcement learning (RL) for searching, and thus suffer from the ex-

pensive computational cost. Thanks to the weight-sharing policy, SANE (Zhao et al., 2021b) avoids heavy retraining processes in RL methods, which leads to superior performance in general. However, this type of work also faces issues from the weight-sharing differential framework, such as bias optimization and instability due to bi-level optimization. In this work, our proposed method NAC focuses on the NAS for GNNs, and we use GraphNAS and SANE as our main baselines.

**Sparse Coding** (Zhang et al., 2015) is a well-studied topic, and its goal is to learn an informative representation as a linear combination from a collection of basis or atom, which is called a *dictionary* collectively. The standard SC formulation is written as follows:

$$\min_{\mathcal{D}, \Gamma} \mathcal{L}(\mathcal{V}, \mathcal{D}, \Gamma) = \|\mathcal{V} - \mathcal{D}\Gamma\|_F^2, \quad (3)$$

$$\text{s.t. } \forall i, \|\Gamma_{\cdot i}\|_0 \leq \tau_0; \forall j, \|\mathcal{D}_{\cdot j}\|_2 = 1,$$

where  $\Gamma$  is the sparse coefficient vector,  $\mathcal{D}$  is the dictionary,  $\tau_0$  is a value to control the sparsity. The first constraint is to encourage a sparse representation on  $\Gamma$  and the second constraint is to normalize the atoms in the dictionary. NAS aims to find a small portion of operators from a large set of operators, making it a sparse coding problem in nature. We, therefore, reformulate the NAS-GNN problem as a sparse coding problem.

### 3. Network Architecture Coding (NAC) for NAS-GNN

To find the best-performing architecture, we introduce a novel NAS-GNN paradigm, Neural Architecture Coding (NAC), based on sparse Coding, which requires no update on the neural weights during training. Technical motivation for NAC stems from the observation that an untrained GNN performs well, and training it might not provide extra benefit. In the remainder of this section, we provide a theoretical analysis of why no-update GNNs are preferred (Section 3.1) and propose a new NAS-GNN paradigm based on our theorems by leveraging untrained GNNs (Section 3.2).

#### 3.1. Analysis of No-update Scheme in GNNs

In the performance analysis of GNNs, the nonlinear activation function  $\phi$  defined in Equation (1) is often skipped. Following (Wu et al., 2019b), we simplify a GNN as  $F = A \dots A(AXW_1)W_2 \dots W_L = A^L X \prod_{l=1}^L W_l$ , where  $A$  is the adjacency matrix,  $W_l$  is the neural weight at the  $l$ -th layer from a total number of layers,  $L$ . We then obtain the final output as  $Out = FW_o$  by passing a linear layer  $W_o$ . At the  $t$ -th iteration, we denote the GNN weight at the  $l$ -th layer as  $W_l(t)$  and the theoretical output layer weight as  $\tilde{W}_o(t)$ .

By initializing the network for all layers, i.e.,  $[W_{1:L}, W_o]$ ,

we aim to train the network to get the optimal weights, denoted as  $[\mathbf{W}_{1:L}^*, \mathbf{W}_o^*]$ , where we assume the optimal weight can be obtained at  $+\infty$ -th iteration:  $\mathbf{W}_l^* = \mathbf{W}_l(+\infty)$ . We use gradient descent as the default optimizer in GNNs.

We now provide our first main theorem to investigate why an untrained GNN model can attain the same performance as the optimal one.

**Theorem 3.1.** *Assume  $\mathbf{W}_l(0)$  is randomly initialized for all  $l \in [1, L]$ , if  $\prod_{l=1}^L \mathbf{W}_l(0)$  is full rank, there must exist a weight matrix for the output layer, i.e.,  $\tilde{\mathbf{W}}_o$ , that makes the final output the same as the one from a well-trained network:*

$$\mathbf{A}^L \mathbf{X} \prod_{l=1}^L \mathbf{W}_l^* \mathbf{W}_o^* = \mathbf{A}^L \mathbf{X} \prod_{l=1}^L \mathbf{W}_l(0) \tilde{\mathbf{W}}_o. \quad (4)$$

Please see Appendix A.2 for more details. This theorem implies that the output layer alone can ensure the optimality of the network even when all previous layers have no updates. Next, we show that training in the standard setting of GNNs is guaranteed to reach the desired weight of the output layer, i.e.  $\tilde{\mathbf{W}}_o$ , under mild conditions. We begin by providing the theorem from (Gunasekar et al., 2018).

**Theorem 3.2** (From (Gunasekar et al., 2018)). *Assume we have the following main conditions: 1) the loss function is either the exponential loss or the log-loss; 2) gradient descent is employed to seek the optimal solution; and 3) data is linearly separable.*

When defining a linear neural network as  $\mathbf{Y} = \mathbf{X} \mathbf{W}_1 \mathbf{W}_2 \dots \mathbf{W}_L = \mathbf{X} \prod_{l=1}^L \mathbf{W}_l = \mathbf{X} \boldsymbol{\beta}$ , we always have the same optimal weight regardless of the number of layers when using gradient descent,

$$\boldsymbol{\beta}^* = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \|\boldsymbol{\beta}\|^2, \quad \text{s.t.} \quad \mathbf{X} \boldsymbol{\beta} \odot \mathbf{s} \geq \mathbf{1}, \quad (5)$$

where  $\mathbf{s}$  is the ground-truth label vectors with elements as 1 or  $-1$ ;  $\mathbf{1}$  denotes a vector of all ones;  $\odot$  denotes the element-wise product.

This theorem suggests that the optimal weights obtained by the gradient descent are the same regardless of how many layers are in the network, which is equivalent to finding a max-margin separating hyperplane.

Following this, we prove that  $\tilde{\mathbf{W}}_o$  can be obtained based on the above conditions:

**Theorem 3.3.** *Assume a GNN model has either the exponential loss or the log-loss, the desired weight  $\tilde{\mathbf{W}}_o$  is secured when updating with gradient descent and initializing  $\prod_{l=1}^L \mathbf{W}_l(0)$  as an orthogonal matrix. Mathematically, we have  $\hat{\mathbf{W}}_o(+\infty) = \hat{\mathbf{W}}_o^* = \tilde{\mathbf{W}}_o$ .*

Please see Appendix A.2 for more details. In summary, our proposed theorems prove that a GNN with randomly

initialized weights can make the final output as good as a well-trained network, where one needs to update the overall network with gradient descent and initialize networks with orthogonal weights. Apart from the optimality, another immediate benefit of this no-update scheme is that the computational cost will be significantly reduced. This is of particular importance when the budget is limited, as the main computational cost of NAS comes from this update.

The above theorems show that one can approximate the optimal output by using the network without training, where we only need to update the final linear layer of the search space. This finding motivates us to learn architectures by taking advantage of randomly initialized networks. We want to emphasize that real-world scenarios can break the optimal conditions, such as the high complexity of data and the early stopping of the training. Even though, our strategy still provides near-optimal performance, if not the best, in experiments.

### 3.2. Architecture Searching via Sparse Coding

In this work, we treat the entire NAS search space of GNNs as a general GNN model where each layer is a mixture of multiple operators. Inspired by Balcilar et al. (2021), we propose a unified expression for each GNN layer in the search space in the following theorem, whose proof is given in Appendix A.3.

**Theorem 3.4.** *Simplifying the complex non-linear functions, each GNN layer in the search space can be unified by  $\sum_k P_l^k(L) \mathbf{H}_l \mathbf{W}_l^k$ , where  $P(\cdot)$  denotes the fixed polynomial function with subscripts indicating individual terms, and superscript  $(l)$  marks the current layer.*

This theorem suggests that all GNN layers are unified as  $\sum_k P_l^k(L) \mathbf{H}_l \mathbf{W}_l^k$ , as the output of each GNN operator shares the same expression. Recall that the non-linear activation functions  $\sigma(\cdot)$  are often neglected in the performance analysis, and even removed in some cases. Following this, we propose the following corollary and provide the details in Appendix A.3.

**Corollary 3.5.** *The search space of GNNs can be unified as  $\sum_{k=0}^K P^k(L) \mathbf{X} \mathbf{W}^k = \mathbf{D} \mathbf{W}$  when removing the activation function, where  $\mathbf{D} = \|\{P^k(L) \mathbf{X}\}$  is the fixed base, and  $\mathbf{W} = \|\{\mathbf{W}^k\}^T$  is the trainable parameters. Here,  $\|$  stands for concatenating a set of matrices horizontally.*

This corollary derives a unified mathematical expression of the search space of GNNs, implying the natural connection between NAS in GNNs and the Sparse Coding.

Building on the above theorems, we present a novel NAS-GNN paradigm in which an untrained GNN model can also yield optimal outputs under mild conditions in the following. This optimality allows us to waive the effort to design complicated models for updating weights and makes us focus on

architecture updates. At a high level, the ultimate goal of NAS is to assign the coefficients of unimportant operators as zeros and the important ones as non-zeros. Towards this, searching the optimal architecture is equivalent to finding the optimal sparse combination of GNN operators, i.e.,  $\alpha$ . Notably, one of the most successful algorithms for this purpose is sparse coding (Olshausen & Field, 1997), which has strong performance guarantees.

*What matters most when we reformulate the NAS problem as a sparse coding problem over a dictionary defined on GNN operators?* According to the research in sparse coding (Aharon et al., 2006; Tropp, 2004; Candes et al., 2006), its performance depends on the quality of the dictionary. In particular, the most desirable property is mutual coherence or orthogonality. High incoherence in the dictionary tends to avoid ambiguity in the solution and improve optimization stability, resulting in optimal or at least sub-optimal solutions. We prove a dictionary that is defined on deep neural networks has built-in orthogonality, as shown in the theorem below.

**Theorem 3.6.** *Let the neural weights of each operator in a deep neural network be an atom and stack them column-wisely, we can guarantee the existence of an orthogonal dictionary.*

Please see Appendix A.5 for more details. This theorem implies that the obtained weights are qualified to be a dictionary due to the built-in high dimensionality in deep neural networks.

Taking inspiration from this, we propose a new NAS paradigm called *neural architecture coding* as follows:

$$\min_{\alpha} \mathcal{L}(\mathbf{y}, f(\mathbf{z})) + \|\mathbf{z} - \mathbf{h}^L\|_F^2 + \rho\|\alpha\|_1, \quad (6)$$

$$\text{where } \begin{cases} \mathbf{h}_v^l = \phi[\mathbf{W}^l \cdot \bar{o}^l(\{\mathbf{h}_u^{l-1}, \forall u \in N(v)\})] \\ \bar{o}^l(x) = \mathbf{o}^l \hat{\alpha}_l = \sum_{k=1}^K \frac{\alpha_{lk}}{\|\alpha_l\|_2} o^{lk}(x) \\ \mathbf{o}^l = [o^{l1}(x), o^{l2}(x), \dots, o^{lK}(x)] \end{cases}$$

where  $\rho$  is the sparsity hyperparameter,  $f(\cdot)$  is a linear function as the final output layer, e.g., MLP,  $\mathbf{h}$  represents all nodes' embeddings. Here, we set the number of GNN operators to  $K$  and the number of hidden layers to  $L$ .  $o^{lk}$  denotes the  $k$ -th operator at the  $l$ -th layer and its output is a vector;  $\alpha_{lk}$  denotes the scalar coefficient for operator  $k$  at  $l$ -th layer and thus  $\alpha_l \in \mathbb{R}^k$ ;  $\mathbf{o}^l$  is the weight vector from all operators at the  $l$ -th layer.  $\mathbf{h}_v^l$  denotes the embedding of the node  $v$  at the  $l$ -th layer, where its neighborhood aggregation process is performed as the weighted summation of the outputs from all operators as  $\bar{o}^l(x)$ . Here,  $\mathbf{y}$  is the node label.

Notice that a standard sparse coding model uses an l2-norm-based objective for optimization, while a cross-entropy objective is preferred as we focus on a classification task. To connect the gap, we prove that optimizing the Cross-Entropy

loss function in GNNs satisfies a sufficient condition for the l2-norm-based objective in sparse coding, as shown in the following theorem.

**Theorem 3.7.** *The gradient from Cross-Entropy loss function is feasible to optimize an l2-norm objective in a gradient descent manner, such that the initial  $z^0$  converges geometrically to the optimal  $z^*$  with a tiny error, e.g.,*

$$\|z^{s+1} - z^*\| \leq (1 - 2\alpha\eta)^s \|z^0 - z^*\| + 2\gamma/\alpha, \quad (7)$$

where  $s$  signifies the current iteration, and  $\gamma, \alpha, \eta$  define the numerical bound for the error.

Please see Appendix A.4 for more details. Therefore, we reformulate the objective in Equation (6) as follows,

$$\min_{\alpha} \mathcal{L}_{CE}(\mathbf{y}, f(\mathbf{z})) + \|\mathbf{z} - \mathbf{h}^L\|_F^2 + \rho\|\alpha\|_1, \quad (8)$$

where  $\mathcal{L}_{CE}$  defines the Cross-Entropy loss function.

**Algorithm 1** The NAC algorithm

---

**Require:** The search space  $\mathcal{A}$ ;  
**Ensure:** The architecture  $\alpha$

Randomly initializing  $\mathbf{W}^l$ , for  $l = 1, \dots, L$ ; set  $\alpha = \mathbf{1}$ ;

- 1: **while**  $t = 1, \dots, T$  **do**
- 2:   Performing the feature aggregation at each layer as  $\bar{o}^l(x) = \mathbf{o}^l \hat{\alpha}_l = \sum_{k=1}^K \frac{\alpha_{lk}}{\|\alpha_l\|_2} o^{lk}(x)$ ;
- 3:   Computing  $\mathbf{h}_v^l = \phi[\mathbf{W}^l \cdot \bar{o}^l(\{\mathbf{h}_u^{l-1}, \forall u \in N(v)\})]$ ;
- 4:   Optimizing  $\alpha$  based on the objective function in Equation (8) w.r.t.  $\alpha$  under the fixed dictionary  $\mathbf{o}$ ;
- 5:   Updating  $\mathbf{W}_o$  based on the objective function in Equation (8) w.r.t.  $\mathbf{W}_o$  under fixed  $\alpha$ ;
- 6: **end while**
- 7: Obtain the final architecture  $\{\alpha^*\}$  from the trained  $\alpha$  via an argmax operation at each layer;

---

It is important to note that, rather than computing the optimal neural weight for  $o^{lk}$  like previous NAS methods explicitly, we use randomly initialized weights instead and focus on optimizing the combination coefficient  $\alpha$ , i.e. architecture parameters. The desired optimal embedding  $\mathbf{z}$  is then used for downstream applications, such as node classification. In practice, we often let  $\mathbf{z} = \mathbf{h}^L$  for computational convenience, which makes the loss of the second term to be zero.

The sparsity regularization  $\|\alpha\|_1$  allows us to rank the importance of the operator directly, alleviating inaccurate gradient estimation such as softmax (Garg et al., 2021). We want to emphasize that this regularization does not break the requirement of using gradient descent. Since the architecture, i.e.,  $\alpha$ , is fixed when updating the output layer weights, we can use gradient descent for updating as required, so the theorems in Section 3.1 hold.

### 3.3. Implementation Details and Complexity Analysis

The overall algorithm is presented at Algorithm 1. Computing NAC consists of two major parts: the forward pass and

the backward pass. Given the search space, the computation of the forward is then fixed and regarded as a constant. Therefore, the computational complexity mainly focuses on the backward pass in the NAC algorithm.

In summary, the algorithmic complexity of NAC is  $O(T * |\alpha|)$ , where  $|\alpha|$  is the size of  $\alpha$ . Please refer to Appendix A.7 for more details about the implementation.

## 4. Experiments

We conduct experiments to address the following issues: (1) *How does NAC perform in comparison to the leading baselines?* (2) *How does the no-update scheme affect other methods?* (3) *Is NAC robust?* Section 4.2-Section 4.4 answer the above questions accordingly.

### 4.1. Experiment Setup

**Datasets.** We performed experiments on transductive learning for node classification (Zhao et al., 2021b). For this setting, we use four benchmark datasets: CiteSeer (Giles et al., 1998), Cora (Kipf & Welling, 2016), PubMed (Sen et al., 2008), and Computers (McAuley et al., 2015). Also, we follow the data partition setting (training/validation/testing) as in (Zhao et al., 2021b). For more details about transductive learning task and introduction of each dataset, please refer to Appendix A.1 for more details.

We also conduct additional experiments on inductive task like PPI Appendix A.11 dataset (Hamilton et al., 2017) and the Open Graph Benchmark datasets (Hu et al., 2020) in Appendix A.11.

**Environment.** For datasets of Citeseer, Cora, and PubMed, we run experiments on a CPU server with a 48-core Intel Xeon Platinum 8260L CPU using PyTorch (Paszke et al., 2019) (version 1.10.2+cpu). Due to the large size of Amazon Computers, we experiment on a GPU server with four NVIDIA 3090 GPUs (24G) with PyTorch (version 1.10.2+gpu). All operators used in the experiments are from the built-in functions of PyG (version 2.0.2) (Fey & Lenssen, 2019).

**Methods.** Following (Zhao et al., 2021b), we compare our NAC with the following strong baselines: 1) Random Search (RS) (Bergstra & Bengio, 2012) and 2) Bayesian Optimization (BO) (Jones et al., 1998), 3) GraphNAS (Gao et al., 2020): a popular reinforcement learning-based (RL) method, 4) GraphNAS-WS (Zhao et al., 2021b): a variant of GraphNAS with weight sharing, and 5) SANE (Zhao et al., 2021b). Among all baselines, the first two are hyperparameter optimization (HPO) methods (Yu & Zhu, 2020). GraphNAS and GraphNAS-WS are two popular methods following the weight sharing paradigm while SANE is the most recent work on automated GNNs using the weight-sharing differen-

tial paradigm, which is the closest one to our work. We also evaluate NAC-updating, i.e., NAC with weight updates, to compare it with the proposed one with no-update scheme.

**Search space.** We select a list of node aggregators (i.e., operators) in GNNs to form the search space. More specifically, they include the following seven aggregators: multi-layer perceptrons (MLP), GCN (Kipf & Welling, 2017b), GAT (Velickovic et al., 2018), GIN (Xu et al., 2019), Geniepath (Liu et al., 2019b), Graphsage (Hamilton et al., 2017), and ChebNet (Defferrard et al., 2016). Note that MLP is a special aggregator with fixed frequency when viewing GNNs from a spectrum perspective (Balcilar et al., 2020). Besides, Graphsage and GAT contain different aggregation functions, i.e., maximum and LSTM. We apply all these variations as operators in the experiments. Note that we do not include the layer aggregators like skip-connection in our search space. The underlying rationale is that recent research (Zela et al., 2020; Chen & Hsieh, 2020; Liang et al., 2019) finds that NAS models converge to skip-connection much faster than other operators due to its parameter-free nature. This introduces biased searching and yields poor generalization. *Our experiments ensure that the search space is the same for all baselines.*

We do not include layer-wise aggregators like JKNNets (Xu et al., 2018) in the main experiment. As we show in Section 3.1, our theory is built upon the classical GNN framework, which equals matrix multiplication. JKNNets cannot be represented in this framework, which contradicts our assumptions.

**Evaluation Metrics.** The proposed approaches aim to boost the ability to learn efficacious architecture for downstream tasks. The classification accuracy and the runtime (i.e., wall-clock time) are two widely used metrics in NAS research to measure the model performance and efficiency, respectively (Zhang et al., 2021a). The reported results are the average of the best results in 4 runs from different random seeds.

We set the runtime of NAC, SANE, and GraphNAS as the time they take to obtain the final architectures with 100 epochs<sup>1</sup>. The runtime of RS and Bayesian is fixed to the time of 100 times trial-and-error processes.

### 4.2. Comparison Results

Results in Table 1 and Figure 1 show that NAC attains superior performance than all baselines regarding both *accuracy* and *efficiency*. More specifically, we observe the following:

- In terms of model performance, our NAC beats all baselines and attains up to 2.83% improvement over the best

<sup>1</sup>The running time for the Computers dataset is measured on a single GPU, which is different from all other three datasets.

Table 1: Experimental results on the compared methods: our NAC attains superior performance in both accuracy (%) and efficiency (in minutes).

	CiteSeer		Cora		PubMed		Computers	
	Accuracy	Time	Accuracy	Time	Accuracy	Time	Accuracy	Time
RS	70.12 $\pm$ 2.36	14.4	71.26 $\pm$ 4.68	30.6	86.75 $\pm$ 0.82	187.8	77.84 $\pm$ 1.35	8.75
BO	70.95 $\pm$ 1.62	18	68.59 $\pm$ 6.66	31.2	87.42 $\pm$ 0.68	189.6	77.46 $\pm$ 2.02	17.65
GraphNAS	68.69 $\pm$ 1.30	253.8	71.26 $\pm$ 4.90	245.4	86.07 $\pm$ 0.51	1363.8	73.97 $\pm$ 1.79	86.37
GraphNAS-WS	65.35 $\pm$ 5.13	80.4	72.14 $\pm$ 2.59	161.4	85.71 $\pm$ 1.05	965.4	72.99 $\pm$ 3.44	42.47
SANE	71.84 $\pm$ 1.33	4.2	84.58 $\pm$ 0.53	10.2	87.55 $\pm$ 0.78	107.4	90.70 $\pm$ 0.89	0.72
NAC	<b>74.62</b> $\pm$ 0.38	<b>1.2</b>	<b>87.41</b> $\pm$ 0.92	<b>1.2</b>	88.04 $\pm$ 1.06	<b>9.0</b>	<b>91.64</b> $\pm$ 0.14	<b>0.23</b>
NAC-updating	74.17 $\pm$ 1.18	4.2	86.62 $\pm$ 1.14	3.6	<b>88.10</b> $\pm$ 0.86	25.8	90.89 $\pm$ 1.10	0.70

baseline, i.e., SANE, while attaining up to 18.8% improvement over the Bayesian method, the best HPO method. Thanks to our non-updating scheme, it prompts the outputs near the optimal to make the selection of the best-performing architecture reliable, as opposed to biased optimization in others. We notice that in one case, NAC is slightly worse than NAC-updating, about 0.06%. This case reminds us that the optimal condition can get compromised due to the high complexity of the data. Still, NAC holds for near-best performance, demonstrating its robustness. Empirical results corroborate our theories in Section 3.1 that non-updating scheme is preferred. The underlying assumption of RL-based and BO-based methods is to obtain an accurate estimation of the distribution of the search space. However, these methods rely on random sampling to perform the estimation, thus have no guarantee of the search quality due to the high complexity in the search space. For instance, RL-based methods are sometimes even worse than RS-based methods, pure randomly sampled, implying the unreliability of such estimation under a limited budget. WS-based methods couple the architecture search and the weight training, so they struggle to train the network to be optimal and obtain a biased network due to oscillations between the two components.

- In terms of model efficiency, our NAC achieves superior performance, around 10 $\times$  faster than SANE and up to 200 $\times$  time faster than GraphNAS. Our non-updating scheme requires nearly no weight update, thus giving NAC an incomparable advantage in reducing the computation cost. All previous NAS methods need to update neural weights, for example, RL-based costs the most due to re-training the network from scratch at each time step, and WS-based costs the least by reusing neural weights.

### 4.3. No-update Scheme at Work

To further validate our no-update scheme, we evaluate its effect on other weight-sharing methods. Since SANE attains the best performance among all baselines, we test the no-update scheme on SANE as SANE\*, i.e., SANE with fixed neural weights. Results in Table 3 show that SANE\*

outperforms the one with updates. This result implies that we can improve the performance of NAS-GNN methods by simply fixing the weights with random initialization. This yields a much lower computational cost in the training.

### 4.4. Analysis of Convergence

A notable benefit of the NAC framework is its guaranteed convergence from the sparse coding perspective. To verify this, Figure 3 offers a convergence comparison between NAC and SANE from the Pubmed dataset by showing the accuracy on the retrained network acquired at each training epoch. NAC shows substantially faster convergence than SANE, which take around 20 and 80 epochs, respectively. Due to the BLO strategy, SANE suffers strong oscillation when optimizing two variables, i.e., the architecture and neural weights, which consumes more epochs to converge.

### 4.5. Ablation Studies

**Initialization:** Recall that in Section 3.1, a prerequisite is the orthogonality on the neural weights matrix. This makes the initialization critical to the final performance. To investigate the effect of initialization, we try three types of initialization methods, including normal, uniform, and orthogonal. Table 4 shows that NAC with orthogonal initialization outperforms other two cases of initialization. The results again confirm our theory.

To get a better understanding of the reason, we visualize the eigenvalues of  $\prod_{l=1}^L \mathbf{W}_l$  on different initialization cases in Figure 2. The results show that the eigenvalues of NACs with normal and uniform initialization decay accordingly. This makes the input feature tends to project to the subspace spanned by the top-ranked eigenvectors, which results in feature distortion and leads to performance degeneration. In contrast, when NAC with orthogonal initialization, the eigenvalues are distributed uniformly, which ensures the spanned space does not bias toward any particular eigenvectors.

In summary, NAC with orthogonal initialization attains the best performance and confirms our theory. In addition, re-

Table 2: Performance comparison between SANE, NAC<sup>+</sup> and NAC.

	CiteSeer Acc(%)	Cora Acc(%)	PubMed Acc(%)	Computers Acc(%)
SANE	71.84 $\pm$ 1.33	84.58 $\pm$ 0.53	87.55 $\pm$ 0.78	90.70 $\pm$ 0.89
NAC <sup>+</sup>	71.76 $\pm$ 2.08	87.27 $\pm$ 1.20	87.59 $\pm$ 0.22	91.07 $\pm$ 0.71
NAC	<b>74.62</b> $\pm$ 0.38	<b>87.41</b> $\pm$ 0.92	<b>88.04</b> $\pm$ 1.06	<b>91.64</b> $\pm$ 0.14

Table 3: Comparison between SANE and SANE\* (w/o. weight updates).

	CiteSeer Acc(%)	Cora Acc(%)	Pubmed Acc(%)	Computers Acc(%)
SANE	71.84 $\pm$ 1.33	84.58 $\pm$ 0.53	87.55 $\pm$ 0.78	90.70 $\pm$ 0.89
SANE*	<b>71.95</b> $\pm$ 1.32	<b>85.46</b> $\pm$ 0.76	<b>88.12</b> $\pm$ 0.35	<b>90.86</b> $\pm$ 0.80

Table 4: Comparison of NAC with different initializations.

	CiteSeer Acc(%)	Cora Acc(%)	PubMed Acc(%)	Computers Acc(%)
Kaiming Normal	71.12 $\pm$ 2.45	86.85 $\pm$ 0.78	87.52 $\pm$ 0.72	91.05 $\pm$ 0.80
Kaiming Uniform	72.06 $\pm$ 2.24	86.99 $\pm$ 1.03	87.68 $\pm$ 0.97	88.52 $\pm$ 3.01
Orthogonal	<b>74.62</b> $\pm$ 0.38	<b>87.41</b> $\pm$ 0.92	<b>88.04</b> $\pm$ 1.06	<b>91.64</b> $\pm$ 0.14

sults from non-orthogonal initializations are still better than other baselines, demonstrating the robustness of NAC.

**The linear output layer:** Recall that for a given linear output layer, there exists the optimal weight, i.e.  $\tilde{W}_o$ , to secure the optimal output according to Theorem 3.1. Nevertheless, obtaining optimal weights in a deep learning algorithm is an open question due to its high dimensionality. Often, in practice, we train the network with early stopping, especially in NAS methods. Therefore, obtaining the optimal weights for the output layer is not possible. *Is training the output layer necessary?* This experiment attempts to answer this by comparing NAC with and without training the output layer. Note that we fix all other layers with random weights. By default, NAC is the one without training any layers in the whole paper, including the output layer. On the contrary, NAC<sup>+</sup> is the one with training the output layer. Here, we set the number of epochs as 100.

In Table 2, we can find that the performance of NAC<sup>+</sup> is in the middle between SANE and NAC. On the one hand, training the linear layer may not be worthwhile because it is impossible to get the desired optimal weights. On the other hand, results with and without training the output layer are always better than SANE, suggesting the superiority of the no-update scheme in general. To further demonstrate this, we implement multiple experiments by training the last layer with different epochs as shown in Appendix A.8.3.

**Sparsity effect:** Our model uses the hyperparameter  $\rho$  to control the sparsity of the architecture parameter  $\alpha$ . We investigate the effect of this hyperparameter by varying its value in a considerably large range, such as [0.001, 10]. We find that the results vary little in a considerably wide range,

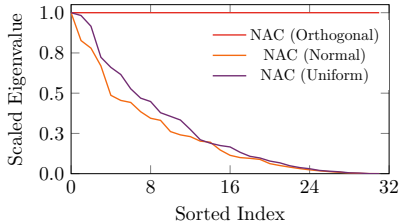


Figure 2: Visualization of the sorted eigenvalues of weights of NAC with different initialization.

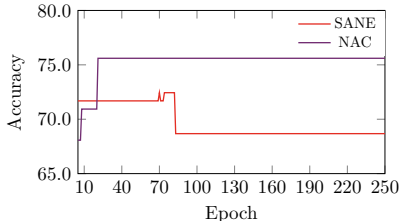


Figure 3: Convergence for SANE and NAC in terms of accuracy. NAC converges much faster than SANE in only around 20 epochs.

this indicates NAC is insensitive to the sparsity hyperparameter.

**Random seeds:** Random seeds often play an importance role in NAS methods as they also affect the initialization significantly. By default, researchers report the average or the best result from different random seeds, this often lead to poor reproducibility. We now present the effect of random seeds on this topic. In particular, we implement the experiment on multiple random seeds and observe a stable performance of NAC. The experimental results show that NAC is robust to both sparse hyperparameter selection and random seeds. Please refer to Appendix A.8.2 for more details.

## 5. Conclusion

We present the first linear complexity NAS algorithm for GNNs, namely NAC. Our NAC is solved by sparse coding and holds no-update scheme for model weights in GNNs because GNNs embed the built-in model linearity and orthogonality for the model weights. Extensive experiments demonstrate that our NAC can achieve higher accuracy (up to 18.8%) and much faster convergence (up to 200 $\times$ ) than SOTA NAS-GNNs baselines.

Several promising directions can be considered. We can further explore more deep neural networks that satisfy the mild condition of our NAC to extent its usability. We can investigate more on the efficiency of different subgradient methods for solving the sparse coding objective. We also intend to investigate how to jointly learn the search space and architectural representation to further enhance the expressive ability of searched architectures.



## References

- Aharon, M., Elad, M., and Bruckstein, A. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE TSP*, 54(11):4311–4322, 2006.
- Arora, S., Ge, R., Ma, T., and Moitra, A. Simple, efficient, and neural algorithms for sparse coding. In *Proc. COLT*, pp. 113–149, 2015.
- Balcilar, M., Renton, G., Héroux, P., Gaüzère, B., Adam, S., and Honeine, P. Analyzing the expressive power of graph neural networks in a spectral perspective. In *Proc. ICLR*, 2020.
- Balcilar, M., Renton, G., Héroux, P., Gaüzère, B., Adam, S., and Honeine, P. Analyzing the expressive power of graph neural networks in a spectral perspective. In *Proc. ICLR*, 2021.
- Bender, G., Kindermans, P.-J., Zoph, B., Vasudevan, V., and Le, Q. Understanding and simplifying one-shot architecture search. In *Proc. ICML*, 2018.
- Bergstra, J. and Bengio, Y. Random search for hyperparameter optimization. *Journal of Machine Learning Research*, 13:281–305, 2012.
- Bi, K., Hu, C., Xie, L., Chen, X., Wei, L., and Tian, Q. Stabilizing darts with amended gradient estimation on architectural parameters. *CoRR*, abs/1910.11831, 2020a.
- Bi, K., Xie, L., Chen, X., Wei, L., and Tian, Q. Gold-nas: Gradual, one-level, differentiable. *CoRR*, abs/2007.03331, 2020b.
- Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. Spectral networks and locally connected networks on graphs. In *Proc. ICLR*, 2014.
- Candes, E., Romberg, J., and Tao, T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory (TIT)*, 52(2):489–509, 2006.
- Chen, W., Gong, X., and Wang, Z. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. *arXiv preprint arXiv:2102.11535*, 2021.
- Chen, X. and Hsieh, C. Stabilizing differentiable architecture search via perturbation-based regularization. In *Proc. ICML*, 2020.
- Defferrard, M., Bresson, X., and Vandergheynst, P. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proc. NIPS*, 2016.
- Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop*, 2019.
- Fischer, H. *A history of the central limit theorem: from classical to modern probability theory*. Springer, 2011.
- Fu, W. J. Penalized regressions: the bridge versus the lasso. *Journal of Computational and Graphical Statistics*, 7(3): 397–416, 1998.
- Gao, Y., Yang, H., Zhang, P., Zhou, C., and Hu, Y. Graph neural architecture search. In *Proc. IJCAI*, 2020.
- Garg, S., Tosatto, S., Pan, Y., White, M., and Mahmood, A. R. An alternate policy gradient estimator for softmax policies. *arXiv preprint arXiv:2112.11622*, 2021.
- Giles, C. L., Bollacker, K. D., and Lawrence, S. Citeseer: An automatic citation indexing system. In *Proc of Conference on Digital Libraries*, pp. 89–98, 1998.
- Gunasekar, S., Lee, J. D., Soudry, D., and Srebro, N. Implicit bias of gradient descent on linear convolutional networks. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Proc. NIPS*, volume 31, 2018.
- Guo, R., Lin, C., Li, C., Tian, K., Sun, M., Sheng, L., and Yan, J. Powering one-shot topological NAS with stabilized share-parameter proxy. In *Proc. ECCV*, 2020a.
- Guo, Z., Zhang, X., Mu, H., Heng, W., Liu, Z., Wei, Y., and Sun, J. Single path one-shot neural architecture search with uniform sampling. In *Proc. ECCV*, 2020b.
- Hamilton, W. L. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020.
- Hamilton, W. L., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Proc. NIPS*, 2017.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. *Proc. NIPS*, 33: 22118–22133, 2020.
- Jones, D. R., Schonlau, M., and Welch, W. J. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4), 1998.
- Jozefowicz, R., Zaremba, W., and Sutskever, I. An empirical exploration of recurrent network architectures. In *Proc. ICML*, 2015.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *Proc. ICLR*, 2017a.

- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *Proc. ICLR*, 2017b.
- Liang, H., Zhang, S., Sun, J., He, X., Huang, W., Zhuang, K., and Li, Z. DARTS+: improved differentiable architecture search with early stopping. *CoRR*, abs/1909.06035, 2019.
- Liu, H., Simonyan, K., and Yang, Y. DARTS: differentiable architecture search. In *Proc. ICLR*, 2019a.
- Liu, Z., Chen, C., Li, L., Zhou, J., Li, X., Song, L., and Qi, Y. Geniepath: Graph neural networks with adaptive receptive paths. In *Proc. AAAI*, 2019b.
- Luo, R., Tian, F., Qin, T., Chen, E., and Liu, T. Neural architecture optimization. In *Proc. NIPS*, 2018.
- McAuley, J., Targett, C., Shi, Q., and Van Den Hengel, A. Image-based recommendations on styles and substitutes. In *Proc. SIGIR*, pp. 43–52, 2015.
- Nguyen, T., Wong, R., and Hegde, C. A provable approach for double-sparse coding. In *Proc. AAAI*, 2018.
- Olshausen, B. A. and Field, D. J. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23):3311–3325, 1997.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *Proc. NIPS*, 32, 2019.
- Qin, Y., Zhang, Z., Wang, X., Zhang, Z., and Zhu, W. Nas-bench-graph: Benchmarking graph neural architecture search. *arXiv preprint arXiv:2206.09166*, 2022.
- Ranstam, J. and Cook, J. Lasso regression. *Journal of British Surgery*, 105(10):1348–1348, 2018.
- Schmidt, M., Fung, G., and Rosales, R. Optimization methods for l1-regularization. *University of British Columbia, Technical Report TR-2009-19*, 2009.
- Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., and Eliassi-Rad, T. Collective classification in network data. *AI Magazine*, 29(3):93, 2008.
- Shevade, S. K. and Keerthi, S. S. A simple and efficient algorithm for gene selection using sparse logistic regression. *Bioinformatics*, 19(17):2246–2253, 2003.
- Shu, Y., Cai, S., Dai, Z., Ooi, B. C., and Low, B. K. H. Nasi: Label-and data-agnostic neural architecture search at initialization. *arXiv preprint arXiv:2109.00817*, 2021.
- Su, X., You, S., Zheng, M., Wang, F., Qian, C., Zhang, C., and Xu, C. K-shot NAS: learnable weight-sharing for NAS with k-shot supernet. In *Proc. ICML*, 2021.
- ter Haar, D. Mathematical foundations of statistical mechanics. a. i. khinchin. *Science*, 110(2865):570–570, 1949.
- Tropp, J. Greed is good: algorithmic results for sparse approximation. *IEEE Transactions on Information Theory (TIT)*, 50(10):2231–2242, 2004.
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *Proc. ICLR*, 2018.
- Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., and Keutzer, K. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proc. CVPR*, 2019a.
- Wu, F., Souza, A., Zhang, T., Fifty, C., Yu, T., and Weinberger, K. Simplifying graph convolutional networks. In *Proc. ICML*, pp. 6861–6871, 2019b.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Yu, P. S. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):4–24, 2021.
- Xie, L. and Yuille, A. L. Genetic CNN. In *Proc. ICCV*, 2017.
- Xie, L., Chen, X., Bi, K., Wei, L., Xu, Y., Wang, L., Chen, Z., Xiao, A., Chang, J., Zhang, X., and Tian, Q. Weight-sharing neural architecture search: A battle to shrink the optimization gap. *ACM Computing Surveys*, 54(9):183:1–183:37, 2022.
- Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K.-i., and Jegelka, S. Representation learning on graphs with jumping knowledge networks. In *Proc. ICML*, volume 80, pp. 5453–5462, 2018.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *Proc. ICML*, 2019.
- You, J., Ying, Z., and Leskovec, J. Design space for graph neural networks. *Proc. NIPS*, 33, 2020.
- Yu, T. and Zhu, H. Hyper-parameter optimization: A review of algorithms and applications. *CoRR*, abs/2003.05689, 2020.
- Zela, A., Elsken, T., Saikia, T., Marrakchi, Y., Brox, T., and Hutter, F. Understanding and robustifying differentiable architecture search. In *Proc. ICLR*, 2020.
- Zhang, X., Huang, Z., Wang, N., Xiang, S., and Pan, C. You only search once: Single shot neural architecture search via direct sparse optimization. *IEEE TPAMI*, 43(9):2891–2904, 2021a.

- Zhang, Z., Xu, Y., Yang, J., Li, X., and Zhang, D. A survey of sparse representation: algorithms and applications. *IEEE Access*, 3:490–530, 2015.
- Zhang, Z., Wang, X., and Zhu, W. Automated machine learning on graphs: A survey. In *Proc. IJCAI*, 2021b.
- Zhao, H., Wei, L., quanming yao, and He, Z. Efficient graph neural architecture search, 2021a. URL <https://openreview.net/forum?id=IjIzIOkK2D6>.
- Zhao, H., Yao, Q., and Tu, W. Search to aggregate neighborhood for graph neural network. In *Proc. ICDE*, 2021b.
- Zhao, Y., Wang, D., Gao, X., Mullins, R., Lio, P., and Jamnik, M. Probabilistic dual network architecture search on graphs. *arXiv preprint arXiv:2003.09676*, 2020.
- Zhou, K., Song, Q., Huang, X., and Hu, X. Auto-gnn: Neural architecture search of graph neural networks. *CoRR*, abs/1909.03184, 2019.
- Zoph, B. and Le, Q. V. Neural architecture search with reinforcement learning. In *Proc. ICLR*, 2017.

## A. Appendix

### A.1. Dataset Details

In the transductive learning task, node features and edges between nodes in the whole datasets were known beforehand. We learn from the already labeled training dataset and then predict the labels of the testing dataset. Three benchmark datasets were used in this setting: CiteSeer (Giles et al., 1998), Cora (Kipf & Welling, 2016) and PubMed (Sen et al., 2008). All of the three benchmark datasets are citation networks. In citation networks, each node represents a work, and each edge shows the relationship between two papers in terms of citations. The datasets contain bag-of-words features for each node, and the goal is to categorize papers into various subjects based on the citation. In addition to the three benchmark datasets, we also employ two another datasets: Amazon Computers (McAuley et al., 2015) and ogbn-arxiv dataset (Hu et al., 2020). Amazon Computers is a subset of the Amazon co-purchase graph, where nodes represent commodities and edges connect them if they are frequently purchased together. Product reviews are encoded as bag-of-word feature vectors in node features, and class labels are assigned based on product category. The ogbn-arxiv dataset (Hu et al., 2020) is a directed graph that represents the citation network of all computer-related papers on the arxiv website. Each node is an arxiv paper, and each edge represents an article citing another article.

In the inductive task, several graphs are used as training data, while other completely unseen graphs are used as validation and test data. For the inductive setting, we use the PPI (Hamilton et al., 2017) dataset as benchmark dataset. The task of PPI dataset is to classify the different protein functions. It consists of 24 graphs, with each representing a different human tissue. Each node has properties such as positional gene sets, motif gene sets, and immunological signatures, and gene ontology sets as labels. Twenty graphs are chosen for training, two graphs for validation, and the remaining for testing. For the inductive task, we use Micro-F1 as the evaluation metric.

The used datasets are concluded in Table 5.

Table 5: The Statistics of Datasets

	Transductive					Inductive
	CiteSeer	Cora	PubMed	Computers	ogbn-arxiv	PPI
#nodes	2708	3327	19717	13752	169343	56944
#edges	5278	4552	44324	245861	1166243	818716
#features	1433	3703	500	767	128	121
#classes	7	6	3	10	40	50

### A.2. Proof of Analysis for No-update Scheme in GNNs

**Theorem A.1.** Assume  $\mathbf{W}_l(0)$  is randomly initialized for all  $l \in [1, L]$ , if  $\prod_{l=1}^L \mathbf{W}_l(0)$  is full rank, there must exist a weight matrix for the output layer, i.e.,  $\tilde{\mathbf{W}}_o$ , that makes the final output the same as the one from a well-trained network:

$$A^L \mathbf{X} \prod_{l=1}^L \mathbf{W}_l^* \mathbf{W}_o^* = A^L \mathbf{X} \prod_{l=1}^L \mathbf{W}_l(0) \tilde{\mathbf{W}}_o. \quad (9)$$

*Proof.* Because  $\prod_{l=1}^L \mathbf{W}_l(0)$  is a full-rank matrix,  $\prod_{l=1}^L \mathbf{W}_l(0)$  is invertible. We can define  $\tilde{\mathbf{W}}_o$  by

$$\tilde{\mathbf{W}}_o \stackrel{\text{def}}{=} \left( \prod_{l=1}^L \mathbf{W}_l(0) \right)^{-1} \prod_{l=1}^L \mathbf{W}_l^* \mathbf{W}_o^*, \quad (10)$$

Hence, by multiplying  $\prod_{l=1}^L \mathbf{W}_l(0)$  and  $A^L \mathbf{X}$  on both sides in Equation (10), we can attain Equation (9).  $\square$

**Theorem A.2.** Assume a GNN model has either the exponential loss or the log-loss, the desired weight  $\tilde{\mathbf{W}}_o$  is secured when updating with gradient descent and initializing  $\prod_{l=1}^L \mathbf{W}_l(0)$  as an orthogonal matrix. Mathematically, we have  $\hat{\mathbf{W}}_o(+\infty) = \hat{\mathbf{W}}_o^* = \tilde{\mathbf{W}}_o$ .

*Proof.* We first frame the simplified GNNs as  $A^L \mathbf{X} \prod_{l=1}^L \mathbf{W}_l = (A^L \mathbf{X}) (\prod_{l=1}^L \mathbf{W}_l)$ . According to Theorem 3.2, we define the corresponding  $\beta$  as  $\beta = \prod_{l=1}^L \mathbf{W}_l$ .

When using the gradient descent to optimize  $\mathbf{W}_l$  and  $\mathbf{W}_o$ , we have the optimized  $\beta$  as  $\beta^* = \prod_{l=1}^L \mathbf{W}_l^* \cdot \mathbf{W}_o^*$ . The obtained  $\beta^*$  can be viewed as the max-margin hyperplane to separate data  $\mathbf{A}^L \mathbf{X}$  according to Theorem 3.2. Omitting the output layer, we let the untrained GNN model as  $\mathbf{A}^L \mathbf{X} \prod_{l=1}^L \mathbf{W}_l(0) = \mathbf{A}^L \mathbf{X} \mathbf{O}$  and assume  $\mathbf{O} = \prod_{l=1}^L \mathbf{W}_l(0)$  is an orthogonal square matrix.

Denoting a complete GNN model as  $\mathbf{A}^L \mathbf{X} \mathbf{O} \mathbf{W}_o$ , the optimized output weight, i.e.  $\hat{\mathbf{W}}_o^*$ , becomes the max-margin separating hyperplane of  $\mathbf{A}^L \mathbf{X} \mathbf{O}$  because it is also trained by gradient descent. Here, we fix the  $\mathbf{W}_l(0)$  during training to maintain its orthogonality. We know that the max-margin hyperplane of any data remains the same if and only if it takes orthogonal transformations, resulting in the following equivalence:  $\mathbf{O} \hat{\mathbf{W}}_o^* = \beta^* = \prod_{l=1}^L \mathbf{W}_l^* \cdot \mathbf{W}_o^*$ . Finally, we get  $\hat{\mathbf{W}}_o^* = \mathbf{O}^{-1} \prod_{l=1}^L \mathbf{W}_l^* \mathbf{W}_o^* = (\prod_{l=1}^L \mathbf{W}_l(0))^{-1} \prod_{l=1}^L \mathbf{W}_l^* \mathbf{W}_o^* = \tilde{\mathbf{W}}_o$ .  $\square$

### A.3. Proof for the Unified Format of Search Space in NAC

**Theorem A.3.** *Simplifying the complex non-linear functions, each GNN layer in the search space can be unified by  $\sum_k \mathbf{P}_l^k(L) \mathbf{H}_l \mathbf{W}_l^k$ , where  $\mathbf{P}(\cdot)$  denotes the fixed polynomial function with subscripts indicating individual terms, and superscript  $(l)$  marks the current layer.*

*Proof.* Without any loss of generality, we omit the lowerscript  $(l)$  of  $\sum_s C^{(s)} \mathbf{H}_l \mathbf{W}_l^{(s)}$  to prove the equivalence between  $\sum_s C^{(s)} \mathbf{H} \mathbf{W}^{(s)}$  and  $\sum_k \mathbf{P}^k(L) \mathbf{H} \mathbf{W}^k$  for arbitrary layer.

According to Table 1 in (Balcilar et al., 2021), we have the frequency support for most GNNs and except for GAT as

$$C^{(s)} = \mathbf{P}^{(s)}(\hat{\mathbf{L}}), \quad (11)$$

where  $\mathbf{P}^{(s)}(\cdot)$  denotes certain polynomial functions and  $\hat{\mathbf{L}}$  is a re-scaling and shifting of the original/normalized/re-normalized Laplacian matrix, which is universally denoted as  $L$  in this proof and its Corollary 3.5.<sup>2</sup> Equivalently, we have:

$$\hat{\mathbf{L}} = \mathbf{P}'_1(L) \quad (12)$$

Putting Equation (12) into Equation (11), it produces:

$$C^{(s)} = \mathbf{P}^{(s)}(\mathbf{P}'_1(L)) \quad (13)$$

$$= \mathbf{P}'^{(s)}(L). \quad (14)$$

And we have:

$$\sum_s C^{(s)} \mathbf{H} \mathbf{W}^{(s)} = \sum_s \mathbf{P}'^{(s)}(L) \mathbf{H} \mathbf{W}^{(s)}. \quad (15)$$

Without loss of generality, suppose  $s = 1$ ,  $C^{(0)} = \mathbf{P}'^{(0)}(L) = \alpha L^2 + \beta L$  and  $C^{(1)} = \mathbf{P}'^{(1)}(L) = \gamma L$ , and it gives:

$$C^{(0)} \mathbf{H} \mathbf{W}^{(0)} + C^{(1)} \mathbf{H} \mathbf{W}^{(1)} = (\alpha L^2 + \beta L) \mathbf{H} \mathbf{W}^{(0)} + \gamma L \mathbf{H} \mathbf{W}^{(1)} \quad (16)$$

$$= L^2 \mathbf{H}(\alpha \mathbf{W}^{(0)}) + L \mathbf{H}(\beta \mathbf{W}^{(0)} + \gamma \mathbf{W}^{(1)}) \quad (17)$$

Replacing  $\alpha \mathbf{W}^{(0)}$  with  $\mathbf{W}^0$  and  $\beta \mathbf{W}^{(0)} + \gamma \mathbf{W}^{(1)}$  with  $\mathbf{W}^1$ , Equation (16) is transformed to  $L^2 \mathbf{H} \mathbf{W}^0 + L \mathbf{H} \mathbf{W}^1$ . Then, assuming  $\mathbf{P}^1 = L^2$  and  $\mathbf{P}^0 = L$ , it produces  $\mathbf{P}^1 \mathbf{H} \mathbf{W}^0 + \mathbf{P}^0 \mathbf{H} \mathbf{W}^1 = \sum_{k=0}^1 \mathbf{P}^k(L) \mathbf{H} \mathbf{W}^k$ .

To this end, we prove the equivalence between  $\sum_s C^{(s)} \mathbf{H} \mathbf{W}^{(s)}$  and  $\sum_k \mathbf{P}^k(L) \mathbf{H} \mathbf{W}^k$ .

Specially, GAT contains complex non-linear functions, and we reduce its expression to  $\mathbf{H}_{(l+1)} = \sigma((\mathbf{H}_l \mathbf{W}_{l+1}^{att} \mathbf{H}_l^T \odot \mathbf{A}) \mathbf{H}_l \mathbf{W}_{l+1})$ . In this expression,  $\mathbf{W}_{att}$  is a attention function, and  $\odot$  denotes element-wise multiplication. Omitting  $\odot$  and  $\sigma$  to obtain the format  $\mathbf{H}_{l+1} = \mathbf{H}_l \mathbf{W}_{l+1}'$ , which is identical to the case when  $\mathbf{P}^{(s)} = \mathbf{I}$  and therefore inclusive in the discussion above.  $\square$

<sup>2</sup>This symbolic simplification does not affect the conclusion of Theorem 3.4 and Corollary 3.5. Because both claims aim to achieve a fixed dictionary format of GNNs, the different expressions of  $L$  do not change the format's nature of being fixed.

**Corollary A.4.** *The Search space of GNNs can be unified as  $\sum_{k=0}^K P^k(\mathbf{L})\mathbf{X}\mathbf{W}^k = \mathbf{D}\mathbf{W}$  when leaving the activation functions, where  $\parallel$  stands for concatenating a set of matrices horizontally,  $\mathbf{D} = \parallel\{P^k(\mathbf{L})\mathbf{X}\}$  is the fixed base, and  $\mathbf{W} = \parallel\{\mathbf{W}^k\}^T$  is the trainable parameters.*

*Proof.* Suppose the node features as  $\mathbf{H}_0 = \mathbf{X}$ , and a two-layer GNN consists of  $P_0^{(S)}, P_1^{(L)}$  and  $\mathbf{W}_0^{(S)}, \mathbf{W}_1^{(L)}$ . Here, we assume  $S, L = 1$  for simplicity.

Based on Theorem 3.4, the first layer outputs  $\mathbf{H}_1 = \sum_{s=0}^2 P_0^{(s)}(\mathbf{L})\mathbf{X}\mathbf{W}_0^{(s)} = P_0^{(0)}(\mathbf{L})\mathbf{X}\mathbf{W}_0^{(0)} + P_0^{(1)}(\mathbf{L})\mathbf{X}\mathbf{W}_0^{(1)}$ . Then, we pass  $\mathbf{H}_1$  to the second layer the below:

$$\begin{aligned}
 H_2 &= \sum_{l=0}^2 P_1^{(l)}(\mathbf{L})\mathbf{H}_1\mathbf{W}_1^{(l)} \\
 &= P_1^{(0)}(\mathbf{L})\mathbf{H}_1\mathbf{W}_1^{(0)} + P_1^{(1)}(\mathbf{L})\mathbf{H}_1\mathbf{W}_1^{(1)} \\
 &= P_1^{(0)}(\mathbf{L})(P_0^{(0)}(\mathbf{L})\mathbf{X}\mathbf{W}_0^{(0)} + P_0^{(1)}(\mathbf{L})\mathbf{X}\mathbf{W}_0^{(1)})\mathbf{W}_1^{(0)} \\
 &\quad + P_1^{(1)}(\mathbf{L})(P_0^{(0)}(\mathbf{L})\mathbf{X}\mathbf{W}_0^{(0)} + P_0^{(1)}(\mathbf{L})\mathbf{X}\mathbf{W}_0^{(1)})\mathbf{W}_1^{(1)} \\
 &= P_{1\times 0}^{(0\times 0)}(\mathbf{L})\mathbf{X}\mathbf{W}_{0\times 1}^{(0\times 0)} + P_{1\times 0}^{(0\times 1)}(\mathbf{L})\mathbf{X}\mathbf{W}_{0\times 1}^{(1\times 0)} \\
 &\quad + P_{1\times 0}^{(1\times 0)}(\mathbf{L})\mathbf{X}\mathbf{W}_{0\times 1}^{(0\times 1)} + P_{1\times 0}^{(1\times 1)}(\mathbf{L})\mathbf{X}\mathbf{W}_{0\times 1}^{(1\times 1)} \\
 &= \sum_k P^k(\mathbf{L})\mathbf{X}\mathbf{W}^k.
 \end{aligned} \tag{18}$$

In Equation (18), we merge the polynomial of a polynomial as  $P_1^{(0)}P_0^{(0)} = P_{1\times 0}^{(0\times 0)}$ . Similarly, we replace  $\mathbf{W}_0^{(0)}\mathbf{W}_1^{(0)}$  with  $\mathbf{W}_{0\times 1}^{(0\times 0)}$ .

Ultimately, we expand Equation (18) by  $\mathbf{D} = \parallel\{P_{1\times 0}^{(0\times 0)}(\mathbf{L})\mathbf{X}, P_{1\times 0}^{(0\times 1)}(\mathbf{L})\mathbf{X}, P_{1\times 0}^{(1\times 0)}(\mathbf{L})\mathbf{X}, P_{1\times 0}^{(1\times 1)}(\mathbf{L})\mathbf{X}\}$  and  $\mathbf{W} = \parallel\{\mathbf{W}_{0\times 1}^{(0\times 0)}, \mathbf{W}_{0\times 1}^{(1\times 0)}, \mathbf{W}_{0\times 1}^{(0\times 1)}, \mathbf{W}_{0\times 1}^{(1\times 1)}\}$ , so that  $\sum_k P_k(\mathbf{L})\mathbf{X}\mathbf{W}_k = \mathbf{D}\mathbf{W}$ .  $\square$

#### A.4. Proof for the Optimization of $l_2$ -norm Objective in NAC

**Theorem A.5.** *The gradient from the Cross-Entropy loss function is feasible to optimize an  $l_2$ -norm objective in a gradient descent manner, such that the initial  $z^0$  converges geometrically to the optimal  $z^*$  with a tiny error, e.g.,*

$$\|z^{s+1} - z^*\| \leq (1 - 2\alpha\eta)^s \|z^0 - z^*\| + 2\gamma/\alpha, \tag{19}$$

where  $s$  signifies the current iteration, and  $\gamma, \alpha, \eta$  define the numerical bound for the error.

*Proof.* Given an iterative algorithm  $\mathcal{A}$  that optimizes for a solution  $z^* \in \mathbb{R}^n$  for a function  $f(z)$ , it yields  $z^{s+1} = z^s - \eta g^s$  in each step  $s$ . Here,  $\eta$  is the step size, and  $g^s$  is the gradient vector. In the following, we give the *sufficient conditions* for such  $g^s$ , quoted from (Nguyen et al., 2018) and (Arora et al., 2015).

**Definition A.6.** A vector  $g^s$  at the  $s$ -th iteration is  $(\alpha, \beta, \gamma_s)$ -correlated with a desired solution  $z^*$  if

$$\langle g^s, z^s - z^* \rangle \geq \alpha \|z^s - z^*\|^2 + \beta \|g^s\|^2 - \gamma_s. \tag{20}$$

Under the convex optimization, if  $f$  is  $2\alpha$ -strong convex and  $1/2\beta$ -smooth, and  $g^s$  is set as  $\nabla_z f(z)$ , then  $g^s$  is  $(\alpha, \beta, 0)$ -correlated with  $z^*$ .

**Theorem A.7.** (Convergence of approximate gradient descent). *Suppose that  $g^s$  satisfies the conditions described in Definition A.6 for  $s = 1, 2, \dots, T$ . Moreover,  $0 < \eta \leq 2\beta$  and  $\gamma = \max_{s=1}^T \gamma_s$ . Then, the following holds for all  $s$ :*

$$\|z^{s+1} - z^*\|^2 \leq (1 - 2\alpha\eta) \|z^s - z^*\|^2 + 2\eta\gamma_s. \tag{21}$$

In particular, the above update converges geometrically to  $z^*$  with an error:

$$\|z^{s+1} - z^*\|^2 \leq (1 - 2\alpha\eta)^s \|z^0 - z^*\|^2 + 2\gamma/\alpha. \quad (22)$$

So, if the cross-entropy loss function satisfies the conditions of convexity and smoothness in Definition A.6, it meets the *sufficient conditions* for convergence.

Recall that the cross-entropy loss is defined as:

$$\begin{aligned} g(W) &= -\frac{1}{m} \sum_{i=1}^m \sum_{j=1}^c Y_{ij} \log\left(\frac{\exp F_i W_j}{\sum_{k=1}^c \exp F_i W_k}\right) \\ &= -\frac{1}{m} \sum_{i=1}^m \sum_{j=1}^c Y_{ij} F_i W_j + \frac{1}{m} \sum_{i=1}^m \log \sum_{k=1}^c \exp F_i W_k \end{aligned} \quad (23)$$

where  $Y$  is one-hot matrix with  $c$  classes. Then, we give the gradient w.r.t  $W_j$ :

$$\nabla_{W_j} g = -\frac{1}{m} \sum_{i=1}^m [Y_{ij} = 1] F_i + \frac{1}{m} \sum_{i=1}^m \frac{\exp F_i W_j}{\sum_{k=1}^c \exp F_i W_k} F_i, \quad (24)$$

where we symbolize the softmax function  $s_i = \frac{\exp F_i W_j}{\sum_{k=1}^c \exp F_i W_k}$ , and thus:

$$\nabla_{W_j} g = \frac{1}{m} \sum_{i=1}^m (s_i - [Y_{ij} = 1]) F_i. \quad (25)$$

**Smoothness:** In the extreme case where all the softmax values are equal, i.e.,  $s_i = \frac{1}{c}$  and  $Y_{ij} = 1$ , we can develop an upper bound as follows,

$$|s_i - [Y_{ij} = 1]| \leq \frac{c-1}{c}. \quad (26)$$

Based on this, we get:

$$|\nabla_{W_j} g| \leq \frac{c-1}{mc} \left| \sum_{i=1}^m F_i \right| \quad (27)$$

This gives us the Lipschitz constant that  $L = \frac{c-1}{mc} \|F\|$ . Sufficiently, there exists  $\beta^*$  that the cross-entropy loss function is  $\beta^*$ -smooth.

**Convexity:** Given  $W_k, W_j, F_i \in \mathbb{R}^d$ , we assume each element sampled from a certain normal distribution. Without loss of generality, we give  $W_k \sim \mathcal{N}(\mu_1, \sigma_1^2)$ ,  $W_k \sim \mathcal{N}(\mu_2, \sigma_2^2)$ , and  $F_i \sim \mathcal{N}(\mu, \sigma^2)$ .

In particular, we take  $F_i$  as an example and the other two follow the same rule:

$$\begin{aligned} \mathbb{E}[F_i^2] &= \mathbb{E}[(F_i - \mu)^2 + 2F_i\mu - \mu^2] \\ &= \sigma^2 + \mu^2, \end{aligned} \quad (28)$$

which equals to  $\frac{1}{d}$ , and thus,  $\mu^2 = \frac{1}{d} - \sigma^2$ . For  $W_k, W_j$ , we have  $\mu_1^2 = \frac{1}{d} - \sigma_1^2$ , and  $\mu_2^2 = \frac{1}{d} - \sigma_2^2$ . For strong convexity, we need to demonstrate there exists  $\alpha \in \mathbb{R}$  supporting  $\langle \nabla_{W_k} g - \nabla_{W_j} g, W_k - W_j \rangle \geq \alpha \langle W_k - W_j, W_k - W_j \rangle$ . By taking the Equation (27), we can analyze the left-hand side as:

$$\begin{aligned} \langle \nabla_{W_k} g - \nabla_{W_j} g, W_k - W_j \rangle &= \frac{1}{m} \sum_{i=1}^m [Y_{ij} \neq Y_{ik}] \langle F_i, W_k - W_j \rangle \\ &\geq \left\langle \frac{1}{m} \sum_{i=1}^m [Y_{ij} \neq Y_{ik}] F_i, W_k - W_j \right\rangle \\ &\geq \frac{2(c-1)}{mc^2} \sum_{i=1}^m \langle F_i, W_k - W_j \rangle. \end{aligned} \quad (29)$$

In Equation (29), we assume  $Y_i$  obeys uniform distribution that  $P(Y_{ik} = 1) = \frac{1}{c}$ , then we replace  $[Y_{ij} \neq Y_{ik}]$  by its average  $\frac{2(c-1)}{c^2}$ , identical to the probability,  $P([Y_{ij} \neq Y_{ik}])$ .

Next, we suppose  $\frac{1}{d} \sum_{h=1}^d dF_i$  is an unbiased estimation of its expectation, and same for  $W_j, W_k$ . Here, we replace  $W_j - W_k$  with a new variable as  $W' \sim \mathcal{N}(\mu_1 - \mu_2, \sigma_1^2 + \sigma_2^2)$ . With the help of Equation (29), we get,

$$\langle \nabla_{W_k} g - \nabla_{W_j} g, W_k - W_j \rangle \geq \frac{2(c-1)}{c^2} d\mu(\mu_1 - \mu_2). \quad (30)$$

By taking these into the right-hand side  $\langle W_k - W_j, W_k - W_j \rangle$ , we have:

$$\begin{aligned} dE[W'^2] &= dE[(W' - (\mu_1 - \mu_2))^2 - (\mu_1 - \mu_2)^2 + 2W'(\mu_1 - \mu_2)] \\ &= d(E[(W' - (\mu_1 - \mu_2))^2] - (\mu_1 - \mu_2)^2 + 2(\mu_1 - \mu_2)E[W']) \\ &= d(\sigma_1^2 + \sigma_2^2 + (\mu_1 - \mu_2)^2) \end{aligned} \quad (31)$$

The solution above is always positive. Finally, we put and in the required inequality, we reach the following inequality :  $\langle \nabla_{W_k} g - \nabla_{W_j} g, W_k - W_j \rangle \geq \alpha \langle W_k - W_j, W_k - W_j \rangle$  only if

$$\alpha \leq \frac{2(c-1)\mu(\mu_1 - \mu_2)}{c^2(\sigma_1^2 + \sigma_2^2 + (\mu_1 - \mu_2)^2)}. \quad (32)$$

Under a mild assumption that  $\alpha^* \in \mathbb{R}$  satisfies Equation (32), we conclude that the cross-entropy loss function is  $\alpha^*$ -strongly convex and  $\beta^*$ -smooth, such that it is  $(\frac{\alpha^*}{2}, \frac{1}{2\beta^*}, 0)$ -correlated with  $z^*$ . Thus, it meets the *sufficient condition* of convergence.  $\square$

### A.5. Proof of Theorem 3.6 for the Dictionary Orthogonality in NAC

**Theorem A.8.** *Let the neural weights of each operator in a deep neural network be an atom and stack them column-wisely, we can guarantee the existence of an orthogonal dictionary.*

*Proof.* Given a dictionary  $\mathbf{H} \in R^{n \times K}$ , where  $n$  is the number of nodes and  $K$  is the number of operators in each layer, we have its mutual coherence computed as follows,

$$\varphi = \max_{\mathbf{h}_i, \mathbf{h}_j \in \mathbf{H}, i \neq j} \left| \left\langle \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2}, \frac{\mathbf{h}_j}{\|\mathbf{h}_j\|_2} \right\rangle \right|, \quad (33)$$

where  $\varphi \in [0, 1]$ , and  $\langle \cdot \rangle$  denotes inner product. Here, each atom,  $\mathbf{h}_i$ , is the weights from an operator. The minimum of  $\varphi$  is 0 and is attained when there is an orthogonal dictionary, while the maximum is 1 and it attained when there are at least two collinear atoms (columns) in a dictionary.

Let  $E_i = \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2}$  and  $E_j = \frac{\mathbf{h}_j}{\|\mathbf{h}_j\|_2}$ , by Central Limit Theorem (Fischer, 2011), we know that  $\langle E_i, E_j \rangle / \sqrt{n}$  converges to a normal distribution, i.e.,

$$\langle E_i, E_j \rangle = \lim_{n \rightarrow \infty} \sqrt{n}Z, \quad (34)$$

where  $Z$  is a standard normal distribution. Consider  $\bar{E}$  as the mean value of all  $\langle E_i, E_j \rangle$ . With weak law of large numbers (a.k.a. Khinchin's law) (ter Haar, 1949), for any positive number  $\varepsilon$ , the probability that sample average  $\bar{E}$  greater than  $\varepsilon$  converges 0 is written as

$$\lim_{n \rightarrow \infty} \Pr(|\bar{E}| \geq \varepsilon) = 0 \quad (35)$$

This implies that the probability that the inner product of  $E_i$  and  $E_j$  is greater than  $\varepsilon$  close to zero when  $n \rightarrow \infty$ . In other words, the probability that  $E_i$  and  $E_j$  are nearly orthogonal goes to 1 when their dimensionality is high. Therefore, the coherence of this dictionary reaches the minimum at a high dimensionality that holds for deep neural networks naturally.  $\square$



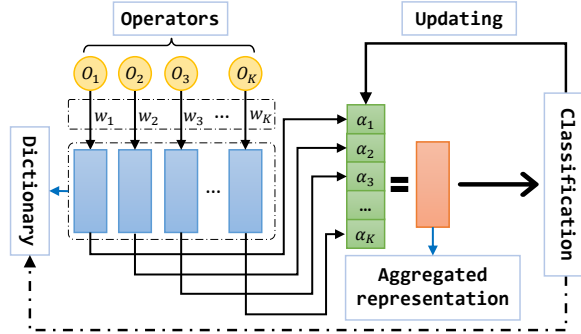


Figure 4: The framework of the proposed NAC in one layer. NAC directly learns the architecture  $\alpha$  with a fixed dictionary. NAC-updating updates the dictionary in training, showing the additional process with a dash line.

---

### Algorithm 2 The NAC algorithm

---

**Require:** The search space  $\mathcal{A}$ ;

**Ensure:** The architecture  $\alpha$

Randomly initializing  $\mathbf{W}^l$ , for  $l = 1, \dots, L$ ; set  $\alpha = \mathbf{1}$ ;

1: **while**  $t = 1, \dots, T$  **do**

2: Performing the feature aggregation at each layer as  $\bar{o}^l(x) = o^l \hat{\alpha}_l = \sum_{k=1}^K \frac{\alpha_{lk}}{\|\alpha_l\|_2} o^{lk}(x)$ ;

3: Computing  $\mathbf{h}_v^l = \phi[\mathbf{W}^l \cdot \bar{o}^l(\{\mathbf{h}_u^{l-1}, \forall u \in N(v)\})]$ ;

4: Optimizing  $\alpha$  based on the objective function in Equation (8) w.r.t.  $\alpha$  under the fixed dictionary  $o$ ;

5: Updating  $\mathbf{W}_o$  based on the objective function in Equation (8) w.r.t.  $\mathbf{W}_o$  under fixed  $\alpha$ ;

6: **end while**

7: Obtain the final architecture  $\{\alpha^*\}$  from the trained  $\alpha$  via an argmax operation at each layer;

---

### A.6. Realization and Details of NAC and NAC-updating

In this work, we propose two realizations of NAC, namely NAC and NAC-updating, where Figure 4 outlines the core idea of both methods in one layer. The objectives in Equation (8) are optimized by gradient descent, and we therefore omit the details but present the basic pipeline in Algorithm 1. The computation of NAC has two major parts: the forward pass and the backward pass. Given the search space, the computation of the forward is then fixed, regarding as a constant. Therefore, the computational complexity mainly focuses on the backward pass in the NAC algorithms.

**NAC.** The main version of our work has no need to update weights, but only to update architectural parameter  $\alpha$  during the training process. Therefore, the algorithmic complexity is as  $O(T * \|\alpha\|)$ , which is a *linear* function w.r.t.  $\alpha$ . The dimension of  $\alpha$  is often small, which makes the model easy to scale to large datasets and high search space. We call it NAC.

**NAC-updating.** When updating weights, similar to DARTS, the complexity is estimated as  $O(T * (\|\alpha\| + \|w\|))$ . The dimension of  $w$  is often much larger than  $\alpha$ , therefore, the complexity is dominated by updating  $w$ , where the complexity is  $O(T * \|w\|)$ . Since the dimension of  $\alpha$  is much smaller than  $w$ , the complexity of NAC is much less than this version.

### A.7. Experimental Implementation Details

**Searching configuration.** In our experiments we adopt 3-layer GNN as the backbone. Unless specified, our experiments follow the same settings for searching architectures as SANE (Zhao et al., 2021b) :

- *Architecture optimizer.* We use Adam for training the architecture parameters  $\alpha$ . We set the learning rate as 0.0003 and the weight decay as 0.001. Also, the  $\beta_1$  and  $\beta_2$  are fixed as 0.5 and 0.999, respectively. All runs a constant schedule for training, such as 100 epochs.
- *Weight optimizer.* We use SGD to update models' parameters, i.e.,  $w$ . The learning rate and SGD momentum are given as 0.025 and 0, respectively, where the learning rate has a cosine decay schedule for 100 epochs. We fix the weight decay value, i.e. set  $\rho_1 = 0.0005$ .
- *Batch size.* For transductive tasks, we adopt in-memory datasets, and the *batch\_size* is fixed as the size of the dataset themselves.

**The configuration for retraining phase.** At the retraining state, we adopt Adam as the optimizer and set the scheduler with cosine decay to adjust the learning rate. The total number of epochs is fixed 400 for all methods for fairness. Please

refer to the setting of SANE (Zhao et al., 2021b) and EGAN (Zhao et al., 2021a) for more details as we follow this in our experiment.

For CiteSeer dataset, we set the initial learning rate as 0.005937 and weight decay as 0.00002007. The configuration for models is as follows:  $hidden\_size = 512$ ,  $dropout = 0.5$ , and using  $ReLU$  as the activation function.

For Cora dataset, we set the initial learning rate as 0.0004150, and weight decay as 0.0001125. In model, we set  $hidden\_size = 256$ ,  $dropout = 0.6$ , and use  $ReLU$  as the activation function.

For PubMed dataset, we set the initial learning rate as 0.002408 and weight decay as 0.00008850. As for the model, we have  $hidden\_size = 64$  and  $dropout = 0.5$ , and use  $ReLU$  as the activation function.

For Amazon dataset, we set the initial learning rate as 0.002111 and weight decay as 0.000331. As for the model, we have  $hidden\_size = 64$  and  $dropout = 0.5$ , and use  $elu$  as the activation function.

For PPI dataset, we set the initial learning rate as 0.00102 and weight decay as 0. As for the model, we have  $hidden\_size = 512$  and  $dropout = 0.5$ , and use  $Relu$  as the activation function.

For ogbn-arXiv data, due to the limitation of computational resources, we use the NAS-Bench-Graph (Qin et al., 2022) to get the performance of searched architecture and make some adjustments to align with the supporting search space in NAS-Bench-Graph.

**Solving  $L^1$  regularization.** The  $L^1$  regularization, also known as **Lasso Regression** (Least Absolute Shrinkage and Selection Operator), adds an absolute value of the magnitude of the coefficient as a penalty term to the loss function (Ranstam & Cook, 2018). Using the  $L^1$  regularization, the coefficient of the less important feature is usually decreased to zero, sparsifying the parameters. It should be noted that since  $\|\alpha\|_1$  is not differentiable at  $\alpha = \mathbf{0}$ , the standard gradient descent approach cannot be used.

Despite the fact that the loss function of the Lasso Regression cannot be differentiated, many approaches to problems of this kind, such as (Schmidt et al., 2009), have been proposed in the literature. These methods can be broadly divided into three groups: constrained optimization methods, unconstrained approximations, and sub-gradient methods.

Since subgradient methods are a natural generalization of gradient descent, this type of methods can be easily implemented in Pytorch’s framework. Lasso Regression can be solved using a variety of subgradient techniques; details on their implementation can be found in (Fu, 1998) and (Shevade & Keerthi, 2003).

**Computational Complexity Estimation of NAC.** The computation of NAC has two major parts: the forward pass and the backward pass. Given the search space, the computation of the forward is then fixed and regarded as a constant. Therefore, the computational complexity mainly focuses on the backward pass in the NAC algorithms.

The main version of our work does not need to update weights, but only to update architectural parameter  $\alpha$  during the training process. Therefore, the algorithmic complexity is as  $O(T * \|\alpha\|)$ , which is a *linear* function w.r.t  $\alpha$ . The dimension of  $\alpha$  is often small, which makes the model easy to scale to large datasets and high search space. When updating weights of the linear layer, the complexity is estimated as  $O(T * (\|\alpha\| + \|\mathbf{W}_o\|))$ . The dimension of  $\mathbf{W}_o$  is a constant number, that equals the number of classes. Therefore, the complexity is almost the same as the main version of NAC, where the complexity is  $O(T * \|\alpha\| + \|\mathbf{W}_o\|)$ .

When updating weights, similar to DARTS, the complexity is estimated as  $O(T * (\|\alpha\| + \|\mathbf{w}\|))$ . The dimension of  $\mathbf{w}$  is often much larger than  $\alpha$ , therefore, the complexity is dominated by updating  $\mathbf{w}$ , where the complexity is  $O(T * \|\mathbf{w}\|)$ . Since the dimension of  $\alpha$  is much smaller than  $\mathbf{w}$ , the complexity of NAC is much less than this type of methods.

**Approximate Architecture Gradient.** Our proposed theorems imply an optimization problem with  $\alpha$  as the upper-level variable and  $\mathbf{W}_o$  as the lower-level variable:

$$\begin{cases} \alpha^* = \underset{\alpha}{\operatorname{argmax}} \mathcal{M}(\mathbf{W}_o^*(\alpha), \alpha) \\ \mathbf{W}_o^*(\alpha) = \underset{\mathbf{W}_o}{\operatorname{argmin}} \mathcal{L}(\alpha, \mathbf{W}_o), \end{cases} \quad (36)$$

Following (Liu et al., 2019a), we can adopt a First-order Approximation to avoid the the expensive inner optimization, which allows us to give the implementation in the Algorithm 1.

A.8. Ablation Studies

A.8.1. THE EFFECT OF SPARSITY

Our model uses the hyperparameter  $\rho$  to control the sparsity of the architecture parameter  $\alpha$ , where a large sparsity is to enforce more elements to be zero. We investigate the effect of this hyperparameter by varying its value in a considerably large range, such as  $[0.001, 10]$ . We first present the accuracy of different sparse setting in Figure 5. We find that the results vary little in a considerably wide range, this indicates our models are insensitive to sparsity hyperparameter in general.

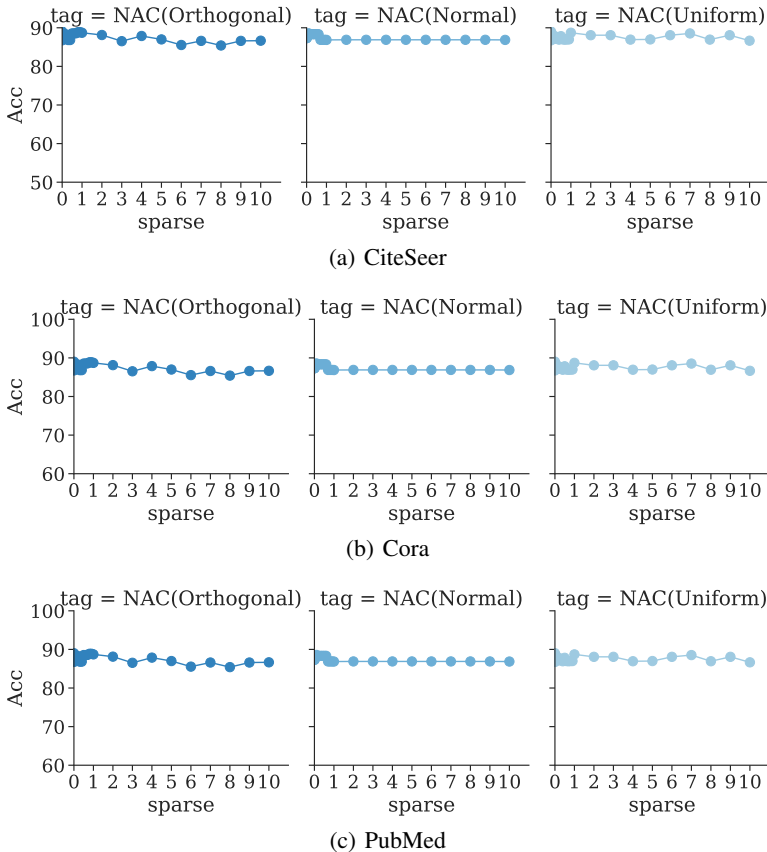


Figure 5: Sensitivity study of the sparsity. Results are with varying sparsity (x-axis) and different initialization NAC methods (i.e., normal, uniform and orthogonal). The variation is small, showing the robustness of our model w.r.t the sparsity.

A.8.2. THE EFFECT OF RANDOM SEEDS

Random seeds often play an importance role in traditional NAS methods as it affects the initialization significantly. People often report the average or the best results under different random seeds, this may lead to poor reproducibility. To the best our knowledge, this is for the first we explicitly demonstrate the effect of random seeds in this subject. We run experiments on several random seeds and report the results of NAC on Pubmed dataset, as shown in Figure 6. In particular, we implement multiple combinations of random seeds and sparsity to observe the variation on performance. Note that we round the values to integer to fit the table. In all these combinations, we have the average and variance as 87.32% and 0.9%, respectively. The average performance is comparable to the best results from all competitive results, which indicates the stability of NAC.

A.8.3. THE EFFECT OF TRAINING ON THE FINAL LINEAR LAYER

Our proposed theorems prove that a GNN with randomly initialized weights can make the final output as good as a well-trained network when initializing networks with orthogonal weights and updating the total network using gradient descent. In practice, we find it difficult to determine at what training epoch the optimal weight parameters can be obtained through training linear layer. We noticed that most of the time, the untrained weights in the initial state can often already exceed the accuracy that can be obtained from the weights after multiple epochs of training the final linear layer, as shown

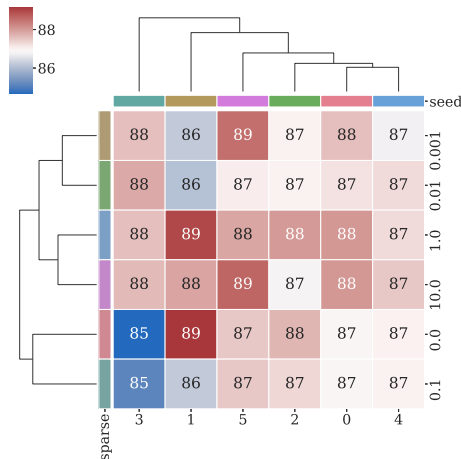


Figure 6: The effects of random seeds of NAC on the accuracy, where x-axis denotes the random seeds and y-axis denotes the sparsity. NAC performs stably with random seeds.

in Figure 7. Therefore, we further omit the training of the final linear layer. It is important to note that this approximation is based on our proposed theorems in which most of the intermediate layers do not require training.

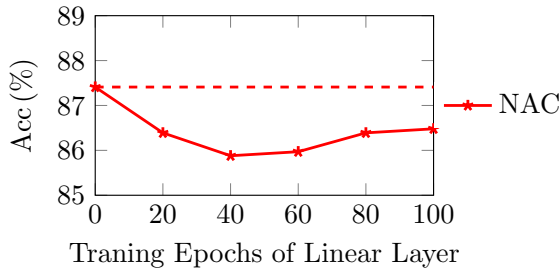


Figure 7: The effects of training final linear layer of NAC on the accuracy, where x-axis denotes the training epochs of the final linear layer and y-axis denotes the averaged accuracy of acquired architecture  $\alpha$  using the corresponding weights.

### A.9. Runtime of Each Method on a Single GPU Server

Apart from the running time on the CPU, we also measure the running time for all methods on a GPU platform, where we use PyTorch (version 1.10.2+gpu) on a GPU server with four NVIDIA 3090 GPUs (24G), as shown in Table Table 6. These results are consistent with the ones in Table 1, demonstrating our advantage in speed.

Table 6: Timing results of the compared methods and our NAC method in one the same GPU server. NAC attains superior performance in efficiency (in seconds).

	CiteSeer	Cora	PubMed	Computers
RS	196.00	328.00	461.00	900.00
BO	225.00	355.00	462.00	878.00
GraphNAS	6193.00	6207.00	6553.00	8969.00
GraphNAS-WS	947.00	1741.00	2325.00	4343.00
SANE	35.00	41.00	43.00	43.00
NAC	<b>14.00</b>	<b>14.00</b>	<b>15.00</b>	<b>14.00</b>
NAC-updating	42.00	31.00	36.00	42.00

### A.10. Performance on Inductive Tasks

To validate the effectiveness of our approach to the Inductive task, we performed a set of experiments on the PPI dataset. The experimental results are concluded in Table 7. The experimental results show that our proposed method effectively outperforms the best baseline method by about 4% in terms of Micro-F1 score. Besides, our method achieves an  $8\times$  speedup in terms of running time than SANE. Also, the non-updating scheme of the NAC approach exceeds the NAC-updating method effectively.

Experiments on the PPI dataset further validate the effectiveness and superiority of our proposed method.

Table 7: Experimental results on the compared methods: our NAC attains superior performance on PPI dataset in both Micro-F1 score (%) and efficiency (in hours).

	PPI(Micro-F1(%))	PPI(Time(h))
SANE	91.01±6.83	2.50
NAC	<b>95.16±0.03</b>	<b>0.31</b>
NAC-updating	94.47±7.09	4.68

### A.11. Performance on ogbn-arXiv

To validate the effectiveness of our approach to the Inductive task, we performed a set of experiments on the ogbn-arXiv dataset. The experimental results are concluded in Table 8. The experimental results show that our proposed method effectively outperforms the best baseline method by about 4% in terms of Micro-F1 score. Besides, our method achieves an 3× speedup in terms of running time than SANE. Also, the non-updating scheme of the NAC approach exceeds the NAC-updating method effectively.

Experiments on the ogbn-arXiv dataset further validate the effectiveness and superiority of our proposed method.

Table 8: Experimental results on the compared methods: our NAC attains superior performance on ogbn-arXiv dataset in Accuracy (%).

	ogbn-arXiv(Accuracy(%))	ogbn-arXiv(Time(min))
SANE	70.94±0.85	1.53
NAC	<b>71.13±0.59</b>	<b>0.50</b>
NAC-updating	71.10±0.47	2.20

### A.12. Searched Architectures for Each Dataset of Our Method

We visualize the searched architectures (top-1) by NAC on different datasets in Figure 8.

- For Citeseer dataset, the searched result is [GAT, GCN, Chebconv];
- For Cora dataset, the searched result is [GIN, GIN, GCN];
- For Pubmed dataset, the searched result is [GCN, GAT\_Linear, Geniepath];
- For Amazon Computers dataset, the searched result is [Geniepath, GCN, SAGE];
- For PPI dataset, the searched result is [Chebconv, GAT\_COS, Chebconv];

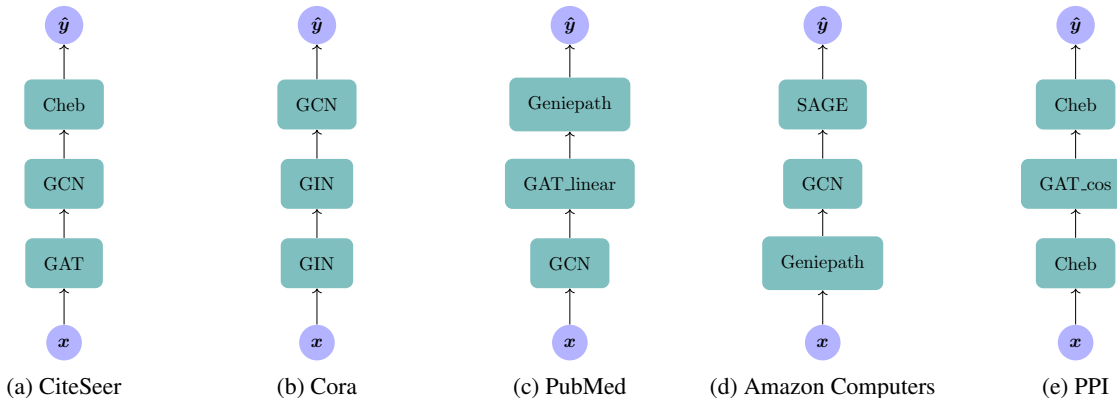


Figure 8: The searched architectures of NAC on benchmark datasets.

### A.13. Related Work

Another line of work (Chen et al., 2021) (Shu et al., 2021) uses neural tangent kernels (NTK) to search the network structure, called training-free NAS and focuses on CNN architectures. These hold strong assumptions when analyzing due to the need

for infinite width of networks, and thus far from reality. In contrast, we do not have such an assumption by taking advantage of the built-in linearity in GNNs to get untrained GNNs to work.