
GLSearch: Maximum Common Subgraph Detection via Learning to Search

Yunsheng Bai^{*1} Derek Xu^{*1} Yizhou Sun¹ Wei Wang¹

Abstract

Detecting the Maximum Common Subgraph (MCS) between two input graphs is fundamental for applications in drug synthesis, malware detection, cloud computing, etc. However, MCS computation is NP-hard, and state-of-the-art MCS solvers rely on heuristic search algorithms which in practice cannot find good solution for large graph pairs given a limited computation budget. We propose GLSEARCH, a Graph Neural Network (GNN) based *learning to search* model. Our model is built upon the branch and bound algorithm, which selects one pair of nodes from the two input graphs to expand at a time. We propose a novel GNN-based Deep Q-Network (DQN) to select the node pair, making the search process much faster. Experiments on synthetic and real-world graph pairs demonstrate that our model learns a search strategy that is able to detect significantly larger common subgraphs than existing MCS solvers given the same computation budget. GLSEARCH can be potentially extended to solve many other combinatorial problems with constraints on graphs.

1. Introduction

Graphs gain increasing attention recently due to their expressive nature in representing real-world data and recent successes in addressing challenging graph tasks via learning, represented by graph neural networks. Among various graph tasks, detecting the largest subgraph that is commonly present in both input graphs, known as Maximum Common Subgraph (MCS) (Bunke & Shearer, 1998) (as shown in Figure 1), is an important yet particularly hard task. MCS naturally encodes the degree of similarity between two graphs, is domain-agnostic, and thus has broad utilities in many domains such as software analysis (Park et al., 2013), graph




Figure 1: For graph pair $(\mathcal{G}_1, \mathcal{G}_2)$ with node labels, the induced connected Maximum Common Subgraph (MCS) is the five-member ring structure highlighted in circle.

database systems (Yan et al., 2005) and cloud computing platforms (Cao et al., 2011). For example, in drug synthesis, finding similar substructures in compounds with similar properties can reduce manual labor (Ehrlich & Rarey, 2011).

MCS detection is NP-hard in its nature and is thus very challenging. The state-of-the-art exact MCS detection algorithms, which use a powerful branch and bound search framework, still run in exponential time in the worst case (Liu et al., 2019). These algorithms aim to provably extract the MCS by exhausting the search space as efficiently as possible. However, in large real-world graphs, exhausting the search space is not computationally tractable. What is worse, they rely on several heuristics on how to explore the search space. For example, MCS_P (McCreesh et al., 2017) uses node degree as its heuristic by choosing high-degree nodes to visit first, but in many cases the true MCS contains low-degree nodes.

Recently, there are some related efforts from the learning community; however, these methods fall short in tackling the constraint posed by the MCS definition that the two extracted subgraphs must be isomorphic to each other. For example, Wang et al. (2019) aims to detect a soft matching matrix between nodes in two input graphs, which, however, cannot be easily transformed into the discrete matched subgraph. Bai et al. (2020b) is the first attempt to use learning based approach to directly output MCS. However, it heavily relies on labeled MCS instances, which requires pre-computation of MCS results by running exact solvers.

In this paper, we present GLSEARCH (Graph Learning to Search), a general framework for MCS detection combining the advantages of search and deep Reinforcement Learning (RL). GLSEARCH learns to search by adopting a Deep

^{*}Equal contribution ¹Department of Computer Science, University of California, Los Angeles, California, USA. Correspondence to: Yunsheng Bai <yba@ucla.edu>, Derek Xu <derekxu@ucla.edu>.

Q-Network (DQN) (Mnih et al., 2015) to replace the node selection heuristics required in state-of-the-art MCS solvers, leading to faster arrival of the optimal solution for an input graph pair, which is particularly useful when applied to large real-world graphs and/or with a limited search budget. Our method proposes a novel representation for states in the DQN framework, and reformulates DQN in a novel way to better capture the effect of different node selections, exploiting the representational power of Graph Neural Networks (GNN).

Besides the novel state representation, another key component of GLSEARCH is the search algorithm. Since the search space is very large especially for large graphs pairs, and DQN policy may choose a bad action during search, the traditional DQN framework tends to stop searching too early, leading to suboptimal solutions. To address this, we adopt a branch-and-bound search framework with backtracking.

Lastly, to enhance the training of DQN, we leverage the search algorithm to not only provide supervised signals in a pre-training stage but also offer guidance during an imitation learning stage.

Experiments on large real graph datasets (that are significantly larger than the datasets adopted by state-of-the-art MCS solvers) demonstrate that GLSEARCH outperforms baseline solvers and machine learning models for graph matching, in terms of effectiveness, by a large margin. Our contributions can be summarized as follows:

- We address the important yet challenging task of Maximum Common Subgraph detection for general-domain input graph pairs and propose GLSEARCH as the solution.
- The key novelty of GLSEARCH is its learning to search framework. With a DQN reformulation trick, it is trained under the reinforcement learning framework to make the best decision at each search step in order to quickly find the best MCS solution during search. The search in turns helps DQN training in a pre-training stage and an imitation learning stage.
- We conduct extensive experiments on medium, large, and million-node real-world graphs to demonstrate the effectiveness of the proposed approach compared against a series of strong baselines in MCS detection and graph matching.

2. Preliminaries and Related Work

2.1. The MCS Detection Problem

We denote a graph as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} and \mathcal{E} denote the vertex and edge set. A node-induced subgraph is defined as $\mathcal{G}_s = (\mathcal{V}_s, \mathcal{E}_s)$ where \mathcal{E}_s preserves all the edges

between nodes in \mathcal{V}_s , i.e. $\forall i, j \in \mathcal{V}_s, (i, j) \in \mathcal{E}_s$ if and only if $(i, j) \in \mathcal{E}$. In this paper, we aim to detect the Maximum Common induced Subgraph (MCS) between an input graph pair, denoted as $\text{MCS}(\mathcal{G}_1, \mathcal{G}_2)$, which is the largest node-induced subgraph contained in both \mathcal{G}_1 and \mathcal{G}_2 . In addition, we require $\text{MCS}(\mathcal{G}_1, \mathcal{G}_2)$ to be a connected subgraph. We allow the nodes of input graphs to be labeled, in which case the labels of nodes in the MCS must match as in Figure 1. Graph isomorphism and subgraph isomorphism can be regarded as two special tasks of MCS: $|\text{MCS}(\mathcal{G}_1, \mathcal{G}_2)| = |\mathcal{V}_1| = |\mathcal{V}_2|$ if \mathcal{G}_1 are isomorphic to \mathcal{G}_2 , $|\text{MCS}(\mathcal{G}_1, \mathcal{G}_2)| = \min(|\mathcal{V}_1|, |\mathcal{V}_2|)$ when \mathcal{G}_1 (or \mathcal{G}_2) is subgraph isomorphic to \mathcal{G}_2 (or \mathcal{G}_1).

2.2. Related Work

Traditional Efforts MCS detection is NP-hard, with existing methods based on constraint programming (Vismara & Valery, 2008; McCreesh et al., 2016), branch and bound (McCreesh et al., 2017; Liu et al., 2019), integer programming (Bahense et al., 2012), conversion to maximum clique detection (Levi, 1973; McCreesh et al., 2016), etc., among which MCSP+RL (Liu et al., 2019) (details presented in Section 2.3) is the state-of-the-art method, which guarantees to find common subgraphs satisfying the isomorphism constraint, but usually cannot extract large common subgraphs when input graphs become large.

Efforts on Learning to Solve Graph Similarity and Matching

There is a growing trend of using machine learning approaches to graph matching and similarity score computation (Zanfir & Sminchisescu, 2018; Wang et al., 2019; Yu et al., 2020; Xu et al., 2019b;a; Bai et al., 2019; 2020a; Li et al., 2019; Ling et al., 2020; Liu et al., 2020; Chen et al., 2020). These methods cannot handle the isomorphism constraint in MCS well, since they were mainly designed for tasks without hard constraints, e.g. finding the similarity score or node-node matching between two graphs supervised by true similarity or matching. Thus, to better satisfy the constraints of MCS, these models need to be embedded into a search framework that uses the scores provided by the models to guide the search for MCS, which will be described next. For example, GW-QAP performs Gromov-Wasserstein discrepancy (Peyré et al., 2016) based optimization and outputs a matching matrix for all node pairs indicating the likelihood of matching (Zanfir & Sminchisescu, 2018). I-PCA performs image matching by outputting a doubly-stochastic matching matrix computed from intermediary Convolution Neural Network features from an input image pair (Wang et al., 2019). Some other methods tackle tasks with hard constraints, such as MCS or GED, but these methods require supervision, which is difficult to obtain (Bai et al., 2020b; Wang et al., 2020). A recent work proposes to use RL to overcome the supervision limitation; however, this approach relies heavily on existing heuristics

to work well (Liu et al., 2019).

Efforts on Learning to Solve NP-hard Graph Problems

Existing works such as Dai et al. (2017), Gasse et al. (2019), and Fan et al. (2020) focus on designing learning based approaches for solving NP-hard tasks on graphs, e.g. Minimum Vertex Cover, Mixed Integer Linear Programming, Network Dismantling, etc., but our problem, Maximum Common Subgraph detection, operates on a pair of input graphs instead of a single graph. Besides, MCS detection requires hard constraint satisfaction, i.e. isomorphism of extracted subgraphs, which is handled by a search algorithm described next.

2.3. Search Algorithms for MCS

In this section, we present the state-of-the-art branch and bound search framework for detecting MCS as shown in Algorithm 1 and Figure 2, which allows the exploration of search space and guarantees the satisfaction of the isomorphism constraint posed by MCS. Thus, it serves as the backbone of our proposed approach. We then discuss several drawbacks in the existing search-based MCS detection algorithms.

Algorithm 1 Branch and Bound for MCS. We highlight in green boxes the two places that will be replaced by GLSEARCH.

```

1: Input: Input graph pair  $\mathcal{G}_1, \mathcal{G}_2$ .
2: Output:  $maxSol$ .
3: Initialize  $s_0 \leftarrow$  empty state.
4: Initialize  $stack \leftarrow$  new Stack( $s_0$ ).
5: Initialize  $maxSol \leftarrow$  empty solution.
6: while  $stack \neq \emptyset$  do
7:    $s_t \leftarrow stack.pop()$ ;
8:    $curSol \leftarrow s_t.getCurSol()$ ;
9:   if  $|curSol| > |maxSol|$  then
10:      $maxSol \leftarrow curSol$ ;
11:   end if
12:    $UB_t \leftarrow |curSol| + overestimate(s_t)$ ;
13:   if  $UB_t \leq |maxSol|$  or  $|s_t.actions| = 0$  then
14:     continue;
15:   end if
16:    $\mathcal{A}_t \leftarrow s_t.actions$ ;
17:    $a_t \leftarrow policy(s_t, \mathcal{A}_t)$ ;
18:    $s_t.actions \leftarrow s_t.actions \setminus \{a_t\}$ ;
19:    $stack.push(s_t)$ ;
20:    $s_{t+1} \leftarrow environment.update(s_t, \mathcal{A}_t)$ ;
21:    $stack.push(s_{t+1})$ ;
22: end while
    
```

MCSP and Its Limitations The basic version, MCSP, is presented in McCreesh et al. (2017) and the more advanced version, MCSP+RL, is proposed in Liu et al. (2019).

The whole search algorithm, outlined in Algorithm 1¹, is a branch-and-bound algorithm that, starting from an empty subgraph, grows the matching subgraph one node pair (between the two graphs) at a time and maintains the best solution found so far. In each search iteration, denote the current search state as s_t consisting of $\mathcal{G}_1, \mathcal{G}_2$, the current matched subgraphs $\mathcal{G}_{1s} = (\mathcal{V}_{1s}, \mathcal{E}_{1s})$ and $\mathcal{G}_{2s} = (\mathcal{V}_{2s}, \mathcal{E}_{2s})$ as well as their node-node mappings. The algorithm tries to select one node pair, (i, j) added to the currently selected subgraphs, where node i is from \mathcal{G}_1 and node j is from \mathcal{G}_2 , as its action, denoted as a_t . It then decides to either continue the search if the solution is promising, or otherwise backtrack to the parent search state, i.e. the current search state is pruned (line 14). Various heuristics on node pair selection policy, denoted as “*policy*” in line 17, are proposed in MCSP and MCSP+RL. For example, in MCSP, nodes of large degrees are selected before small-degree nodes.

At each search state, in order to determine whether the solution is promising or not, an upper bound of the size of the MCS, “ UB_t ” in line 12 is computed. A concept of “bidomain” is introduced to facilitate its estimation. Bidomains partition the unselected nodes, $(V_1 \setminus V_{1s}) \cup (V_2 \setminus V_{2s})$, into equivalent classes, $\mathcal{D} = \{D_k\}_{k=0}^{2^{|V_{1s}|-1}}$, where the k -th bidomain D_k consists of two sets of nodes, $D_k = \langle \mathcal{V}'_{k1}, \mathcal{V}'_{k2} \rangle$ such that nodes in \mathcal{V}'_{k1} and nodes in \mathcal{V}'_{k2} connect to subgraph nodes, in \mathcal{V}_{1s} and \mathcal{V}_{2s} , that match to each other. Figure 3 shows an example with three bidomains. Due to the subgraph isomorphism constraint posed by MCS, only nodes in \mathcal{V}'_{k1} can match to \mathcal{V}'_{k2} and vice versa. Since we require the MCS to be connected subgraphs, we differentiate bidomains $\mathcal{D}^{(c)} = \{D_k\}_{k=1}^{2^{|V_{1s}|-1}}$ that are connected (adjacent) to \mathcal{G}_{1s} and \mathcal{G}_{2s} (e.g. bidomain “01” and “10” in Figure 3) from the single bidomain D_0 disconnected (unconnected) from \mathcal{G}_{1s} and \mathcal{G}_{2s} (e.g. bidomain “00” in Figure 3) The candidate node pairs to select from, i.e. the action space “ \mathcal{A}_t ”, consists of all node pairs in all connected bidomains, $\mathcal{D}^{(c)}$. This also guarantees the extracted subgraphs at each state are isomorphic to each other.

To estimate the upper bound, it is noteworthy that each bidomain can contribute at most $\min(|\mathcal{V}'_{k1}|, |\mathcal{V}'_{k2}|)$ nodes to the future best solution. The upper bound can therefore be estimated as $\sum_{D_k \in \mathcal{D}} \min(|\mathcal{V}'_{k1}|, |\mathcal{V}'_{k2}|)$, which is the “ $overestimate(s_t)$ ” function in line 12. This upper bound computation is consistently used for all the methods in the paper. The major difference is and the node pair selection policy and backtracking method, i.e. lines 17 and 7.

As mentioned previously, MCSP adopts a heuristic that selects node pairs with the largest degree as its policy. The most severe limitation of MCSP is that the node-degree-

¹The original algorithm is recursive. To highlight our novelty, we rewrite into an equivalent iterative version.

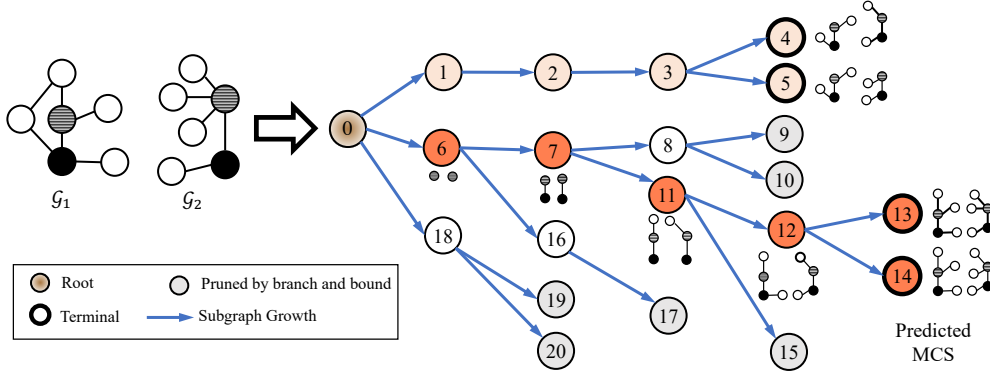


Figure 2: An illustration of the branch and bound search algorithm for MCS detection (Algorithm 1). Each node represents one state (s_t), with node id reflecting the order states are visited, and each edge represents an action (a_t) of selecting a new node pair. The search is essentially depth-first with pruning by the upper bound check. Our model learns the node pair selection strategy, i.e. which state to visit first. Notice, if the policy had visited state 6 before state 1, a large solution may be found in less iterations. When the search completes or a pre-defined search iteration budget is used up, the best solution (largest subgraphs) will be returned, which in this example is state 13 (and 14).




Figure 3: An example to illustrate the concept of bidomains. We denote the bidomain of each non-subgraph node by its bitstring, where each bit denotes whether there is a connection (“1”) or not (“0”) to a particularly node in subgraph \mathcal{V}_{1s} or the node it maps to in subgraph \mathcal{V}_{2s} (both circled in red). For example, each node in the “10” bidomain is connected to the top “C” node in the subgraph and disconnected to the bottom “C” node. In this case, there are 3 bidomains denoted as “00” (D_0), “01” (D_1), and “10” (D_2). By definition, the bidomain denoted with all zeros, e.g. “00” in this case, is called the disconnected bidomain. Notice, the bidomain corresponding to “11” (D_3) is empty.

based heuristic is not adaptive to the complex real-world graph structures. MCSP always backtracks to the parent search state; however, this leads to suboptimal solutions, as the wrong action could have occurred earlier in search.

MCSP+RL and Its Limitations MCSP+RL improves MCSP by replacing the node pair selection policy with a value function for each node or node pair. Their goal is to minimize the search tree size so that a search tree leaf can be reached as early as possible. Specifically, MCSP+RL aims to reduce the UB_t to make it tighter, so that more pruning (line 14) can happen in subsequent search steps, resulting in a smaller search tree (search space). To achieve that goal, they design the reward function for each state-action pair as the reduction (or reduction rate) of search space by selecting

that node pair. The value function maintains a score for each node (or node pair), which is initialized to 0 and updated during search. In each step during search, the policy is to select the node pair with the largest score.

We identify two limitations of MCSP+RL: (1) Since the reward definition is defined using a heuristic and there is no training stage, for each new graph pair, the scores must be re-initialized and the policy has to be re-updated. (2) The fact that the scores for each node (or node pair) are 0 at the beginning of search has another problem: MCSP+RL breaks ties using node degrees, essentially degenerating to the same policy as MCSP initially for each graph pair, which is also verified by our experiments as shown in Figure 4 where MCSP and MCSP+RL perform the same.

Besides, a common issue of MCSP and MCSP+RL is that, during the search, they can enter a bad locally optimal search state and get “stuck” without finding a better (larger) solution, $maxSol$, for many iterations, as shown in the flat line segments of Figure 4.

3. Proposed Method

In this section, we present our RL based MCS detection method, GLSEARCH. The rest of Section 3 is organized as follows. Section 3.1 presents a high-level overview of how to leverage Deep Q-Network (DQN) (Mnih et al., 2015) for search, including the basic definitions of state, action, reward, etc., and how DQN can address the various issues of search methods for MCS described previously. Section 3.2 describes the details of how to leverage DQN for search, focusing on how to effectively design representation learning for DQN for the task of MCS detection. Section 3.3 explains how to effectively train the DQN with the help of search, i.e. how search can in turn help DQN training.

Table 1: Table of Notations.

| Notation | Description |
|--|--|
| $\mathcal{G}_1, \mathcal{G}_2$ | The input pair of graphs. |
| $\mathcal{G}_{1s}, \mathcal{G}_{2s}$ | The selected subgraph of some given state. |
| $\mathcal{V}_x, \mathcal{E}_x$ | Set of nodes and edges for graph \mathcal{G}_x . |
| s_t | The search state at iteration t . |
| a_t | The action selected at iteration t . |
| \mathcal{D} | Set of all bidomains, $\{D_k\}_{k=0}^{2^{ \mathcal{V}_{1s} -1}}$. |
| $\mathcal{D}^{(c)}$ | Set of connected bidomains, $\{D_k\}_{k=1}^{2^{ \mathcal{V}_{1s} -1}}$. |
| D_0 | The disconnected bidomain. |
| $\mathcal{V}'_{k1}, \mathcal{V}'_{k2}$ | Set of nodes in \mathcal{G}_1 and \mathcal{G}_2 in D_k . |
| h_i, h_j | Embeddings of nodes $i \in \mathcal{V}_1$ and $j \in \mathcal{V}_2$. |
| h_{s1}, h_{s2} | Embeddings of subgraphs \mathcal{G}_{1s} and \mathcal{G}_{2s} . |
| $h_{\mathcal{G}_1}, h_{\mathcal{G}_2}$ | Embeddings of graph \mathcal{G}_1 and subgraph \mathcal{G}_2 . |
| $h_{D^{(c)}}$ | Embedding for connected bidomains, $D^{(c)}$ |
| h_{D_0} | Embedding for disconnected bidomains, D_0 |

3.1. Leveraging DQN for Search: Overview

GLSEARCH enables graph representation learning techniques to tackle the hard isomorphism constraint posed by MCS and uses deep Q-learning to select node pairs smartly in each search state. GLSEARCH represents states and actions in continuous embeddings, and maps (s_t, a_t) to a score $Q(s_t, a_t)$ via a DQN which consists of a Graph Neural Network encoder and learnable components to project the representations into the final score. GLSEARCH is trained on a set of diverse small and medium-sized graphs, and once trained, can be applied to any new graph pair.

Unlike MCSP and MCSP+RL, which aim to reduce the search tree size, the aim of our agent is to directly maximize the common subgraph size, allowing large common subgraph to be found even on very large graph pairs.

State s_t consists of the (1) current selected subgraphs, (2) the node-node mappings between the nodes in the selected subgraphs, and (3) the input graphs. We include the node-node mappings as part of the state definition since node-node mappings can be used to derive the bidomain partitioning as illustrated in Figure 3, which constrains the node pairs that can be selected in future, and thus affects the future common subgraph size. Action a_t is defined as a node pair to select. For GLSEARCH, given our goal, the immediate reward for transitioning from one state to any next state is defined as $r_t = +1$ since one new node pair is selected, so that $Q(s_t, a_t)$ captures the largest common subgraph size starting at s_t by performing a_t .

GLSEARCH is trained to find large common subgraphs quickly, but due to the large action space of large graph pairs, our model may still be susceptible to the local optimum without increasing $maxSol$ as described in Section 2.3. Thus, when this occurs, we utilize additional information stored in the search tree to backtrack to a state that will

most likely improve $maxSol$. We find that in practice, states with a large action space, \mathcal{A}_t , tend to include more high-quality unexplored actions. Hence, if the best solution found so far does not increase² for a pre-defined number of iterations, then in the next iteration, instead of popping from the stack³, we find the state with the largest action space, and visit it. We refer to this improved search methodology as **promise-based search**. More details can be found in the supplementary material.

3.2. Search Policy Learning via GNN-based DQN

Since the action space can be large for MCS, we leverage the representation learning capacity of continuous representations for DQN design. At state s_t , for each action a_t , our DQN predicts a $Q(s_t, a_t)$ representing the remaining future reward after selecting action $a_t = (i, j)$ where $i \in \mathcal{V}_1$ and $j \in \mathcal{V}_2$, which intuitively corresponds to the largest number of nodes that will be eventually selected starting from the action edge (s_t, a_t) as shown in tree in Figure 2.

Based on the above insights, one can design a simple DQN leveraging the representation learning power of Graph Neural Networks (GNN) such as Kipf & Welling (2016) and Velickovic et al. (2018) by passing \mathcal{G}_1 and \mathcal{G}_2 to a GNN to obtain one embedding per node, $\{h_i | \forall i \in \mathcal{V}_1\}$ and $\{h_j | \forall j \in \mathcal{V}_2\}$. Denote CONCAT as concatenation, READOUT as a readout operation that aggregates node-level embeddings into subgraph embeddings h_{s1} and h_{s2} , and whole-graph embeddings $h_{\mathcal{G}_1}$ and $h_{\mathcal{G}_2}$. A state can then be represented as $h_{s_t} = \text{CONCAT}(h_{\mathcal{G}_1}, h_{\mathcal{G}_2}, h_{s1}, h_{s2})$. An action can be represented as $h_{a_t} = \text{CONCAT}(h_i, h_j)$. The Q function would then be designed as:

$$\begin{aligned} Q(s_t, a_t) &= \text{MLP}(\text{CONCAT}(h_{s_t}, h_{a_t})) \\ &= \text{MLP}(\text{CONCAT}(h_{\mathcal{G}_1}, h_{\mathcal{G}_2}, h_{s1}, h_{s2}, h_i, h_j)). \end{aligned} \quad (1)$$

However, there are several flaws to this simple design of Q function:

- (A) h_i and h_j , generated by typical GNNs, encode only *local* neighborhood information, but $Q(s_t, a_t)$ should capture the *long-term* effect of adding (i, j) . What is worse, different node pairs have different embeddings, but their immediate rewards are always $+1$, a constant, in MCS, making differentiating the quality of different actions even more difficult.
- (B) Swapping the order of \mathcal{G}_1 and \mathcal{G}_2 should not cause $Q(s_t, a_t)$ to change, but concatenating embeddings from the two graphs causes the DQN to be sensitive to their ordering.

²If the search does not enter line 10 of Algorithm 1.

³Line 7 of Algorithm 1.

(C) Lastly, how to effectively leverage the node-node mappings between \mathcal{G}_{1s} and \mathcal{G}_{2s} , an important part of the state definition as explained in Section 3.1, for predicting $Q(s_t, a_t)$ remains a challenge.

To address these issues, we propose the following improvements over the simple DQN design.

Factoring out Action In order to maximally reflect the effect of adding node pair (i, j) to \mathcal{G}_{1s} and \mathcal{G}_{2s} , we reformulate the optimal Q score, $Q^*(s_t, a_t)$, as $r_t + \gamma V^*(s_{t+1}) = 1 + \gamma V^*(s_{t+1})$ (using the fact that $r_t = +1$) in MCS, where V is the value function, and γ is the discount factor. Then, in order to compute the effect of a_t , we can compute the value associated with s_{t+1} which does not depend on a_t and avoids the use of local \mathbf{h}_i and \mathbf{h}_j . In this case, we can rely on our state embedding to capture global information and amplify differences between different actions by looking at the states they will arrive.

Interaction between Input Graphs To resolve the graph symmetry issue, we first construct the interaction between the embeddings from two graphs, i.e. $\text{INTERACT}(\mathbf{h}_{x1}, \mathbf{h}_{x2})$, where \mathbf{h}_{x1} and \mathbf{h}_{x2} represent any embedding from \mathcal{G}_1 and \mathcal{G}_2 respectively, and $\text{INTERACT}(\cdot)$ is any commutative function to combine the two embeddings. Instead of summation, We define a CNN-based $\text{INTERACT}(\cdot)$ operator, defined in Section B.3 of the Supplementary Material, since cross-graph interaction is important for the MCS task. This interacted embedding is later concatenated with other useful representations and fed into a final MLP to compute the Q score.

Bidomain Representations Bidomains are derived from node-node mappings and partition the rest of \mathcal{G}_1 and \mathcal{G}_2 , which is a more useful signal for predicting the future reward. In fact, as described in Section 2.3, bidomains have been adopted in search-based MCS solvers to estimate the upper bound. Here, we require the harder prediction of $Q(s_t, a_t)$ for which we propose to also use the representation of bidomains to amplify the differences in different states. Denote \mathbf{h}_{D_k} as the representation for bidomain $D_k = \langle \mathcal{V}'_{k1}, \mathcal{V}'_{k2} \rangle$. Similar to computing the graph-level and subgraph-level embeddings, we compute \mathbf{h}_{D_k} as

$$\mathbf{h}_{D_k} = \text{INTERACT}\left(\text{READOUT}(\{\mathbf{h}_i | i \in \mathcal{V}'_{k1}\}), \text{READOUT}(\{\mathbf{h}_j | j \in \mathcal{V}'_{k2}\})\right). \quad (2)$$

Given all the bidomain embeddings, we compute a single representation for all the connected bidomains, $\mathcal{D}^{(c)}$, $\mathbf{h}_{\mathcal{D}^{(c)}} = \text{READOUT}(\{\mathbf{h}_{D_k} | k \in \mathcal{D}^{(c)}\})$. Notice, since bidomains are a function of node-node mappings, the mapping is implicitly accounted for by this representation. In fact, bidomain representations are better at capturing difference in states than node-node mappings. Across different actions, the node-node mappings only differ by a single node-node pair,

but the bidomains can differ drastically as they capture the effect of matching certain nodes, rather than the mappings themselves. Our final DQN has the form:

$$Q(s_t, a_t) = 1 + \gamma \text{MLP}\left(\text{CONCAT}\left(\text{INTERACT}(\mathbf{h}_{\mathcal{G}_1}, \mathbf{h}_{\mathcal{G}_2}), \text{INTERACT}(\mathbf{h}_{s1}, \mathbf{h}_{s2}), \mathbf{h}_{\mathcal{D}^{(c)}}, \mathbf{h}_{D_0}\right)\right). \quad (3)$$

3.3. Leveraging Search for DQN Training

For large graph pairs, the action space can be quite large. Thus, to enhance the training of our DQN, before the standard training of DQN (Mnih et al., 2013), we pre-train DQN and guide its exploration with expert trajectories supplied by the search algorithm.

For the pre-training stage, we run the search to completion on small graph pairs (thus, the exact MCS solution is found), and use a supervised mse loss function to replace the DQN loss function. The overall loss function is $(y_t - Q(s_t, a_t))^2$ where y_t the target for iteration t and $Q(s_t, a_t)$ is the predicted score. In this case, y_t denotes the remaining size of the largest common subgraph starting from s_t to its leaf node in the current branch of the search tree.

For larger graph pairs though, finding the true target becomes too slow. In that case, after pre-training, we enter the imitation learning stage where we follow the expert trajectories provided by MCSP instead of its own predicted $Q(s_t, a_t)$, to incorporate more trustworthy policy decisions into the training signal. More details can be found in the supplementary material.

4. Experiments

We evaluate GLSEARCH against two state-of-the-art exact MCS detection algorithms and a series of graph matching methods from various domains. We conduct experiments on 7 hundred-node medium-sized synthetic and real-world graph datasets, 8 thousand-node large real-world graph datasets, and 2 million-node very large real-world datasets, whose details can be found in the supplementary material. Among the different baseline models, we find no consistent trend. This indicates the difficulty of our task, as existing methods cannot find a consistent policy that guarantees good performance on datasets from different domains. Our model can substantially outperform the baselines, highlighting the significance of our contributions to learning for search.

4.1. Baseline Methods

There are two groups of methods: Exact MCS algorithms including MCSP and MCSP+RL, learning based graph matching models including GW-QAP, I-PCA, and NEURALMCS.

Table 2: Results on medium graphs. Each synthetic dataset consists of 50 randomly generated pairs labeled as “(generation algorithm)-(number of nodes in each graph)”. “BA”, “ER”, and “WS” refer to the Barabási-Albert (BA) (Barabási & Albert, 1999), the Erdős-Rényi (ER) (Gilbert, 1959), and the Watts–Strogatz (WS) (Watts & Strogatz, 1998) algorithms, respectively. NCI109 consists of 100 chemical compound graph pairs whose average graph size is 28.73. We show the ratio of the (average) size of the subgraphs found by each method with respect to the best result on that dataset.

| Method | BA-50 | BA-100 | ER-50 | ER-100 | WS-50 | WS-100 | NCI109 |
|--------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| MCSP | 0.913 | 0.892 | 0.842 | 0.896 | 0.905 | 0.856 | 0.948 |
| MCSP+RL | 0.923 | 0.857 | 0.844 | 0.877 | 0.913 | 0.875 | 0.948 |
| GW-QAP | 0.945 | 0.887 | 0.855 | 0.925 | 0.916 | 0.898 | 0.966 |
| I-PCA | 0.899 | 0.863 | 0.848 | 0.923 | 0.879 | 0.852 | 0.951 |
| NEURALMCS | 0.908 | 0.889 | 0.846 | 0.906 | 0.889 | 0.865 | 0.954 |
| GLSEARCH-RAND | 0.995 | 0.987 | 0.920 | 0.978 | 0.967 | 0.931 | 0.989 |
| GLSEARCH | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 |
| BEST SOLUTION SIZE | 19.12 | 34.38 | 26.56 | 37.64 | 29.48 | 55.56 | 10.48 |

Table 3: Results on real-world large graph pairs. Each dataset consists of one large real graph pair ($\mathcal{G}_1, \mathcal{G}_2$ may not be isomorphic, but $\mathcal{G}_{1s}, \mathcal{G}_{2s}$ are isomorphic guaranteed by search). Below each dataset name, we show its size $\min(|\mathcal{V}_1|, |\mathcal{V}_2|)$ to indicate these pairs are significantly larger than the ones in Table 2. Consistent with Table 2, we show the ratio of the subgraph sizes.

| Method | ROAD 652 | DBEN 1945 | DBZH 1907 | DBPD 1907 | ENRO 3369 | COPR 3518 | CIRC 4275 | HPPI 2152 |
|--------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| MCSP | 0.374 | 0.815 | 0.797 | 0.722 | 0.694 | 0.684 | 0.498 | 0.864 |
| MCSP+RL | 0.771 | 0.699 | 0.589 | 0.434 | 0.742 | 0.674 | 0.583 | 0.787 |
| GW-QAP | 0.305 | 0.929 | 0.855 | 0.808 | 0.711 | 0.860 | 0.354 | 0.834 |
| I-PCA | 0.267 | 0.551 | 0.589 | 0.607 | 0.650 | 0.707 | 0.203 | 0.762 |
| NEURALMCS | 0.977 | 0.785 | 0.616 | 0.620 | 0.737 | 0.742 | 0.561 | 0.785 |
| GLSEARCH-RAND | 0.641 | 0.762 | 0.658 | 0.639 | 0.814 | 0.755 | 0.603 | 0.814 |
| GLSEARCH | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 |
| BEST SOLUTION SIZE | 131 | 508 | 482 | 521 | 543 | 791 | 3515 | 404 |

All the methods either originally use or are adapted to use the branch and bound search framework in Section 2.3 with differences in node pair selection policy and training strategies. During testing, we apply the trained model on all testing graph pairs. We give a budget of 500 and 7500 search iterations for medium-size and large graph pairs. For each of the two million-node graph pairs, since the true MCS is much larger, we run each method for 50 minutes and plot the subgraph size growth across time. Due to the search algorithm, all the methods can find exact MCS solutions given long enough budget, albeit an unrealistic assumption in practice for large graph pairs.

To validate the usefulness of the learned DQN, we compare GLSEARCH, our full model, with a randomly initialized model, GLSEARCH-RAND, which replaces the output of our DQN with a completely random scalar.

4.2. Hyperparameter Settings

For I-PCA, NEURALMCS and GLSEARCH, we utilize 3 layers of Graph Attention Networks (GAT) (Velickovic et al., 2018) each with 64 dimensions for the embeddings. The initial node embedding is encoded using the local degree profile (Cai & Wang, 2018). We use $\text{ELU}(x) = \alpha(\exp(x)-1)$ for $x \leq 0$ and x for $x > 0$ as our activation function

where $\alpha = 1$. We run all experiments with Intel i7-6800K CPU and one Nvidia Titan GPU. For DQN, we use MLP layers to project concatenated embeddings to a scalar. We use SUM followed by an MLP for READOUT and IDCONV+MAXPOOL followed by an MLP for INTERACT. Further details can be found in supplementary material. For training, we set the learning rate to 0.001, the number of training iterations to 10000, and use the Adam optimizer (Kingma & Ba, 2015). The models were implemented with the PyTorch and PyTorch Geometric libraries (Fey & Lenssen, 2019).

4.3. Results

The key property of GLSEARCH is its ability to find the best solution in the fewest number of search iterations. As shown in Table 2, our model outperforms baselines in terms of size of extracted subgraphs on all medium-sized synthetic graph datasets and the chemical compound dataset NCI109.

Regarding large real-world graphs, as shown in Table 3, our model outperforms baselines in terms of the size of the extracted subgraphs on all datasets. The exact solvers rely on heuristics for node selection, and consistently find much smaller subgraphs compared to our results.

Compared with learning based graph matching models,

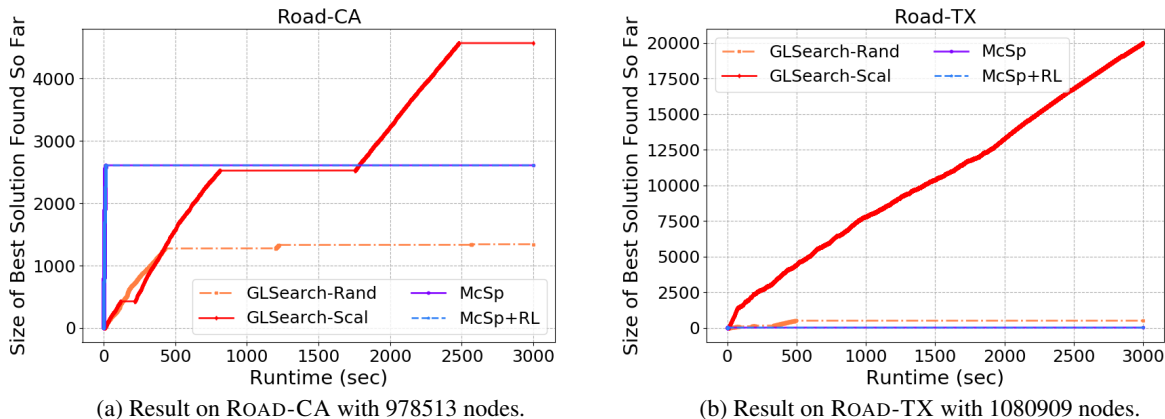


Figure 4: Comparison of the best solution sizes of different methods on two million-node graph pairs, ROAD-CA and ROAD-TX. GW-QAP, I-PCA, and NEURALMCS encounter memory error on these graph pairs due to their computation of a quadratic node-node matching matrix

Table 4: Ablation study on large real-world datasets. We demonstrate our Q function design choices indeed solve the various shortcomings presented in Section 3.2, through better representation learning.

| Method | ROAD | DBEN | DBZH | DBPD | ENRO | CoPR | CIRC | HPPI |
|--------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| GLSEARCH (no h_G) | 0.977 | 0.878 | 0.925 | 0.845 | 0.860 | 0.987 | 0.980 | 0.960 |
| GLSEARCH (no h_s) | 1.000 | 0.874 | 0.894 | 0.869 | 0.928 | 1.000 | 0.801 | 0.913 |
| GLSEARCH (no h_{D_c}) | 0.803 | 0.780 | 0.687 | 0.818 | 0.740 | 0.804 | 0.505 | 0.849 |
| GLSEARCH (no h_{D_0}) | 0.576 | 0.856 | 0.782 | 0.768 | 0.823 | 0.932 | 0.323 | 0.938 |
| GLSEARCH (SUM interact) | 0.902 | 0.913 | 0.963 | 0.885 | 0.899 | 0.957 | 1.000 | 0.948 |
| GLSEARCH (unfactored) | 0.447 | 0.807 | 0.712 | 0.582 | 0.816 | 0.816 | 0.512 | 0.861 |
| GLSEARCH (unfactored-i) | 0.500 | 0.789 | 0.741 | 0.772 | 0.748 | 0.825 | 0.902 | 0.864 |
| GLSEARCH | 0.992 | 1.000 | 1.000 | 1.000 | 1.000 | 0.990 | 0.881 | 1.000 |
| BEST SOLUTION SIZE | 132 | 508 | 482 | 521 | 543 | 799 | 3989 | 404 |

GLSEARCH is the only model which learns a reward that is dependent on both state and action, i.e. $Q(s_t, a_t)$. GW-QAP, I-PCA, and NEURALMCS essentially pre-compute the matching scores for all the node pairs in the input graphs, and therefore at each search step, the scores cannot adapt to the particular state, i.e. the matching scores only depend on $\mathcal{G}_1, \mathcal{G}_2$. Notice our state representation includes $\mathcal{G}_1, \mathcal{G}_2$ as well, hence GLSEARCH has more representational power than baselines. Trained under a reinforcement learning framework guided by search, GLSEARCH also performs the best among learning based baselines.

4.4. Million-Node Graph Pairs

GLSEARCH can scale to very large graph pairs, the limit of which is only bounded by the scalability of the GNN embedding step. To demonstrate this, we run GLSEARCH on million-node real-world graph datasets, ROAD-CA and ROAD-TX. To fit the model onto our GPU resources, we construct a lighter version of GLSEARCH, called GLSEARCH-SCAL, which reduces the GAT encoder dimensions from 64 to 16.

As shown in Figure 4, GLSEARCH significantly outperforms baseline solvers on the two million-node real-world datasets. On ROAD-TX, MCS and MCS+RL perform poorly (getting “stuck” in local optimum as pointed out in Section 2.3) while GLSEARCH continues to find larger and larger common subgraph after 50 minutes.

4.5. Ablation Study

To evaluate the effectiveness of different components proposed in our DQN model, we run ablation studies on the 8 large real world datasets.

We first measure the importance of each embedding vector fed to our DQN module, as described by Equation 3. We remove each embedding vector (specifically: $h_G = \text{INTERACT}(h_{G_1}, h_{G_2})$, $h_s = \text{INTERACT}(h_{s1}, h_{s2})$, h_{D_c} , and h_{D_0}) individually from the DQN model and retrain the model under the same training settings. Table 4 is consistent with our conclusion that every embedding vector used by GLSEARCH is critical in capturing the search state’s representation. Furthermore, we find leveraging bidomain representations is very beneficial to our model.

We next measure the importance of interaction to address the symmetry issue of the MCS calculation, where input graph pairs must be order insensitive. We first test the necessity of using more complex interaction functions, by replacing our 1DCONV+MAXPOOL interaction with simple SUM for interaction (still followed by an MLP). As shown in Table 4, we see that simpler interaction functions may not be powerful enough to encode the interaction between 2 graphs. Particularly, this suggests that interaction is quite important to model performance.

Finally, we measure the importance of factoring out actions from our DQN model. We test this with 2 models. The first utilizes Equation 1 to encode the Q-value, which we refer to as GLSEARCH (unfactored). Since Equation 1 also suffers from the issue of graph symmetry, we adapt this model to use the same interaction function as GLSEARCH to construct 3 order-invariant embeddings $h_G = \text{INTERACT}(h_{G_1}, h_{G_2})$, $h_s = \text{INTERACT}(h_{s1}, h_{s2})$, $h_a = \text{INTERACT}(h_i, h_j)$ to concatenate and pass to the final MLP layer in Equation 1. We refer to this model as GLSEARCH (unfactored-i). Our results show that without factoring out the action, our performance is comparable to or worse than MCSP, indicating the significant performance boost introduced by factoring out the action.

5. Insights and Contributions of GLSEARCH

To Search community on MCS detection The major challenge that prevents existing search algorithms from extracting large common subgraphs for large input graph pairs is that the focus of these algorithms is on reducing the **overall** search space rather than making smarter node pair selections in **each** search step. By improving the order it searches candidate solutions, GLSEARCH can quickly find **better** MCS candidates, without much (or any) backtracking and pruning, than state-of-the-art search algorithms.

To General Learning community GLSEARCH tackles the hard constraint that the subgraphs must be isomorphic to each other by only choosing actions from connected bidomains. However, there is an additional advantage of introducing the bidomain concept: Bidomains partition the rest of the input graphs into different regions where future actions can be selected. Thus, properly encoding of the bidomains gives more information about hard constraints to the DQN, which improves performance as experimentally verified by the ablation studies. More generally, this shows that learning components can be further enriched by incorporating knowledge on tackling hard constraints of an NP-hard task (e.g. bidomain) into their model design.

To Graph Deep Learning community Although various works have pointed out and analyzed the limitation of GNN’s expressive power (Xu et al., 2019c; Balcilar et al.,

2021), for particular tasks such as MCS detection, GNNs can still be used if augmented properly. We aim to predict a Q score for a state-action pair, using a DQN with two components, one component computing the local node embeddings using several layers of GNN, the other component combining embeddings at a different granularity, i.e. embeddings at the subgraph, whole-graph, and bidomain levels, to produce the final score. Overall our DQN design adopts a similar general principle as Position-Aware GNN (You et al., 2019), which allows a regular GNN to absorb information from non-local nodes (called “anchor” nodes which are randomly selected nodes globally).

To Reinforcement Learning community We are aware of efforts in the RL community to tackle NP-hard problems, but they either focus on non-graph tasks, such as Knapsack (Bello et al., 2017) and Job Shop Scheduling (Soloabal et al., 2020; Dai et al., 2019), or address single-graph NP-hard tasks without hard constraints, such as Minimum Vertex Cover (Dai et al., 2017), Graph Exploration (Dai et al., 2019), and Network Dismantling (Fan et al., 2020). Fundamentally different from these works, MCS detection requires a graph pair as input, and we show how to properly encode such an input. The subgraph isomorphism constraint of the task also sets us apart from the aforementioned graph tasks which we tackle via fully taking advantage of a key property of the task, i.e. bidomain, while in contrast, Soloabal et al. (2020) relies on penalty signals generated from constraint dissatisfaction in order to guide the agent to achieve feasible solutions for non-graph tasks.

6. Conclusion

We believe the interplay of search and learning is a promising research direction, and take a step towards bridging the gap by tackling the NP-hard challenging task, Maximum Common Subgraph detection. We have proposed a reinforcement learning method which unifies search and deep Q-learning into a single framework. By using the search to train our carefully designed DQN, the DQN provides better node selection policy for search to find large common subgraph solutions faster, which is experimentally verified on synthetic and real-world large graph pairs. In the future, we will explore the adaptation of our framework which combines learning with search to other constrained combinatorial problems, e.g. Maximum Clique Detection.

7. Acknowledgements

This work was partially supported by NSF IIS- 2031187, NSF DGE-1829071, NSF III-1705169, NSF CAREER Award 1741634, NSF 1937599, NIH R35-HL135772, NIH/NIBIB R01 EB027650, DARPA HR00112090027, Okawa Foundation Grant, and Amazon Research Awards.

References

- Bahiense, L., Manić, G., Piva, B., and De Souza, C. C. The maximum common edge subgraph problem: A polyhedral investigation. *Discrete Applied Mathematics*, 160 (18):2523–2541, 2012.
- Bai, Y., Ding, H., Bian, S., Chen, T., Sun, Y., and Wang, W. Simgnn: A neural network approach to fast graph similarity computation. *WSDM*, 2019.
- Bai, Y., Ding, H., Gu, K., Sun, Y., and Wang, W. Learning-based efficient graph similarity computation via multi-scale convolutional set matching. *AAAI*, 2020a.
- Bai, Y., Xu, D., Gu, K., Wu, X., Marinovic, A., Ro, C., Sun, Y., and Wang, W. Neural maximum common subgraph detection with guided subgraph extraction, 2020b. URL <https://openreview.net/forum?id=BJgcwh4FwS>.
- Balcilar, M., Renton, G., Héroux, P., Gaüzère, B., Adam, S., and Honeine, P. Analyzing the expressive power of graph neural networks in a spectral perspective. In *ICLR*, 2021. URL <https://openreview.net/forum?id=-qh0M9XWxnv>.
- Barabási, A.-L. and Albert, R. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- Bello, I., Pham, H., Le, Q. V., Norouzi, M., and Bengio, S. Neural combinatorial optimization with reinforcement learning. *ICLR*, 2017.
- Bunke, H. and Shearer, K. A graph distance metric based on the maximal common subgraph. *Pattern recognition letters*, 19(3-4):255–259, 1998.
- Cai, C. and Wang, Y. A simple yet effective baseline for non-attributed graph classification. *arXiv preprint arXiv:1811.03508*, 2018.
- Cao, N., Yang, Z., Wang, C., Ren, K., and Lou, W. Privacy-preserving query over encrypted graph-structured data in cloud computing. In *2011 31st International Conference on Distributed Computing Systems*, pp. 393–402. IEEE, 2011.
- Chen, Z., Chen, L., Villar, S., and Bruna, J. Can graph neural networks count substructures? *arXiv preprint arXiv:2002.04025*, 2020.
- Dai, H., Khalil, E. B., Zhang, Y., Dilkina, B., and Song, L. Learning combinatorial optimization algorithms over graphs. *NeurIPS*, 2017.
- Dai, H., Li, Y., Wang, C., Singh, R., Huang, P.-S., and Kohli, P. Learning transferable graph exploration. *NeurIPS*, 2019.
- Ehrlich, H.-C. and Rarey, M. Maximum common subgraph isomorphism algorithms and their applications in molecular science: a review. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 1(1):68–79, 2011.
- Fan, C., Zeng, L., Sun, Y., and Liu, Y.-Y. Finding key players in complex networks through deep reinforcement learning. *Nature Machine Intelligence*, 2(6):317–324, 2020.
- Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- Gasse, M., Chételat, D., Ferroni, N., Charlin, L., and Lodi, A. Exact combinatorial optimization with graph convolutional neural networks. *arXiv preprint arXiv:1906.01629*, 2019.
- Gilbert, E. N. Random graphs. *The Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *ICLR*, 2015.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *ICLR*, 2016.
- Levi, G. A note on the derivation of maximal common subgraphs of two directed or undirected graphs. *Calcolo*, 9(4):341, 1973.
- Li, Y., Gu, C., Dullien, T., Vinyals, O., and Kohli, P. Graph matching networks for learning the similarity of graph structured objects. *ICML*, 2019.
- Ling, X., Wu, L., Wang, S., Ma, T., Xu, F., Wu, C., and Ji, S. Hierarchical graph matching networks for deep graph similarity learning, 2020. URL <https://openreview.net/forum?id=rkeqnlrtDH>.
- Liu, X., Pan, H., He, M., Song, Y., Jiang, X., and Shang, L. Neural subgraph isomorphism counting. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1959–1969, 2020.
- Liu, Y.-l., Li, C.-m., Jiang, H., and He, K. A learning based branch and bound for maximum common subgraph problems. *IJCAI*, 2019.
- McCreesh, C., Ndiaye, S. N., Prosser, P., and Solnon, C. Clique and constraint models for maximum common (connected) subgraph problems. In *International Conference on Principles and Practice of Constraint Programming*, pp. 350–368. Springer, 2016.

- McCreesh, C., Prosser, P., and Trimble, J. A partitioning algorithm for maximum common subgraph problems. 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *NeurIPS Deep Learning Workshop 2013*, 2013.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- Park, Y., Reeves, D. S., and Stamp, M. Deriving common malware behavior through graph clustering. *Computers & Security*, 39:419–430, 2013.
- Peyré, G., Cuturi, M., and Solomon, J. Gromov-wasserstein averaging of kernel and distance matrices. In *ICML*, pp. 2664–2672, 2016.
- Solozabal, R., Ceberio, J., and Takáč, M. Constrained combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:2006.11984*, 2020.
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. *ICLR*, 2018.
- Vismara, P. and Valery, B. Finding maximum common connected subgraphs using clique detection or constraint satisfaction algorithms. In *International Conference on Modelling, Computation and Optimization in Information Systems and Management Sciences*, pp. 358–368. Springer, 2008.
- Wang, R., Yan, J., and Yang, X. Learning combinatorial embedding networks for deep graph matching. *ICCV*, 2019.
- Wang, R., Zhang, T., Yu, T., Yan, J., and Yang, X. Combinatorial learning of graph edit distance via dynamic embedding. *arXiv preprint arXiv:2011.15039*, 2020.
- Watts, D. J. and Strogatz, S. H. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440, 1998.
- Xu, H., Luo, D., and Carin, L. Scalable gromov-wasserstein learning for graph partitioning and matching. In *NeurIPS*, pp. 3046–3056, 2019a.
- Xu, H., Luo, D., Zha, H., and Carin, L. Gromov-wasserstein learning for graph matching and node embedding. *ICML*, 2019b.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? *ICLR*, 2019c.
- Yan, X., Yu, P. S., and Han, J. Substructure similarity search in graph databases. In *SIGMOD*, pp. 766–777. ACM, 2005.
- You, J., Ying, R., and Leskovec, J. Position-aware graph neural networks. In *ICML*, pp. 7134–7143. PMLR, 2019.
- Yu, T., Wang, R., Yan, J., and Li, B. Learning deep graph matching with channel-independent embedding and hungarian attention. In *ICLR*, 2020. URL <https://openreview.net/forum?id=rJgBd2NYPH>.
- Zanfir, A. and Sminchisescu, C. Deep learning of graph matching. In *CVPR*, pp. 2684–2693, 2018.

GLSearch: Maximum Common Subgraph Detection via Learning to Search

Yunsheng Bai^{*1} Derek Xu^{*1} Yizhou Sun¹ Wei Wang¹

A. Dataset Description

This section describes the datasets used for evaluating our model and baselines. Section B.1 describes the dataset we use for training GLSEARCH as well as baseline learning based graph matching models.

We use the following real-world datasets for evaluation:

- NCI109: It is a collection of small-sized chemical compounds (Wale et al., 2008) whose nodes are labeled indicating atom type. We form 100 graph pairs from the dataset whose average graph size (number of nodes) is 28.73.
- ROAD: The graph is a road network of California whose nodes indicate intersections and endpoints and edges represent the roads connecting the intersections and endpoints (Leskovec et al., 2009). The graph contains 1965206 nodes, from which we randomly sample a connected subgraph of around 0.05% nodes twice to generate two subgraphs for the graph pair $\mathcal{G}_1 = (\mathcal{V}_1, \mathcal{E}_1)$ and $\mathcal{G}_2 = (\mathcal{V}_2, \mathcal{E}_2)$.
- DBEN, DBZH, and DBPD: It is a dataset originally used in a work on cross-lingual entity alignment (Sun et al., 2017). The dataset contains pairs of DBpedia knowledge graphs in different languages. For DBEN, we use the English knowledge graph and sample 10% nodes twice to generate two graphs for our task. For DBZH, we sample around 10% nodes from the knowledge graph in Chinese. For DBPD, we sample once from the English graph to get \mathcal{G}_1 and sample once from the Chinese graph to get \mathcal{G}_2 . Note that although the nodes have features, we do not use them because our task is more about graph structural matching rather than node semantic meanings, and leave the incorporation of continuous node initial representations as future work.
- ENRO: The graph is an email communication network whose nodes represent email addresses and (undirected) edges represent at least one email sent between the addresses (Klimt & Yang, 2004). From the total 36692 nodes, we sample around 10% nodes to generate the graph pair.
- COPR: An Amazon computer product network whose nodes represent goods and edges represent two goods frequently purchased together (Shchur et al., 2018). The graph contains 703655 nodes from which we sample around 0.5% to get the pair we use.
- CIRC: This is a graph pair where each graph is a circuit diagram whose nodes represent devices/wires and edges represent the connecting relations between devices and wires. In other words, each node is either a device or a wire, and the entire graph is bipartite. The two graphs given are known to be **isomorphic**¹ and we do not perform any sampling. Nodes have labels about the type of the device/wire. In real world, the successful matching of circuit layout diagrams is an essential process in circuit design verification.
- HPPI: It is a human protein-protein interaction network whose nodes represent proteins and edges represent physical interaction between proteins in a human cell (Agrawal et al., 2018). From the 21557 nodes, we sample around 10% nodes to generate the pair used in experiments.
- ROAD-CA: We use the same road network of California as the ROAD dataset, but this time, from the 1965206 nodes, we randomly sample a connected subgraph of around 50.0% nodes twice to generate two subgraphs for the graph pair.
- ROAD-TX: Similar to ROAD-CA, but the road network is in Texas (Leskovec et al., 2009). The graph contains 1379917 nodes, from which we randomly sample a connected subgraph of around 80% nodes twice to generate two subgraphs for the graph pair.

^{*}Equal contribution ¹Department of Computer Science, University of California, Los Angeles, California, USA. Correspondence to: Yunsheng Bai <yba@ucla.edu>, Derek Xu <derekxu@ucla.edu>.

The details of all the graph pairs can be found in Table 1.

¹Section E.3 shows results on synthetic datasets where the MCS size lower bound known.

Table 1: Details of real-world graph pairs used in evaluating the performance of baseline methods and GLSEARCH.

| Name | Description | $ \mathcal{V}_1 $ | $ \mathcal{V}_2 $ | $ \mathcal{E}_1 $ | $ \mathcal{E}_2 $ | $\frac{ \mathcal{E}_1 }{ \mathcal{V}_1 }$ | $\frac{ \mathcal{E}_2 }{ \mathcal{V}_2 }$ |
|----------------|-------------------------------------|-------------------|-------------------|-------------------|-------------------|---|---|
| ROAD | Road Network | 1114 | 652 | 1454 | 822 | 1.305 | 1.261 |
| DBEN | Knowledge Graph | 1945 | 1945 | 6242 | 5851 | 3.209 | 3.008 |
| DBZH | Knowledge Graph | 1907 | 1907 | 4856 | 4948 | 2.546 | 2.595 |
| DBPD | Knowledge Graph | 1945 | 1907 | 6242 | 4856 | 3.209 | 2.546 |
| ENRO | Email Communication Network | 3369 | 3369 | 46399 | 50637 | 13.772 | 15.030 |
| COPR | Product Co-purchasing Network | 3518 | 3518 | 56028 | 40633 | 15.926 | 11.550 |
| CIRC | Circuit Layout Diagram | 4275 | 4275 | 6128 | 6128 | 1.433 | 1.433 |
| HPPI | Protein-Protein Interaction Network | 2152 | 2152 | 54910 | 54132 | 25.516 | 25.154 |
| ROAD-CA | Road Network | 978513 | 978513 | 1404115 | 1366917 | 1.435 | 1.397 |
| ROAD-TX | Road Network | 1080909 | 1080909 | 1503531 | 1507440 | 1.391 | 1.395 |

For synthetic datasets, we generate graph pairs using the Barabási-Albert (BA) (Barabási & Albert, 1999) algorithm (edge density set to 5), the Erdős-Rényi (ER) (Gilbert, 1959) algorithm (edge density set to 0.08), and the Watts–Strogatz (WS) (Watts & Strogatz, 1998) algorithm (rewiring probability set to 0.2 and ring density set to 4), respectively.

B. Details on DQN and Training GLSEARCH

B.1. Training Data Preparation: Curriculum Learning

Curriculum learning (Bengio et al., 2009) is a strategy for training machine learning models whose core idea is to train a model first using “easy” examples before moving on to using “hard” ones. Our goal is to train a general model for MCS detection task which works well on general testing graph pairs from different domains. Therefore, we employ the idea of curriculum learning in training our GLSEARCH. More specifically, we prepare the training graph pairs in the following way:

- Curriculum 1: The first curriculum consists of the easiest graph pairs that are small: (1) We sample 30 graph pairs from AIDS (Zeng et al., 2009), a chemical compound dataset usually for graph similarity computation (Bai et al., 2019) where each graph has less than or equal to 10 nodes; (2) We sample 30 graph pairs from LINUX (Wang et al., 2012), another dataset commonly used for graph matching consisting of small program dependency graphs generated from Linux kernel; (3) So far we have 60 real-world graph pairs. We then generate 60 graph pairs using popular graph generation algorithms. Specifically, we generate 20 graph pairs using the BA algorithm, 20 graph pairs using the ER algorithm, and 20 graph pairs using the Watts–Strogatz WS algorithm, respectively. Details of the graphs can be found in Table 2. In summary, the first curriculum contains 120 graph pairs in total.
- Curriculum 2: After the first curriculum, each next

curriculum contains graphs that are larger and harder to match than the previous curriculum. For the second curriculum, we sample 30 graph pairs from PTC (Shrivastava & Li, 2014), a collection of chemical compounds, 30 graph pairs from IMDB (Yanardag & Vishwanathan, 2015), a collection of ego-networks of movie actors/actresses, and generate 20 graph pairs again using the BA, ER, and WS algorithms but with larger graph sizes.

- Curriculum 3: For the third curriculum, we sample 30 graph pairs from MUTAG (Debnath et al., 1991), a collection of chemical compounds, 30 graph pairs from REDDIT (Yanardag & Vishwanathan, 2015), a collection of ego-networks corresponding to online discussion threads, and generate 20 even larger graph pairs using the BA, ER, and WS algorithms.
- Curriculum 4: For the last curriculum, we sample 30 graph pairs from WEB (Riesen & Bunke, 2008), a collection of text document graphs, 30 graph pairs from MCSPLAIN-CONNECTED (McCreech et al., 2017), a collection of synthetic graph pairs adopted by MCSP, and generate 20 graph pairs again using BA, ER, and WS algorithms but with larger graph sizes.

For each curriculum, we train the model for 2500 iterations before moving on to the next, resulting in 10000 training iterations in total.

B.2. Training Techniques and Details

B.2.1. STAGE 1: PRE-TRAINING

For the first 1250 iterations, we pre-train our DQN with the supervised true target y_t obtained as follows:

- For each graph pair, we run the complete search, i.e. we do not perform any pruning for unpromising states. The entire search space is explored, and the future

Table 2: Training graph details. For synthetic graphs, “ed”, “p”, and “rd” represent edge density, rewiring probability, and ring density, respectively.

| Curriculum | Data Source | # Pairs |
|--------------|--------------------|---------|
| Curriculum 1 | AIDS | 30 |
| | LINUX | 30 |
| | BA:n=16,ed=5 | 20 |
| | ER:n=14,ed=0.14 | 20 |
| | WS:n=18,p=0.2,rd=2 | 20 |
| Curriculum 2 | PTC | 30 |
| | IMDB | 30 |
| | BA:n=32,ed=4 | 20 |
| | ER:n=30,ed=0.12 | 20 |
| | WS:n=34,p=0.2,rd=2 | 20 |
| Curriculum 3 | MUTAG | 30 |
| | REDDIT | 30 |
| | BA:n=48,ed=4 | 20 |
| | ER:n=46,ed=0.1 | 20 |
| | WS:n=50,p=0.2,rd=4 | 20 |
| Curriculum 4 | WEB | 30 |
| | MCSPLAIN-CONNECTED | 30 |
| | BA:n=62,ed=3 | 20 |
| | ER:n=64,ed=0.08 | 20 |
| | WS:n=66,p=0.2,rd=4 | 20 |

reward for every action can be found by finding the longest path starting from the action to a terminal state. Since graphs are small in the initial stage, such complete search can be affordable. Using Figure 2 in the main text as an example, for the action that causes state 0 to transition to state 6, the longest path is 0, 6, 7, 11, 12, 13 (or 0, 6, 7, 11, 12, 14).

- Given the longest path found for each action, we then compute $y_t = 1 + \gamma + \gamma^2 + \dots + \gamma^{(L-1)}$, where γ is the discount factor set to 1.0, L is the length of the longest path. In the example above, $y_t = 5$, intuitively meaning that at state 0, for the action that leads to state 6, in future the best solution will have 5 more nodes. In contrast, the action 0 to 1 has $y_t = 4$, meaning the action 0 to 6 is more preferred.
- Given the true target computed for each action, we run the mini-batch gradient descents over the mse loss $(y_t - Q(s_t, a_t))^2$, where the batch size (number of sampled actions) is set to 32.

B.2.2. STAGE 2: IMITATION LEARNING AND STAGE 3

For stage 2 (2500 iterations) and stage 3 (6250 iterations), we train the DQN using the framework proposed in Mnih et al. (2013). The difference is that in stage 2, instead of allowing the model to use its own predicted $Q(s_t, a_t)$

at each state, we let the model make a decision using the heuristics by the MCSP algorithm, which serves as an expert providing trajectories in stage 2. We aim to outperform MCSP eventually after training using our own predicted $Q(s_t, a_t)$ in stage 3.

Here we describe the procedure of the training process. In each training iteration, we sample a graph pair from the current curriculum for which we run the DQN multiple times until a terminal state is reached to collect all the transitions, i.e. 4-tuples in the form of (s_t, a_t, r_t, s_{t+1}) where r_t is 1 and $y_t = 1 + \gamma \max_{a'} Q(s_{t+1}, a')$, and store them into a global experience replay buffer, a queue that maintains the most recent L 4-tuples. In our calculations, $L = 1024$. Afterwards, at the end of the iteration, the agent gets updated by performing the mini-batch gradient descents over the mse loss $(y_t - Q(s_t, a_t))^2$, where the batch size (number of sampled transitions from the replay buffer) is set to 32.

To stabilize our training, we adopt a target network which is a copy of the DQN network and use it for computing $\max_{a'} \gamma Q(s_{t+1}, a')$. This target network is synchronized with the DQN periodically, in every 100 iterations.

Since at the beginning of stage 3, the Q approximation may still be unsatisfactory, and random behavior may be better, we adopt the epsilon-greedy method by switching between random policy and Q policy using a probability hyperparameter ϵ . Thus, the decision is made as $\arg \max q$ where q is the our predicted $Q(s_t, a_t)$ for all possible actions (i, j) with $1 - \epsilon$ probability; With ϵ probability, the decision is random. This probability is tuned to decay slowly as the agent learns to play the game, eventually stabilizing at a fixed probability. We set the starting epsilon to 0.1 decaying to 0.01.

B.3. DQN Parameter Details

In experiments, we use SUM followed by an MLP for READOUT and 1DCONV+MAXPOOL followed by an MLP for INTERACT. Specifically, the MLP has 2 layers down-projecting the node embeddings from 64 to 32 dimensions. Notice that different types of embeddings require different MLPs, e.g. the MLP used for aggregating and generating graph-level embeddings is different from the MLP used for aggregating and generating subgraph-level embeddings.

For 1DCONV+MAXPOOL, we apply a 1-dimensional convolutional neural network to each one of two embeddings being interacted, followed by performing max pooling across each dimension in the two embeddings before feeding into an MLP to generate the final interacted embedding. Specifically, the 1DCONV contains a filter of size 3 and stride being 1. The MLP afterwards is again a 2-layer MLP projecting the dimension to 32. As shown in the main text, such learnable interaction operator brings performance gain compared

to simple summation based interaction.

The final MLP takes four components, $\mathbf{h}_G = \text{INTERACT}(\mathbf{h}_{G_1}, \mathbf{h}_{G_2})$, $\mathbf{h}_s = \text{INTERACT}(\mathbf{h}_{s1}, \mathbf{h}_{s2})$, \mathbf{h}_{Dc} , and \mathbf{h}_{D0} , each with dimension 32. It consists of 7 layers down-projecting the 128-dimensional input (32×4) to a scalar as the predicted q score. For every MLP used in experiments, all the layers except the last use the $\text{ELU}(x)$ activation function. An exception is the final MLP, whose last layer uses the $\text{ELU}(x) + 1$ as the activation function to ensure positive q output.

A subtle point to notice is the necessity of using either non-linear readout or nonlinear interaction for generating the bidomain representation. Otherwise, if both operators are a simple summation, the representation for all the connected bidomains (\mathbf{h}_{Dc}) is essentially the global summation of all nodes in all the connected bidomains. In other words, the nonlinearity of MLP in the readout operation or the interaction operator allows our model to capture the bidomain partitioning information in \mathbf{h}_{Dc} .

C. Notes on Search

C.1. Comparison with MCSP and MCSP+RL

The key idea of our model is that under a limited search budget, by exploring the most promising node pairs first, search can reach a larger common subgraph solution faster. In other words, for small graph pairs, all baseline models would obtain the exact MCS result as long as the search algorithm runs to complete, i.e. the stack is eventually empty, meaning no more actions to select and no more states to backtrack to (all states have been visited and fully expanded to all possible next states).

However, for large graph pairs, the task is NP-hard, and the complete search becomes nearly impossible. Exceptions exist though: For example, if the pruning condition based on the upper bound estimation is powerful enough to prune many states, the search may finish in relatively few iterations. However, we observe that the state-of-the-art solvers, MCSP and MCSP+RL, cannot finish completely for all the graph pairs used in testing. Instead of trying to improve the upper bound estimation to be more exact, in this paper, our goal is to learn a better node pair selection policy, to replace the heuristics used by baseline solvers.

Notice that our focus on node pair selection policy instead of upper bound estimation implies that a better selection policy would mean the search can quickly find a larger solution and update its best solution found so far $maxSol$. This not only mean when the search budget is used up, the result returned is larger, but also mean that for subsequent iterations (before the iteration limit is reached), more states would be pruned by checking $UB_t \leq |maxSol|$, thus further helping the

search. In summary, in our framework, the upper bound computation strategy remains unchanged, yet the successful node pair selection policy benefits the search in two major ways.

Since we use MCSP in the imitation learning stage of training our DQN, and compare with MCSP and MCSP+RL in the main text, we describe their node pair selection heuristics. For MCSP, when entering a new state, it first selects the node with the largest node degree in G_1 , and then enumerates through all the nodes in G_2 in descending order of node degrees. In the original implementation provided by MCSP, this is achieved by recursive function calls. After all the nodes in G_2 are visited, i.e. the depth-first search of all the node pairs (i, j) finishes where i is the largest-degree node in G_1 and j is every node in G_2 , the algorithm selects the second largest-degree node in G_1 , and repeats the enumeration of nodes in G_2 . After all node pairs are exhausted, the function returns and the algorithm essentially backtracks to the parent state. If the current state is the root node in the search tree, the search is complete and the exact MCS is returned. However, as noted earlier, for large graph pairs it is almost impractical to search exhaustively and a budget on the amount of search conducted has to be applied. Thus, which node pairs to visit first matters a lot for successfully extracting a large solution for large input graphs. However, as seen in Figures 7 and 8, in many cases the true MCS does not contain large-degree nodes, since large-degree nodes tend to form more complicated subgraphs which are harder to match in the other input graph compared to simpler subgraphs like a chain. Thus, by visiting large-degree nodes first, MCSP may not always yield a large solution fast.

In contrast, MCSP+RL maintains a promising score for each node and iteratively updates the scores as search visits more states. The update formula is based on the reduction of upper bound for search, where upper bound in an overestimation of future subgraph size. As search makes progress, the scores are updated in each iteration, and nodes which cause large reduction in upper bound computation get large reward. This has the limitation that for each new graph pair, the scores associated with each node must be re-initialized to 0 and re-learned, since there is no neural network and the only learnable parameters are the scores for each node. At the beginning of search, all scores are initialized to 0, and the search has to break the tie using another heuristic, while once trained, our GLSEARCH can be applied to any new testing pair, and at the beginning of search, the learned parameters in GLSEARCH starts to benefit the search. In other words, the whole design of MCSP+RL can be regarded as a search framework with shallow learning (without neural networks or training via back-propagation). GLSEARCH is the first model to use deep learning for node pair selection.

Another limitation of MCSP+RL is that the scores main-

tained for nodes reflect the potential ability of a node to reduce the upper bound for future iterations in the current search, which is indirect as the MCS aims to find the largest common subgraph, not the reduction of upper bound. Moreover, the upper bound itself is an overestimation of future subgraph size, which may or may not be close enough to the actual best future subgraph size. In contrast, we aim to predict the q score for actions which directly reflect the best future subgraph size. Overall, the lack of deep learning ability causes MCSP+RL not only to re-estimate the scores for the nodes for each new graph pair, but also to resort to the upper bound heuristic for updating the scores.

To ensure the budget on search iterations is applied consistently for all the models evaluated in the main text, we adapt the original recursive implementation of MCSP in C++ to an iterative implementation in Python so that all the models compare with each other in the same programming language and the same search backbone algorithm. To be specific, we check the iteration count at the beginning of every search iteration and early stop the search if the pre-defined budget is reached.

C.2. Tree vs Sequence

At this point, having illustrated the differences between GLSEARCH and MCSP and MCSP+RL, it is worth clarifying whether GLSEARCH search yields a tree or sequence in different stages. For training stage 1 and 2, as described in Section B.2, our model is just randomly initialized and not well trained, so the pre-training and imitation learning stages use the policy of MCSP instead of using its own predicted q scores. In stage 1, the complete search is performed to provide maximum amount of supervised signals, i.e. y_t , but in stage 2, we start using the RL training framework, i.e. experience replay buffer, target network, etc, so we run the agent multiple times until a terminal state is reached, corresponding to a sequence in a tree, which starts from the root node and ends at a leaf node. For stages 2 and 3, since sequences are generated instead of trees, for each sequence, the upper bound check always passes, because the pruning only happens when backtracking is allowed, i.e. a tree is formed. To see this more clearly, recall the pruning only happens if $UB_t \leq |maxSol|$ in Algorithm 1 in the main text. However, during the sequence generation process, $maxSol$ keeps increasing by one each time a new state is reached. Since $UB_t \leftarrow |curSol| + \text{overestimate}(s_t)$, $UB_t > |curSol|$ for non-terminal s_t , and $|curSol| = |maxSol|$, and thus $UB_t > |maxSol|$, and thus the pruning never happens.

At the beginning of stage 2, the sequences we collect are usually not long, since the policy is the same as MCSP in stage 2. However, in stage 3, we start using our predicted $Q(s_t, a_t)$ to get such sequences, and at the end of training, when we apply GLSEARCH to testing pairs during inference,

as shown in the main text, we perform better than MCSP and all the other baselines. In inference, we completely rely on our predicted q scores as the policy, and a search tree is yielded, although the tree is not complete since the graph pairs are large and a search budget is reached.

C.3. Terminal Conditions

This section discusses on how a terminal state, i.e. a leaf node in the search tree, is determined. Notice our definition of bidomain (equivalence class) does not include node labels, and in each iteration, we allow the matching between nodes in the same bidomain with the same node label. We consider the node labels as additional pruning on each bidomain, i.e. we further only allow nodes with the same label to match within bidomain when considering actions to be fed into DQN. Suppose \mathcal{G}_1 and \mathcal{G}_2 are connected graphs. There are two cases:

- Case 1: Nodes are unlabeled (or equivalently, all the nodes have the same label). The terminal condition is that there is no non-empty connected (adjacent) bidomains. For example, there are still some adjacent bidomains, but for each bidomain $\langle \mathcal{V}'_{k_1}, \mathcal{V}'_{k_2} \rangle$, at least one of \mathcal{V}'_{k_1} and \mathcal{V}'_{k_2} is empty (containing no nodes), so there is no nodes to match in each bidomain. Examples are states 3 and 6 in Figure 1.
- Case 2: Nodes are labeled. For the terminal condition, there may be still some non-empty connected bidomains, but the node labels do not match causing no more node pairs to select from. For example, one bidomain contains \mathcal{V}'_{k_1} with C and N as node labels and \mathcal{V}'_{k_2} with H as node label. Then essentially there is no more node pairs left.

C.4. Promise-based Search: Improving Search with Backtracking

Unlike MCSP or MCSP+RL, GLSEARCH is optimized to find the largest common subgraph in one try, not to prune the search space. This is because, in practice, even with advanced pruning techniques, it is not practical to exhaust the entire search space for large graphs. As a consequence, GLSEARCH may still fall into local solutions if it follows the branch-and-bound algorithm. Thus, GLSEARCH improves upon MCSP search by backtracking to an earlier state with the most promise of finding a larger subgraph when the current best solution is not improved upon within a fixed number of iterations. Typically, when the action space is larger, there is more potential for equally good or better actions to exist outside of the one selected, thus a state's action space size is equated to its promise.

In implementation, GLSEARCH keep track of promise by maintaining a priority queue of states, where a state's pri-

ority is given by its action space size, in parallel with the search stack. Thus, whenever GLSEARCH suspects the model is in a local solution, the next state popped on line 7 of Algorithm 1 in the main text will be the state with the largest action space from the priority queue, instead of the top state in the search stack. GLSEARCH detects when the model is in a local solution by keeping track of the largest subgraph found since the last time the priority queue was popped (or since the start of search). If this local best solution is not improved within a fixed number of iterations, the model knows it is in a local solution. In practice, we set this number to 3.

C.5. Notes on Equivalent States and Multiple Ground Truths

It is well known that for the MCS detection task, there can be multiple ground truth solutions with the same subgraph size. For example, in Figure 1, both states 3 and 6 correspond to the same subgraph size but the states 3 and 6 are different due to their different node-node mappings. Our model maintains the node-node mappings for each state, and therefore states 3 and 6 would be reached as different states. It is important to note that for large graph pairs, reaching both states 3 and 6 usually does not happen, since the search would first reach state 3 and need many iterations to backtrack to state 0 and further many iterations to reach state 6.

There is an even more subtle point in Figure 1. Suppose the search first matches node a to node 1, denoted as state 1, then matches node b to node 2, leading to state 2. After several iterations, it backtracks to state 0 and chooses to match node b to node 2, denoted as state x, then matches node a to node 1, denoted as state y. Although states 1 and x are different, but states 2 and y are equivalent, since both states 2 and y have the same node-node mapping, i.e. a to 1 and b to 2. Thus, the search maintains an additional set of visited states and at each time a new state is reached, a checking is performed to avoid revisiting the same state twice.

The node-node mapping is an important component of the definition of state, not only because it differentiates the otherwise equivalent states, but also because different node-node mappings can lead to different final states and future reward (and thus must be considered by the design of DQN). Suppose the bidomain “01” in state 1 contains more than two nodes, i.e. there are many nodes connected to b in \mathcal{G}_1 besides node d and many nodes connected to node 2 in \mathcal{G}_2 besides node 3. Then state 2 is an intuitively more preferred state compared to state 5, since the matching of node b to node 2 allows more node pairs to be matched to each other in future, thus a larger action space in state 2. The value associated with state 2 thus should be larger than state 5.

C.6. Dealing with Large Action Space

For large graph pairs, the successful detection of MCS not only depend on the design of our critical DQN component as well as its training, but also rely on techniques which prune the large action space at each state.

The bidomain partitioning idea has been outlined in the main text which effectively reduces the action space size by only matching nodes in the same bidomain. However, for extremely large and dense graphs, the bidomains may not split the rest of the graphs enough and the action space may still be too large. For example, consider two fully connected graphs \mathcal{G}_1 and \mathcal{G}_2 , i.e. for every two nodes there is an edge. Then initially, there is no nodes selected, and there is only one bidomain consisting of all the node pairs. What is worse, at any state, there is always only one bidomain consisting of all the node pairs in the remaining subgraphs. Therefore, to reduce the action space further, we only compute the q scores for N_d nodes at most in each state. Specifically, we first sort the candidate bidomains by their size in ascending order, and select the first N_b small bidomains. Next, we sort the nodes in each bidomain by degree in descending order and select the first N_d/N_b nodes with large node degrees. For our experiments, $N_b = 1$ and $N_d = 20$ in GLSEARCH; $N_b = 1$ and $N_d = 3$ in GLSEARCH-SCAL. We find, in practice, these settings do not drastically alter performance.

In future, additional techniques for pruning the action space can be explored. For example, instead of bounding the computation to be N_d , we may perform a hierarchical graph matching by first running a clustering algorithm and then matching clusters which will be bounded by the number of clusters. Another possible direction is to learn an additional Q function learning which node is more promising instead of which node pair, i.e. $Q(s_t, a_t^{(i)})$ for node i and $Q(s_t, a_t^{(j)})$ for node j in the action $a_t = (i, j)$. We suppose such additional Q function may bring further performance gain.

C.7. Analysis of Time Complexity

Overall the branch-and-bound search has exponential worst-case time complexity due to the NP-hard nature of exact MCS detection, and our goal is to use additional overhead per search iteration to make “smarter” decision each iteration so that we can find a larger common subgraph faster (in less iterations *and* real running time). Per iteration, our model requires the neural network operations to compute a Q score instead of simply using a degree heuristic which is $\mathcal{O}(1)$. Here we analyze the time complexity of these neural operations:

- To compute the node embeddings, the complexity is the same as the GNN model, which in our case is $\mathcal{O}(|\mathcal{V}| + |\mathcal{E}|)$ for GAT (since nodes must aggregate embeddings from neighbors and attention scores must be

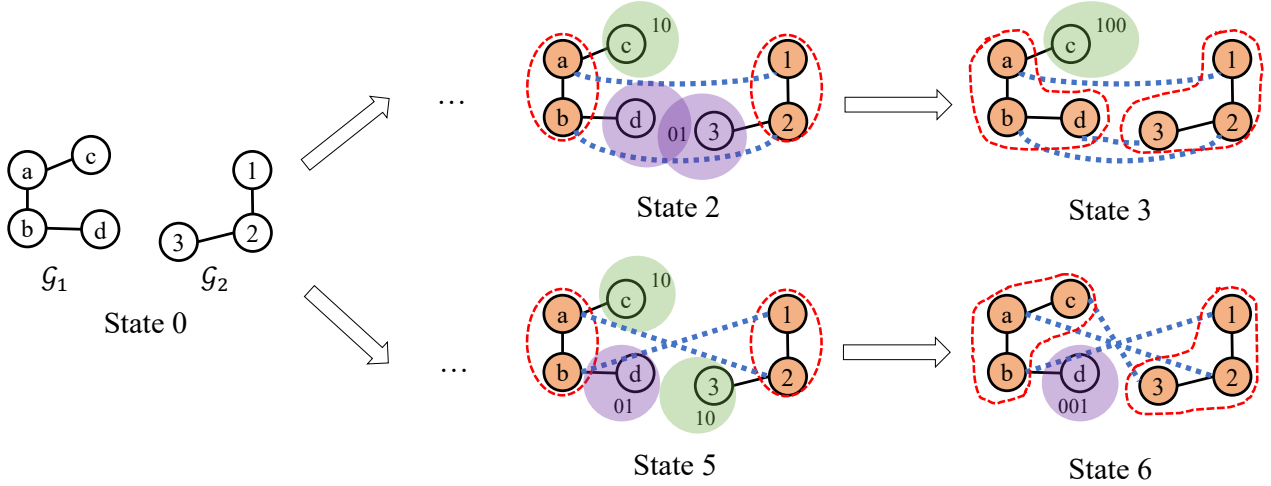


Figure 1: An example illustrating the idea of equivalent states. It is important to note that states 2, 3, 5, 6 are different since their node-node mappings are different. However, the solutions derived from both states 3 and 6 have the same subgraph size, 3. In other words, there can be multiple ways to arrive at the same solution size, with different underlying sequential processes to reach the final states.

computed for each edge). Notice the node embeddings are computed by local neighborhood aggregation, and will not be updated in search, and therefore we compute the node embeddings only once at the beginning of search, and can be cached for efficiency.

- At each iteration, to compute a Q score for a state-action pair, we run Equation (3) (in the main text) which requires computing the whole-graph, subgraph, and bidomain embeddings. Overall the time complexity for each state-action pair is $\mathcal{O}(|\mathcal{V}| - |\mathcal{V}_s|)$ where \mathcal{V}_s is the number of nodes in the currently matched subgraph. The whole-graph embeddings do not change across search, so they only need to be computed once at the beginning. The subgraph embeddings can be maintained incrementally, i.e. adding new node embeddings as search grows the subgraph. The bidomain embeddings are computed via a series of READOUT and INTERACT operations (Equation (2)): For READOUT: We use summation followed by MLP so the runtime is $\mathcal{O}(|\mathcal{V}| - |\mathcal{V}_s|)$; For INTERACT: We use a 1D CNN followed by MLP which depends on the embedding dimension set to a constant, and does not depend on the number of nodes in the input graphs.

Overall the time complexity for each iteration is $\mathcal{O}(N_d^2(|\mathcal{V}| - |\mathcal{V}_s|))$.

D. Baseline Description and Comparison

For all the models used in experiments, we evaluate their performance under the same search framework, i.e. with consistent search iteration counting, upper bound estima-

tion, etc. MCS and MCS+RL use heuristics to select node pairs, which is ineffective as shown in the main text and has been described in Section C.1. Therefore, this Section focuses on the comparison with the rest baselines, i.e. GW-QAP (Xu et al., 2019), I-PCA (Wang et al., 2019), and NEURALMCS (Bai et al., 2020).

GW-QAP performs Gromov-Wasserstein discrepancy (Peyré et al., 2016) based optimization for each graph pair and outputs a matching matrix \mathbf{Y} for all node pairs indicating the likelihood of matching which is treated the same way as our q scores, i.e. at each search iteration we index into \mathbf{Y} to select a node pair. I-PCA and NEURALMCS also output a matching matrix but require supervised training, and thus are trained using the same training data graph pairs as our GLSEARCH but with different loss functions and training signals. During testing, we apply the trained model on all testing graph pairs. For medium-size synthetic and real-world testing graph pairs, each method is given a budget of 500 search iterations. For large real-world graph pairs, each method is given a budget of 7500 search iterations. For million-node real-world graph pairs, each method is given a budget of 50 minutes. These budgets were chosen based on when the models’ performances stabilized. We then describe each method in more details.

D.1. GW-QAP

GW-QAP is a state-of-the-art graph matching model for general graph matching. The task is not about MCS specifically, but instead about matching two graphs with its own criterion based on the Gromov-Wasserstein discrepancy (Peyré et al., 2016). Therefore, we suppose the matching matrix \mathbf{Y} generated for each graph pair can be used as a guidance

for which node pairs should be visited first. In other words, we pre-compute the matching scores for all the node pairs before the search starts, and in iteration, we look up the matching matrix and treat the score as the q score for action selection. I-PCA and NEURALMCS essentially compute a matching matrix too, and it is worth mentioning that all the three methods cannot learn a score based on both the state and the action. They can be regarded as generating the matching scores based on the whole graphs only without being conditioned and dynamically updated on states and actions.

D.2. I-PCA

I-PCA is a state-of-the-art image matching model, where each image is turned into a graph with techniques such as Delaunay triangulation (Lee & Schachter, 1980). It utilizes similarity scores and normalization to perform graph matching. We adapt the model to our task by replacing these layers with 3 GAT layers, consistent with GLSEARCH. As the loss is designed for matching image-derived graphs, we alter their loss functions to binary cross entropy loss similar to NEURALMCS which will be detailed below.

D.3. NEURALMCS

NEURALMCS is proposed for MCS detection with similar idea from I-PCA that a matching matrix is generated for each graph pair using GNNs. However, they both require the supervised training signals, i.e. the complete search for training graph pairs must be done to guide the update of I-PCA and NEURALMCS. In contrast, GLSEARCH is trained under the RL framework which does not require the complete search (in stage 2 and 3, only sequences are generated as detailed in Section C.2). This has the benefit of exploring the large action space in a “smarter” way and eventually allows our model to outperform I-PCA and NEURALMCS. In implementation, the complete search is not possible for large training graph pairs, so instead we apply a search budget and use the best solution found so far to guide the training of I-PCA and NEURALMCS.

Regarding the subgraph extraction strategy, for all the baselines, we use the same branch and bound algorithm, which is the state-of-the-art search designed for MCS (McCreesh et al., 2017). However, as mentioned in Section C.4, only our model is equipped with the ability to backtrack in a principled way. The main text shows the performance gain to GLSEARCH brought by the backtracking ability.

E. More Results with Analysis

E.1. Best Solution Sizes across Time

Figure 2 shows that under the budget we set, for the large real-world graph pairs, all the methods reach a “steady state” where the best solution found so far no longer grows. This means the search continues but the search cannot find a larger solution, illustrating the fact that search gets “stuck”. In theory, given infinitely many iterations, all the models will eventually find the true MCS, which is the largest, but since the task is NP-hard, the search space can be exponential in the worst case (subject to pruning but in all the testing graph pairs the search has not finished yet), such budget has to be applied to search. At the end of the budget iterations, though, all the models have made “mistakes” (visiting unpromising states) and in order to find even larger common subgraphs, the search needs to backtrack potentially many times to fix those “mistakes” (backtracking to a very early state).

Admittedly, there is always the possibility that for more iterations, some baseline method may find a larger solution. Besides, in each iteration GLSEARCH does take more running time overhead as shown in Table 3. However, the point of our model is to quickly find a larger solution in as few iterations as possible, not to find a large solution given too many iterations. In other words, the goal of GLSEARCH is to be **smart** enough to **quickly** find a large solution instead of purely finding a good solution. Figure 2 shows that GLSEARCH not only finds solutions larger than baselines *when* the iteration budget is reached but also finds larger solutions faster than baselines *before* the budget is reached.

In addition to reaching a larger solution in less iterations, GLSEARCH also reaches better solutions with respect to runtime, as shown in Figure 3.² Notice, GLSEARCH finds the same large subgraph in 10 minutes as in 7500 iterations. Although per iteration, it is slower than MCS^P per iteration, GLSEARCH finds a larger solution in usually less than a minute. Moreover, our implementation can be further optimized. Note, we adapted all baselines to run on Python for fair comparison, and made sure the search iteration counting is consistent across all baselines and the results shown in the main text. At this stage, our main goal is to explore the idea of “*learning to search*”, which has been experimentally verified to be a promising direction of research, and leave the efforts of implementation optimization using various techniques as future focus.

²In general, we refer to the running time results to show the performance gains of GLSEARCH and GLSEARCH-SCAL; however, to focus solely on the effects of different actions selected by different policies on the final common subgraph, we also compare by iterations.




Figure 2: For each method, we maintain the best solution found so far in each iteration during the search process. We plot the size of the largest extracted common subgraphs found so far vs search iteration count for all the methods across all the datasets. The larger the subgraph size, the better (“smarter”) the model in terms of quickly finding a large MCS solution under limited budget for large graphs.

GLSearch: Maximum Common Subgraph Detection via Learning to Search

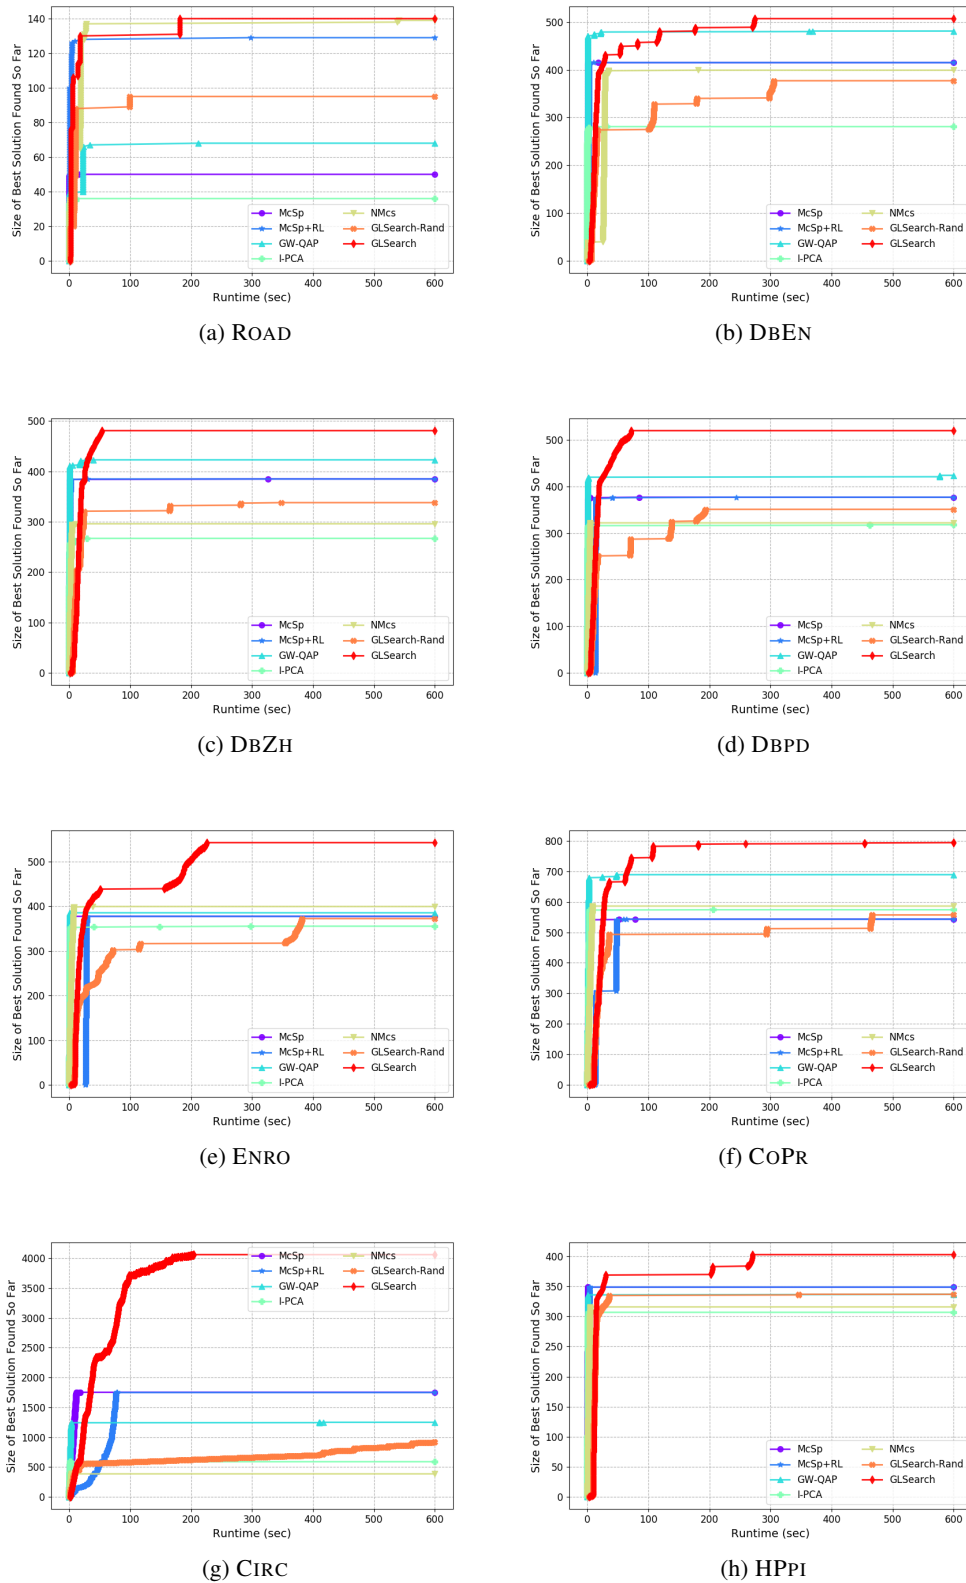


Figure 3: For each method, we maintain the best solution found so far in each iteration during the search process. We plot the size of the largest extracted common subgraphs found so far vs the real running time for all the methods across all the datasets. The larger the subgraph size, the better (“smarter”) the model in terms of quickly finding a large MCS solution under limited budget for large graphs.

Table 3: Average running time per iteration (msec).

| Method | ROAD | DBEN | DBZH | DBPD | ENRO | CoPR | CIRC | HPPI |
|---------------|--------|--------|--------|--------|---------|--------|---------|--------|
| MCSP | 2.040 | 10.724 | 1.415 | 0.974 | 1.722 | 2.891 | 1.776 | 0.498 |
| MCSP+RL | 0.894 | 6.834 | 2.103 | 1.247 | 2.166 | 3.107 | 2.080 | 0.559 |
| GW-QAP | 0.548 | 0.834 | 0.546 | 4.692 | 1.041 | 3.419 | 1.550 | 0.546 |
| I-PCA | 1.152 | 1.797 | 0.967 | 0.897 | 1.739 | 2.725 | 3.792 | 0.636 |
| NEURALMCS | 2.394 | 4.172 | 4.648 | 5.667 | 9.610 | 9.788 | 7.471 | 15.742 |
| GLSEARCH-RAND | 17.392 | 66.418 | 67.342 | 67.946 | 163.005 | 71.972 | 655.447 | 83.488 |
| GLSEARCH | 8.132 | 66.552 | 71.409 | 96.262 | 135.087 | 51.181 | 37.377 | 60.509 |

E.2. Additional Ablation Study

Table 4 shows that pre-training and imitation learning benefit the performance under four out of the eight datasets. On ENRO and HPPI, without pre-training, our model performs better, which may be attributed to the fact that they are dense graphs (Table 1) while the training graphs used in stage 1 are relatively small and sparse (Section B.1).

Table 5 shows that the promised-based search improves the performance under four out of the eight datasets. For the other four datasets, the performance does not change, indicating that the backtracking to an earlier promising state based on the DQN output at least does not hurt the performance. In the cases like ENRO and HPPI, whose average node degrees are large, the promise-based search improves the performance by a large amount, showing the usefulness of the proposed strategy.

E.3. Results on Graph Pairs with Known MCS Size Lower Bound

To better understand the quality of subgraphs found by GLSEARCH, we construct new datasets with known lower bound MCS sizes. This is accomplished through generating 2 new graphs that share a common subgraph from the existing large real-world graph datasets. For each real-world graph, \mathcal{G}_0 , we randomly extract 3 different subgraphs of the same size, \mathcal{S}_0 , \mathcal{S}_1 , and \mathcal{S}_2 , by running breadth first search from 3 different starting nodes and extracting the explored induced subgraph. To construct a new graph pair, $(\mathcal{G}_1, \mathcal{G}_2)$, we form \mathcal{G}_1 by connecting \mathcal{S}_0 to \mathcal{S}_1 with 20 random edges and form \mathcal{G}_2 by connecting \mathcal{S}_0 to \mathcal{S}_2 with 20 random edges. Thus, connections between \mathcal{S}_0 nodes are the same in both \mathcal{G}_1 and \mathcal{G}_2 , but connections between $\mathcal{G}_1 \setminus \mathcal{S}_0$ and $\mathcal{G}_2 \setminus \mathcal{S}_0$ are different. Notice, the lower bound of MCS size in these new datasets would be $|\mathcal{S}_0|$, and we name the new dataset by adding ‘ss’ to the parent dataset’s name.

The results on CIRC and these datasets (Table 6) show no MCS method guarantees to always detect a solution that is as large as the known MCS solution/lower bound. This suggests the difficulty of the task itself. In practice, though, GLSEARCH is still preferred compared to baselines due to its better performance.

E.4. Result Visualization on ROAD-CA and ROAD-TX

For the largest two graph pairs, ROAD-CA and ROAD-TX, in order to clearly see and compare the subgraph growth across time of GLSEARCH-SCAL and MCSP, we perform the following visualization: At every 1000 or 2000 iterations, we plot the graph pair and highlight the matched subgraphs. As shown in Figures 4 and 5, from the left to the right, the growing of the extracted common subgraphs can be seen.

In order to render the best visualization, the following techniques are used: (1) Since the input graphs are two large, we only plot the extracted subgraphs and the remaining graphs around the extracted subgraphs by performing breadth-first search starting from the extracted subgraphs (gray color); (2) For the matched subgraphs, to clearly see the node-node mapping, we ensure the node layout positions are the same across \mathcal{G}_1 (top) and \mathcal{G}_2 (bottom), and use colors for the matched subgraph nodes to indicate the node-node mapping³.

It is noteworthy that on ROAD-TX, MCSP only finds 30 nodes as shown in Figure 5. By examining the plot carefully, we can see that this is caused by poor choice of actions that make further selections of node pairs impossible, i.e. there is no more node pairs to choose in action space, reaching a terminal condition⁴. As shown in Figure 4 in the main text, MCSP spends the rest of the time backtracking and exploring the rest of the search space, but given the exponentially growing search space size, MCSP cannot easily leave the local optimum. Given infinitely long running time, however, all methods under the branch and bound algorithm will eventually find the exact MCS solution, which is an impractical assumption in real world. This illustrates the necessity of making smart node pair selection choices through the search instead of relying on heuristics. As a fact, MCSP+RL also finds 30 nodes, since it degenerates

³The nodes of ROAD-CA and ROAD-TX are unlabeled and the colors are only used to highlight node-node mapping.

⁴We aim to find induced common subgraphs, meaning that if a new node pair is selected, all the edges between the new nodes and the existing nodes must also be included, and in this example, any new node pair would lead to the resulting subgraphs **not** isomorphic to each other.

Table 4: Contribution of pre-training and imitation learning to the performance of GLSEARCH. “no-sup” denotes the removal of the pre-training stage (The first 3750 iterations: IL; The last 6250 iterations: Normal QDN training); “no IL” denotes the removal of the imitation learning stage (The first 3750 iterations: pre-training; The last 6250 iterations: normal DQN training); “no sup; no IL” indicates the entire training (10000 iterations) is normal DQN training.

| Method | ROAD | DBEN | DBZH | DBPD | ENRO | CoPr | CIRC | HPPI |
|--------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| GLSEARCH (no sup) | 0.557 | 0.957 | 0.946 | 0.904 | 1.000 | 0.999 | 1.000 | 1.000 |
| GLSEARCH (no IL) | 1.000 | 0.933 | 0.965 | 0.887 | 0.357 | 0.875 | 0.666 | 0.632 |
| GLSEARCH (no sup; no IL) | 0.678 | 0.907 | 0.896 | 0.837 | 0.401 | 0.949 | 0.949 | 0.651 |
| GLSEARCH | 0.879 | 1.000 | 1.000 | 1.000 | 0.412 | 1.000 | 0.855 | 0.688 |
| BEST SOLUTION SIZE | 149 | 486 | 465 | 471 | 1318 | 790 | 4112 | 587 |

Table 5: Contribution of promise-based search (Section C.4) to the performance of GLSEARCH. “no promise” denotes that the search does not use the proposed promise-based search, i.e. it does not backtrack to an earlier state if the search makes no progress after a certain amount of iterations, and instead, it continues the regular branch and bound search.

| Method | ROAD | DBEN | DBZH | DBPD | ENRO | CoPr | CIRC | HPPI |
|-----------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| GLSEARCH (no promise) | 0.879 | 1.000 | 1.000 | 1.000 | 0.412 | 1.000 | 0.855 | 0.688 |
| GLSEARCH | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 |
| BEST SOLUTION SIZE | 131 | 508 | 482 | 521 | 543 | 791 | 3515 | 404 |

to the same policy as MCSP at the beginning of each new graph pair as mentioned in the main text. We verify that its found subgraphs are indeed exactly the same as MCSP.

Another insight from Figure 5 is that, the initial node pair selection made by MCSP, which is highlighted as “nodes with high degrees” in the plot, misleads MCSP into eventually finding a very small solution. In contrast, GLSEARCH-SCAL uses embeddings at subgraph, whole-graph, and bidomain levels to make a decision at each step, which capture the network structure better than just the node degree information used by MCSP. A real-world analogy can be road networks in downtown areas which tend to be grid structures versus road networks in rural areas which tend to be less rigid. GLSEARCH-SCAL takes graph structure into account and matches one downtown area with another downtown area, so the resulting matched subgraphs can be very large, while MCSP is misled by its heuristic into matching two high-degree nodes⁵ at the beginning of search, but unfortunately the two areas do not have similar road network structures and eventually the matched subgraphs are small. It must be noted, however, that this analogy is only a high-level hypothesis for this specific case, and in general, for more complicated graph structures, the actual decisions made by GLSEARCH-SCAL usually cannot be easily explained using such an analogy.

F. Extensions of GLSEARCH

GLSEARCH can be extended for a flurry of other MCS definitions, e.g. approximate MCS, MCS for weighted and

⁵The node degree of both nodes is 12 and we verify they are the highest-degree nodes in \mathcal{G}_1 and \mathcal{G}_2 .

directed graphs, etc. via a moderate amount of change to the search and learning components. In this section, we briefly outline what could be done for these tasks.

For approximate MCS detection, the bidomain constraint must be relaxed. One method of relaxing this constraint is to allow sets of nodes belonging to different but similar bidomains to match to each other. For instance, nodes in \mathcal{G}_1 from the bidomain of bitstring “00110” could map with nodes in \mathcal{G}_2 from the bidomain of bitstring “00111”, since they are only 1 hamming distance away. Such relaxations as this can be made stricter or looser based on the application. The difference would be the search framework, thus the learning part of GLSEARCH can largely stay the same.

Regarding MCS for graphs with non-negative edge weights, assuming our task is to maximize the sum of edge weights in the MCS, instead of defining $r_t = 1$, we can alter the reward function to be the difference of the sum of edge weights before and after selecting a node pair $r_t = \sum_{e \in S_t^{(u,v)}} w(e) - \sum_{e \in S_t} w(e)$ where S_t is the edges of currently selected subgraph, $S_t^{(u,v)}$ is the edges of the subgraph after adding node pair, (u, v) , and $w(\cdot)$ is a function that takes an edge and returns its weight. As the cumulative sum of rewards at step T is the sum of edge weights $\sum_{t \in [1, \dots, T]} r_t = \sum_{e \in S_T} w(e)$ and reinforcement learning aims to maximize the cumulative sum of rewards, we can adapt GLSEARCH to optimize for MCS problems with weighted edges.

Regarding MCS for directed graphs, the bidomain constraint may be altered such that every bit in the bidomain string representations now has 3 states: ‘0’ for disconnected, ‘1’ for connected by in-edge, and ‘2’ for connected by out-edge.

Table 6: Results on graph pairs with a common core subgraph (lower bound of MCS), with a fixed runtime of 10 minutes.

| Method | ROAD-ss | DBEN-ss | DBZH-ss | ENRO-ss | COPR-ss | HPPI-ss |
|-------------------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | 444 | 778 | 762 | 1346 | 1406 | 860 |
| MCSP | 0.588 | 0.466 | 0.544 | 0.216 | 1.000 | 0.233 |
| MCSP+RL | 0.588 | 0.466 | 0.544 | 0.214 | 1.000 | 0.233 |
| GLSEARCH | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 |
| BEST SOLUTION SIZE | 188 | 389 | 350 | 673 | 703 | 430 |
| CORE (LOWER BOUND) SIZE | 222 | 389 | 381 | 673 | 703 | 430 |

By considering the inward/outward direction of a bitstring, we can guarantee the isomorphism of directed graphs. In this case, the search framework would only differ in how bidomains are partitioned. The learning part of GLSEARCH would stay the same for this application.

Regarding MCS for chemical compounds, we are aware of works that aim to tackle MCS in the chemoinformatics domain (Schietgat et al., 2013; Duesbury et al., 2018), and we do admit that the current definition of MCS may not satisfy constraints in the chemoinformatics domain, e.g. Figure 6. However, since we aim to design a general solver, we believe our current work has strong potential to be extended in the future with domain-specific constraints.

More generally, we believe that there are many more extensions to GLSEARCH in addition to the ones listed, such as disconnected MCS, network alignment, or subgraph extraction. Further exploration of these are to be done as future efforts.

G. Additional Result Visualization

We plot the testing graph pairs and the results of MCSP and GLSEARCH in this Section using a software called Gephi (Bastian et al., 2009). For all the figures except Figure 19, we use two colors for nodes, one for the selected subgraphs by the model, the other for the remaining subgraphs that cannot be further matched within the search budget. When plotting, we use larger circle size for nodes with larger degrees.


In general, GLSEARCH is a less interpretable but more powerful method compared to heuristic baselines. That said, GLSEARCH presents some insights that may be useful for producing new hand-crafted heuristics.

GLSEARCH identifies “smart” nodes which can lead to larger common subgraphs faster. For example, in the road networks (ROAD), as in Figure 7 and 8, our learned policy selects nodes with smaller degrees which allow for easier matching. The common subgraphs in road networks are most likely long chains, where nodes tend to have low degrees. In contrast MCSP always chooses high-degree nodes first leading to smaller extracted subgraphs.

GLSEARCH identifies “smart” *matching* of nodes which can lead to larger common subgraphs faster. For example, in the circuit graph (CIRC), we find 3 high-degree nodes that, when correctly matched, greatly reduces the matching difficulty of remaining nodes (see Figure 19 and 20). Upon further analysis, MCSP incorrectly matches the 3 high degree nodes (matching high degree node to low degree node). This happens when matching high-degree node correctly would break the isomorphism constraint (due to the current selected subgraph being incorrectly matched). GLSEARCH conscientiously adds node pairs so that it will always be able to match the 3 high degree nodes correctly.

We believe two aspects of GLSEARCH design lead to this phenomenon. First, GLSEARCH encodes neighborhood structures that are k-hop away. MCSP only looks at a single node and not its relationship with k-hop neighbors. Second, GLSEARCH considers scores on the node-node pair granularity, thus it will only match nodes with similar local neighborhoods. MCSP only considers scores on the node granularity, potentially matching 2 nodes with dissimilar neighborhoods together.

From these insights, one can potentially design a heuristic to first detect highly valuable nodes and guide a policy which prioritizes the matching of these critical nodes, or create better heuristics that consider not only uses the features of a single node but also the similarity between the 2 nodes being matched.



Subgraphs found by GLSearch-Scal

Time →

Subgraphs found by McSp






Figure 4: Visualization of subgraphs found by GLSEARCH-SCAL and MCSP on ROAD-CA. The subgraphs found by each method grows across time (until the budget of 50 minutes is reached), and the sizes of the subgraphs are denoted at the bottom of each figure. MCSP only finds 2609 nodes as shown in Figure 4 (a) in the main text, and the solution does not increase further.



Subgraphs found by GLSearch-Scal



Subgraphs found by McSp

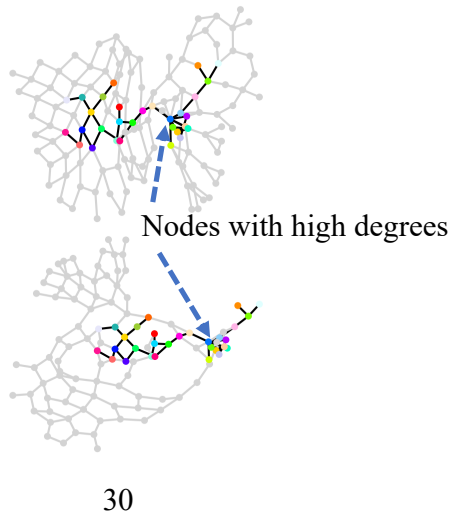


Figure 5: Visualization of subgraphs found by GLSEARCH-SCAL and MCSP on ROAD-TX. The subgraphs found by each method grows across time (until the budget of 50 minutes is reached), and the sizes of the subgraphs are denoted at the bottom of each figure. MCSP only finds 30 nodes as shown in Figure 4 (b) in the main text, and the solution does not increase further.




Figure 6: Visualization of 5 sampled graph pairs with the MCS results by GLSEARCH on NCI109. Each chemical compound node has its label indicated in the plot. Extracted subgraphs are highlighted in green.




Figure 7: Visualization of MCS result on ROAD. Extracted subgraphs are highlighted in green.




Figure 8: Visualization of GLSEARCH result on ROAD. Extracted subgraphs are highlighted in green.




Figure 9: Visualization of MCSP result on DBEN. Extracted subgraphs are highlighted in blue.




Figure 10: Visualization of GLSEARCH result on DBEN. Extracted subgraphs are highlighted in blue.




Figure 11: Visualization of MCS result on DBZH. Extracted subgraphs are highlighted in pink.




Figure 12: Visualization of GLSEARCH result on DBZH. Extracted subgraphs are highlighted in pink.




Figure 13: Visualization of MCSF result on DBPD. Extracted subgraphs are highlighted in purple.




Figure 14: Visualization of GLSEARCH result on DBPD. Extracted subgraphs are highlighted in purple.




Figure 15: Visualization of MCSP result on ENRO. Extracted subgraphs are highlighted in blue.




Figure 16: Visualization of GLSEARCH result on ENRO. Extracted subgraphs are highlighted in blue.




Figure 17: Visualization of MCSP result on CoPR. Extracted subgraphs are highlighted in blue.




Figure 18: Visualization of GLSEARCH result on CoPR. Extracted subgraphs are highlighted in blue.




Figure 19: Visualization of the original graph pair of CIRC. The two graphs are in fact isomorphic. Different colors denote different node labels. There are 6 node labels in total: M (71.67%), null (10.41%), PY (9.1%), NY (8.23%), N (0.37%), and P (0.21%).




Figure 20: Visualization of MCS SP result on CIRC. Extracted subgraphs are highlighted in yellow.




Figure 21: Visualization of GLSEARCH result on CIRC. Extracted subgraphs are highlighted in yellow.




Figure 22: Visualization of MCS result on HPPI. Extracted subgraphs are highlighted in cyan.




Figure 23: Visualization of GLSEARCH result on HPPI. Extracted subgraphs are highlighted in cyan.

References

- Agrawal, M., Zitnik, M., Leskovec, J., et al. Large-scale analysis of disease pathways in the human interactome. In *PSB*, pp. 111–122. World Scientific, 2018.
- Bai, Y., Ding, H., Bian, S., Chen, T., Sun, Y., and Wang, W. Simgnn: A neural network approach to fast graph similarity computation. *WSDM*, 2019.
- Bai, Y., Xu, D., Gu, K., Wu, X., Marinovic, A., Ro, C., Sun, Y., and Wang, W. Neural maximum common subgraph detection with guided subgraph extraction, 2020. URL <https://openreview.net/forum?id=BJgcwh4FwS>.
- Barabási, A.-L. and Albert, R. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- Bastian, M., Heymann, S., and Jacomy, M. Gephi: an open source software for exploring and manipulating networks. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 3, 2009.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *ICML*, pp. 41–48, 2009.
- Debnath, A. K., Lopez de Compadre, R. L., Debnath, G., Shusterman, A. J., and Hansch, C. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of medicinal chemistry*, 34(2):786–797, 1991.
- Duesbury, E., Holliday, J., and Willett, P. Comparison of maximum common subgraph isomorphism algorithms for the alignment of 2d chemical structures. *ChemMedChem*, 13(6):588–598, 2018.
- Gilbert, E. N. Random graphs. *The Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.
- Klimt, B. and Yang, Y. Introducing the enron corpus. In *CEAS*, 2004.
- Lee, D.-T. and Schachter, B. J. Two algorithms for constructing a delaunay triangulation. *International Journal of Computer & Information Sciences*, 9(3):219–242, 1980.
- Leskovec, J., Lang, K. J., Dasgupta, A., and Mahoney, M. W. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1):29–123, 2009.
- McCreesh, C., Prosser, P., and Trimble, J. A partitioning algorithm for maximum common subgraph problems. 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *NeurIPS Deep Learning Workshop 2013*, 2013.
- Peyré, G., Cuturi, M., and Solomon, J. Gromov-wasserstein averaging of kernel and distance matrices. In *ICML*, pp. 2664–2672, 2016.
- Riesen, K. and Bunke, H. Iam graph database repository for graph based pattern recognition and machine learning. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*, pp. 287–297. Springer, 2008.
- Schiethat, L., Ramon, J., and Bruynooghe, M. A polynomial-time maximum common subgraph algorithm for outerplanar graphs and its application to chemoinformatics. *Annals of Mathematics and Artificial Intelligence*, 69(4):343–376, 2013.
- Shchur, O., Mumme, M., Bojchevski, A., and Günnemann, S. Pitfalls of graph neural network evaluation. *Relational Representation Learning Workshop (R2L 2018), NeurIPS 2018*, 2018.
- Shrivastava, A. and Li, P. A new space for comparing graphs. In *Proceedings of the 2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pp. 62–71. IEEE Press, 2014.
- Sun, Z., Hu, W., and Li, C. Cross-lingual entity alignment via joint attribute-preserving embedding. In *International Semantic Web Conference*, pp. 628–644. Springer, 2017.
- Wale, N., Watson, I. A., and Karypis, G. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14(3):347–375, 2008.
- Wang, R., Yan, J., and Yang, X. Learning combinatorial embedding networks for deep graph matching. *ICCV*, 2019.
- Wang, X., Ding, X., Tung, A. K., Ying, S., and Jin, H. An efficient graph indexing method. In *ICDE*, pp. 210–221. IEEE, 2012.
- Watts, D. J. and Strogatz, S. H. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440, 1998.
- Xu, H., Luo, D., and Carin, L. Scalable gromov-wasserstein learning for graph partitioning and matching. In *NeurIPS*, pp. 3046–3056, 2019.
- Yanardag, P. and Vishwanathan, S. Deep graph kernels. In *SIGKDD*, pp. 1365–1374. ACM, 2015.

Zeng, Z., Tung, A. K., Wang, J., Feng, J., and Zhou, L.
Comparing stars: On approximating graph edit distance.
PVLDB, 2(1):25–36, 2009.