

# GRASP: Simple yet Effective Graph Similarity Predictions

Haoran Zheng<sup>1\*</sup>, Jieming Shi<sup>2†</sup>, and Renchi Yang<sup>1</sup>

<sup>1</sup>Hong Kong Baptist University, Hong Kong, China

<sup>2</sup>The Hong Kong Polytechnic University, Hong Kong, China

cshrzeng@comp.hkbu.edu.hk, jieming.shi@polyu.edu.hk, renchi@hkbu.edu.hk

## Abstract

Graph similarity computation (GSC) is to calculate the similarity between one pair of graphs, which is a fundamental problem with fruitful applications in the graph community. In GSC, graph edit distance (GED) and maximum common subgraph (MCS) are two important similarity metrics, both of which are NP-hard to compute. Instead of calculating the exact values, recent solutions resort to leveraging graph neural networks (GNNs) to learn data-driven models for the estimation of GED and MCS. Most of them are built on components involving node-level interactions crossing graphs, which engender vast computation overhead but are of little avail in effectiveness. In the paper, we present GRASP, a simple yet effective GSC approach for GED and MCS prediction. GRASP achieves high result efficacy through several key instruments: enhanced node features via positional encoding and a GNN model augmented by a gating mechanism, residual connections, as well as multi-scale pooling. Theoretically, GRASP can surpass the 1-WL test, indicating its high expressiveness. Empirically, extensive experiments comparing GRASP against 10 competitors on multiple widely adopted benchmark datasets showcase the superiority of GRASP over prior arts in terms of both effectiveness and efficiency.

**Code** — <https://github.com/HaoranZ99/GraSP>

**Extended version** — <https://arxiv.org/abs/2412.09968>

## Introduction

Graphs are an omnipresent data structure to model complex relationships between objects in nature and society, such as social networks, transportation networks, and biological networks. Graph similarity calculation (hereinafter GSC) is a fundamental operation for graphs and finds widespread use in applications, e.g., ranking documents (Lee, Jin, and Jain 2008), disease prediction (Borgwardt and Kriegel 2007), and code similarity detection (Li et al. 2019). In GSC, *graph edit distance* (GED) (Bunke and Allermann 1983) and *maximum common subgraph* (MCS) (Bunke and Shearer 1998) are two popular similarity metrics. However, the exact computations of GED and MCS are both NP-hard (Zeng et al. 2009; Liu

et al. 2020). As revealed in (Blumenthal and Gamper 2020), calculating the exact GED between graphs containing over just 16 nodes is infeasible.

Early attempts have been made to trade accuracy for efficiency via combinatorial approximation (Neuhaus, Riesen, and Bunke 2006; Riesen and Bunke 2009; Fankhauser, Riesen, and Bunke 2011). However, such methods still incur sub-exponential or cubic asymptotic complexities, rendering them impractical for sizable graphs or a large collection of graphs. In recent years, inspired by the powerful capabilities of graph neural networks (GNNs) in capturing graph structures, a series of GNN-based data-driven models (Bai et al. 2019; Li et al. 2019; Bai et al. 2020; Ling et al. 2021; Tan et al. 2023; Piao et al. 2023) have been developed for GSC prediction. At a high level, most of them usually encompass three components, including the node/graph embedding module, the cross-graph node-level interaction module, as well as the similarity computation module. Amid them, modeling the node-level interactions that cross graphs leads to a quadratic time complexity in offline training or online inference, significantly impeding the efficiency and scalability of such models. Efforts have been made to alleviate this (Qin et al. 2021; Zhuo and Tan 2022; Ranjan et al. 2022), but there is still significant space for improvement.

In this paper, to achieve superior efficacy, we present GRASP (short for GRAph Similarity Prediction), a simple but effective approach for GED/MCS prediction. Notably, GRASP can achieve efficient training and inference while offering superb performance in GED/MCS predictions. The design recipes for GRASP include several major ingredients. GRASP first utilizes positional encoding to inject global topological patterns in graphs for enhanced node features. The augmented node features are then fed into a carefully crafted graph embedding architecture, which comprises an optimized GNN backbone for learning node representations and additionally, a multi-scale pooling trick that capitalizes on the merits of different pooling schemes to transform the node representations into high-caliber graph-level embeddings. As such, the GED/MCS prediction can be directly solved with such graph embeddings, without acquiring all pair-wise node interactions for node embeddings or cross-level similarities required in previous methods. Furthermore, compared to conventional GNN models (Kipf and Welling 2017; Xu et al. 2019) with expressive power bounded by the

\*Work partially done at the Hong Kong Polytechnic University.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

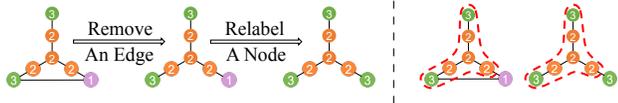


Figure 1: GED and MCS examples from AIDS700nef dataset. Left: GED is 2 and right: MCS is 6.

1-order Weisfeiler-Lehman test (1-WL test) (Weisfeiler and Leman 1968), our analysis manifests that GRASP obtains higher expressiveness by surpassing the 1-WL test, theoretically corroborating its practical effectiveness in accurate GSC. We experimentally evaluate GRASP against 10 competitors over 4 real datasets in GED and MCS prediction tasks under various settings. The empirical results exhibit that GRASP can consistently achieve superior GED/MCS estimation performance over all the baselines while retaining high efficiency.

In summary, our contributions are as follows:

- We propose a new approach GRASP for GSC, which offers remarkable performance by virtue of several designs, including techniques for feature enhancement via positional encoding and multi-scale pooling over node representations.
- We theoretically analyze and reveal that GRASP can generate expressive representations that are beyond the 1-WL graph isomorphism test while achieving efficient time complexity with respect to the size of a graph pair.
- The superior performance of GRASP, in terms of efficiency and effectiveness, is validated against various competitors over multiple benchmark datasets.

## Preliminaries

A graph  $\mathcal{G}$  consists of a set of nodes  $\mathcal{V}$  and a set of edges  $\mathcal{E}$ , i.e.,  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . The number of nodes and edges are  $|\mathcal{V}|$  and  $|\mathcal{E}|$ , respectively. Given a graph database  $\mathcal{D}$  containing a collection of graphs, we aim to build an end-to-end model that accurately estimates the similarity values of graph pairs. The model is designed to predict multiple similarity/distance metrics, including Graph Edit Distance (GED) and Maximum Common Subgraph (MCS).

**Definition 1** (Graph Edit Distance (Bunke and Allermann 1983)). *GED is the edit distance between two graphs  $\mathcal{G}_1$  and  $\mathcal{G}_2$ , i.e., the minimum number of edit operations to convert  $\mathcal{G}_1$  to  $\mathcal{G}_2$ , where edit operations include adding/removing a node, adding/removing an edge, and relabeling a node.*

**Definition 2** (Maximum Common Subgraph (Bunke and Shearer 1998)). *MCS is an isomorphic graph to the largest subgraph shared by two graphs  $\mathcal{G}_1$  and  $\mathcal{G}_2$ .*

Following the choice of (Bai et al. 2020), the MCS is connected. Figure 1 shows GED and MCS examples. Let  $m$  be the number of labels, and we can define the node feature matrix  $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times m}$  of a graph  $\mathcal{G}$ .

During the training phase, a training sample consisting of a set of graph pair  $(\mathcal{G}_1, \mathcal{G}_2) \in \mathcal{D} \times \mathcal{D}$  with ground-truth similarity value  $s(\mathcal{G}_1, \mathcal{G}_2)$ . All training samples are used to train a prediction model. In the inference phase, the model predicts the similarity/distance of unseen graph pairs.

## Our Approach: GRASP

The architecture of GRASP is illustrated in Figure 2. In the node feature pre-processing, we first enhance node features by concatenating positional encoding of the node. This enhancement considers global graph topological features. Second, the node embedding module preserves the enhanced node features via an RGGC backbone into node embeddings. Third, within the graph embedding module, we devise a multi-scale node embedding pooling technique to obtain graph embeddings. Lastly, the prediction module predicts the similarity/distance between two graph embeddings.

### Enhanced Node Features via Positional Encoding

Previous works like (Bai et al. 2019, 2020; Ranjan et al. 2022) obtain features using the one-hot encoding of node labels. Specifically, every node  $i$  has  $\mathbf{x}_i \in \mathbb{R}^m$  to represent its label, where  $m$  is the number of labels in the graph database. Then  $\mathbf{x}_i$  is transformed by an MLP to  $\boldsymbol{\mu}_i \in \mathbb{R}^d$  (Ranjan et al. 2022), which is then used as the input of GNN layers. However, only using node labels as node features to be the input of message-passing graph neural networks (MPGNNs) cannot pass the 1-WL graph isomorphism test. Specifically, MPGNNs update node representations based on message passing and by aggregating neighbor information, a node  $i$  in a graph can be represented as:

$$\mathbf{h}_i^{(\ell)} = f\left(\mathbf{h}_i^{(\ell-1)}, \left\{\mathbf{h}_j^{(\ell-1)} \mid j \in \mathcal{N}(i)\right\}\right), \quad (1)$$

where  $\ell$  represents  $\ell$ -th layer of the stack of the MPGNN layers,  $\mathcal{N}(i)$  denotes the neighbors of node  $i$ .

It has been shown by (Xu et al. 2019) that MPGNNs can perform up to, but not beyond, the level of the 1-WL test. In general, the 1-WL test can effectively distinguish non-isomorphic graphs. However, since MPGNNs only aggregate information from their direct neighbors, they fail to capture higher-order structural properties and the global topology of the graph, thus limiting the expressive capability of MPGNNs. We take one case that MPGNNs fail to distinguish (Sato 2020), as shown in Figure 3. MPGNNs produce the same set of node embeddings for these two non-isomorphic graphs. Thus, the same set of node embeddings causes a model to give incorrect predictions that the GED is 0, which apparently should not be 0.

To mitigate this issue and enhance the expressiveness of our method, we propose to enhance node features with positional encoding that can exceed the expressive capacity of the 1-WL test. Specifically, we use the random walk method for position encoding, RWPE, which has been empirically proven to work by (Dwivedi et al. 2022). Unlike the distance encoding method proposed by (Li et al. 2020) that uses random walks to learn the distance of all pairs of nodes in the graph, RWPE uses only the probability of a node landing on itself, i.e., the diagonal part of the random walk matrix. Given the random walk length  $k$ , we can use RWPE to pre-compute the positional features  $\mathbf{p}_{i_{\text{init}}} \in \mathbb{R}^k$  of a node  $i$  in the graph, denoted as  $\mathbf{p}_{i_{\text{init}}} = \left\{\mathbf{RW}_{ii}^{\ell} \mid \ell = 1, 2, \dots, k\right\}$ , where  $\mathbf{RW} = \mathbf{AD}^{-1}$  is the random walk matrix obtained by adjacency matrix  $\mathbf{A}$  and diagonal degree matrix  $\mathbf{D}$  of a graph. Then we transform the RWPE  $\mathbf{p}_{i_{\text{init}}}$  into  $\mathbf{p}_i \in \mathbb{R}^d$  by an MLP.

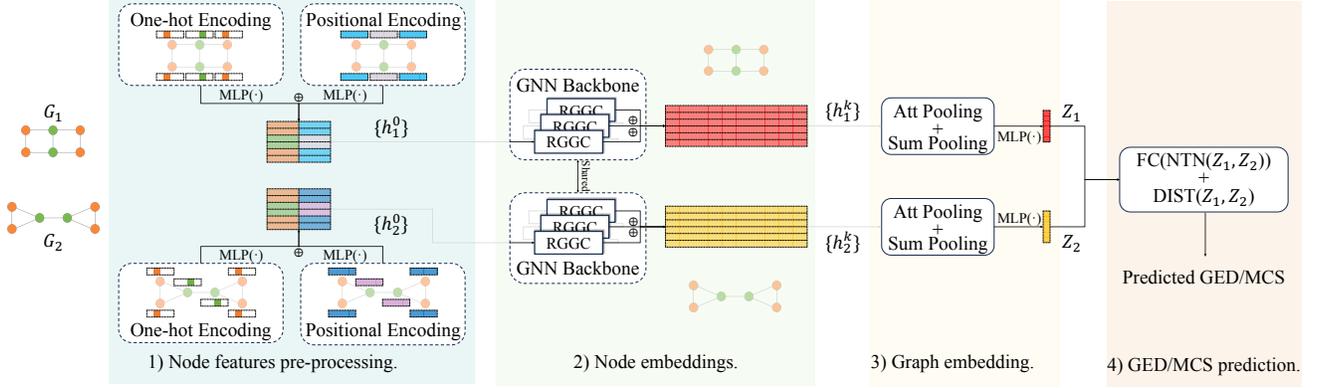


Figure 2: The architecture of GRASP.

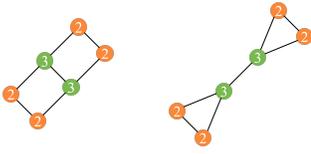


Figure 3: An example that MPGNNs fail to distinguish.

We concatenate the transformed positional encoding  $\mathbf{p}_i$  and the encoding  $\mu_i$  of node  $i$  to get its enhanced representation  $\mathbf{h}_i^{(0)} \in \mathbb{R}^{2d}$ , to be the input of GNN backbone,

$$\mathbf{h}_i^{(0)} = \text{CONCAT}(\mu_i, \mathbf{p}_i). \quad (2)$$

### Multi-Scale Pooling on RGGC

The enhanced node features  $\mathbf{h}_i^{(0)}$  obtained in previous subsection are then fed into a graph neural network consisting of  $n$  layers of ResGatedGraphConv (RGGC layers) (Bresson and Laurent 2017) to learn the hidden representations of the nodes. The RGGC layers can leverage a gating mechanism where gating units learn to control information flow through the network. Moreover, residual connections are used in RGGC layers to help with gradient flow during training. With these two techniques incorporated, the RGGC layers can learn complex patterns in graph data, and therefore, are more powerful and versatile than basic graph convolutional layers like GCNs and GINs. At the  $\ell$ -th layer, the node representation  $\mathbf{h}_i^{(\ell)}$  is

$$\mathbf{h}_i^{(\ell)} = \mathbf{h}_i^{(\ell-1)} + \text{ReLU}(\mathbf{W}_S \mathbf{h}_i^{(\ell-1)} + \sum_{j \in \mathcal{N}(i)} \eta_{i,j} \odot \mathbf{W}_N \mathbf{h}_j^{(\ell-1)}), \quad (3)$$

where  $\mathbf{W}_S \in \mathbb{R}^{2d \times 2d}$  and  $\mathbf{W}_N \in \mathbb{R}^{2d \times 2d}$  are learnable weight matrices,  $\odot$  is the Hadamard point-wise multiplication operator and the gate  $\eta_{i,j}$  is defined as  $\eta_{i,j} = \sigma(\mathbf{W}_{GS} \mathbf{h}_i^{(\ell-1)} + \mathbf{W}_{GN} \mathbf{h}_j^{(\ell-1)})$ , where  $\mathbf{W}_{GS} \in \mathbb{R}^{2d \times 2d}$ ,  $\mathbf{W}_{GN} \in \mathbb{R}^{2d \times 2d}$  are learnable weight matrices and with  $\sigma$

as an activation function. In (Bresson and Laurent 2017), the sigmoid function is chosen as the activation function so that the gate  $\eta_{i,j}$  can learn the weight controlling how important the information from node  $j$  to node  $i$ .

We concatenate the node representations obtained in all layers to preserve the information of multi-hop neighbors better. The concatenated representation of node  $i$  after  $n$  layers is  $\mathbf{h}_i \in \mathbb{R}^{2(n+1)d}$ ,

$$\mathbf{h}_i = \text{CONCAT} \left( \left\{ \mathbf{h}_i^{(\ell)} \mid \ell = 0, 1, \dots, n \right\} \right). \quad (4)$$

Note that our task is to estimate the similarity between two graphs. Therefore, we need to generate graph-level representations based on the node representations above. We design a multi-scale pooling technique that considers both attention pooling and summation pooling. Attention pooling (Bai et al. 2019) assigns weights to each node according to its importance, and then pools the node embeddings using a weighted sum based on the attention weights. Denote the resulting graph embedding as  $\mathbf{z}_{\text{att}} \in \mathbb{R}^{2nd}$ . We get

$$\mathbf{z}_{\text{att}} = \sum_{i=1}^{|\mathcal{V}|} \sigma(\mathbf{h}_i^T \tanh(\frac{1}{|\mathcal{V}|} \mathbf{W}_A \sum_{j=1}^{|\mathcal{V}|} \mathbf{h}_j)) \mathbf{h}_i,$$

where  $\mathbf{W}_A \in \mathbb{R}^{[2(n+1)d] \times [2(n+1)d]}$  is learnable and  $\sigma$  is a sigmoid function. Attention pooling employs an attention mechanism and therefore one merit is that it is globally context-aware. Summation pooling  $\mathbf{z}_{\text{sum}} \in \mathbb{R}^{2(n+1)d}$  simply sums the node embeddings, i.e.,  $\mathbf{z}_{\text{sum}} = \sum_{i=1}^{|\mathcal{V}|} \mathbf{h}_i$ . Summation pooling has the advantage of simplicity and thus is more computationally efficient.

We observe that both of the above pooling methods have some drawbacks: summation pooling treats all nodes equally, which may not be optimal; and attention in attention pooling runs the risk of overfitting on specific nodes. Therefore, in our multi-scale pooling, we mix these two pooling methods and let the model learn to trade off the two pooling operations, thus reducing the drawbacks of the two pooling methods. We get the combined graph embedding  $\mathbf{z} \in \mathbb{R}^{2(n+1)d}$ :

$$\mathbf{z}_{\text{combined}} = \mathbf{a} \mathbf{z}_{\text{att}} + (1 - \mathbf{a}) \mathbf{z}_{\text{sum}}, \quad (5)$$

where  $\mathbf{a} \in \mathbb{R}^{2(n+1)d}$  is a vector that can be learned.

Similar to the node feature pre-processing, we pass the graph embedding that has gone through the pooling layer via an MLP to adjust its dimension for subsequent processing, and finally, we get the graph embedding  $\mathbf{z} \in \mathbb{R}^d$ .

### GED and MCS Prediction Objectives

After obtaining the graph embeddings  $\mathbf{z}_1$  and  $\mathbf{z}_2$  of two graphs  $\mathcal{G}_1$  and  $\mathcal{G}_2$ , we explain how to obtain predicted GED values and the design of training objective. The way to get MCS estimation and training objectives naturally follows.

To get a predicted GED, an intuitive idea is to compute the Euclidean distance between  $\mathbf{z}_1$  and  $\mathbf{z}_2$ ,

$$\text{dist}(\mathbf{z}_1, \mathbf{z}_2) = \|\mathbf{z}_1 - \mathbf{z}_2\|_2. \quad (6)$$

Moreover, inspired by (Zhuo and Tan 2022), we introduce the Neural Tensor Network (NTN) (Socher et al. 2013) as a multi-headed weighted cosine similarity function to compute the interaction value of the two graph embeddings in the capacity of a bias value as a complement to Eq. (6). NTN is a powerful method for quantifying relationships between representations. We denote the interaction value of embeddings  $\mathbf{z}_1$  and  $\mathbf{z}_2$  as:

$$\text{interact}(\mathbf{z}_1, \mathbf{z}_2) = \text{MLP}(\text{ReLU}(\mathbf{z}_1^T \mathbf{W}_I^{[1:t]} \mathbf{z}_2 + \mathbf{W}_C \text{CONCAT}(\mathbf{z}_1, \mathbf{z}_2) + \mathbf{b})), \quad (7)$$

where  $\mathbf{W}_I^{[1:t]} \in \mathbb{R}^{d \times d \times t}$  is a learnable weight tensor,  $\mathbf{W}_C$  is a learnable weight matrix,  $\mathbf{b} \in \mathbb{R}^t$  is a bias vector,  $t$  is the hyperparameter controlling the NTN output and  $\text{MLP}(\cdot)$  is a fully connected neural network that maps the similarity from  $\mathbb{R}^t$  to  $\mathbb{R}$ .

Finally, our predicted GED value is

$$\text{GED}(\mathcal{G}_1, \mathcal{G}_2) = \beta \cdot \text{dist}(\mathbf{z}_1, \mathbf{z}_2) + (1 - \beta) \cdot \text{interact}(\mathbf{z}_1, \mathbf{z}_2), \quad (8)$$

where  $\beta$  is a scalar that can be learned.

Then we adopt the mean squared error between our predicted GED value  $\text{GED}(\mathcal{G}_1, \mathcal{G}_2)$  and ground-truth GED value  $\text{GED}^*(\mathcal{G}_1, \mathcal{G}_2)$ , and have the loss function:

$$\mathcal{L} = \frac{1}{T} \sum_{(\mathcal{G}_1, \mathcal{G}_2) \in \mathcal{D} \times \mathcal{D}} \text{MSE}(\text{GED}(\mathcal{G}_1, \mathcal{G}_2), \text{GED}^*(\mathcal{G}_1, \mathcal{G}_2)), \quad (9)$$

where  $T$  is the number of training graph pairs in a graph database  $\mathcal{D}$  and value  $\text{GED}^*(\mathcal{G}_1, \mathcal{G}_2)$  is the ground-truth GED between graph  $\mathcal{G}_1$  and  $\mathcal{G}_2$ .

Eq. (8) can be generalized to other similarity metrics like MCS. Recall that MCS is a similarity measuring the largest common subgraph of two graphs. Therefore, we can consider the output of interaction, i.e., Eq. (7) as the similarity of the two graph embeddings and further consider Eq. (6) as a bias value. Therefore, we keep the right-hand side of Eq. (8) unchanged and change the left-hand side to  $\text{MCS}(\mathcal{G}_1, \mathcal{G}_2)$ . The loss function in Eq. (9) follows for MCS.

### Analysis of GRASP

We first prove that GRASP achieves high expressiveness and can pass the 1-WL test, and then analyze the complexity of GRASP that is linear to the number of nodes in a graph pair.

We use the following Proposition 1 to formally state that our method can outperform the 1-WL test. The proof is provided in the extended version. We utilize the proof in (Xu et al. 2019) that the representation ability of the 1-WL test is equivalent to that of standard MPGNNs in the graph isomorphism problem.

**Proposition 1.** *Given a pair of non-isomorphic graphs  $\mathcal{G}_1$  and  $\mathcal{G}_2$  that cannot be discriminated by the 1-WL test, and the two graphs have different sets of initial position encodings, then GRASP can generate different graph representations for the two non-isomorphic graphs.*

The preconditions of Proposition 1 include that two graphs should have different initial position encodings. According to the definition of RWPE, nodes on non-isomorphic graphs generally get different sets of RWPEs when  $k$  is sufficiently large (Dwivedi et al. 2022). Thus, RWPE satisfies this precondition that the sets of positional encodings are different. Therefore, graph embeddings that are more discriminative than the 1-WL test are obtained. These powerful graph embeddings can benefit our graph similarity prediction task when predicting GED and MCS objectives.

**Complexity Analysis.** The inference time complexity of GRASP is linear to the number of nodes of the graph pair. Our node feature preprocessing module requires downscaling the dimensionality of the node features from  $\mathbb{R}^m$  to  $\mathbb{R}^d$ , resulting in a time complexity of  $O(md|\mathcal{V}|)$ . The positional encoding module contains a random walk positional encoding pre-computation and an MLP. The pre-computation of random walk positional encoding takes  $O(k|\mathcal{V}|^2)$  due to the sparse matrix multiplication. This pre-computation is performed only once when the number of steps  $k$  for the random walker is determined. Since it can be reused across both training and inference stages, it is regarded as a preprocessing technique and excluded from the complexity analysis. The MLP( $\cdot$ ) takes  $O(kd|\mathcal{V}|)$  time. The node embedding module contains  $n$  layers of RGGC with a time complexity of  $O(n|\mathcal{E}|)$ . The multi-scale pooling module contains attention pooling and summation pooling, both with time complexity of  $O(nd|\mathcal{V}|)$ . In the phase of generating the final graph embedding, we downscale the dimensionality of the graph embedding from  $\mathbb{R}^{2nd}$  to  $\mathbb{R}^d$  with a time complexity of  $O(nd^2)$ . In the similarity prediction module, the time complexity of NTN interaction is  $O(d^2t)$ , where  $t$  is the dimension of NTN output, and the time complexity of Euclidean distance calculation is  $O(d)$  and the total time is  $O(d^2t)$ . Hence, the time complexity of GRASP for predictions is  $O(d|\mathcal{V}|(m+k+n) + nd^2 + d^2t)$ , which is linear to the number of nodes of the graph pair.

## Experiments

We evaluate GRASP against competitors on GED and MCS prediction tasks on well-adopted benchmarking datasets.

### Experiment Setup

**Data.** We conduct experiments on four real-world datasets, including AIDS700nef, LINUX, IMDBMulti (Bai et al. 2019) and PTC (Bai et al. 2020). The statistics and descriptions of the four datasets can be found in the extended ver-

	AIDS700nef					IMDBMulti				
	MAE ↓	$\rho$ ↑	$\tau$ ↑	P@10 ↑	P@20 ↑	MAE ↓	$\rho$ ↑	$\tau$ ↑	P@10 ↑	P@20 ↑
SIMGNN	0.852±0.033	0.838±0.005	0.659±0.006	0.489±0.027	0.600±0.018	5.836±0.439	0.924±0.003	0.811±0.006	0.802±0.033	0.819±0.021
GRAPHSIM	0.979±0.019	0.808±0.012	0.654±0.008	0.382±0.021	0.502±0.013	7.907±0.670	0.738±0.012	0.621±0.009	0.578±0.028	0.588±0.023
GMN	0.881±0.010	0.831±0.003	0.650±0.004	0.461±0.010	0.573±0.006	5.101±0.860	0.925±0.006	0.814±0.008	0.830±0.018	0.855±0.010
MGMN	0.805±0.006	0.863±0.001	0.727±0.002	0.553±0.010	0.637±0.004	12.876±0.728	0.771±0.029	0.669±0.023	0.307±0.033	0.353±0.024
H2MN	0.731±0.014	0.875±0.002	0.740±0.003	0.552±0.008	0.652±0.010	9.735±1.542	0.909±0.010	0.809±0.012	0.746±0.057	0.783±0.032
EGSC	0.667±0.010	0.889±0.001	0.720±0.002	0.662±0.008	0.735±0.004	4.397±0.101	0.940±0.001	0.847±0.002	0.861±0.009	0.874±0.007
ERIC	0.765±0.016	0.876±0.003	0.746±0.005	0.592±0.015	0.681±0.015	5.487±0.929	0.934±0.003	<u>0.862±0.006</u>	0.837±0.014	0.836±0.012
GREED	0.731±0.013	0.886±0.005	0.758±0.006	0.617±0.013	0.709±0.015	4.418±0.285	0.930±0.003	0.857±0.004	0.837±0.010	0.850±0.003
NA-GSL	0.765±0.014	0.881±0.001	0.753±0.002	0.562±0.023	0.665±0.010	7.588±0.975	0.900±0.009	0.814±0.009	0.544±0.049	0.611±0.033
GEDGNN	0.724±0.012	0.885±0.002	0.757±0.003	0.604±0.011	0.685±0.012	8.585±1.494	0.902±0.006	0.834±0.014	0.748±0.008	0.771±0.026
GRASP	<b>0.640±0.013</b>	<b>0.917±0.002</b>	<b>0.804±0.003</b>	<b>0.741±0.009</b>	<b>0.800±0.002</b>	<b>3.966±0.064</b>	<b>0.942±0.001</b>	<b>0.874±0.003</b>	<b>0.863±0.011</b>	<b>0.876±0.014</b>

	LINUX					PTC				
	MAE ↓	$\rho$ ↑	$\tau$ ↑	P@10 ↑	P@20 ↑	MAE ↓	$\rho$ ↑	$\tau$ ↑	P@10 ↑	P@20 ↑
SIMGNN	0.269±0.027	0.939±0.002	0.792±0.004	0.947±0.012	0.959±0.016	4.200±0.395	0.930±0.006	0.821±0.008	0.465±0.028	0.628±0.031
GRAPHSIM	0.212±0.033	0.962±0.003	0.889±0.003	0.970±0.010	0.982±0.007	5.003±0.491	0.850±0.009	0.721±0.011	0.408±0.008	0.521±0.011
GMN	0.293±0.047	0.937±0.001	0.788±0.002	0.906±0.022	0.932±0.012	4.353±1.051	0.941±0.005	0.836±0.008	0.558±0.041	0.675±0.028
MGMN	0.368±0.035	0.958±0.003	0.883±0.005	0.907±0.042	0.931±0.017	3.839±0.374	0.920±0.004	0.802±0.007	0.503±0.027	0.629±0.017
H2MN	0.909±0.523	0.961±0.004	0.874±0.006	0.951±0.016	0.953±0.016	5.391±2.151	0.918±0.012	0.809±0.017	0.435±0.047	0.562±0.035
EGSC	0.083±0.019	0.948±0.003	0.811±0.003	0.981±0.009	0.988±0.009	3.743±0.203	0.938±0.002	0.835±0.003	0.577±0.012	0.695±0.008
ERIC	<u>0.080±0.011</u>	0.971±0.002	0.906±0.003	0.977±0.005	0.983±0.006	3.908±0.099	0.940±0.002	<u>0.842±0.002</u>	0.555±0.010	0.670±0.011
GREED	0.342±0.006	0.963±0.001	0.884±0.001	0.964±0.004	0.970±0.003	3.934±0.132	0.911±0.004	0.797±0.005	0.432±0.015	0.545±0.017
NA-GSL	0.132±0.023	0.973±0.001	0.914±0.001	<b>0.987±0.005</b>	0.989±0.003			OOM		
GEDGNN	0.108±0.015	0.971±0.003	0.906±0.003	0.966±0.011	0.975±0.008	4.164±0.274	0.935±0.002	0.839±0.002	0.543±0.025	0.659±0.008
GRASP	<b>0.042±0.003</b>	<b>0.975±0.001</b>	<b>0.922±0.001</b>	<u>0.983±0.003</u>	<b>0.992±0.002</b>	<b>3.556±0.080</b>	<b>0.952±0.002</b>	<b>0.861±0.003</b>	<b>0.612±0.019</b>	<b>0.707±0.010</b>

Table 1: Effectiveness results on GED predictions with standard deviation. **Bold**: best, Underline: runner-up.

sion. We split training, validation, and testing data with a ratio of 6:2:2 for all datasets and all methods by following the setting in (Bai et al. 2019) and (Bai et al. 2020). For small datasets AIDS and LINUX, we use A\* to calculate ground-truth GEDs. For IMDBMulti and PTC, we follow the way in (Bai et al. 2019) and use the minimum of the results of three approximation methods Beam (Neuhaus, Riesen, and Bunke 2006), Hungarian (Riesen and Bunke 2009) and VJ (Fankhauser, Riesen, and Bunke 2011) to be ground-truth GED. We use the MCSPLIT (McCreech, Prosser, and Trimble 2017) algorithm to calculate the ground-truth MCS.

**Baseline Methods.** We compare with ERIC (Zhuo and Tan 2022), GREED (Ranjan et al. 2022), NA-GSL (Tan et al. 2023), GEDGNN (Piao et al. 2023), H2MN (Zhang et al. 2021), EGSC (Qin et al. 2021), SIMGNN (Bai et al. 2019), GMN (Li et al. 2019), GRAPHSIM (Bai et al. 2020), and MGMN (Ling et al. 2021). We use their official code provided by the authors and because some methods are designed to predict graph similarity scores, i.e., exponentially normalized GEDs instead of original GED, we modify the regression head part of their models to align with experiment settings. Moreover, since some methods are only implemented for GED, we extend their official code for the MCS objective. GENN-A\* (Wang et al. 2021) cannot scale to large data, IMDBMulti and PTC, and thus it is omitted.

**Hyperparameters.** In our method, we use a search range of  $\{w/o, 8, 16, 24, 32\}$  for the step size  $k$  of the RWPE,  $\{4, 6, 8, 10, 12\}$  for the number of layers  $\ell$  of the GNN backbone, and  $\{16, 32, 64, 128, 256\}$  for the dimensionality  $d$  of the node hidden representations and also the final graph embedding. Our full hyperparameter settings on the four datasets can be found in the extended version. For competitors, we

conduct experiments according to their hyperparameter settings reported by their works.

**Evaluation Metrics.** We compare the performance of all methods by Mean Absolute Error (MAE) between the predicted and ground-truth GED and MCS scores, Spearman’s Rank Correlation Coefficient ( $\rho$ ) (Spearman 1987), Kendall’s Rank Correlation Coefficient ( $\tau$ ) (Kendall 1938), and Precision at 10 and 20 (P@10 and 20). The lower MAE proves that the model performs better; the higher the latter four, the better the model performs. Note that all these metrics are evaluated over ground-truth GED and MCS scores, instead of normalized graph similarity scores.

All experiments are conducted on a Linux machine with a CPU Intel(R) Xeon(R) Gold 6226R CPU 2.90GHz, and GPU model NVIDIA GeForce RTX 3090.

## Effectiveness

Tables 1 and 2 report the overall effectiveness of all methods on all datasets for GED and MCS predictions, respectively. OOM means out of GPU memory. In Table 1, for GED predictions, our method GRASP outperforms all methods by all metrics on AIDS700nef, IMDBMulti and PTC datasets, and by 4 out of 5 metrics on LINUX dataset. For example, On AIDS700nef, our method GRASP achieves high P@10 0.741, significantly outperforming the best competitor performance 0.662 of EGSC by 11.9% relative improvement. On LINUX, our method GRASP significantly reduces MAE to 0.042, compared with the best competitor ERIC with 0.080 MAE. In Table 2 for MCS predictions, our method GRASP outperforms all methods by all metrics on all four datasets. For example, On PTC, GRASP can achieve 0.762 for  $\tau$  metric, which is much higher than the runner-

	AIDS700nef					IMDBMulti				
	MAE ↓	$\rho$ ↑	$\tau$ ↑	P@10 ↑	P@20 ↑	MAE ↓	$\rho$ ↑	$\tau$ ↑	P@10 ↑	P@20 ↑
SIMGNN	0.572±0.019	0.796±0.004	0.615±0.004	0.610±0.014	0.662±0.014	0.284±0.086	0.783±0.005	0.618±0.006	0.889±0.014	0.918±0.016
GRAPHSIM	0.442±0.014	0.802±0.009	0.679±0.006	0.529±0.019	0.596±0.018	0.786±0.077	0.764±0.005	0.671±0.008	0.787±0.019	0.791±0.023
GMN	0.747±0.023	0.701±0.002	0.552±0.001	0.421±0.013	0.489±0.005	0.274±0.038	0.782±0.003	0.617±0.003	0.910±0.012	0.917±0.011
MGMN	0.510±0.011	0.841±0.007	0.723±0.007	0.658±0.018	0.686±0.010	0.758±0.068	0.770±0.002	0.686±0.003	0.884±0.025	0.845±0.023
H2MN	0.537±0.013	0.838±0.006	0.713±0.008	0.578±0.014	0.639±0.015	0.299±0.027	0.872±0.001	0.772±0.005	0.903±0.015	0.912±0.009
EGSC	0.335±0.006	0.874±0.006	0.706±0.004	0.798±0.003	0.826±0.007	0.147±0.012	0.792±0.007	0.636±0.005	0.945±0.005	0.956±0.005
ERIC	0.382±0.006	0.895±0.008	0.788±0.008	0.747±0.004	0.789±0.025	0.183±0.017	0.879±0.002	0.796±0.002	0.919±0.012	0.938±0.008
GREED	0.458±0.026	0.890±0.003	0.781±0.004	0.747±0.015	0.790±0.008	0.251±0.042	0.876±0.002	0.791±0.003	0.897±0.014	0.920±0.011
NA-GSL	0.513±0.010	0.850±0.002	0.731±0.002	0.638±0.012	0.698±0.009	0.251±0.039	0.877±0.002	0.794±0.003	0.902±0.016	0.908±0.015
GEDGNN	0.439±0.019	0.890±0.005	0.780±0.003	0.718±0.009	0.779±0.010	0.348±0.017	0.872±0.003	0.781±0.014	0.883±0.026	0.889±0.022
GRASP	<b>0.330±0.007</b>	<b>0.923±0.002</b>	<b>0.822±0.002</b>	<b>0.848±0.007</b>	<b>0.872±0.005</b>	<b>0.134±0.018</b>	<b>0.881±0.001</b>	<b>0.797±0.002</b>	<b>0.947±0.007</b>	<b>0.957±0.004</b>

	LINUX					PTC				
	MAE ↓	$\rho$ ↑	$\tau$ ↑	P@10 ↑	P@20 ↑	MAE ↓	$\rho$ ↑	$\tau$ ↑	P@10 ↑	P@20 ↑
SIMGNN	0.138±0.016	0.690±0.010	0.518±0.007	0.888±0.064	0.933±0.043	1.391±0.060	0.793±0.010	0.628±0.010	0.492±0.015	0.593±0.015
GRAPHSIM	0.287±0.034	0.788±0.017	0.699±0.004	0.810±0.021	0.823±0.034	1.789±0.071	0.794±0.012	0.648±0.014	0.448±0.025	0.531±0.028
GMN	0.187±0.016	0.678±0.007	0.506±0.007	0.876±0.038	0.904±0.023	1.610±0.075	0.735±0.015	0.571±0.011	0.462±0.014	0.539±0.015
MGMN	0.176±0.027	0.806±0.024	0.710±0.027	0.832±0.104	0.888±0.054	1.441±0.152	0.829±0.020	0.695±0.024	0.530±0.037	0.610±0.027
H2MN	0.196±0.027	0.803±0.002	0.690±0.003	0.901±0.060	0.936±0.037	1.360±0.122	0.816±0.013	0.683±0.013	0.516±0.007	0.617±0.017
EGSC	0.049±0.007	0.700±0.008	0.528±0.005	0.988±0.009	0.991±0.006	1.170±0.127	0.846±0.012	0.690±0.012	0.603±0.016	0.694±0.014
ERIC	0.037±0.009	0.812±0.001	0.715±0.001	0.991±0.005	0.995±0.002	1.406±0.119	0.817±0.007	0.694±0.007	0.570±0.018	0.647±0.021
GREED	0.101±0.021	0.811±0.001	0.711±0.001	0.978±0.007	0.988±0.006	1.306±0.086	0.824±0.008	0.697±0.009	0.578±0.007	0.654±0.012
NA-GSL	0.085±0.021	0.812±0.001	0.715±0.001	0.985±0.006	0.991±0.002			OOM		
GEDGNN	0.089±0.014	0.809±0.005	0.711±0.003	0.943±0.028	0.969±0.011	1.493±0.113	0.811±0.015	0.700±0.017	0.535±0.012	0.630±0.016
GRASP	<b>0.037±0.005</b>	<b>0.813±0.001</b>	<b>0.716±0.001</b>	<b>0.995±0.006</b>	<b>0.996±0.002</b>	<b>1.162±0.028</b>	<b>0.887±0.002</b>	<b>0.762±0.003</b>	<b>0.665±0.005</b>	<b>0.729±0.006</b>

Table 2: Effectiveness results on MCS predictions with standard deviation. **Bold**: best, Underline: runner-up.

	MAE	$\rho$	$\tau$	P@10	P@20
GRASP	<b>0.640</b>	<b>0.917</b>	<b>0.804</b>	<b>0.741</b>	<b>0.800</b>
GRASP (GIN)	0.677	0.908	0.790	0.713	0.774
GRASP (GCN)	0.654	0.911	0.795	0.721	0.783
w/o pe	<u>0.651</u>	<u>0.915</u>	<u>0.800</u>	0.726	0.793
w/o att	0.662	0.912	0.797	0.732	0.792
w/o sum	0.666	0.913	0.797	<u>0.735</u>	<u>0.797</u>
w/o NTN	0.662	0.911	0.793	0.718	0.782

Table 3: Ablation study on AIDS700nef under GED metric.

up GEDGNN with 0.7. Moreover, the P@10 of GRASP is 0.665 while the best competitor EGSC has 0.603.

The overall results in Tables 1 and 2 demonstrate the superior power of GRASP with simple but effective designs, including positional encoding enhanced node features and multi-scale pooling on RGGC backbone. Further, for intuitive results, interested readers may refer to the extended version for the absolute GED error heatmaps produced by GRASP, and recent baselines.

**Ablation study.** We first compare GRASP with RGGC backbone over GRASP with GCN and GIN backbones. We also ablate the positional encoding, the attention pooling and summation pooling in the multi-scale pooling, and the NTN, denoted as w/o pe, w/o att, w/o sum, and w/o NTN respectively. The results on AIDS700nef for GED are reported in Table 3. Observe that GRASP obtains the best performance on all metrics than all its ablated versions, which proves the effectiveness of all our proposed components in GRASP. In particular, with only either attention or summation pooling, the performance is inferior to GRASP with the proposed multi-scale pooling technique that hybrids both pool-

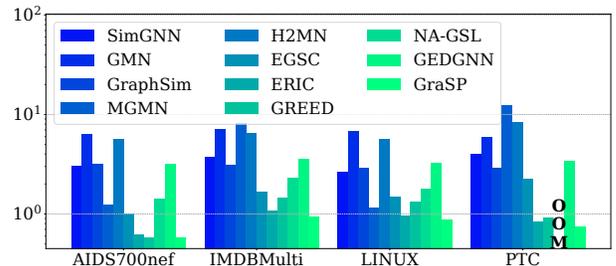


Figure 4: Inference time in second(s) per 10k pairs.

ing techniques, which validates the rationale for designing the technique. Similar results for the other three datasets are presented in the extended version.

## Efficiency

We report the inference time in log scale per 10k graph pairs of all approaches on every dataset in Figure 4. We omit the results when a certain approach is OOM. Our method GRASP is the fastest method on all datasets to complete the inference. The efficiency of GRASP is due to the following designs. First, GRASP does not need the expensive cross-graph node-level interactions that are usually adopted in existing methods. Second, the positional encoding used in GRASP to enhance node features can be precomputed and reused for the efficiency of online inference. Third, all the technical components in GRASP are designed to make the complicated simple and effective for graph similarity predictions, resulting in efficient performance. The inference time

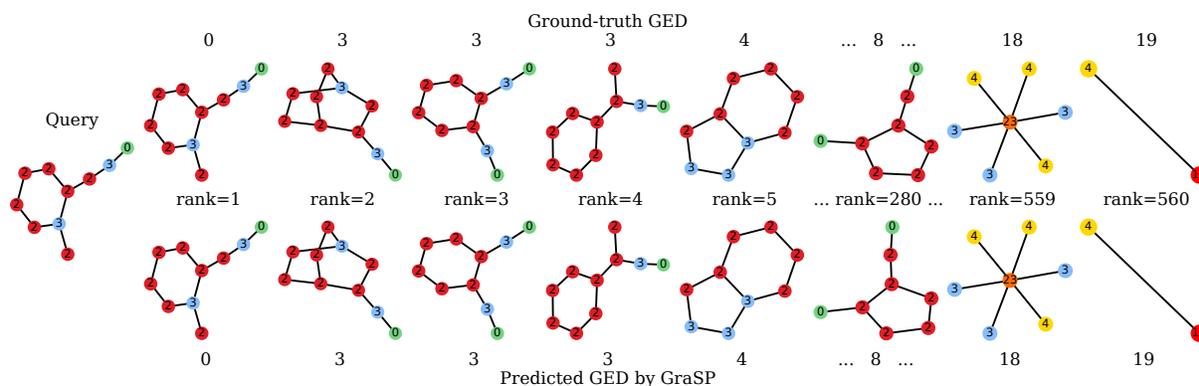


Figure 5: A ranking case study of GED prediction on AIDS700nef.

of SIMGNN, GMN, GRAPHSIM, MGMN, NA-GSL, and GEDGNN is longer, which is due to the expensive cross-graph node-level interactions that these models explicitly perform during inference. EGSC, ERIC, and GREED are relatively faster but do not exceed our method.

### Case Study

We conduct a case study of GRASP over drug discovery, where graph similarity search often has an essential role (Ranu et al. 2011). We consider the ranking of returned graphs w.r.t. a query graph. Figure 5, shows a case of ranking under the GED metric on the AIDS700nef dataset. Given a query graph on the left side of Figure 5, the first row shows the ground-truth ranking of graphs with GED to the query, and the second row shows the graphs ranked by predicted GED from GRASP. As shown, GRASP can accurately predict the GED values and rank the graphs with similar structures to the top, and the ranking from 1 to 560 is almost the same as the ground truth. We provide the case studies on the other three datasets in the extended version.

### Related Work

For GED, recent works (Kim, Choi, and Li 2019; Kim 2020; Chang et al. 2020) compress the search space for faster filtering and verification, but remain inefficient. For MCS, recent work (McCreesh, Prosser, and Trimble 2017) introduces a branch and bound algorithm to reduce memory and computational demands. Approximation methods like Beam (Neuhaus, Riesen, and Bunke 2006), Hungarian (Riesen and Bunke 2009), and VJ (Fankhauser, Riesen, and Bunke 2011) employ heuristic search, trading precision for reduced complexity but with sub-exponential or cubical costs. (Wang et al. 2021) combines heuristic and learning methods but lacks scalability (Ranjan et al. 2022).

In recent years, many learning-based methods have emerged for accurate and fast graph similarity prediction. SIMGNN (Bai et al. 2019) uses a Neural Tensor Network (NTN) to capture graph-level interaction and histogram features for node-level interactions. GMN (Li et al. 2019) encodes cross-graph node matching similarities into embeddings. GRAPHSIM (Bai et al. 2020) directly captures multi-

scale node interactions via embeddings, bypassing graph-level embedding computation. MGMN (Ling et al. 2021) introduces a node-graph matching layer to capture interactions across levels. H2MN (Zhang et al. 2021) transforms graphs into hypergraphs and matches subgraphs pooled from hyperedges. ERIC (Zhuo and Tan 2022) employs a soft matching module during training, removed during inference for speed. NA-GSL (Tan et al. 2023) uses a self-attention module for node embeddings and a cross-graph attention module for similarity matrices. GEDGNN (Piao et al. 2023) applies cross-graph node matching with a graph-matching loss and a post-processing step to generate a graph edit path. These methods involve cross-graph node-level interactions, resulting in quadratic time costs. However, the cross-graph node-level interaction module may not be necessary, and simple but effective designs can already achieve superior performance. EGSC (Qin et al. 2021) omit this cross-graph node-level interaction module and uses knowledge distillation but sacrifices accuracy on unseen graphs. GREED (Ranjan et al. 2022) proposes the concept of graph pair-independent embeddings, but experiments reveal room for further improvement. Our model includes a novel embedding structure that incorporates a positional encoding technique to promote the expressiveness power of the node and graph embeddings over the 1-WL test, as well as a new multi-scale pooling technique, to improve performance.

### Conclusion

We present GRASP, a simple but effective method for accurate predictions on GED and MCS. To make the complicated simple, we design a series of rational and effective techniques in GRASP to achieve superior performance. We propose techniques to enhance node features via positional encoding, employ a robust graph neural network, and develop a multi-scale pooling technique. We theoretically prove that GRASP is expressive and passes the 1-WL test, with efficient time complexity linear to the size of a graph pair. In extensive experiments, GRASP is versatile in predicting GED and MCS scores accurately on real-world data, often outperforming existing methods by a significant margin.

## Acknowledgments

The work described in this paper was supported by grants from the Research Grants Council of Hong Kong Special Administrative Region, China (No. PolyU 25201221, PolyU 15205224). Jieming Shi is supported by NSFC No. 62202404, Otto Poon Charitable Foundation Smart Cities Research Institute (SCRI) P0051036-P0050643. This work is supported by Tencent Technology Co., Ltd. P0048511. Renchi Yang is supported by the NSFC Young Scientists Fund (No. 62302414) and the Hong Kong RGC ECS grant (No. 22202623).

## References

- Bai, Y.; Ding, H.; Bian, S.; Chen, T.; Sun, Y.; and Wang, W. 2019. SimGNN: A Neural Network Approach to Fast Graph Similarity Computation. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 384–392.
- Bai, Y.; Ding, H.; Gu, K.; Sun, Y.; and Wang, W. 2020. Learning-Based Efficient Graph Similarity Computation via Multi-Scale Convolutional Set Matching. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, 3219–3226.
- Blumenthal, D. B.; and Gamper, J. 2020. On the exact computation of the graph edit distance. *Pattern Recognit. Lett.*, 134: 46–57.
- Borgwardt, K. M.; and Kriegel, H. 2007. Graph Kernels For Disease Outcome Prediction From Protein-Protein Interaction Networks. In *Proceedings of the Pacific Symposium*, 4–15.
- Bresson, X.; and Laurent, T. 2017. Residual Gated Graph ConvNets. *arXiv preprint arXiv:1711.07553*.
- Bunke, H.; and Allermann, G. 1983. Inexact graph matching for structural pattern recognition. *Pattern Recognit. Lett.*, 1(4): 245–253.
- Bunke, H.; and Shearer, K. 1998. A graph distance metric based on the maximal common subgraph. *Pattern Recognit. Lett.*, 19(3-4): 255–259.
- Chang, L.; Feng, X.; Lin, X.; Qin, L.; Zhang, W.; and Ouyang, D. 2020. Speeding Up GED Verification for Graph Similarity Search. In *36th IEEE International Conference on Data Engineering*, 793–804.
- Dwivedi, V. P.; Luu, A. T.; Laurent, T.; Bengio, Y.; and Bresson, X. 2022. Graph Neural Networks with Learnable Structural and Positional Representations. In *The Tenth International Conference on Learning Representations*.
- Fankhauser, S.; Riesen, K.; and Bunke, H. 2011. Speeding Up Graph Edit Distance Computation through Fast Bipartite Matching. In *Graph-Based Representations in Pattern Recognition - 8th IAPR-TC-15 International Workshop*, volume 6658 of *Lecture Notes in Computer Science*, 102–111.
- Kendall, M. G. 1938. A New Measure of Rank Correlation. *Biometrika*, 30(1/2): 81–93.
- Kim, J. 2020. Nass: A New Approach to Graph Similarity Search. *arXiv preprint arXiv:2004.01124*.
- Kim, J.; Choi, D.; and Li, C. 2019. Inves: Incremental Partitioning-Based Verification for Graph Similarity Search. In *Advances in Database Technology - 22nd International Conference on Extending Database Technology*, 229–240.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations*.
- Lee, J.; Jin, R.; and Jain, A. K. 2008. Rank-based distance metric learning: An application to image retrieval. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.
- Li, P.; Wang, Y.; Wang, H.; and Leskovec, J. 2020. Distance Encoding: Design Provably More Powerful Neural Networks for Graph Representation Learning. In *Annual Conference on Neural Information Processing Systems*.
- Li, Y.; Gu, C.; Dullien, T.; Vinyals, O.; and Kohli, P. 2019. Graph Matching Networks for Learning the Similarity of Graph Structured Objects. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 3835–3845.
- Ling, X.; Wu, L.; Wang, S.; Ma, T.; Xu, F.; Liu, A. X.; Wu, C.; and Ji, S. 2021. Multilevel graph matching networks for deep graph similarity learning. *IEEE Trans. Neural Networks Learn. Syst.*, 34(2): 799–813.
- Liu, Y.; Li, C.; Jiang, H.; and He, K. 2020. A Learning Based Branch and Bound for Maximum Common Subgraph Related Problems. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2392–2399.
- McCreech, C.; Prosser, P.; and Trimble, J. 2017. A Partitioning Algorithm for Maximum Common Subgraph Problems. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 712–719.
- Neuhaus, M.; Riesen, K.; and Bunke, H. 2006. Fast Suboptimal Algorithms for the Computation of Graph Edit Distance. In *Structural, Syntactic, and Statistical Pattern Recognition, Joint IAPR International Workshops*, volume 4109 of *Lecture Notes in Computer Science*, 163–172.
- Piao, C.; Xu, T.; Sun, X.; Rong, Y.; Zhao, K.; and Cheng, H. 2023. Computing Graph Edit Distance via Neural Graph Matching. *Proc. VLDB Endow.*, 16(8): 1817–1829.
- Qin, C.; Zhao, H.; Wang, L.; Wang, H.; Zhang, Y.; and Fu, Y. 2021. Slow Learning and Fast Inference: Efficient Graph Similarity Computation via Knowledge Distillation. In *Annual Conference on Neural Information Processing Systems*, 14110–14121.
- Ranjan, R.; Grover, S.; Medya, S.; Chakaravarthy, V. T.; Sabharwal, Y.; and Ranu, S. 2022. GREED: A Neural Framework for Learning Graph Distance Functions. In *Annual Conference on Neural Information Processing Systems*.
- Ranu, S.; Calhoun, B. T.; Singh, A. K.; and Swamidass, S. J. 2011. Probabilistic Substructure Mining From Small-Molecule Screens. *Molecular Informatics*, 30(9): 809–815.
- Riesen, K.; and Bunke, H. 2009. Approximate graph edit distance computation by means of bipartite graph matching. *Image Vis. Comput.*, 27(7): 950–959.

- Sato, R. 2020. A Survey on The Expressive Power of Graph Neural Networks. *arXiv preprint arXiv:2003.04078*.
- Socher, R.; Chen, D.; Manning, C. D.; and Ng, A. Y. 2013. Reasoning With Neural Tensor Networks for Knowledge Base Completion. In *27th Annual Conference on Neural Information Processing Systems*, 926–934.
- Spearman, C. 1987. The Proof and Measurement of Association between Two Things. *The American Journal of Psychology*, 100(3/4): 441–471.
- Tan, W.; Gao, X.; Li, Y.; Wen, G.; Cao, P.; Yang, J.; Li, W.; and Zaiane, O. R. 2023. Exploring attention mechanism for graph similarity learning. *Knowl. Based Syst.*, 276: 110739.
- Wang, R.; Zhang, T.; Yu, T.; Yan, J.; and Yang, X. 2021. Combinatorial Learning of Graph Edit Distance via Dynamic Embedding. In *IEEE Conference on Computer Vision and Pattern Recognition*, 5241–5250.
- Weisfeiler, B.; and Leman, A. 1968. The reduction of a graph to canonical form and the algebra which appears therein. *Nauchno-Tekhnicheskaya Informatsia*, 12–16.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *7th International Conference on Learning Representations*.
- Zeng, Z.; Tung, A. K. H.; Wang, J.; Feng, J.; and Zhou, L. 2009. Comparing Stars: On Approximating Graph Edit Distance. *Proc. VLDB Endow.*, 2(1): 25–36.
- Zhang, Z.; Bu, J.; Ester, M.; Li, Z.; Yao, C.; Yu, Z.; and Wang, C. 2021. H2MN: Graph Similarity Learning with Hierarchical Hypergraph Matching Networks. In *The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2274–2284.
- Zhuo, W.; and Tan, G. 2022. Efficient Graph Similarity Computation with Alignment Regularization. In *Annual Conference on Neural Information Processing Systems*.