
PACE: A Parallelizable Computation Encoder for Directed Acyclic Graphs

Zehao Dong^{*1} Muhan Zhang^{*23} Fuhai Li⁴ Yixin Chen¹

Abstract

Optimization of directed acyclic graph (DAG) structures has many applications, such as neural architecture search (NAS) and probabilistic graphical model learning. Encoding DAGs into real vectors is a dominant component in most neural-network-based DAG optimization frameworks. Currently, most DAG encoders use an asynchronous message passing scheme which sequentially processes nodes according to the dependency between nodes in a DAG. That is, a node must not be processed until all its predecessors are processed. As a result, they are inherently not parallelizable. In this work, we propose a Parallelizable Attention-based Computation structure Encoder (PACE) that processes nodes simultaneously and encodes DAGs in parallel. We demonstrate the superiority of PACE through encoder-dependent optimization subroutines that search the optimal DAG structure based on the learned DAG embeddings. Experiments show that PACE not only improves the effectiveness over previous sequential DAG encoders with a significantly boosted training and inference speed, but also generates smooth latent (DAG encoding) spaces that are beneficial to downstream optimization subroutines. Our source code is available at <https://github.com/zehao-dong/PACE>.

1. Introduction

Directed acyclic graphs (DAGs) are ubiquitous in various real-world problems including neural architecture search (Elsken et al., 2019; Wen et al., 2020), source code

modeling (Allamanis et al., 2018), structure learning of Bayesian networks (Koller & Friedman, 2009; Zhang et al., 2019), etc. One key challenge in DAG optimization problems is the difficulty to use gradient strategies to quickly adjust the structure of a DAG towards the right direction due to the absence of gradient information. Some earlier works propose to directly optimize the discrete DAG structure through black-box optimization techniques such as reinforcement learning (Zoph & Le, 2016), evolutionary algorithms (Liu et al., 2017), and Bayesian optimization (Kandasamy et al., 2018), which are inherently less efficient. A more recent approach is to encode DAGs into some continuous space for searching, and various DAG encoders have been developed. In general, these DAG encoding schemes fall into two categories: structure-aware encoding scheme (Ying et al., 2019; Wen et al., 2020; Shi et al., 2020) and computation-aware (performance-aware) encoding scheme (Zhang et al., 2019; Thost & Chen, 2021).

Due to the superior graph representation learning ability, graph neural networks (GNNs) have broadly achieved state-of-art performance on various graph learning tasks, such as node classification (Velickovic et al., 2018; Hamilton et al., 2017), graph classification (Xu et al., 2019; Zhang et al., 2018b; Duvenaud et al., 2015), link prediction (Zhang & Chen, 2018), and hyperlink prediction (Zhang et al., 2018a). Basically, GNNs follow the message passing scheme (Gilmer et al., 2017) where each node aggregates node features from its one-hop neighborhood repeatedly to update its own feature, and this aggregation happens at all nodes simultaneously. However, Thost & Chen (2021) suggests that such a framework cannot exploit the inductive bias of the computation dependency defined by DAGs, thus failing to generate a smooth encoding space beneficial to downstream optimization and prediction routines.

Hence, in order to model the dependency between nodes in DAGs, various GNNs specifically designed for encoding DAGs, such as D-VAE (Zhang et al., 2019) and DAGNN (Thost & Chen, 2021), are developed to inject the computation dependency between nodes into the representation learning process. Instead of updating node features simultaneously, these DAG encoders are constructed upon a gated recurrent unit (GRU) and will not update the representation of a node until all of its predecessors are updated. Such an asynchronous message passing scheme actually

^{*}Equal contribution ¹Department of Computer Science & Engineering, Washington University in St. Louis, St. Louis, USA ²Institute for Artificial Intelligence, Peking University, Beijing, China ³Beijing Institute for General Artificial Intelligence, Beijing, China ⁴Institute for Informatics and Department of Pediatrics, Washington University in St. Louis, St. Louis, USA. Correspondence to: Yixin Chen <chen@cse.wustl.edu>.

simulates how a real computation is performed along the DAG—the message passing order respects the computation dependency defined by the DAG, thus better exploiting the inductive bias. One way to achieve this is to perform message passing sequentially following a topological ordering of the nodes. However, one key limitation of such DAG encoders is that the encoding process is inherently sequential, precluding processing all nodes in parallel. Although DAGNN proposes a topological batching trick to accelerate the training speed by partitioning nodes into disjoint batches where nodes within a batch can be processed in parallel, the time complexity is still lower-bounded by the longest path (diameter) of the DAG and the fundamental constraint of the sequential computation nature still remains.

Numerous efforts have been made to reduce the sequential computation cost in the sequence modeling literature (Cho et al., 2014; Wu et al., 2016). For instance, ConvS2S (Gehring et al., 2017) utilizes convolutional layers as building blocks to compute output representations at different positions, while Transformer (Vaswani et al., 2017) proposes to inject the position (order) information into the model through positional encoding, and then the dependency between positions can be captured through the attention mechanism (Bahdanau et al., 2015; Gehring et al., 2017) in parallel instead of resorting to recurrent neural networks. However, the success of these techniques relies on the inherent linear order of symbols in the input/output sequences that automatically characterizes the dependency between symbols. That is, the dependency between symbols is fully captured by their positions in the sequence. Such a condition is not satisfied by nodes in a DAG since each node can have **multiple parents** instead of only one like in plain sequences, and the dependency between nodes forms a (strong) partial order rather than a linear order. Thus, previous parallel sequence modeling methods would fail in the representation learning of DAGs.

In this paper, we propose a novel Parallelizable Computation Encoder, PACE, to improve the efficiency over existing GRU-based DAG encoders. In order to borrow the power of Transformer for sequence modeling to DAG modeling problems, we need to design a *positional encoding* scheme specifically for DAGs which can fully capture the dependency between nodes in a DAG before applying the pairwise self-attention mechanism. To achieve this, we propose a GNN-based dag2seq framework which is proved to injectively map DAGs to sequences of node embeddings. This means, we are able to fully recover the DAG structure from these produced node embeddings, the same as the positional encoding in the original Transformer. After that, a Transformer encoder (with mask operation) is applied to the node embedding sequence to simultaneously learn representations for all nodes in the DAG through the self-attention mechanism. This way, PACE incorporates the relational

inductive bias (Battaglia et al., 2018; Xu et al., 2020) carried by DAG structures into the encoding process, while eschewing the recurrence in previous works thus greatly improving the parallelization and encoding efficiency.

To demonstrate the superiority of the proposed PACE model, we evaluate PACE against current state-of-art DAG encoders, GNN-based graph encoders, and recent (undirected) graph Transformers. Massive experiments show that PACE not only outperforms competitive baselines but also significantly boosts the training and inference speed through parallelization.

2. Backgrounds

2.1. Parallelizable Sequence Models

Encoding the complexities and nuances of sequences plays a central role in various machine learning tasks, including sentiment classification (Medhat et al., 2014), speech recognition (Abdel-Hamid et al., 2014), and other natural language processing (NLP) tasks (Khan et al., 2016). For many years, sequential models, such as recurrent neural networks (RNNs) (Medsker & Jain, 2001), were the primal way to solve the sequence encoding problem. These models are computationally expansive due to the sequential encoding process. Hence, many parallelizable sequence models are proposed, including Transformer (Vaswani et al., 2017), BERT (Devlin et al., 2019), etc. Our proposed PACE model is built upon the Transformer (encoder) architecture.

Transformer (Vaswani et al., 2017) is arguably the earliest translation model that solves the sequence-to-sequence task without using sequence-aligned RNNs or convolutional architectures. Transformer relates information from different positions in the sequence through the (masked) self-attention mechanism to encode/decode the representation of items in the input/output sequence, and incorporates the order of the items into the encoding/decoding process through a positional encoding mechanism. Briefly, the positional encoding mechanism is an injective function $f_{pe} : \mathbb{N} \rightarrow \mathbb{R}^d$ that represents the positions (i.e. indices) of items in the sequence as d -dimensional vectors. Hence, it consistently outputs a unique encoding for each position in the sequence. The Transformer architecture is inherently parallelizable and can capture the long-term dependency with ease, thus is broadly applied to sequence modeling tasks in following works (Dai et al., 2019; Al-Rfou et al., 2019; Devlin et al., 2019; Lewis et al., 2020).

2.2. DAG Encoding Problem

A directed acyclic graph (DAG) G is represented as a pair (V, E) with $V = \{v_1, v_2, \dots, v_n\}$ denoting the set of nodes and $E \in V \times V$ denoting the set of directed edges. A DAG often carries a computation. We use \mathcal{O} to denote an (com-

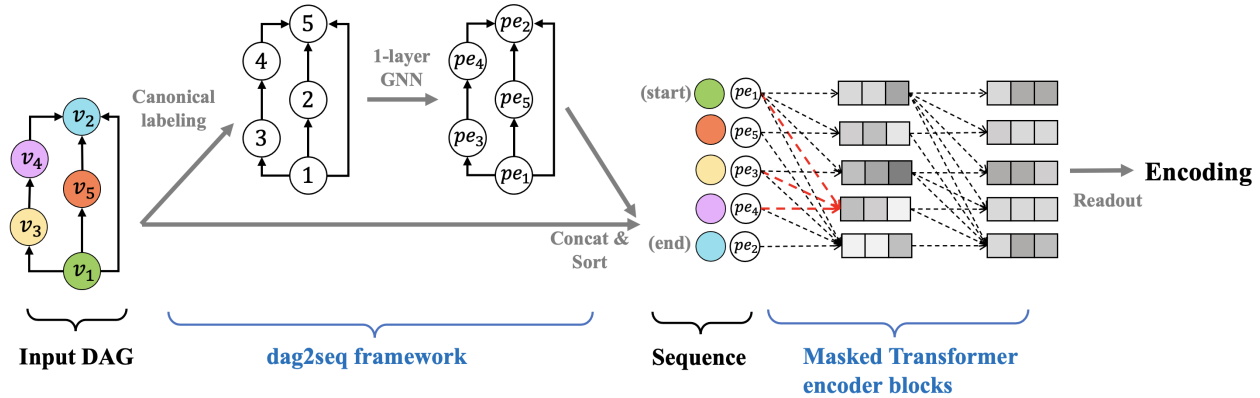


Figure 1. Illustration of PACE. The input DAG is first injectively represented as a sequence through the dag2seq framework, and then the sequence is fed into multiple stacked masked Transformer encoder blocks. The operations of nodes (i.e. node types) are visualized as colors, and nodes in the sequence is sorted according to the canonical label generated in the dag2seq framework.

putational) operation dictionary. For instance, the operation dictionary \mathcal{O} for NAS-Bench-101 dataset contains five operations: “Input”, “Output”, “ 3×3 convolution”, “ 1×1 convolution”, and “ 3×3 max-pool”. Let $G = (V, E)$ be a DAG whose nodes represent operations in \mathcal{O} . Then the DAG G represents a **computation structure** in which dependencies between operations are determined by directed edges in E . Hence, isomorphic DAGs define the same computation structure. Then the objective of the DAG encoding problem is to develop encoders that can generate embeddings to distinguish the computation structures defined by DAGs in the encoding space.

DAGs $G = (V, E)$ show a close relation with partial order. For any two nodes $v_i, v_j \in V$, let \preceq be a binary relation such that $v_i \preceq v_j$ if and only if there is a directed path from v_i to v_j , then the binary relation \preceq defines a partial order on the node

set V . Based on the partial order, sequential DAG encoders, such as D-VAE (Zhang et al., 2019) and DAGNN (Thost & Chen, 2021), use GRU (Cho et al., 2014) to recursively encode nodes in the input DAG, where a node is not encoded until all of its predecessors (those nodes with a \preceq relation with it) are encoded. Because of the possibly very long dependency chains, these sequential DAG encoders inherently share the same limitations as RNNs, such as the slow training and inference speed and the difficulty to capture long-term dependencies.

To address these limitations, it is intuitive to generalize




Figure 2. Illustration of the ambiguity when applying sequence models (such as Transformer) to the DAG encoding problem.

Transformer to the DAG encoding problem, as Transformer brings undoubtedly a huge improvement over the RNN-based sequence models. However, Figure 2 illustrates that it can be ambiguous how to represent DAGs as sequences due to the complex topological structure of DAGs. Let $o : V \rightarrow \mathcal{O}$ be the function that maps each node in G to an operation in \mathcal{O} , and f_{pe} be the original positional encoding function in Transformer. One intuitive way to linearize a DAG into a sequence is to sort its nodes with a topological order, which is also used in GRU-based DAG encoders. However, the topological orders of nodes in a DAG are often not unique. For instance, graph G_1 in Figure 2 has two different topological orders: v_1, v_2, v_3, v_4, v_5 and v_2, v_1, v_3, v_4, v_5 , hence resulting in two different node sequences Seq_1 and Seq_2 . Note that this does not hurt existing GRU-based DAG encoders as they are invariant to which topological order is used.

To avoid the ambiguity of topological order, another approach is to sort nodes according to a canonical order (i.e., node index in the canonical form of the DAG) as suggested by (Niepert et al., 2016). For example, let v_1, v_2, v_3, v_4, v_5 be the nodes sorted by a canonical order. Then, we can linearize a DAG into a sequence $(o(v_1), f_{pe}(v_1)), (o(v_2), f_{pe}(v_2)), \dots, (o(v_n), f_{pe}(v_n))$, similar to how Transformer represents a sentence. Although canonical order guarantees the generated sequence is unique for the same DAG, different DAGs might have the same sequence. For instance, DAGs G_1 and G_2 in Figure 2 are not isomorphic, yet they will be represented as the same sequence Seq_2 . This is because the positional encoding function f_{pe} only captures node positions in the sequence but fails to encode the node dependencies, causing a significant structural information loss.

In summary, it is not straightforward to generalize Transformer to DAGs. We need to design a smart linearization

method for DAGs which guarantees that the same DAG is transformed to the same sequence, while different DAGs are always transformed to different sequences, in order to achieve a lossless transformation.

3. The PACE Model

In this section, we describe the proposed **Parallelizable Computation Encoder (PACE)**. The key component of PACE is a **dag2seq** framework, which leverages graph canonization procedure and a graph neural network to transform DAGs into sequences while preserving the distinctiveness of non-isomorphic DAGs, so that the DAG encoding problem can be addressed efficiently by applying Transformer to the sequences.

3.1. The Dag2seq Framework

Here we describe the proposed dag2seq framework. Let $G = (V, E, o)$ be a labeled DAG (node labels are operations in \mathcal{O}), where $V = \{1, 2, \dots, n\}$ is the finite set of nodes, E is the set of directed edges, and $o : V \rightarrow \mathcal{O}$ is a function that associates to each node an operation in \mathcal{O} . We denote the *canonical form* of G as $C(G) = (V^C, E^C, o^C)$, which assigns to each labeled graph G an isomorphic labeled graph $C(G)$ that is a unique representation of its isomorphism class. That is, all labeled graphs isomorphic to G will have the same canonical form $C(G)$. Since G and $C(G)$ are isomorphic, there exists a bijection $\pi : V \rightarrow V^C$ (note that $V^C = \{1, 2, \dots, n\}$) between the node sets such that $o^C(\pi(i)) = o(i)$ for all $i \in V$ and $(\pi(i), \pi(j)) \in E^C$ if and only if $(i, j) \in E$. The graph canonization process assigns a new index $\pi(i)$ to each node i . Based on the new indices $\pi(i)$, the proposed dag2seq computes the *positional encoding* of node $j \in V$ (denoted as p_j) as follows

$$a_j = \text{Agg}(\{\pi(i), (i, j) \in E\}) \quad (1)$$

$$p_j = \text{Combine}(\pi(j), a_j) \quad (2)$$

where functions *Agg* and *Combine* follow the same definition as Graph Isomorphism Network (GIN) (Xu et al., 2019). Note that $\pi(i), i \in V$ are discrete variables, and function (1) and (2) take their one-hot encodings as input in practice. From Equations (1) and (2), we can see that dag2seq uses the canonical indices as node features and applies a one-layer injective GNN to obtain the positional encoding of each node. Note that $(i, j) \in E$ is equivalent to $(\pi(i), \pi(j)) \in E^C$. Hence, Equations (1) and (2) can also be interpreted as applying a one-layer injective GNN on the canonical graph $C(G)$ with the true node indices of $C(G)$ as node features. For notation convenience, we use $\pi^{-1} : V^C \rightarrow V$ to denote the inverse function of π . Then Theorem 3.1 describes how dag2seq generates sequences that uniquely represent DAGs.

Theorem 3.1. *Let $G = (V, E, o)$ be a labeled DAG,*

*and p_1, p_2, \dots, p_n be the positional encodings generated by dag2seq. If functions *Agg* and *Combine* are injective, then the sequence $(o(\pi^{-1}(1)), p_{\pi^{-1}(1)}), (o(\pi^{-1}(2)), p_{\pi^{-1}(2)}), \dots, (o(\pi^{-1}(n)), p_{\pi^{-1}(n)})$ injectively encodes the computation structures defined by DAGs.*

We prove Theorem 3.1 in Appendix A. Theorem 3.1 provides a guarantee that node sequences generated by dag2seq injectively encode the original DAGs which fully preserve the node type information as well as structure information of the original DAGs. In other words, two labeled DAGs will be encoded into the same sequence if and only if they are isomorphic (i.e. their computation structures are the same). Then advanced parallelizable encoders (such as the Transformer encoder) for sequence modeling can be applied to the DAG encoding problem to improve the efficiency of the encoding process, thus facilitating the downstream optimization and prediction routines. Intuitively, the canonical form provides a unified node indexing for isomorphic DAGs which initially may have different node orderings, while the one-layer injective GNN encodes the direct predecessors of each node into its positional encoding. Then it is not difficult to see that from the canonical indices of nodes and their direct predecessors as well as the node types we can fully recover the original DAG. The one-layer GNN is parallelizable, in contrast to the sequential GNNs used in previous works (Zhang et al., 2019; Thost & Chen, 2021).

It is also worth discussing the complexity of graph canonization. The graph canonization problem is theoretically at least as computationally hard as the graph isomorphism problem, which is in NP but not known to be solvable in polynomial time nor to be NP-complete. However, getting the canonical form of graphs is not too difficult in practice, thanks to the well-known graph canonization tools such as Nauty (McKay & Piperno, 2014). Empirically, such tools usually return the canonical form of a reasonable-sized graph in seconds. Theoretically, Nauty has an average time complexity of $O(n)$, and polynomial-time graph canonization algorithms also exist for graphs of bounded degrees. We also found in our experiments that graph canonization adds a negligible overhead.

3.2. The Transformer in PACE

With the dag2seq framework to injectively map DAGs to sequences, we next develop the attention-based parallelizable encoder which consists of K stacked Transformer (encoder) blocks.

For each item $(o(\pi^{-1}(i)), p_{\pi^{-1}(i)})$ in the output sequence of dag2seq, $o(\pi^{-1}(i))$ provides the operation information, while the positional encoding $p_{\pi^{-1}(i)}$ contains the structural information. Hence, PACE concatenates the trainable embedding of operation $o(\pi^{-1}(i))$ and the positional encoding

$p_{\pi^{-1}(i)}$ as the embedding e_i of item i in the sequence.

$$e_i = \text{Concat}(\text{Emb}(o(\pi^{-1}(i))), p_{\pi^{-1}(i)}) \quad (3)$$

Then the sequence e_1, e_2, \dots, e_n is fed into the first Transformer encoder block. Each transformer encoder block performs the multi-head self-attention mechanism (Vaswani et al., 2017) to update the embedding of each item in the sequence. We provide details of the multi-head self-attention mechanism in Appendix C.

In the original transformer encoder blocks, the attention operations are not masked. In other words, for any two items i and j in the sequence, the embedding of item i will be used to update the embedding of item j . Such unmasked operation is reasonable as the positional encodings $\{p_{\pi^{-1}(i)}, \forall i\}$ have already encoded the partial order between nodes. However, due to the complex dependencies a DAG may encode, fully relying on the positional encodings to capture such dependencies might not be enough. Therefore, we introduce a masked attention operation that helps better capture the dependency between nodes in practice. The masked attention operation can be specified through a binary mask matrix M such that $M_{i,j} = \text{False}$ if there exists a path from i to j in $C(G)$ and $M_{i,j} = \text{True}$ otherwise.

In this mask matrix M , element $M_{i,j} = \text{True}$ indicates that the effect of item i in updating the embedding of item j will be masked out. When $M_{i,j} = \text{True}$, there is no (directed) path from node i to node j in the canonical form $C(G)$. In other words, we only allow j 's predecessors in $C(G)$ (or equivalently, $\pi^{-1}(j)$'s predecessors in G) to participate in the updating of j 's embedding. Such a masking operation has two benefits: 1) the structure information of the DAG is strengthened in the masked self-attention, and 2) the partial order between operations in the computation structure is exploited, which aligns with the logic of a real computation in the sense that the operation at some node does not depend on its successor operations. We also empirically verify the effectiveness of the masked attention operation in the ablation study. The mask matrix M can be efficiently computed through the DFS algorithm or the Floyd algorithm, which we describe in Appendix B.

Together with the proposed dag2seq framework, Transformer in PACE helps reduce the complexity of covering the entire DAG's diameter significantly. The positional encoding (from dag2seq) of each node encodes its canonical index and its incoming neighbors. Since the masked self-attention aggregates messages from all predecessors of a node, this node will know all its predecessors and their connection structures in one step, which provides us all the information needed to compute graph metrics such as shortest path distance to this node with only one message passing step. In contrast, ordinary GNNs require at least $O(\text{shortest path distance})$ message passing steps to capture

the dependency/attention between two nodes. For the last node to capture the entire graph structure, GNNs will require $O(\text{diameter})$ message passing steps. Therefore, PACE allows using much shallower layers than GNNs to encode DAGs (especially when the DAG has a long diameter) and increases the speed significantly.

3.3. Training Methodology

We design two schemes to train the PACE model. One is to train a variational autoencoder (VAE) for DAGs that is able to encode and decode/generate DAGs into and from a latent space, like D-VAE. On the other hand, since the proposed dag2seq transforms the DAG encoding problem to the sequence encoding problem, pre-training (self-supervised learning) techniques in NLP are also suitable for training PACE.

Without loss of generality, we assume that there is a single output node in each DAG that has no successor. If not, we can add a virtual output node and add directed edges from all nodes whose out-degree is 0 to the output node. When PACE is trained in a VAE architecture, we use a common trick in standard Transformers by assuming that there are at most N nodes in the input DAG. If a DAG has $n < N$ nodes, we pad $N - n$ end symbols to the end of the sequence generated by dag2seq. Then PACE reads out the DAG encoding by concatenating the learned embeddings of the N symbols. When PACE is trained in a pre-training architecture, similar to the sentiment classification task in BERT, PACE takes the learned embedding of the output node as the DAG encoding.

Training PACE in a VAE architecture. In the PACE-VAE architecture, we take PACE as the encoder, and connect the output of PACE with two fully connected (FC) layers to predict the mean and variance of the approximated posterior distribution in the evidence lower bound (ELBO) (Kingma & Welling, 2013). The decoder of PACE-VAE consists of $K = 3$ standard Transformer decoder blocks of dimension d_k . Given the latent vector z to decode, the decoder uses a FC layer to reconstruct a vector of dimension $N \times d_k$. The vector is then reshaped to a matrix Z of shape (N, d_k) , which plays the same role as the "memory" matrix in a standard Transformer in NLP. Hence, the encoder-decoder attention layer in each Transformer decoder block takes the decoded matrix Z as the "Key" matrix and "Value" matrix, and uses the output from the previous self-attention layer as the "Query" matrix. Similar to the standard Transformer, during the generation, the decoder of PACE-VAE sequentially predicts nodes in G according to the learnt canonical order, and this process is ended until a special symbol is predicted indicating the decoding process is completed. For each generated node i in the decoding process, we do a softmax to select the operation of the node, and use a binary

classifier to predict the existence of an edge between node i and any node $j < i$. We describe details about PACE-VAE and its parallelizable training framework in Appendix D.

Training PACE in a pre-training architecture. After converting DAGs to sequences by the proposed dag2seq framework, PACE is essentially a sequence modeling encoder. (Yan et al., 2020) validates that the pre-training architecture in NLP that generates embeddings without using accuracies can better preserve the local structural relationship in the latent space. As such, PACE can also take the masked language modeling (MLM) (Devlin et al., 2019; Yan et al., 2021) objective for (pre-)training to capture the locality information of the computation structure defined by DAGs. For each input DAG, we randomly select 20% nodes for masking and prediction, where 80% of them are replaced with the `[MASK]` token and the remaining nodes are unchanged. The output embeddings are used to predict the original node operations $o(\pi^{-1}(i))$, and we train PACE by minimizing the cross-entropy loss of the predicted node operations and the true node operations.

4. Comparison to Related Works

Despite the great success of Transformers in modeling sequential data such as natural languages and images (Carion et al., 2020; Dosovitskiy et al., 2020), there has been no work generalizing Transformers to another important type of sequential data, namely DAGs. Existing DAG encoders mostly adopt an RNN-style encoding scheme to sequentially process nodes in a DAG. S-VAE (Bowman et al., 2016) takes as input the sequence of node strings which consist of the node type as well as the adjacency vector of each node, and then applies a GRU-based RNN to the topologically sorted node sequence to encode a DAG. To encode the semantics of the computation carried by a DAG, D-VAE (Zhang et al., 2019) proposes an DAG-style message passing framework that sequentially updates the DAG encoding following a topological order of nodes. For each node in a DAG, D-VAE takes a gated summation aggregator to combine information from all its direct predecessors, and then a GRU is used to update the node embedding based on the aggregated information and the node type. Similar to the encoder of D-VAE, Thost & Chen (2021) proposes DAGNN which also sequentially aggregates node features. It differs from D-VAE in the sense that the aggregator is constructed through the attention mechanism and it comes with a layer notion in the DAG encoding process. However, the sequential nature of these RNN-style DAG encoders precludes their parallelizability. In contrast, PACE adopts a dag2seq framework that encodes the dependencies between nodes in the positional encodings, and then applies a Transformer to the node sequence to encode a DAG parallelly.

Various previous works (e.g., S-VAE and GraphRNN) have

proposed to linearize graphs to node sequences by adopting different node ordering techniques. S-VAE topologically sorts nodes, where each node feature is the concatenation of the one-hot encoding of the node type and a 0/1 vector indicating whether (directed) edges exist from previous nodes to itself. On the other hand, GraphRNN takes the BFS ordering and utilizes the fixed M -dimensional vector to represent the node connectivity in the BFS queue. The above methods share a major limitation—both topological ordering and BFS ordering are not unique. That is, the same DAG can still be represented as different sequences, thus generating different DAG embeddings. Consequently, similar DAGs might be encoded to vastly different vectors in the DAG encoding space, which makes the encoding space unsmooth and the DAG optimization problem difficult. In contrast, our dag2seq framework injectively represents DAGs as sequences, which guarantees that the same DAG is always represented as the same sequence, and different DAGs are always represented as different sequences. This is crucial for the downstream DAG optimization problems.

Concurrent to our work, a large body of works are proposed to apply Transformer to general graph data. GT (Dwivedi & Bresson, 2020) and SAN (Kreuzer et al., 2021) propose to use the Laplacian positional encodings (Laplacian PEs) (Dwivedi & Bresson, 2020) to encode the relative positions of nodes in graphs, instead of the original positional encoding framework for sequences. On the other hand, GraphiT (Mialon et al., 2021) and Graphormer (Ying et al., 2021) incorporate the (relative) position of nodes in graphs in the attention mechanism. The difference is that GraphiT uses graph kernels to adjust the attention scores, while Graphormer utilizes the shortest path between nodes to do so. Furthermore, (Jain et al., 2021) applies a Transformer module after a standard GNN module to capture the both local and long-term context. However, all previous graph Transformers have problems when applying to DAGs. Laplacian PEs are constructed from the k smallest non-trivial eigenvectors of the graph Laplacian, which is inherently not suitable for DAGs because graph Laplacian is only defined for undirected graphs. Other relative-position-based approaches such as graph kernels and shortest path are symmetric in nature and do not take the absolute (canonical) index of each node, thus are not injective. In contrast, our positional encoding framework, dag2seq, is DAG-friendly, and injectively converts DAGs into sequences, which preserves the full structure and type information of DAGs so that we can safely treat DAG encoding as sequence encoding.

5. Experiments

In this section, we conduct experiments on popular DAG encoding datasets to validate the effectiveness and efficiency of the proposed PACE model against state-of-art GRU-

Table 1. Predictive performance on NA and BN.

| Evaluation Metric | NA | | BN | |
|-------------------|----------------------|----------------------|----------------------|----------------------|
| | RMSE ↓ | Pearson’s r ↑ | RMSE ↓ | Pearson’s r ↑ |
| PACE (our model) | 0.254 ± 0.002 | 0.964 ± 0.001 | 0.115 ± 0.004 | 0.994 ± 0.001 |
| DAGNN | 0.264 ± 0.004 | 0.964 ± 0.001 | 0.122 ± 0.004 | 0.991 ± 0.000 |
| D-VAE | 0.384 ± 0.002 | 0.920 ± 0.001 | 0.281 ± 0.004 | 0.964 ± 0.001 |
| S-VAE | 0.478 ± 0.002 | 0.873 ± 0.001 | 0.499 ± 0.006 | 0.873 ± 0.002 |
| GraphRNN | 0.726 ± 0.002 | 0.669 ± 0.001 | 0.779 ± 0.007 | 0.634 ± 0.001 |
| DeepGMG | 0.478 ± 0.002 | 0.873 ± 0.001 | 0.843 ± 0.007 | 0.555 ± 0.003 |
| GCN | 0.832 ± 0.001 | 0.527 ± 0.001 | 0.599 ± 0.006 | 0.809 ± 0.002 |
| gated-GCN | 0.416 ± 0.002 | 0.891 ± 0.001 | 0.311 ± 0.003 | 0.953 ± 0.002 |
| GT | 0.329 ± 0.001 | 0.942 ± 0.001 | 0.166 ± 0.003 | 0.987 ± 0.000 |
| SAN | 0.311 ± 0.003 | 0.950 ± 0.001 | 0.158 ± 0.005 | 0.989 ± 0.001 |
| Graphormer | 0.352 ± 0.002 | 0.936 ± 0.001 | 0.181 ± 0.004 | 0.971 ± 0.001 |
| GraphiT | 0.299 ± 0.002 | 0.955 ± 0.001 | 0.142 ± 0.005 | 0.990 ± 0.003 |

Table 2. Downstream search performance on NAS101 and NAS301.

| Search Method | NAS101 (Regret) | | NAS301 (Acc) | |
|------------------|----------------------|----------------------|-----------------------|-----------------------|
| | DNGO (%) ↓ | DNGO-LS (%) ↓ | DNGO (%) ↑ | DNGO-LS (%) ↑ |
| PACE (our model) | 0.391 ± 0.241 | 0.278 ± 0.178 | 94.507 ± 0.165 | 94.547 ± 0.145 |
| DAGNN | 0.445 ± 0.224 | 0.448 ± 0.127 | 94.445 ± 0.219 | 94.433 ± 0.156 |
| D-VAE | 0.439 ± 0.203 | 0.430 ± 0.222 | 94.453 ± 0.148 | 94.428 ± 0.131 |
| S-VAE | 0.458 ± 0.175 | 0.451 ± 0.225 | 94.332 ± 0.183 | 94.371 ± 0.203 |
| GIN | 0.593 ± 0.177 | 0.518 ± 0.201 | 94.451 ± 0.224 | 94.411 ± 0.198 |
| GAT | 0.597 ± 0.269 | 0.509 ± 0.187 | 94.430 ± 0.171 | 94.421 ± 0.202 |
| GCN | 0.627 ± 0.161 | 0.538 ± 0.233 | 94.448 ± 0.149 | 94.404 ± 0.160 |
| gated-GCN | 0.451 ± 0.177 | 0.428 ± 0.168 | 94.461 ± 0.170 | 94.421 ± 0.164 |
| GT | 0.573 ± 0.276 | 0.460 ± 0.148 | 94.421 ± 0.174 | 94.533 ± 0.139 |
| SAN | 0.390 ± 0.250 | 0.291 ± 0.166 | 94.446 ± 0.191 | 94.501 ± 0.143 |
| Graphormer | 0.429 ± 0.302 | 0.314 ± 0.182 | 94.477 ± 0.142 | 94.551 ± 0.137 |
| GraphiT | 0.407 ± 0.233 | 0.307 ± 0.181 | 94.482 ± 0.156 | 94.489 ± 0.138 |

based DAG encoders, general-purpose GNN-based graph encoders, and recent (undirected) graph Transformers.

5.1. Datasets and Metrics

NA and BN. The dataset NA consists of approximately 19K neural architectures generated by the software ENAS (Pham et al., 2018). Each architecture has its pre-computed weight-sharing (WS) accuracy on CIFAR-10 (Krizhevsky et al., 2009) and includes 8 nodes. The dataset BN consists of 200K Bayesian networks randomly generated by the bnlearn package (Scutari, 2010). Each Bayesian network has 8 nodes and is associated with a Bayesian Information Criterion (BIC) score that measures the architecture performance on dataset Asia (Lauritzen & Spiegelhalter, 1988). Following the experimental settings used in (Zhang et al., 2019), PACE is evaluated under a VAE architecture, and we take 90% NA/BN data as the training set and hold out the rest for testing. To make a fair comparison, we evaluate the quality of DAG encoders by measuring the predictive performance as well as the downstream search performance. Briefly, a sparse Gaussian process (SGP) regression model

(Snelson & Ghahramani, 2005) is trained to predict the DAG performance from its encoding, and we use rooted mean square error (RMSE) and Pearson correlation (Pearson’s r) as metrics to evaluate the predictive performance. When evaluating the downstream search performance, we perform Bayesian optimization (BO) in the DAG encoding space based on the SGP regression model, and compare the performance of the best searched architecture.

NAS101 and NAS301. NAS101 (NAS-Bench-101) and NAS301 (NAS-Bench-301) are two well-known neural architecture search (NAS) benchmark datasets. NAS101 (Ying et al., 2019) consists of approximately 420K neural architectures with pre-computed validation and test accuracies on CIFAR-10, where each architecture has up to 7 nodes and 9 edges. NAS301 (Siems et al., 2020) is a surrogate benchmark for DARTS (Liu et al., 2018b). Following (Liu et al., 2018a; Yan et al., 2021), we randomly sample 1M neural architectures, where each architecture contains at most 15 nodes. On NAS101 and NAS301, the PACE model is trained with the pre-training architecture using the MLM objective. Since better predictive performance

of DAG encoders always facilitates the downstream search, we implement two popular BO-based downstream search methods, DNGO (Snoek et al., 2015) and DNGO-LS (Yan et al., 2021), and compare the downstream search performance of different encoders. For NAS101, following the original work of (Ying et al., 2019), we take the regret as the evaluation metric, where the regret is the difference between the test accuracy of the (offline) best neural architecture and that of the best searched neural architecture (after 20 rounds). For NAS301, since we do not have an oracle for the (offline) best neural architecture, we use the test accuracy of the best searched neural architecture, instead.

OGBG-CODE2. Encoding ASTs (abstract syntax trees) is another application area for DAG-based processing. Hence, we also evaluate PACE on dataset OGBG-CODE2 (Hu et al., 2020) against state-of-art GRU-based DAG encoders (i.e. D-VAE and DAGNN). Dataset OGBG-CODE2 contains approximately 450K Python functions parsed into DAGs. These DAGs in average contain more than 120 nodes, and the largest DAG has more than 30000 nodes. In the experiment, PACE takes the learnt representation of the output node as the whole-graph embedding for the token prediction task (TOK), and the test F1 score is used as the evaluation metric.

5.2. Baselines and Model Configuration

We benchmark the proposed PACE with 1) recent (undirected) graph Transformers (i.e. GT (Dwivedi & Bresson, 2020), SAN (Kreuzer et al., 2021), GraphiT (Mialon et al., 2021) and Graphormer (Ying et al., 2021)); 2) DAG encoders (i.e. GraphRNN (You et al., 2018), DeepGMG (Li et al., 2018), S-VAE, D-VAE and DAGNN); and 3) GNN-based graph encoders (i.e. gated GCN (Bresson & Laurent, 2017), GCN (Kipf & Welling, 2016), GIN (Xu et al., 2019), GAT (Velickovic et al., 2018)). In the experiments, PACE uses 3 Transformer encoder blocks to boost the training and inference speed. The dimension of the embedding layer that maps node types to embeddings is 64. The output dimension of the 1-layer GNN in dag2seq is also 64. On NA and BN, we concatenate the positional encodings and node type embeddings as the node features fed into the first Transformer encoder block. On NAS101 and NAS301, we use the summation of positional encodings and node type embeddings, instead. All the experiments are done on NVIDIA Tesla P100 12GB GPUs.

5.3. Experimental Results

NA and BN: Table 1. In the experiment, PACE achieves the smallest RMSE and the largest Pearson’s r on both NA and BN, indicating that PACE generates the smoothest encoding space with respect to the computation structure defined by DAGs. In addition, the improvement is more significant

on NA. One possible reason is that DAGs in NA always have a Hamiltonian path which introduces the long-term dependencies between nodes. Since the Transformer-style attention in PACE can better capture the long-term dependency than GRU-based DAG encoders, PACE can preserve the similarity of DAGs better in the learnt DAG encoding space. Furthermore, Appendix E and Appendix G compare the downstream search performance and generation performance, respectively. The results show that PACE still achieves the best performance, and the observation suggests that PACE has superior DAG-encoding ability than competitive encoders.

NAS101 and NAS301: Table 2. In the experiment, PACE significantly outperforms other baselines. When applying DNGO-IS search on NAS301, PACE achieves the second best performance. In other cases, PACE has the best downstream search performance. Similar to PACE, DAGNN and GAT also use the attention mechanism to model the dependencies (relations) between nodes. However, the attention mechanism in these encoders is only applied to nodes and their direct predecessors (DAGNN) or adjacent nodes (GAT), hence making it hard to capture the long-term dependencies between nodes. On the contrary, PACE allows all predecessor nodes to participate in the attention mechanism while encoding the neighbors through positional encoding, which enables PACE to learn both the long-term dependencies and short-term dependencies with ease, thus effectively capturing the similarity of computation structure defined by DAGs. In addition, Laplacian PEs is dependent on the factorization of the graph Laplacian matrix, hence techniques based on Laplacian PEs (i.e. GT, SAN, and gated-GCN) are inherently not suitable for DAGs, and their performance are not comparable to the proposed DAG-friendly PACE model.

OGBG-CODE2: Table 3. The results interestingly reveal that PACE effectively learn the long-term dependencies between nodes in a DAG. In this experiment, the training process always takes a batch size of 16 and a training epoch of 30, thus we use the average training and inference time (per epoch) to measure the encoding speed. Experimental results show that PACE significantly improves the predictive performance over the state-of-art GRU-based DAG encoders (i.e. DAGNN and D-VAE), and PACE only requires about $\frac{1}{3}$ training time and $\frac{1}{3}$ inference time. More details about the experiment are provided in Appendix F.

Table 3. OGBG-CODE2: PACE versus state-of-art DAG encoders

| Methods | Test F1 score \uparrow | Training time (min) \downarrow | Inference time (min) \downarrow |
|---------|---------------------------------------|------------------------------------|-----------------------------------|
| PACE | 0.1779 \pm 0.0021 | 40.71 \pm 0.19 | 2.11 \pm 0.22 |
| DAGNN | 0.1751 \pm 0.0049 | 166.19 \pm 3.52 | 6.77 \pm 0.36 |
| D-VAE | 0.1596 \pm 0.0041 | 134.54 \pm 3.06 | 6.08 \pm 0.16 |

Table 4. Ablation study.

| Model configuration | NA | | BN | | NAS101 (Regret) | | NAS301 (Acc) | |
|------------------------------|--------------|---------------|--------------|---------------|-----------------|---------------|---------------|---------------|
| | RMSE ↓ | Pearson's r ↑ | RMSE ↓ | Pearson's r ↑ | DNGO (%) ↓ | DNGO-LS (%) ↓ | DNGO (%) ↑ | DNGO-LS (%) ↑ |
| Model1: dag2seq & Mask | 0.254 | 0.964 | 0.115 | 0.9942 | 0.391 | 0.278 | 94.507 | 94.547 |
| Model2: dag2seq & No Mask | 0.255 | 0.967 | 0.119 | 0.9941 | 0.479 | 0.368 | 94.501 | 94.505 |
| Model3: No dag2seq & No mask | 0.981 | 0.001 | 0.368 | 0.9318 | 0.600 | 0.498 | 94.467 | 94.401 |

Table 5. Comparison of different DAG-to-sequence methods + Transformer

| | NA | | BN | | NAS101 (regret) | | NAS301 (acc) | | ogbg-code2 |
|------------------------|--------------|---------------|--------------|---------------|-----------------|---------------|---------------|---------------|-----------------|
| | RMSE ↓ | Pearson's r ↑ | RMSE ↓ | Pearson's r ↑ | DNGO (%) ↓ | DNGO-LS (%) ↓ | DNGO (%) ↑ | DNGO-LS (%) ↑ | Test F1 score ↑ |
| PACE | 0.254 | 0.964 | 0.115 | 0.994 | 0.391 | 0.278 | 94.507 | 94.547 | 0.1779 |
| S-VAE (Transformer) | 0.392 | 0.903 | 0.417 | 0.901 | 0.439 | 0.386 | 94.361 | 94.384 | 0.1481 |
| GraphRNN (Transformer) | 0.406 | 0.896 | 0.431 | 0.889 | 0.427 | 0.371 | 94.410 | 94.392 | 0.1463 |

5.4. Computational Cost

A key advantage of PACE is the parallelizable DAG encoding process, so we compare the computational cost of PACE to GRU-based DAG encoders (D-VAE and DAGNN). We use a single GPU for each experiment, and **Figure 3** shows our results. In this experiment, PACE is trained with the VAE architecture on datasets NA and BN, but with BERT-like objective (the pre-training architecture) on datasets NAS101 and NAS031. On the other hand, D-VAE and DAGNN are always trained with a VAE architecture. Then, to make a fair comparison, we compare the total training time and the average inference time per epoch to evaluate the computational cost of each method. **Figure 3** shows that PACE requires about $\frac{1}{3}$ total training time and about $\frac{1}{3}$ average inference time compared to GRU-based DAG encoders (i.e. D-VAE and DAGNN). Hence, PACE significantly boosts the DAG encoding speed.

5.5. Ablation Study

In the ablation study, we demonstrate the effectiveness of our proposed dag2seq (positional encoding) framework and the attention mask in PACE. From **Table 4**, we have the following observations: 1) In general, PACE trained with attention mask outperforms the one without attention mask, indicating that the attention mask helps better capture the inductive bias of DAGs. Nevertheless, even without attention mask, PACE still performs relatively well because the dag2seq framework also captures the node dependencies. 2) We also find that the dag2seq framework is vital for the performance of PACE. Without dag2seq, PACE (without mask) almost completely fails at the NA dataset. This verifies the importance of dag2seq for solving the ambiguity issue of DAG encoding illustrated in **Figure 2**.

Comparison of different DAG-to-sequence methods To further decouple the contribution of dag2seq from that of the Transformer, we replace the RNN encoders in the baselines S-VAE and GraphRNN with Transformers, and compare them with our PACE model (dag2seq + Transformer). GraphRNN is originally designed as a pure genera-

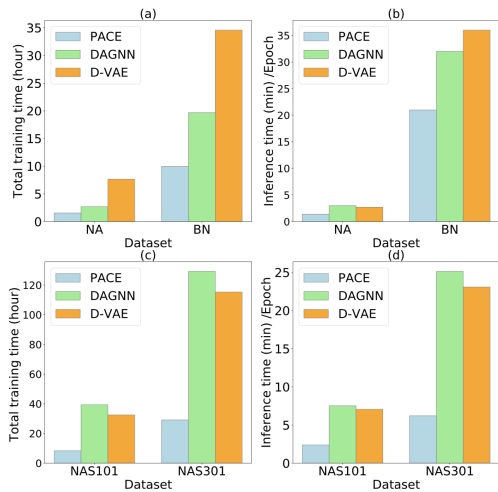


Figure 3. Computational cost.

tive model, and thus does not have a DAG encoder. However, following the idea of S-VAE, we generate the node feature sequence by replacing the topological ordering in S-VAE with the BFS ordering as in GraphRNN, and replacing the 0/1 vector in S-VAE with the fixed M -dimensional vector. **Table 5** shows the results. PACE still demonstrates significant performance advantages, illustrating the importance of dag2seq for DAG encoding.

6. Conclusion

In this paper, we have proposed PACE, a novel DAG encoder based on Transformer. Unlike traditional GRU-based DAG encoders which sequentially encode DAG nodes, PACE is fully parallelizable and effectively learn the long-range dependencies of node pairs in a DAG. PACE incorporates the strong relational inductive bias through a node-dependency-aware positional encoding framework, dag2seq, and a masked self-attention mechanism. Experiments demonstrate that PACE not only generates smooth latent (DAG encoding) space beneficial to the downstream prediction and optimization routines, but also significantly boosts the encoding speed.

References

- Abdel-Hamid, O., Mohamed, A.-r., Jiang, H., Deng, L., Penn, G., and Yu, D. Convolutional neural networks for speech recognition. *IEEE/ACM Transactions on audio, speech, and language processing*, 22(10):1533–1545, 2014.
- Al-Rfou, R., Choe, D., Constant, N., Guo, M., and Jones, L. Character-level language modeling with deeper self-attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3159–3166, 2019.
- Allamanis, M., Barr, E. T., Devanbu, P., and Sutton, C. A survey of machine learning for big code and naturalness. *ACM Computing Surveys (CSUR)*, 51(4):1–37, 2018.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *International Conference on Learning Representations (ICLR)*, 2015.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Bowman, S., Vilnis, L., Vinyals, O., Dai, A., Jozefowicz, R., and Bengio, S. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pp. 10–21, 2016.
- Bresson, X. and Laurent, T. Residual gated graph convnets. *arXiv preprint arXiv:1711.07553*, 2017.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., and Zagoruyko, S. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pp. 213–229. Springer, 2020.
- Cho, K., van Merriënboer, B., Gülçehre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*, 2014.
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J. G., Le, Q., and Salakhutdinov, R. Transformer-xl: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2978–2988, 2019.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2019.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Duvenaud, D., Maclaurin, D., Aguilera-Iparraguirre, J., Gómez-Bombarelli, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. *Advances in Neural Information Processing Systems*, 2015:2224–2232, 2015.
- Dwivedi, V. P. and Bresson, X. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699*, 2020.
- Elsken, T., Metzen, J. H., Hutter, F., et al. Neural architecture search: A survey. *J. Mach. Learn. Res.*, 20(55):1–21, 2019.
- Gehring, J., Auli, M., Grangier, D., Yarats, D., and Dauphin, Y. N. Convolutional sequence to sequence learning. In *International Conference on Machine Learning*, pp. 1243–1252. PMLR, 2017.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pp. 1263–1272. PMLR, 2017.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7ebea9-Paper.pdf>.
- 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. *arXiv preprint arXiv:2005.00687*, 2020.
- Jain, P., Wu, Z., Wright, M., Mirhoseini, A., Gonzalez, J. E., and Stoica, I. Representing long-range context for graph neural networks with global attention. *Advances in Neural Information Processing Systems*, 34, 2021.
- Kandasamy, K., Neiswanger, W., Schneider, J., Póczos, B., and Xing, E. P. Neural architecture search with bayesian optimisation and optimal transport. In *NeurIPS*, 2018.

- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., and Tang, P. T. P. On large-batch training for deep learning: Generalization gap and sharp minima. *International Conference on Learning Representations (ICLR)*, 2017.
- Khan, W., Daud, A., Nasir, J. A., and Amjad, T. A survey on the state-of-the-art machine learning models in the context of nlp. *Kuwait journal of Science*, 43(4), 2016.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Koller, D. and Friedman, N. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Kreuzer, D., Beaini, D., Hamilton, W. L., Létourneau, V., and Tossou, P. Rethinking graph transformers with spectral attention. *arXiv preprint arXiv:2106.03893*, 2021.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Lauritzen, S. L. and Spiegelhalter, D. J. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 50(2):157–194, 1988.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7871–7880, 2020.
- Li, Y., Vinyals, O., Dyer, C., Pascanu, R., and Battaglia, P. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.
- Liu, C., Zoph, B., Neumann, M., Shlens, J., Hua, W., Li, L.-J., Fei-Fei, L., Yuille, A., Huang, J., and Murphy, K. Progressive neural architecture search. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 19–34, 2018a.
- Liu, H., Simonyan, K., Vinyals, O., Fernando, C., and Kavukcuoglu, K. Hierarchical representations for efficient architecture search. *arXiv preprint arXiv:1711.00436*, 2017.
- Liu, H., Simonyan, K., and Yang, Y. Darts: Differentiable architecture search. In *International Conference on Learning Representations*, 2018b.
- McKay, B. D. and Piperno, A. Practical graph isomorphism. ii. *Journal of symbolic computation*, 60:94–112, 2014.
- Medhat, W., Hassan, A., and Korashy, H. Sentiment analysis algorithms and applications: A survey. *Ain Shams engineering journal*, 5(4):1093–1113, 2014.
- Medsker, L. R. and Jain, L. Recurrent neural networks. *Design and Applications*, 5:64–67, 2001.
- Mialon, G., Chen, D., Selosse, M., and Mairal, J. Graphit: Encoding graph structure in transformers. *arXiv preprint arXiv:2106.05667*, 2021.
- Niepert, M., Ahmed, M., and Kutzkov, K. Learning convolutional neural networks for graphs. In *International conference on machine learning*, pp. 2014–2023. PMLR, 2016.
- Pham, H., Guan, M., Zoph, B., Le, Q., and Dean, J. Efficient neural architecture search via parameters sharing. In *International Conference on Machine Learning*, pp. 4095–4104. PMLR, 2018.
- Scutari, M. Learning bayesian networks with the bnlearn r package. *Journal of Statistical Software*, 35(i03), 2010.
- Shi, H., Pi, R., Xu, H., Li, Z., Kwok, J., and Zhang, T. Bridging the gap between sample-based and one-shot neural architecture search with bonas. *Advances in Neural Information Processing Systems*, 33, 2020.
- Siems, J., Zimmer, L., Zela, A., Lukasik, J., Keuper, M., and Hutter, F. Nas-bench-301 and the case for surrogate benchmarks for neural architecture search. *arXiv preprint arXiv:2008.09777*, 2020.
- Snelson, E. and Ghahramani, Z. Sparse gaussian processes using pseudo-inputs. *Advances in neural information processing systems*, 18:1257–1264, 2005.
- Snoek, J., Rippel, O., Swersky, K., Kiros, R., Satish, N., Sundaram, N., Patwary, M., Prabhat, M., and Adams, R. Scalable bayesian optimization using deep neural networks. In *International conference on machine learning*, pp. 2171–2180. PMLR, 2015.
- Thost, V. and Chen, J. Directed acyclic graph neural networks. *ArXiv*, abs/2101.07965, 2021.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 6000–6010, 2017.
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. *ArXiv*, abs/1710.10903, 2018.

- Vinyals, O., Bengio, S., and Kudlur, M. Order matters: Sequence to sequence for sets. *arXiv preprint arXiv:1511.06391*, 2015.
- Wen, W., Liu, H., Chen, Y., Li, H., Bender, G., and Kirdermans, P.-J. Neural predictor for neural architecture search. In *European Conference on Computer Vision*, pp. 660–676. Springer, 2020.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ryGs6iA5Km>.
- Xu, K., Li, J., Zhang, M., Du, S. S., Kawarabayashi, K.-i., and Jegelka, S. What can neural networks reason about? *ICLR 2020*, 2020.
- Yan, S., Zheng, Y., Ao, W., Zeng, X., and Zhang, M. Does unsupervised architecture representation learning help neural architecture search? *Advances in Neural Information Processing Systems*, 33, 2020.
- Yan, S., Song, K., Liu, F., and Zhang, M. Cate: Computation-aware neural architecture encoding with transformers. *arXiv preprint arXiv:2102.07108*, 2021.
- Ying, C., Klein, A., Christiansen, E., Real, E., Murphy, K., and Hutter, F. Nas-bench-101: Towards reproducible neural architecture search. In *International Conference on Machine Learning*, pp. 7105–7114. PMLR, 2019.
- Ying, C., Cai, T., Luo, S., Zheng, S., Ke, G., He, D., Shen, Y., and Liu, T.-Y. Do transformers really perform bad for graph representation? *arXiv preprint arXiv:2106.05234*, 2021.
- You, J., Ying, R., Ren, X., Hamilton, W., and Leskovec, J. Graphrnn: Generating realistic graphs with deep autoregressive models. In *International Conference on Machine Learning*, pp. 5708–5717. PMLR, 2018.
- You, J., Gomes-Selman, J., Ying, R., and Leskovec, J. Identity-aware graph neural networks. *arXiv preprint arXiv:2101.10320*, 2021.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- Zhang, M. and Chen, Y. Link prediction based on graph neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 5171–5181, 2018.
- Zhang, M. and Li, P. Nested graph neural networks. *Advances in Neural Information Processing Systems*, 34, 2021.
- Zhang, M., Cui, Z., Jiang, S., and Chen, Y. Beyond link prediction: Predicting hyperlinks in adjacency space. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018a.
- Zhang, M., Cui, Z., Neumann, M., and Chen, Y. An end-to-end deep learning architecture for graph classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018b.
- Zhang, M., Jiang, S., Cui, Z., Garnett, R., and Chen, Y. D-vae: A variational autoencoder for directed acyclic graphs. *Advances in neural information processing systems*, 32, 2019.
- Zoph, B. and Le, Q. V. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.

A. Proof of Theorem 3.1

Let $G_1 = (V_1, E_1, o_1)$ and $G_2 = (V_2, E_2, o_2)$ be two labelled graphs, then G_1 and G_2 are isomorphic (i.e. G_1 and G_2 represent the same computation structure.) if and only if their canonical forms are identical, i.e. $C(G_1) = C(G_2)$. Note that this equation means equality between the canonical forms, not isomorphism. Let $C(G_1) = (V_1^C, E_1^C, o_1^C)$ and $C(G_2) = (V_2^C, E_2^C, o_2^C)$. As discussed in section 3.1, there exists bijections $\pi_1 : V_1 \rightarrow V_1^C$ and $\pi_2 : V_2 \rightarrow V_2^C$, and we use $\pi_1^{-1} : V_1^C \rightarrow V_1$ and $\pi_2^{-1} : V_2^C \rightarrow V_2$ to denote their inverse functions. Then, we have $C(G_1) = C(G_2)$ if and only if (1) $o_1^C(i) = o_2^C(i)$ for $\forall i$; and (2) $(i, j) \in E_1^C \Leftrightarrow (i, j) \in E_2^C$ for $\forall i, j$.

Next, we will prove Theorem 3.1 by equivalently showing that the sequence $(o(\pi^{-1}(1)), p_{\pi^{-1}(1)}), (o(\pi^{-1}(2)), p_{\pi^{-1}(2)}), \dots, (o(\pi^{-1}(n)), p_{\pi^{-1}(n)})$ can guarantee the distinctness of canonical forms $C(G)$. For the notation convenience, let function $f(\pi(j), \{\pi(i), (i, j) \in E\}) = \text{Combine}(\pi(j), \text{Agg}(\{\pi(i), (i, j) \in E\}))$ be the composition of functions Agg and Combine, then it is injective if and only if both Agg and Combine are injective. Furthermore, we use Seq_1 to denote the sequence $(o_1(\pi_1^{-1}(1)), p_{\pi_1^{-1}(1)}), (o_1(\pi_1^{-1}(2)), p_{\pi_1^{-1}(2)}), \dots, (o_1(\pi_1^{-1}(n)), p_{\pi_1^{-1}(n)})$, and Seq_2 to denote the sequence $(o_2(\pi_2^{-1}(1)), p_{\pi_2^{-1}(1)}), (o_2(\pi_2^{-1}(2)), p_{\pi_2^{-1}(2)}), \dots, (o_2(\pi_2^{-1}(n)), p_{\pi_2^{-1}(n)})$.

So far, we know $C(G_1) \neq C(G_2) \Leftrightarrow$ there (1) exists i such that $o_1^C(i) \neq o_2^C(i)$, or (2) exists i, j such that $(i, j) \in E_1^C$ but $(i, j) \notin E_2^C$ (equivalently, $(i, j) \notin E_1^C$ but $(i, j) \in E_2^C$).

Now, let's prove $C(G_1) \neq C(G_2) \Rightarrow Seq_1 \neq Seq_2$.

- **(1) For the first case**, since π_1, π_2 are the bijections that map G_1 and G_2 to their canonical forms, then we have:

$$\begin{aligned} o_1(\pi_1^{-1}(i)) &= o_1^C(\pi_1(\pi_1^{-1}(i))) \\ &= o_1^C(i) \\ o_2(\pi_2^{-1}(i)) &= o_2^C(\pi_2(\pi_2^{-1}(i))) \\ &= o_2^C(i) \end{aligned}$$

Since $o_1^C(i) \neq o_2^C(i)$, then we get $o_1(\pi_1^{-1}(i)) \neq o_2(\pi_2^{-1}(i))$, indicating that $Seq_1 \neq Seq_2$.

- **(2) For the second case**, according to the definition of canonical form, we know that $(\pi_1^{-1}(i), \pi_1^{-1}(j)) \in E_1 \Leftrightarrow (i, j) \in E_1^C$ (similarly, $(\pi_2^{-1}(i), \pi_2^{-1}(j)) \in E_2 \Leftrightarrow (i, j) \in E_2^C$). As such, we get:

$$\begin{aligned} p_{\pi_1^{-1}(j)} &= f(\pi_1(\pi_1^{-1}(j)), \{\pi_1(\pi_1^{-1}(s)), (\pi_1^{-1}(s), \pi_1^{-1}(j)) \in E_1\}) \\ &= f(j, \{s, (s, j) \in E_1^C\}) \\ p_{\pi_2^{-1}(j)} &= f(\pi_2(\pi_2^{-1}(j)), \{\pi_2(\pi_2^{-1}(s)), (\pi_2^{-1}(s), \pi_2^{-1}(j)) \in E_1\}) \\ &= f(j, \{s, (s, j) \in E_2^C\}) \end{aligned}$$

Then, since $(i, j) \in E_1^C$ but $(i, j) \notin E_2^C$, we have $\{s, (s, j) \in E_1^C\} \neq \{s, (s, j) \in E_2^C\}$. Since function f is injective, then we have $p_{\pi_1^{-1}(j)} \neq p_{\pi_2^{-1}(j)}$. Hence, $Seq_1 \neq Seq_2$.

In the end, let's prove the other direction, i.e. $Seq_1 \neq Seq_2 \Rightarrow C(G_1) \neq C(G_2)$. When $Seq_1 \neq Seq_2$, there (1) exists i such that $o_1(\pi_1^{-1}(i)) \neq o_2(\pi_2^{-1}(i))$, or (2) exists j such that $p_{\pi_1^{-1}(j)} \neq p_{\pi_2^{-1}(j)}$.

- **(1) For the first case**, according to the previous analysis, we have:

$$\begin{aligned} o_1^C(i) &= o_1(\pi_1^{-1}(i)) \\ o_2^C(i) &= o_2(\pi_2^{-1}(i)) \end{aligned}$$

Hence, we can get $o_1^C(i) \neq o_2^C(i)$, which indicates $C(G_1) \neq C(G_2)$.

- **(2) For the second case**, according to the previous analysis, we get:

$$\begin{aligned} p_{\pi_1^{-1}(j)} &= f(j, \{s, (s, j) \in E_1^C\}) \\ p_{\pi_2^{-1}(j)} &= f(j, \{s, (s, j) \in E_2^C\}) \end{aligned}$$

Since f is injective, $p_{\pi_1^{-1}(j)} \neq p_{\pi_2^{-1}(j)}$ implies that $\{s, (s, j) \in E_1^C\} \neq \{s, (s, j) \in E_2^C\}$. Then, there must exist i such that $(i, j) \in E_1^C$ but $(i, j) \notin E_2^C$ (or $(i, j) \notin E_1^C$ but $(i, j) \in E_2^C$). Henceforth, we get $C(G_1) \neq C(G_2)$.

B. Mask Matrix

Here we provide two potential ways to get the mask matrix in PACE. Following the same notation as the main paper, we use $C(G) = (V^C, E^C, o^C)$ to denote the canonical form of the input DAG G .

DFS Algorithm This algorithm takes the canonical form $C(G)$ as input and performs the DFS (depth first search) algorithm on the graph to explore all the nodes of the graph. Before we start the depth first search, we traverse all edges in E^C to find all direct-successors of each node i , and then put them in a set $S(i)$ for $i \in V^C$. Then, for each node i , we perform the DFS to get a dependent set $D(i)$, and we have $M_{j,i} = \text{False}$ if and only if $j \in D(i)$.

Algorithm 1 DFS Algorithm

Input: $D(i) = \{\}$; $Visited = [\text{False for } j \in V^C]$; a source (start) node i , $T = [i]$ (T is a stack).
 $Visited[i] = \text{True}$.
while $|T| > 0$. **do**
 $j = T[-1]$.
 delete j from T
 for k in $S(j)$ **do**
 if $Visited[k] = \text{False}$ **then**
 put k in $D(i)$
 $Visited[k] = \text{True}$
 put k in T
 end if
 end for
end while

Floyd Algorithm The Floyd algorithm is originally proposed for finding shortest paths in directed weighted graphs. Here, we initialize the edge weights to be 1, and implement the Floyd algorithm to find the distance $dist(i, j)$ (i.e. length of the shortest directed path) between each node pair i, j in $C(G)$. Then we have $M_{i,j} = \text{False}$ if and only if $dist(i, j) \neq \infty$.

Algorithm 2 Floyd Algorithm

Input: $dist(i, j) = 1$ if $(i, j) \in E^C$ else ∞
for $i \in V^C$ **do**
 for $j \in V^C$ **do**
 for $k \in V^C$ **do**
 if $dist(j, k) > dist(j, i) + dist(i, k)$. **then**
 $dist(j, k) = dist(j, i) + dist(i, k)$.
 end if
 end for
 end for
end for

C. Multi-Head Self-Attention Mechanism

Here we introduce the multi-head (masked) self-attention mechanism in the Transformer encoder block of PACE. For the notation convenience, we use H_k to denote the output representation of the k th Transformer encoder block, and use H_0 to denote the input (i.e. the representation of the sequence generated by dag2seq) to the first Transformer encoder block. Furthermore, we denote the number of heads in the self-attention mechanism as h , and the embedding dimension (of each

item in the sequence) as d . Then the Transformer encoder blocks update representation H_k as following.

$$H_k^j = \text{softmax}\left(\frac{Q_k^j (K_k^j)^T}{d}\right) V_k^j \quad \text{for } j = 1, 2, \dots, h \quad (4)$$

$$H_{k+1} = \text{feed-forward}\left(\|_{j=1}^h H_k^j\right) \quad (5)$$

where $Q_k^j = H_k W_{k,q}^j$, $K_k^j = H_k W_{k,k}^j$, $V_k^j = H_k W_{k,v}^j$ are the query matrix, key matrix and value matrix, respectively (i.e. $W_{k,q}^j, W_{k,k}^j, W_{k,v}^j$ are trainable parameter matrices); $\|$ represents the concatenation operation; Feed-forward is a one-layer MLP. When we introduce the mask operation into the Transformer encoder block, let M be the mask matrix from the Floyd algorithm or the BFS algorithm, then we use following equation to replace equation 5 in the Transformer encoder block.

$$H_k^j = \text{softmax}\left(\frac{Q_k^j (K_k^j)^T + -\infty * M}{d}\right) V_k^j \quad \text{for } j = 1, 2, \dots, h \quad (6)$$

D. Details about PACE in the VAE Architecture

In the section, we describe the decoder of PACE-VAE. Figure 4 illustrates the overall architecture. In the main paper, we have introduced how PACE maps input DAGs to the latent space, here we focus on the decoder of PACE-VAE.

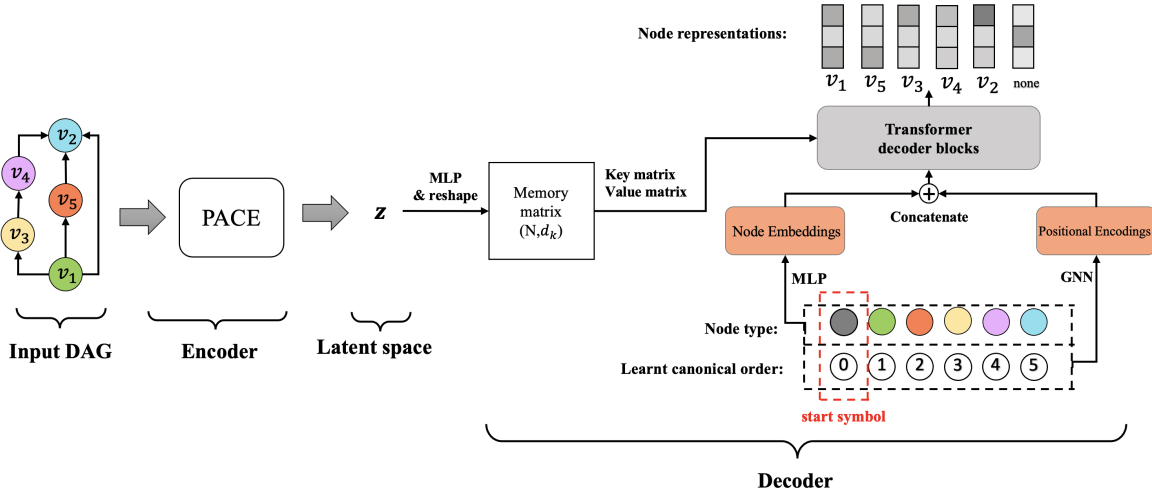


Figure 4. The illustration of PACE in the VAE architecture (PACE-VAE)

Similar to PACE, the decoder is constructed upon the Transformer decoder block. Each Transformer decoder block consists of a masked multi-head self-attention layer (i.e. Equation 6), a multi-head attention layer (i.e. Equation 4 except that the key matrix and value matrix are computed from points z in the latent space), and a feed-forward layer (i.e. Equation 5). The decoder takes a MLP as the embedding layer to generate node type embeddings as PACE. In analogous to the dag2seq framework in PACE, the decoder also uses a GNN to generate the positional encoding based on the learnt canonical order of nodes. The node embeddings and positional encodings are concatenated and then fed into multiple stacked Transformer decoder blocks to predict the node representations, which is used to predict the node types and the existence of edges. In analogous to the standard Transformer decoder, the decoder performs the shift right trick (i.e. the i th output node representation corresponds to the $i + 1$ th node in the sequence) and adds a start symbol node (i.e. the black node in Figure 4) at the beginning of the node sequence. Specifically, the canonical label of the start symbol node is different from any possible canonical label in the dag2seq framework to distinguish its position. For instance, DAG in the searching space contains at most N nodes, then the canonical order of the start symbol node can be 0 or $N + 1$. Let o_i denotes the output representation of node i in the sequence, then it is used to predict the type of node $i + 1$ in the sequence through a MLP. Similarly, for any $j < i$, we use another MLP, which takes the concatenation of o_j and o_i as input, to predict the existence of an directed edge from node $j + 1$ to node $i + 1$ in the sequence. Note that the canonical order can be generated from the topological sort by breaking ties using canonicalization tools, such as Nauty. Thus, for each node i in the sequence, any

dependent node j of this node must be arranged in a prior position in the sequence (i.e. $j < i$). In the end, based on these predictions (node representations), we can perform the teacher forcing to train the VAE.

Although the ordering of the output sequence can have a significant impact on the performance in the sequence-to-sequence model (Vinyals et al., 2015), it is not the same case for the decoder of PACE-VAE. Let x_1, x_2, \dots, x_n be an output sequence in the sequence-to-sequence model, and v_1, v_2, \dots, v_n an output node sequence in the decoder of PACE-VAE. Then, for any i, j such that $i < j$, the sequence-to-sequence model knows that x_j is arranged in a later position in the output sequence than x_i , in other words, x_j is dependent on x_i . However, in the decoder of PACE-VAE, the existence of an edge between v_i and v_j is predicted by the decoder itself, hence, any topological order is suitable in the decoder of PACE-VAE, and we select (topological) canonical order to facilitate the teacher forcing.

E. Downstream Search Performance on NA and BN

Table 6. Downstream search performance on NA and BN.

| Model | PACE | DAGNN | D-VAE | S-VAE |
|-------------------------------|------------------|------------------|-----------|-----------|
| (NA) Test accuracy \uparrow | 95.08 | 95.06 | 94.80 | 92.79 |
| (BN) BIC score \uparrow | -11107.29 | -11107.29 | -11125.75 | -11125.77 |




Figure 5. Best architectures on NA and BN detected by PACE.

In this section, we compare the downstream search performance on datasets NA and BN. As Table 6 shows, PACE detects architectures with the best performance on both NA and BN. Compared with GRU-based DAG encoders, the Transformer-style attention in PACE has superior ability for learning the long-term dependency. Hence, PACE’s latent embeddings are more informative about the computation structure, which favors using a predictive model in its latent space to guide the downstream search routines.

We also visualize the optimal architectures detected by Bayesian optimization (over the latent DAG encoding space generated by PACE) on datasets NA and BN. Figure 5 illustrates our results. On dataset BN, we find that the detected optimal Bayesian network structure is almost the same as the ground truth (Figure 2 in (Lauritzen & Spiegelhalter, 1988)), except that there is another directed edge from node A (visit to Asia ?) to node T (Tuberculosis) in the ground truth.

F. Experiment Details and Further Discussion on OGBG-CODE2

Our paper mainly focuses on the DAG optimization problem. Since searching the optimal neural architectures and Bayesian network structures is essentially typical DAG optimization task, we choose datasets NA, NAS101, NAS301 (neural architecture search benchmark) and BN (Bayesian network structure learning benchmark) for evaluation. However, encoding ASTs (Abstract Syntax Trees) can be another standard application area for DAG-based processing, then we implement additional experiments on dataset OGBG-CODE2 to evaluate PACE against GRU-based DAG encoders. On average, DAGs in dataset OGBG-CODE2 contain more than 120 nodes, while the largest DAG contains more than 30,000 nodes, and the depth of tree can reach 275. Hence, sequentially encoding DAGs in OGBG-CODE2 can be extremely costly, and PACE can take a relative larger number of (self-) attention layers (6 layers) in this setting. On dataset OGBG-CODE2, it’s conventionally to use a AST node encoder that incorporates edge features to generate continuous node embeddings. Similarly, PACE uses a GNN as AST node encoder to generate continuous node embeddings, and then dag2seq in PACE

approximately takes the DFS order as inputs.

Training method: As discussed in the main paper, PACE takes the learnt representation of the output node as the graph embedding. On dataset OGBG-CODE2, PACE is not trained in the VAE version nor in the BERT version. Instead, PACE is directly trained on the TOK (token prediction task) in a supervised version based on the graph embedding. Thus, all methods are trained with a single objective and then evaluated on that objective, so that the test F1 score can determine how easily the models can learn.

In general cases, PACE takes Floyd algorithm 2 (see Appendix B) to get the attention mask. However, Floyd algorithm has a time complexity of $O(N^3)$, where N is the number of nodes in a DAG. Then, the data preprocessing might be time-consuming when applied to large-scale graphs. Hence, on dataset OGBG-CODE2, we propose a backtracking algorithm based on the tree structure of ASTs. The backtracking algorithm has a time complexity of $O(N)$ when the depths of nodes in ASTs are bounded. Then we can significantly reduce the preprocessing time to obtain the attention mask. In algorithm 3, we use $AST = (V, E)$ to denote an abstract syntax tree, where nodes in V are sorted according to the DFS order. In all experiments, we include the computation time of finding the canonical form of DAGs in the training and inference time of PACE.

Algorithm 3 Backtracking Algorithm

Input: backtracking dictionary $B = \{\}$, $AST = (V, E)$, function $f : V \rightarrow \mathbb{N}$ that maps nodes to the depth, Attention mask M whose elements are set to be True.

```

for  $v \in V$  do
   $j = f(v)$ 
   $B(j) = v$ 
  if  $j > 0$ . then
    for  $k$  in 1 to  $j$  do
       $M_{B(j), B(k)} = \text{False}$ 
    end for
  end if
end for

```

G. Reconstruction Accuracy and Generation Performance

Table 7. Recon. accuracy, valid prior, uniqueness, novelty and overall (ave) performance %

| Methods | NA | | | | | BN | | | | |
|----------|---------------------|------------------|-------------------|------------------|--------------------|---------------------|------------------|-------------------|------------------|--------------------|
| | Accuracy \uparrow | Valid \uparrow | Unique \uparrow | Novel \uparrow | Overall \uparrow | Accuracy \uparrow | Valid \uparrow | Unique \uparrow | Novel \uparrow | Overall \uparrow |
| PACE | 99.97 | 98.16 | 57.77 | 100.00 | 88.98 | 99.99 | 99.96 | 45.10 | 98.50 | 85.88 |
| DAGNN | 99.97 | 99.98 | 37.36 | 100.00 | 84.33 | 99.96 | 99.89 | 37.61 | 98.16 | 83.91 |
| D-VAE | 99.96 | 100.00 | 37.26 | 100.00 | 84.31 | 99.94 | 98.84 | 38.98 | 98.01 | 83.94 |
| S-VAE | 99.98 | 100.00 | 37.03 | 99.99 | 84.25 | 99.99 | 100.00 | 35.51 | 99.70 | 83.80 |
| GraphRNN | 99.85 | 99.84 | 29.77 | 100.00 | 82.37 | 96.71 | 100.00 | 27.30 | 98.57 | 80.65 |
| GCN | 5.42 | 99.37 | 41.48 | 100.00 | 61.57 | 99.07 | 99.89 | 30.53 | 98.26 | 81.94 |

Models parameterized with neural networks contribute to the inductive biases of the deep generative models (Zhang et al., 2016; Keskar et al., 2017). Thus, the quality of DAG encoders can also be characterized by the reconstruction accuracy (Accuracy) and the generation performance. (i.e. the proportions of valid/ unique/ novel architectures in generated DAGs)

The reconstruction accuracy, prior validity, uniqueness and novelty are calculated in the same way as (Zhang et al., 2019). Empirical results are presented in Table 7, and we take the average of these four metrics to characterize the overall performance of the deep generative model (i.e. VAE), which also measures the quality of the DAG encoder. We find that PACE performs similarly well in reconstruction accuracy, prior validity and novelty with D-VAE, DAGNN and S-VAE, while significantly improving the uniqueness. Hence, PACE achieves the best overall performance and generates more diverse DAG architectures.

H. Extension of Theorem 3.1

The expressive power of message passing Graph Neural Networks (GNNs) is upper bounded by Weisfeiler-Lehman (1-WL) isomorphism test. To address this limitation, subgraph-based techniques (Zhang & Li, 2021; You et al., 2021) are applied to general message passing GNNs to extend their expressiveness beyond the 1-WL test by learning the representation of each node based on the rooted subgraph around it. Thanks to the plug-and-play framework of these subgraph-based techniques, they share advantages of simplicity and efficiency in various graph learning tasks.

Following this idea, the proposed dag2seq framework can also be extended to a subgraph-based version. Given a DAG $G = (V, E)$ and a node i , we use G_i^h to denote the height- h rooted subgraph of node i , which is the subgraph induced from G by the nodes within h hops of node i . Let f be a function that maps subgraphs to vectors, then we have the following corollary:

Corollary H.1. *Let $G = (V, E)$ be a DAG, and p_1, p_2, \dots, p_n be the positional encodings generated by dag2seq. If functions f , Agg and Combine are injective, then the sequence $(f(G_{\pi^{-1}(1)}^h), p_{\pi^{-1}(1)}), (f(G_{\pi^{-1}(2)}^h), p_{\pi^{-1}(2)}), \dots, (f(G_{\pi^{-1}(n)}^h), p_{\pi^{-1}(n)})$ can injectively represent DAG G .*

As the function f is injective, we can prove corollary H.1 by simply replacing $o(\pi^{-1}(i))$ in the proof A by $f(G_{\pi^{-1}(i)}^h)$ for $\forall i$. Corollary H.1 shows that the proposed dag2seq framework can be equipped with advancements in GNN literature. When G has a tree structure (e.g. OGBG-CODE), dag2seq can simply take message passing GNN as the function f to injectively represent DAGs as sequences. On the other hand, for general DAGs, the corollary indicates that the more expressive GNNs, such as IDGNN (You et al., 2021) and NGNN (Zhang & Li, 2021), will generate sequences with higher quality.