

# SIREN: SUBGRAPH ISOMORPHISM VIA REINFORCEMENT ENHANCED GRAPH NEURAL NETWORKS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

The subgraph isomorphism problem comprises two distinct objectives: (a) Existence determination: Verifying whether an input graph contains a subgraph isomorphic to another input graph; and (b) Complete solution enumeration: Outputting the exhaustive set of all isomorphic mappings when they exist. Solving this problem serves as a fundamental requirement for numerous application domains. However, as an NP-complete problem, existing mainstream solvers primarily rely on heuristic techniques, demonstrating limited efficiency when handling large-scale input graphs. To address this challenge, we propose SIREN - a graph neural network enhanced with deep reinforcement learning for subgraph isomorphism resolution. SIREN establishes graph embeddings through partial order-aware GNNs, while employing Deep Q-Networks with bidomain-based pruning to accelerate the graph matching process. Experimental results on real-world datasets demonstrate that SIREN achieves 100% precision with modest computational time, outperforming AI-based approximate matching methods. Compared to state-of-the-art exact solvers, SIREN delivers  $\sim 52\%$  faster execution than leading AI approaches and  $\sim 21\%$  acceleration over top heuristic methods.

## 1 INTRODUCTION

The representation of structured data using graphs has evolved over decades as a foundational methodology in data modeling West et al. (2001). Meanwhile, graph-based algorithms have found widespread adoption for analyzing complex relational patterns across scientific and industrial domains Dijkstra (2022); Kruskal (1956). Within graph theory, the subgraph isomorphism problem represents a core computational challenge Nguyen et al. (2022); Zhang et al. (2024). Specifically, it requires determining whether a given graph contains a structurally isomorphic subgraph of another graph, while also deriving the explicit node correspondence when such isomorphism exists.

The subgraph isomorphism problem has garnered significant attention across pivotal domains due to its critical role in enabling structural pattern analysis. In biomolecular sciences, it underpins the determination of structural compatibility between molecules and proteins for drug discovery and protein interaction studies Balaban (1985); Bonnici et al. (2013). Within semantic web technologies, it facilitates efficient Resource Description Framework (RDF) query processing to traverse complex knowledge graphs Kim et al. (2015). Social network analytics leverages subgraph isomorphism detection to generate personalized recommendation systems through dynamic community subpattern mining Rong et al. (2018). Furthermore, in domain-specific computing architectures, this capability proves essential for optimizing loop mapping schemes in reconfigurable computing systems, where topological constraints demand rigorous subgraph matching Hamzeh et al. (2012). These cross-disciplinary applications collectively demonstrate the problem’s fundamental importance in modern computational paradigms.

The subgraph isomorphism problem, known to be NP-complete Conte et al. (2004), is classically tackled using heuristic methods, which can be broadly classified into three categories: (1) **Tree search algorithms**, ranging from classical ones (Ullmann Ullmann (1976), VF-series Carletti et al. (2017)) to modern extensions (RM Sun et al. (2020), VEQ Kim et al. (2021), CaLiG Yang et al. (2023)), that employ depth-first search combined with pruning techniques; (2) **Constraint programming frameworks**, which model the problem as a constraint satisfaction problem (CSP) using integer linear programming (ILP), SAT, or other formalisms (e.g., McGregor McGregor (1979), Solnon

Solnon (2010), Zampelli Zampelli et al. (2010)); (3) **Graph indexing approaches**, often inspired by database systems (e.g., GraphQL He & Singh (2008), QuickSI Shang et al. (2008), GADDI Zhang et al. (2009)), that use precomputed structural signatures and inverted indices to accelerate filtering. Despite the use of advanced pruning strategies, state-of-the-art heuristic solvers still suffer from exponential worst-case time complexity, which often limits their practicality on large real-world graphs.

Recent advances in artificial intelligence, particularly graph neural networks (GNNs) designed for graph isomorphism analysis Xu et al. (2019), have spurred new learning-based approaches for subgraph isomorphism. Representative methods include IsoNet Roy et al. (2022), streaming models Duong et al. (2021), GMN Li et al. (2019), GNN-PE Ye et al. (2024b), EinsMatch Ramachandran et al. (2024), and SubMDSE Raj et al. (2025). Furthermore, AI techniques have been extended to address related problems such as subgraph alignment Bainson et al. (2024), graph edit distance computation Piao et al. (2023), and maximum common subgraph identification Bai et al. (2021).

However, current learning-based methods exhibit notable limitations: they often yield probabilistic approximations rather than exact solutions Ramachandran et al. (2024), or are confined to very small query graphs (e.g.,  $\leq 10$  nodes) Ye et al. (2024b). A more fundamental constraint is their heavy reliance on solver-generated labels produced by traditional algorithms for supervised training, which introduces computational bottlenecks and restricts practical deployment.

To address the limitations of existing approaches, we present **SIREN** (Subgraph Isomorphism via Reinforcement-Enhanced Graph Neural Networks). Our framework integrates a Deep Q-Network (DQN) Volodymyr et al. (2019) with bidomain-based pruning to autonomously discover optimal node selection heuristics, which are critical components in state-of-the-art subgraph isomorphism solvers. Complementing this, SIREN employs a pretrained graph neural network grounded in partial order relation learning to hierarchically encode subgraph structural dependencies. While maintaining **probabilistic completeness**, experimental results demonstrate that SIREN outperforms all machine learning-based methods in accuracy for approximate matching tasks. Simultaneously, it achieves significant efficiency gains over both heuristic approaches and AI-based methods when enumerating complete solution sets. The primary contributions of this paper are as follows:

1. **DQN-GNN integration with probabilistic completeness:** We propose a novel DQN-GNN integration method with bidomain-based pruning that efficiently solves subgraph isomorphism problems while guaranteeing **probabilistic completeness**.
2. **Partial-order-aware pretrained GNN:** We introduce a partial-order-aware GNN pretraining strategy that enhances substructure relationship understanding.
3. **Unified framework:** Our framework simultaneously addresses both existence determination and complete solution enumeration for subgraph isomorphism.
4. **Superior precision and efficiency:** Experimental results on real-world datasets demonstrate that SIREN achieves 100% precision with modest computational time, outperforming AI-based approximate matching methods. Compared to state-of-the-art exact solvers, SIREN delivers  $\sim 52\%$  faster execution than leading AI approaches and  $\sim 21\%$  acceleration over top heuristic methods.

## 2 PRELIMINARIES

**1) Graph:** A graph can be formally defined as  $G = (V, E)$ , where  $V = \{v_1, \dots, v_n\}$  denotes the finite set of vertices and  $E \subseteq V \times V$  denotes the edge set. Each edge  $e_{ij} \in E$  connects two vertices  $v_i, v_j$  in  $V$ .

**2) Graph Isomorphism:** Two graphs  $G_1 = (V_1, E_1)$  and  $G_2 = (V_2, E_2)$  are isomorphic ( $G_1 \cong G_2$ ) if there exists a bijective mapping  $f : V_1 \xrightarrow{\sim} V_2$  such that:

$$\forall v_i, v_j \in V_1, (v_i, v_j) \in E_1 \iff (f(v_i), f(v_j)) \in E_2 \quad (1)$$

**3) Subgraph:** Let  $G = (V, E)$  be a graph. A subgraph  $G' = (V', E')$  of  $G$  is defined as  $G' \sqsubseteq G$ , which satisfies:

$$V' \subseteq V \quad \text{and} \quad E' \subseteq \{(u, v) \in E \mid u, v \in V'\} \quad (2)$$

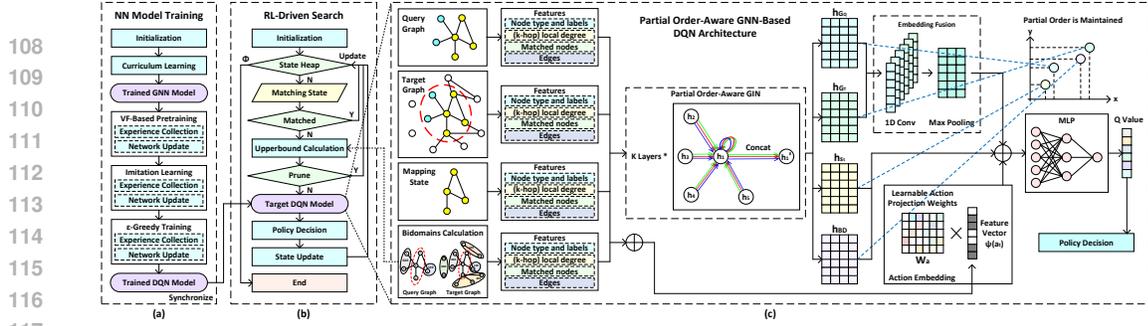


Figure 1: Overview of the framework of SIREN. (a) NN Training Framework, (b) Reinforcement Learning-Driven Search Framework, (c) Partial Order-Aware GNN-based DQN Architecture

**4) Subgraph Isomorphism Problem:** The subgraph isomorphism problem involves determining whether there exists a subgraph  $G' \subseteq G_T$  such that  $G_Q \cong G'$ , where  $G_Q$  denotes the query graph and  $G_T$  denotes the target graph. Formally:

$$\exists G' = (V', E') \subseteq G_T \quad \text{s.t.} \quad G_Q \cong G' \quad (3)$$

If such subgraphs exist, the solution set  $\mathcal{S}_G$  comprises all valid subgraphs:

$$\mathcal{S}_G = \{G' \subseteq G_T \mid G_Q \cong G'\} \quad (4)$$

### 3 THE SIREN METHOD

In this section, we provide a detailed description of the SIREN framework, which stands for Subgraph Isomorphism via Reinforcement-Enhanced Graph Neural Networks. Section 3.1 presents our DQN-based **probabilistic complete** framework for subgraph isomorphism, and Section 3.2 introduces our partial order relation-aware GNN architecture.

#### 3.1 PROBABILISTIC COMPLETE FRAMEWORK FOR SUBGRAPH ISOMORPHISM

SIREN addresses the subgraph isomorphism problem by integrating graph neural networks with reinforcement learning. This approach formulates the problem as a Markov Decision Process (MDP) Puterman (1990) and employs a DQN-based framework Volodymyr et al. (2019) to enhance traditional tree-search heuristics Carletti et al. (2017). Our search framework guarantees the completeness of SIREN, with a detailed proof provided in Section A.4.

**1) DQN-Enhanced Search Framework.** The DQN-enhanced search framework of SIREN, depicted in Figure 1(b), maintains a state heap  $\mathcal{ST}$  to store feasible states. At each decision step, the state  $s_t$  with the maximal action space cardinality  $|\mathcal{A}_t|$  is selected as the current state, prioritizing branches with higher combinatorial potential. At each step, the agent either adds a new node pair to the current partial matching or backtracks from a previous decision. The selection of candidate nodes is optimized using a partial order-aware Graph Isomorphism Network (GIN) integrated within the DQN. The search process terminates when all possible isomorphisms are found or non-existence is proven. SIREN can also be configured to terminate upon discovering the first feasible solution, which is suitable for applications requiring only a single valid subgraph matching.

**2) DQN Framework.** As illustrated in Figure 1(c), the framework utilizes continuous embedding representations to encode states  $s_t$  and actions  $a_t$ . These representations are processed by a DQN architecture that consists of a partial-order aware GNN encoder and learnable projection modules. The DQN maps state-action pairs  $(s_t, a_t)$  to Q-value estimates  $Q(s_t, a_t)$ , thereby enabling data-driven policy optimization.

**3) State Representation.** A state  $s_t$  comprises the node-node mapping  $\mathcal{M}_t$  between the selected subgraphs, the input graphs themselves, and the bidomain information corresponding to the current mapping  $\mathcal{M}_t$ . The features provided to the GNN model for state  $s_t$  include node types/labels, (k-hop) local degree profiles, matching status indicators, and edge information of the graphs.

**4) Action Space.** At each step, the agent selects (1) a node pair  $(v_q, v_t) \in \mathcal{C}_t$  that maintains topological consistency with  $\mathcal{M}_t$ , and (2) a special  $\langle \text{terminate} \rangle$  action to prune unpromising branches,

where  $\mathcal{C}_t$  denotes the candidate node pairs for expansion. The action space size is dynamic:  $|\mathcal{C}_t| + 1$  (including termination).

**5) Action Embedding.** SIREN incorporates an action embedding optimization mechanism. For each feasible node pair  $(v_q, v_t)$ , the associated node features are projected into a low-dimensional embedding space via a learnable matrix  $\mathbf{W}_a$ . This embedding captures the essential value characteristics of the action, which in turn refines the Q-value computation.

**6) Reward Function.** The reward function of SIREN is designed to guide the learning process through a dense reward mechanism:

- A base reward of  $+\alpha \cdot 1/|V_Q|$  for each valid node pair matching (default  $\alpha = 0.5$ ), where  $|V_Q|$  is the number of nodes in the query graph.
- An incremental structural reward of  $+\beta \cdot \Delta E/|E_T|$  for each newly matched edge (default  $\beta = 0.5$ ), where  $\Delta E$  denotes the increase in the number of correctly matched edges and  $|E_T|$  is the total number of edges in the target graph.
- A sparse success reward of  $+\gamma \cdot |V_Q|$  (default  $\gamma = 0.05$ ) upon finding a complete isomorphism.

**7) Bidomain-Based Embeddings and Pruning.** SIREN leverages Bidomains Bai et al. (2021) to enhance its search process in two key ways: by providing additional graph embeddings that inform the Q-value computation, and by facilitating efficient pruning through estimates of the exploration upper bound. For state  $s_t = \mathcal{M}_t = (\mathcal{M}_Q, \mathcal{M}_T)$  with matched sets  $\mathcal{M}_Q \subseteq V(G_Q)$ ,  $\mathcal{M}_T \subseteq V(G_T)$ , the  $k$ -th bidomain  $\mathcal{B}_k$  is defined as:

$$\mathcal{B}_k = \langle \mathcal{V}_k^Q, \mathcal{V}_k^T \rangle \quad (5)$$

where  $\mathcal{V}_k^Q \subseteq V(G_Q)$  and  $\mathcal{V}_k^T \subseteq V(G_T)$  satisfy:

$$\forall u \in \mathcal{V}_k^Q, v \in \mathcal{V}_k^T : \text{adj}(u, \mathcal{M}_Q) \equiv \text{adj}(v, \mathcal{M}_T) \quad (6)$$

denoting identical connectivity patterns to matched sets  $\mathcal{M}_Q \subseteq V(G_Q)$  and  $\mathcal{M}_T \subseteq V(G_T)$ . A detailed description of the bidomain technique is provided in Section A.8.

**8) Embedding Fusion.** We perform embedding fusion on the embeddings  $\mathbf{h}_{G_Q}$  and  $\mathbf{h}_{G_T}$  generated by the GNN for the query graph and target graph, respectively. This alignment operation involves processing each embedding through a 1D convolutional layer followed by a pooling layer, and subsequently merging the resulting representations.

**9) Action-Value Function.** For state transition, candidate actions  $\mathcal{A}_t = \{(u, v) \in \mathcal{C}_t\}$  are evaluated by the DQN’s action value function  $Q_\theta(s_t, a_t)$ , where  $\mathbf{h}$  denotes the graph embeddings generated with our trained GNN model:

$$\begin{aligned} Q_\theta(s_t, a_t) &= \mathcal{F} \left( \underbrace{\text{GNN}_{\text{enc}}(G_Q, G_T, s_t, \mathcal{B}_{s_t})}_{\text{state embedding}} \oplus \underbrace{\mathbf{W}_a \psi(a_t)}_{\text{action embedding}} \right) \\ &= \mathcal{F}(\mathbf{h}_{s_t}, \text{Fuse}(\mathbf{h}_{G_Q}, \mathbf{h}_{G_T}), \mathbf{h}_{BD}, a_t) \end{aligned} \quad (7)$$

where  $\psi(a_t)$  is the feature vector of action  $a_t$ , and  $\mathbf{W}_a$  denotes learnable action projection weights.

**10)  $\epsilon$ -Greedy Action Selection.** The agent selects an action  $a_t$  at each timestep using an  $\epsilon$ -greedy policy based on the current Q-value estimates:

$$a_t^* = \begin{cases} \arg \max_{a_t \in \mathcal{A}_t} (Q_\theta(s_t, a_t)) & \text{with probability } 1 - \epsilon \\ \text{random action} & \text{with probability } \epsilon \end{cases}$$

The updated state  $s_{t+1} = \mathcal{M}_t \cup a_t^*$  is then pushed back onto  $\mathcal{S}$ .

**11) Action Space Partition.** For large-scale graphs where the number of candidate matching actions (node pairs) becomes prohibitively large, we partition the action space into chunks and compute Q-values separately for each chunk. This approach prevents GPU memory overflow while maintaining computational efficiency.

The training process of our DQN model consists of three consecutive phases: pretraining, imitation learning, and reinforcement learning, which is detailed in Section A.5. To analyze the impact of different reinforcement learning paradigms, we conducted comparisons between DQN and Proximal Policy Optimization (PPO) Schulman et al. (2017), as detailed in Section A.11.

### 3.2 PARTIAL ORDER-AWARE GNN ARCHITECTURE

The subgraph relationship satisfies the properties of reflexivity, transitivity, and antisymmetry, thereby establishing it as a partial order on the set of graphs. Formally:

- **Reflexivity:**  $G \sqsubseteq G, G \cong G$ .
- **Transitivity:**  $(G_1 \sqsubseteq G_2) \wedge (G_2 \sqsubseteq G_3) \implies G_1 \sqsubseteq G_3$ .
- **Antisymmetry:**  $(G_1 \sqsubseteq G_2) \wedge (G_2 \sqsubseteq G_1) \iff G_1 = G_2$ .

The formal proof of these properties is provided in Section A.1. Moreover, The intersection of the set of  $G_1$ 's subgraphs and the set of  $G_2$ 's subgraphs contains all common subgraphs of  $G_1$  and  $G_2$  Lou et al. (2020). Therefore, SIREN employs a partial order-aware GNN to enforce that the learned graph embeddings preserve the isomorphic partial order relationships within the embedding space. This geometric constraint ensures that for any two graphs  $G_Q$  and  $G_T$  with embeddings  $\mathbf{h}_{G_Q} = (\mathbf{h}_{G_Q}^1, \dots, \mathbf{h}_{G_Q}^D)$  and  $\mathbf{h}_{G_T} = (\mathbf{h}_{G_T}^1, \dots, \mathbf{h}_{G_T}^D)$ :

$$\forall i \in 1, \dots, D, \mathbf{h}_{G_Q}^i \leq \mathbf{h}_{G_T}^i \iff G_1 \sqsubseteq G_2 \quad (8)$$

The proof that our GNN model satisfies the partial order relationships described in Equation 8 is provided in Section A.3. To preserve the partial order relationship of the subgraphs, we use the max margin loss to train our GNN model. Within each minibatch, we define  $P = \{(q, t) \in \mathcal{MB} \mid \mathcal{N}(q) \sqsubseteq \mathcal{N}(t)\}$  as the set of positive pairs where the neighborhood subgraph of query node  $q$  is isomorphic to a subgraph of target node  $t$ 's neighborhood, and  $N = \mathcal{MB} \setminus P$  as the negative pairs violating this structural constraint. The loss function  $\mathcal{L}$  then operates on these sets to enforce geometric consistency in the embedding space:

$$\mathcal{L}(\mathbf{h}_q, \mathbf{h}_t) = \sum_{\mathbf{h}_q, \mathbf{h}_t \in P} E(\mathbf{h}_q, \mathbf{h}_t) + \sum_{\mathbf{h}_q, \mathbf{h}_t \in N} \max(0, \alpha - E(\mathbf{h}_q, \mathbf{h}_t)) \quad (9)$$

where

$$E(\mathbf{h}_q, \mathbf{h}_t) = \|\max(0, \mathbf{h}_q - \mathbf{h}_t)\|_2^2 \quad (10)$$

We employed an improved GIN model Xu et al. (2019), incorporating multi-scale feature fusion techniques to generate graph embedding vectors. For layer  $l \in \{0, 1, \dots, L-1\}$ , the embedding of node  $v$  is computed as:

$$\mathbf{h}_v^{(l+1)} = \text{MLP}^{(l)}((1 + \epsilon^{(l)}) \cdot \mathbf{h}_v^{(l)} + \sum_{u \in \mathcal{N}(v)} \mathbf{h}_u^{(l)}) \quad (11)$$

where  $\epsilon^{(l)} \in \mathbb{R}$  is a learnable scalar parameter,  $\mathcal{N}(v)$  denotes the neighborhood of node  $v$ , and  $\text{MLP}^{(l)}$  denotes the multi-layer perceptron with LeakyReLU Maas et al. (2013) activation  $\sigma$ . Let  $\gamma_l$  be the learnable hierarchical weight coefficient, the graph embedding is obtained by concatenating sum-pooled features across layers:

$$h_G = \sum_{l=0}^L \gamma_l \cdot \sum_{v \in V} h_v^{(l)} \quad (12)$$

The training protocol of our partial order-aware GNN model is detailed in Section A.6.

## 4 EXPERIMENTS

To evaluate the effectiveness and efficiency of SIREN, we compared SIREN with 19 state-of-the-art neural network-based methods and heuristic methods. The experiments are conducted real large graph datasets from TUDataset Morris et al. (2020) and on synthetic graph datasets.

### 4.1 BASELINE METHODS

We compared SIREN with 10 state-of-the-art neural network-based methods for subgraph isomorphism, including SimGNN Bai et al. (2019), GraphSim Bai et al. (2020), GEDGNN Piao et al. (2023), GOTSim Doan et al. (2021), ERIC Zhuo & Tan (2022), NeuroMatch Lou et al. (2020), GMN Li et al.

Table 1: Evaluation of Mean Average Precision (MAP) on real-world graph pairs.  $K$  = layers,  $|ND|$  = node state dimensions. Each dataset contains 300 query graphs and 800 target graphs.

Method	$K$	$ ND $	AIDS	MUTAG	PTC-FM	PTC-FR	PTC-MM	PTC-MR
Average $ V_Q $			11.61	12.91	11.73	11.81	11.80	11.87
Max $ V_Q $			15	15	15	15	15	15
Average $ V_T $			18.50	18.41	18.30	18.32	18.36	18.32
Max $ V_T $			20	20	20	20	20	20
SIMGNN	3	10	0.326 ± 0.019	0.303 ± 0.012	0.416 ± 0.015	0.355 ± 0.015	0.358 ± 0.015	0.308 ± 0.017
GRAPHSIM	3	10	0.173 ± 0.007	0.182 ± 0.008	0.231 ± 0.011	0.165 ± 0.007	0.2 ± 0.009	0.216 ± 0.013
GEDGNN	3	10	0.340 ± 0.015	0.605 ± 0.029	0.437 ± 0.013	0.497 ± 0.018	0.509 ± 0.018	0.309 ± 0.009
GOTSIM	5	10	0.336 ± 0.017	0.387 ± 0.018	0.459 ± 0.017	0.361 ± 0.013	0.417 ± 0.017	0.430 ± 0.017
ERIC	5	10	0.512 ± 0.022	0.558 ± 0.027	0.624 ± 0.019	0.572 ± 0.021	0.573 ± 0.02	0.639 ± 0.018
GMN-MATCH	3	10	0.609 ± 0.02	0.693 ± 0.026	0.686 ± 0.018	0.667 ± 0.021	0.627 ± 0.02	0.683 ± 0.017
NEUROMATCH	3	10	0.454 ± 0.025	0.583 ± 0.027	0.622 ± 0.019	0.572 ± 0.023	0.522 ± 0.019	0.565 ± 0.02
ISONET	3	10	0.704 ± 0.021	0.733 ± 0.023	0.782 ± 0.017	0.734 ± 0.02	0.758 ± 0.016	0.764 ± 0.015
SUBMDSE-LATE	5	10	0.712 ± 0.018	0.721 ± 0.025	0.793 ± 0.016	0.744 ± 0.019	0.758 ± 0.015	0.782 ± 0.014
SUBMDSE-EARLY	5	10	0.817 ± 0.017	0.837 ± 0.02	0.887 ± 0.012	0.854 ± 0.013	0.849 ± 0.012	0.864 ± 0.011
SIREN-MINI	3	10	<b>1.000 ± 0.000</b>					
Avg. Runtime (ms)			41.3	33.4	35.5	43.7	34.8	44.2
SIREN	8	64	<b>1.000 ± 0.000</b>					
Avg. Runtime (ms)			50.5	47.3	42.5	56.3	52.4	52.7
Improvement			18.3%	16.3%	11.3%	14.6%	15.1%	13.6%

(2019), IsoNet Roy et al. (2022), SubMDSE Raj et al. (2025) and GNN-PE Ye et al. (2024b), as well as 9 heuristic state-of-the-art methods for solving the subgraph isomorphism problem, including RI Bonnici et al. (2013), VF2++ Jüttner & Madarasi (2018), GraphQL He & Singh (2008), QuickSI Shang et al. (2008), VF3 Carletti et al. (2017), RM Sun et al. (2020), VEQ Kim et al. (2021), CaLiG Yang et al. (2023), and DPIso Han et al. (2019).

Existing neural network-based subgraph isomorphism methods can be divided into two categories: exact matching and approximate prediction. Exact matching methods return the precise solution set for the subgraph isomorphism problem, while approximate prediction methods provide a quick assessment of whether two graphs satisfy the subgraph isomorphism relationship. Among these methods, GNN-PE is an exact matching method, while the other 7 methods are approximate prediction methods. To ensure a fair comparison, we evaluated the solving speed against GNN-PE and heuristic methods (while the accuracy of all methods is inherently 1.000), and compared the solving accuracy with the other 7 methods, using the same real-world datasets and sampling settings as in the original papers.

## 4.2 EXPERIMENTAL SETTINGS

In SIREN, we utilize 8 layers of Graph Isomorphism Networks (GIN) Xu et al. (2019) each with 64 dimensions for the embeddings. For DQN, we use MLP layers to project concatenated embeddings to a scalar. The discount factor  $\gamma$  of the DQN is set to 1.0, and the learning rate of the DQN and the GIN is set to 0.001. The models are trained using the Adam optimizer Kingma (2014). The learning rate is annealed with a cosine annealer with restarts every 100 epochs. The DQN is trained by 10000 iterations. Prior to DQN training, we conduct a 50000-epochs supervised pre-training of the GNN model to generate geometrically consistent partial order-preserving graph embeddings. The training data is generated by randomly sampling neighborhoods from large real-world graph datasets Morris et al. (2020), while being regenerated every 50 epochs. To ensure fairness across diverse tasks, we trained two distinct GNN configurations (in SIREN-Mini and SIREN). Detailed parameter specifications are provided in Table 1. The experiments were conducted on a Ubuntu server equipped with a 128-Core Intel(R) Xeon(R) Gold 5218 CPU running at 2.30 GHz and 256 GB of memory, along with 4× Nvidia Ampere A800 GPU. SIREN were implemented with the PyTorch and PyTorch Geometric libraries Fey & Lenssen (2019).

## 4.3 EFFECTIVENESS OF SIREN

To compare SIREN with approximate prediction deep learning methods, we selected the same datasets as SubMDSE Raj et al. (2025) from TUDataset Morris et al. (2020) for testing, i.e., PTC-FR, PTC-FM, PTC-MM, PTC-MR, MUTAG, and AIDS. We compared SIREN with 9 state-of-the-art neural methods. The experimental results indicate that, without limiting the search time, SIREN achieved correct solutions for all datasets, with the longest time not exceeding 2 seconds ( $\sim 1.8s$ ). Compared

Table 2: Ablation Study of SIREN by Mean Average Precision (MAP) on Real-World Graph Pairs.

Method	PTC-FR	PTC-FM	PTC-MM	PTC-MR	MUTAG	AIDS
SIREN-DQN	0.880	0.893	0.961	0.828	0.952	0.904
v.s. IsoNet	<b>+2.3%</b>	-0.8%	<b>+5.5%</b>	-6.6%	<b>+3.2%</b>	<b>+1.0%</b>
SIREN-GNN	0.766	0.796	0.886	0.738	0.861	0.842
v.s. IsoNet	-9.1%	-10.5%	-2.0%	-15.6%	-5.9%	-5.2%
SIREN	<b>1.000</b>	<b>1.000</b>	<b>1.000</b>	<b>1.000</b>	<b>1.000</b>	<b>1.000</b>

Table 3: Runtime on Large-Scale Dense Real-World Graphs (Unit: s), OOM: Out of Memory.

	1000	2000	3000	4000	5000	10000
nodes of $G_T$	1000	2000	3000	4000	5000	10000
nodes of $G_Q$	200	400	600	800	1000	2000
Edge density	0.4	0.3	0.2	0.2	0.2	0.2
VF3	12151.6	16351.2	4292.6	$> 10^5$	$> 10^5$	$> 10^5$
VF3P	1187.8	1923.6	636.2	24058.0	51020.4	$> 10^5$
VEQ	860.9	OOM	OOM	OOM	OOM	OOM
GNN-PE	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$
GLSEARCH	7534.5	$> 10^5$	4682.1	$> 10^5$	$> 10^5$	$> 10^5$
SIREN	<b>642.4</b>	<b>943.2</b>	<b>486.7</b>	<b>16842.4</b>	<b>34719.5</b>	$> 10^5$

to the most accurate neural method SubMDSE-Early Raj et al. (2025), SIREN improved Mean Average Precision by an average of 14.9%. The experimental results show that SIREN outperforms approximate prediction deep learning methods in terms of effectiveness. These experimental results align with the theory, as SIREN guarantees that, given enough time, it can obtain an exact solution or prove that no valid solution exists. Furthermore, these results demonstrate that SIREN is a more effective method for general real-world datasets, as it provides acceptable running times and more accurate results.

#### 4.4 EFFICIENCY AND THROUGHPUT RATE

To compare SIREN with exact matching deep learning methods and heuristic methods, we selected the same datasets as GNN-PE Roy et al. (2022) from TUDataset Morris et al. (2020) for testing, i.e., Yeast, Human, HPRD, WordNet, DBLP, Youtube, and US Patents. Statistics of these real graphs are summarized in Section A.10. We divided the dataset according to the method used in GNN-PE Roy et al. (2022), where the size of  $|V(q)|$  is set to 8. Figure 2 shows the efficiency test results on large-scale datasets, which indicate that our SIREN method outperforms all existing heuristic methods across all test data, and it also surpasses GNN-PE in the majority of cases. Only for the DBLP dataset does SIREN ( $\sim 0.19$ s) perform slightly slower than GNN-PE ( $\sim 0.17$ s), with almost no difference. On average, compared to RI Bonnici et al. (2013), VF Jüttner & Madarasi (2018), GraphQL He & Singh (2008), QuickSI Shang et al. (2008), GNN-PE Roy et al. (2022), VF3 Carletti et al. (2017), RM Sun et al. (2020), VEQ Kim et al. (2021), CaLiG Yang et al. (2023), and DPIso Han et al. (2019), SIREN achieves speedups of  $36.5\times$ ,  $29.0\times$ ,  $171.2\times$ ,  $25.7\times$ ,  $52.1\%$ ,  $14.4\times$ ,  $92.4\%$ ,  $21.3\%$ ,  $7.70\times$ , and  $55.3\%$ , respectively.

Figure 3 compares the throughput, measured in generated matched embeddings per second (EPS), of SIREN against other methods on large-scale datasets. As shown, SIREN achieves a higher throughput than all baseline methods, outperforming the state-of-the-art approach VEQ by  $1.046\times$ .

#### 4.5 SCALABILITY OF SIREN

To evaluate the scalability of SIREN, we compared SIREN with GNN-PE Roy et al. (2022), GLSearch Bai et al. (2021), VEQ Kim et al. (2021), and the state-of-the-art CPU-parallelized subgraph isomorphism method VF3p Carletti et al. (2017) on synthetic large adversarial dense graph dataset by VF3 Carletti et al. (2017). As shown in Table 3, experimental results on large graphs demonstrate that despite the server CPU’s higher theoretical FP32 compute capacity (CPU: 140 TFLOPS vs. GPU: 40 TFLOPS), SIREN still outperforms traditional methods. On average, compared to

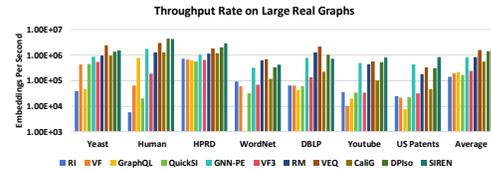
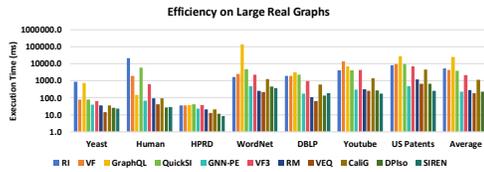


Figure 2: Efficiency of different methods. Figure 3: Throughput Rate of different methods. GNN-PE, GLSearch, VF3, VF3p and VEQ, SIREN achieves speedups of  $8.32\times$ ,  $4.82\times$ ,  $3.34\times$ ,  $47.0\%$ , and  $34.0\%$ , respectively (excluding cases exceeding  $10^5$  seconds or OOM). Our experiments further demonstrate that VEQ encounters out-of-memory failures under a 256GB memory constraint—already substantial for graph processing—when handling graphs exceeding 2000 nodes with edge density  $\rho > 0.3$ . This limitation stems from VEQ’s significantly higher space complexity compared to SIREN, preventing it from scaling to large dense graphs.

#### 4.6 ABLATION STUDY

Table 2 presents the ablation study results for two modules in SIREN: DQN and GNN. Here, SIREN-DQN refers to the results obtained using only the DQN + GNN structure, where GNN is randomly initialized without any pre-training method. SIREN-GNN indicates the scenario where the DQN structure is not used, and the NeuralMatch method Lou et al. (2020) is employed to directly output estimates for the subgraph isomorphism problem. The experimental results demonstrate that both the lack of GNN pre-training and the absence of the DQN model lead to a decline in solution quality.

## 5 RELATED WORK

### 5.1 TRADITIONAL APPROACHES

The subgraph isomorphism problem is known to be NP-complete and has traditionally been tackled using heuristic methods. These approaches can be broadly categorized as follows:

**1) Tree Search:** Classical algorithms such as Ullmann’s algorithm Ullmann (1976) and the VF-series Cordella et al. (2001; 2004); Carletti et al. (2017) employ depth-first search combined with pruning heuristics (e.g., degree and label filters). VF3 Carletti et al. (2017) introduces state-space precomputation and look-ahead pruning to improve efficiency, particularly for large dense graphs. More recent methods, including RM Sun et al. (2020), VEQ Kim et al. (2021), and CaLiG Yang et al. (2023), further accelerate the resolution process.

**2) Constraint Programming:** These approaches model the problem as a Constraint Satisfaction Problem (CSP), using integer linear programming (ILP), Boolean satisfiability (SAT), or related formalisms. Variables typically represent query nodes, with domains consisting of candidate target nodes, and constraints encode structural requirements. Methods by McGregor McGregor (1979), Solnon Solnon (2010), and Zampelli Zampelli et al. (2010) employ arc consistency techniques to iteratively prune the solution space.

**3) Graph Indexing:** Methods inspired by database systems, such as GraphQL He & Singh (2008), QuickSI Shang et al. (2008), and GADDI Zhang et al. (2009), leverage precomputed structural signatures and inverted indices to enable efficient pre-match filtering and early termination.

**Limitations:** Despite their practical utility, all these methods exhibit exponential worst-case time complexity—typically  $\mathcal{O}(n^k)$  for a query pattern of size  $k$ —and are sensitive to label noise and graph density. Indexing techniques also involve significant memory overhead, often reaching  $\mathcal{O}(m^d)$  for depth- $d$  neighborhood features.

### 5.2 NEURAL APPROACHES FOR SUBGRAPH ISOMORPHISM

Recent neural network-based approaches for subgraph isomorphism can be divided into two main categories:

**1) Exact Solvers:** These methods aim to provably determine the existence of subgraph isomorphisms and recover corresponding node mappings. Representative techniques include subgraph index

432 embeddings Duong et al. (2021) and path dominance embeddings Ye et al. (2024b), which encode  
433 structural hierarchies to facilitate exact matching.

434 **2) Approximate Heuristics:** These methods prioritize scalability and efficiency at the cost of  
435 completeness, often relying on probabilistic or learned similarity measures. Examples include  
436 similarity-based networks Li et al. (2019), geometric embeddings Lou et al. (2020), [D2Match Liu  
437 et al. \(2023\)](#) and IsoNet Roy et al. (2022). More recent advances include EinsMatch Ramachandran  
438 et al. (2024), which introduces iterative alignment refinement, and SubMDSE Raj et al. (2025), which  
439 explores multifaceted design spaces for improved performance.

440 **Limitations:** Exact solvers typically depend on expensive solver-generated labels for supervised  
441 training, leading to high annotation and computational costs. On the other hand, approximate  
442 methods often lack theoretical guarantees and tend to exhibit reduced accuracy when applied to strict  
443 isomorphism problems.

### 444 5.3 NEURAL APPROACHES FOR RELATED PROBLEMS

445  
446  
447 **1) Graph Isomorphism:** While graph isomorphism can be solved in quasi-polynomial time, i.e.,  
448  $O(e^{(\log n)^c})$  Babai (2016), it is often studied to characterize the expressive power of graph neural  
449 networks. Methods such as GIN Xu et al. (2019)—which simulates the Weisfeiler-Leman (WL) test  
450 Leman & Weisfeiler (1968)—and Graph Transformers Yun et al. (2019); Lee et al. (2024) have been  
451 developed, but they generally lack the precision required for subgraph isomorphism tasks.

452  
453 **2) Maximum Common Subgraph (MCS):** Subgraph isomorphism is a special case of MCS in which  
454 the common subgraph must be isomorphic to the query graph  $G_Q$ . Neural solvers such as GLSearch  
455 Bai et al. (2021) and MCSP+RL Liu et al. (2019) have been proposed for MCS, but they often fail to  
456 fully leverage the topological structure of the query graph, resulting in suboptimal efficiency.

457 **3) Subgraph Alignment:** This problem involves determining whether a query graph  $G_Q$  is isomor-  
458 phic to an induced subgraph of a target graph  $G_T$ . A recent spectral-based method Bainson et al.  
459 (2024) has been proposed to address it. Although subgraph alignment is a special case of subgraph  
460 isomorphism, the two are not identical; clarifying their distinctions is necessary to prevent confusion  
461 arising from terminological overlap.

462 **4) Graph Alignment:** Graph alignment aims to find a bijective mapping between the nodes of two  
463 graphs such that structural differences after mapping are minimized. FUGAL Bommakanti et al.  
464 (2024) introduces a learning-based approach for this task, while a differentiable top-k-method Wang  
465 et al. (2023) addresses partial graph matching, a related relaxation.

466 **5) Graph Edit Distance (GED):** Subgraph isomorphism can be reduced to GED by setting infinite  
467 substitution costs. However, general neural solvers for GED Raveaux (2021); Piao et al. (2023) are  
468 not well-suited for exact isomorphism due to their flexible cost models and broader objective.

469  
470 **6) Large Language Model (LLM)-Based Methods:** Recent efforts such as ThinkOnGraph Sun  
471 et al. (2024) and GraphGPT Tang et al. (2024) focus primarily on attributed graphs. Although  
472 InstructGLM Ye et al. (2024a) encodes structural information through prompting, it—like other  
473 LLM-based approaches—has not yet shown effectiveness for exact combinatorial isomorphism  
474 problems.

## 475 6 CONCLUSION

476  
477  
478 The subgraph isomorphism problem is a challenging NP-complete problem with wide applications  
479 across various fields. In this paper, we introduce SIREN, an RL-enhanced GNN for subgraph  
480 isomorphism. Through our proposed DQN-based reinforcement learning framework and the GNN  
481 model based on partial order relations, we can improve the candidate node selection process in solving  
482 the subgraph isomorphism problem. Experiments on real datasets show that SIREN can effectively  
483 accelerate the solving of the subgraph isomorphism problem and enhance solution quality. Future  
484 work includes further improvements to the reinforcement learning framework and GNN, testing on  
485 more large-scale real datasets, and extending similar methods to other NP problems, such as the  
maximum clique problem.

## REFERENCES

- 486  
487  
488 László Babai. Graph isomorphism in quasipolynomial time. In *Proceedings of the forty-eighth*  
489 *annual ACM symposium on Theory of Computing*, pp. 684–697, 2016.
- 490 Yunsheng Bai, Hao Ding, Song Bian, Ting Chen, Yizhou Sun, and Wei Wang. Simgnn: A neural  
491 network approach to fast graph similarity computation. In *Proceedings of the twelfth ACM*  
492 *international conference on web search and data mining*, pp. 384–392, 2019.
- 493 Yunsheng Bai, Hao Ding, Ken Gu, Yizhou Sun, and Wei Wang. Learning-based efficient graph  
494 similarity computation via multi-scale convolutional set matching. In *Proceedings of the AAAI*  
495 *conference on artificial intelligence*, volume 34, pp. 3219–3226, 2020.
- 496  
497 Yunsheng Bai, Derek Xu, Yizhou Sun, and Wei Wang. Gsearch: Maximum common subgraph  
498 detection via learning to search. In *International Conference on Machine Learning*, pp. 588–598.  
499 PMLR, 2021.
- 500 Ama Bembua Bainsan, Judith Hermanns, Petros Petsinis, Niklas Aavad, Casper Dam Larsen, Tiarnan  
501 Swayne, Amit Boyarski, Davide Mottin, Alex M. Bronstein, and Panagiotis Karras. Spectral  
502 subgraph localization. In Soledad Villar and Benjamin Chamberlain (eds.), *Proceedings of the Sec-*  
503 *ond Learning on Graphs Conference*, volume 231 of *Proceedings of Machine Learning Research*,  
504 pp. 7:1–7:11. PMLR, 27–30 Nov 2024. URL [https://proceedings.mlr.press/v231/](https://proceedings.mlr.press/v231/bainsan24a.html)  
505 [bainsan24a.html](https://proceedings.mlr.press/v231/bainsan24a.html).
- 506 Alexandru T Balaban. Applications of graph theory in chemistry. *Journal of chemical information*  
507 *and computer sciences*, 25(3):334–343, 1985.
- 508  
509 Luc Rey Bellet. Ergodic properties of markov processes. In *Open Quantum Systems II: The Markovian*  
510 *Approach*, pp. 1–39. Springer, 2006.
- 511 Aditya Bommakanti, Harshith Reddy Vonteri, Konstantinos Skitsas, Sayan Ranu, Davide Mottin,  
512 and Panagiotis Karras. Fugal: Feature-fortified unrestricted graph alignment. In A. Globerson,  
513 L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Ad-*  
514 *vances in Neural Information Processing Systems*, volume 37, pp. 19523–19546. Curran Asso-  
515 ciates, Inc., 2024. URL [https://proceedings.neurips.cc/paper\\_files/paper/](https://proceedings.neurips.cc/paper_files/paper/2024/file/22b111819c74453837899689166c4cf9-Paper-Conference.pdf)  
516 [2024/file/22b111819c74453837899689166c4cf9-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/22b111819c74453837899689166c4cf9-Paper-Conference.pdf).
- 517 Vincenzo Bonnici, Rosalba Giugno, Alfredo Pulvirenti, Dennis Shasha, and Alfredo Ferro. A  
518 subgraph isomorphism algorithm and its application to biochemical data. *BMC bioinformatics*, 14:  
519 1–13, 2013.
- 520  
521 Vincenzo Carletti, Pasquale Foggia, Alessia Saggese, and Mario Vento. Challenging the time  
522 complexity of exact subgraph isomorphism for huge and dense graphs with vf3. *IEEE transactions*  
523 *on pattern analysis and machine intelligence*, 40(4):804–818, 2017.
- 524 Donatello Conte, Pasquale Foggia, Carlo Sansone, and Mario Vento. Thirty years of graph matching  
525 in pattern recognition. *International journal of pattern recognition and artificial intelligence*, 18  
526 (03):265–298, 2004.
- 527  
528 Luigi P Cordella, Pasquale Foggia, Carlo Sansone, and Mario Vento. A (sub) graph isomorphism algo-  
529 rithm for matching large graphs. *IEEE transactions on pattern analysis and machine intelligence*,  
530 26(10):1367–1372, 2004.
- 531 Luigi Pietro Cordella, Pasquale Foggia, Carlo Sansone, Mario Vento, et al. An improved algorithm  
532 for matching large graphs. In *3rd IAPR-TC15 workshop on graph-based representations in pattern*  
533 *recognition*, pp. 149–159. Citeseer, 2001.
- 534 Edsger W Dijkstra. A note on two problems in connexion with graphs. In *Edsger Wybe Dijkstra: his*  
535 *life, work, and legacy*, pp. 287–290. 2022.
- 536  
537 Khoa D Doan, Saurav Manchanda, Suchismit Mahapatra, and Chandan K Reddy. Interpretable graph  
538 similarity computation via differentiable optimal alignment of node embeddings. In *Proceedings*  
539 *of the 44th international ACM SIGIR conference on research and development in information*  
*retrieval*, pp. 665–674, 2021.

- 540 Chi Thang Duong, Trung Dung Hoang, Hongzhi Yin, Matthias Weidlich, Quoc Viet Hung Nguyen,  
541 and Karl Aberer. Efficient streaming subgraph isomorphism with graph neural networks. *Proceed-*  
542 *ings of the VLDB Endowment*, 14(5):730–742, 2021.
- 543 Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv*  
544 *preprint arXiv:1903.02428*, 2019.
- 545 Mahdi Hamzeh, Aviral Shrivastava, and Sarma Vrudhula. Epimap: Using epimorphism to map  
546 applications on cgras. In *Proceedings of the 49th Annual Design Automation Conference*, pp.  
547 1284–1291, 2012.
- 548 Myoungji Han, Hyunjoon Kim, Geonmo Gu, Kunsoo Park, and Wook-Shin Han. Efficient subgraph  
549 matching: Harmonizing dynamic programming, adaptive matching order, and failing set together.  
550 In *Proceedings of the 2019 international conference on management of data*, pp. 1429–1446, 2019.
- 551 Huahai He and Ambuj K Singh. Graphs-at-a-time: query language and access methods for graph  
552 databases. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of*  
553 *data*, pp. 405–418, 2008.
- 554 Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology*  
555 *and distribution*, pp. 492–518. Springer, 1992.
- 556 Alpár Jüttner and Péter Madarasi. Vf2++—an improved subgraph isomorphism algorithm. *Discrete*  
557 *Applied Mathematics*, 242:69–81, 2018.
- 558 Hyunjoon Kim, Yunyoung Choi, Kunsoo Park, Xuemin Lin, Seok-Hee Hong, and Wook-Shin  
559 Han. Versatile equivalences: Speeding up subgraph query processing and subgraph matching. In  
560 *Proceedings of the 2021 international conference on management of data*, pp. 925–937, 2021.
- 561 Jinha Kim, Hyungyu Shin, Wook-Shin Han, Sungpack Hong, and Hassan Chafi. Taming subgraph  
562 isomorphism for rdf query processing. *arXiv preprint arXiv:1506.01973*, 2015.
- 563 Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,  
564 2014.
- 565 Joseph B Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem.  
566 *Proceedings of the American Mathematical society*, 7(1):48–50, 1956.
- 567 O-Joun Lee et al. Transitivity-preserving graph representation learning for bridging local connectivity  
568 and role-based similarity. In *Proceedings of the AAAI Conference on Artificial Intelligence*,  
569 volume 38, pp. 12456–12465, 2024.
- 570 Andrei Leman and Boris Weisfeiler. A reduction of a graph to a canonical form and an algebra arising  
571 during this reduction. *Nauchno-Technicheskaya Informatsiya*, 2(9):12–16, 1968.
- 572 Giorgio Levi. A note on the derivation of maximal common subgraphs of two directed or undirected  
573 graphs. *Calcolo*, 9(4):341–352, 1973.
- 574 Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. Graph matching networks  
575 for learning the similarity of graph structured objects. In *International conference on machine*  
576 *learning*, pp. 3835–3845. PMLR, 2019.
- 577 Xuanzhou Liu, Lin Zhang, Jiaqi Sun, Yujiu Yang, and Haiqin Yang. D2match: leveraging deep learn-  
578 ing and degeneracy for subgraph matching. In *Proceedings of the 40th International Conference*  
579 *on Machine Learning, ICML’23*. JMLR.org, 2023.
- 580 Yan-li Liu, Chu-min Li, Hua Jiang, and Kun He. A learning based branch and bound for maximum  
581 common subgraph problems. *arXiv preprint arXiv:1905.05840*, 2019.
- 582 Zhaoyu Lou, Jiakuan You, Chengtao Wen, Arquimedes Canedo, Jure Leskovec, et al. Neural subgraph  
583 matching. *arXiv preprint arXiv:2007.03092*, 2020.
- 584 Andrew L Maas, Awni Y Hannun, Andrew Y Ng, et al. Rectifier nonlinearities improve neural  
585 network acoustic models. In *Proc. icml*, volume 30, pp. 3. Atlanta, GA, 2013.

- 594 James J McGregor. Relational consistency algorithms and their application in finding subgraph and  
595 graph isomorphisms. *Information Sciences*, 19(3):229–250, 1979.
- 596
- 597 Volodymyr Mnih. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*,  
598 2013.
- 599 Christopher Morris, Nils M Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion  
600 Neumann. Tudataset: A collection of benchmark datasets for learning with graphs. *arXiv preprint*  
601 *arXiv:2007.08663*, 2020.
- 602
- 603 Lam BQ Nguyen, Ivan Zelinka, Vaclav Snasel, Loan TT Nguyen, and Bay Vo. Subgraph mining in a  
604 large graph: A review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*,  
605 12(4):e1454, 2022.
- 606 Chengzhi Piao, Tingyang Xu, Xiangguo Sun, Yu Rong, Kangfei Zhao, and Hong Cheng. Computing  
607 graph edit distance via neural graph matching. *Proc. VLDB Endow.*, 16(8):1817–1829, April 2023.  
608 ISSN 2150-8097. doi: 10.14778/3594512.3594514. URL [https://doi.org/10.14778/  
609 3594512.3594514](https://doi.org/10.14778/3594512.3594514).
- 610 Martin L Puterman. Markov decision processes. *Handbooks in operations research and management*  
611 *science*, 2:331–434, 1990.
- 612
- 613 Vaibhav Raj, Indradyumna Roy, Ashwin Ramachandran, Soumen Chakrabarti, and Abir De. Charting  
614 the design space of neural graph representations for subgraph matching. In *The Thirteenth*  
615 *International Conference on Learning Representations*, 2025. URL [https://openreview.  
616 net/forum?id=5pd78GmXC6](https://openreview.net/forum?id=5pd78GmXC6).
- 617 Ashwin Ramachandran, Vaibhav Raj, Indradyumna Roy, Soumen Chakrabarti, and Abir De. It-  
618 eratively refined early interaction alignment for subgraph matching based graph retrieval. In  
619 A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.),  
620 *Advances in Neural Information Processing Systems*, volume 37, pp. 77593–77629. Curran Asso-  
621 ciates, Inc., 2024. URL [https://proceedings.neurips.cc/paper\\_files/paper/  
622 2024/file/8dc8a46f3981224217d32eb3f8362998-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/8dc8a46f3981224217d32eb3f8362998-Paper-Conference.pdf).
- 623 Romain Raveaux. On the unification of the graph edit distance and graph matching problems.  
624 *Pattern Recognition Letters*, 145:240–246, 2021. ISSN 0167-8655. doi: [https://doi.org/10.1016/  
625 j.patrec.2021.02.014](https://doi.org/10.1016/j.patrec.2021.02.014). URL [https://www.sciencedirect.com/science/article/  
626 pii/S0167865521000763](https://www.sciencedirect.com/science/article/pii/S0167865521000763).
- 627
- 628 Huan Rong, Tinghuai Ma, Meili Tang, and Jie Cao. A novel subgraph k+-isomorphism method in  
629 social network based on graph similarity detection. *Soft Computing*, 22(8):2583–2601, 2018.
- 630
- 631 Indradyumna Roy, Venkata Sai Baba Reddy Velugoti, Soumen Chakrabarti, and Abir De. Interpretable  
632 neural subgraph matching for graph retrieval. In *Proceedings of the AAAI conference on artificial*  
633 *intelligence*, volume 36, pp. 8115–8123, 2022.
- 634 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy  
635 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 636
- 637 Haichuan Shang, Ying Zhang, Xuemin Lin, and Jeffrey Xu Yu. Taming verification hardness: an  
638 efficient algorithm for testing subgraph isomorphism. *Proceedings of the VLDB Endowment*, 1(1):  
639 364–375, 2008.
- 640 Christine Solnon. Alldifferent-based filtering for subgraph isomorphism. *Artificial Intelligence*, 174  
641 (12-13):850–864, 2010.
- 642 Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel Ni,  
643 Heung-Yeung Shum, and Jian Guo. Think-on-graph: Deep and responsible reasoning of large  
644 language model on knowledge graph. In *The Twelfth International Conference on Learning*  
645 *Representations*, 2024. URL <https://openreview.net/forum?id=nnV01PvbTv>.
- 646
- 647 Shixuan Sun, Xibo Sun, Yulin Che, Qiong Luo, and Bingsheng He. Rapidmatch: A holistic approach  
to subgraph query processing. *Proceedings of the VLDB Endowment*, 14(2):176–188, 2020.

- 648 Jiabin Tang, Yuhao Yang, Wei Wei, Lei Shi, Lixin Su, Suqi Cheng, Dawei Yin, and Chao Huang.  
649 Graphgpt: Graph instruction tuning for large language models. In *Proceedings of the 47th*  
650 *International ACM SIGIR Conference on Research and Development in Information Retrieval*,  
651 SIGIR '24, pp. 491–500, New York, NY, USA, 2024. Association for Computing Machinery.  
652 ISBN 9798400704314. doi: 10.1145/3626772.3657775. URL [https://doi.org/10.1145/  
653 3626772.3657775](https://doi.org/10.1145/3626772.3657775).
- 654 Julian R Ullmann. An algorithm for subgraph isomorphism. *Journal of the ACM (JACM)*, 23(1):  
655 31–42, 1976.
- 656 Mnih Volodymyr et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):  
657 529–533, 2019.
- 658 Runzhong Wang, Ziao Guo, Shaofei Jiang, Xiaokang Yang, and Junchi Yan. Deep learning of partial  
659 graph matching via differentiable top-k. In *2023 IEEE/CVF Conference on Computer Vision and*  
660 *Pattern Recognition (CVPR)*, pp. 6272–6281, 2023. doi: 10.1109/CVPR52729.2023.00607.
- 661 Douglas Brent West et al. *Introduction to graph theory*, volume 2. Prentice hall Upper Saddle River,  
662 2001.
- 663 Keyulu Xu et al. How powerful are graph neural networks? In *7th International Conference*  
664 *on Learning Representations, ICLR 2019, New Orleans, LA, USA*, 2019. URL [https://  
665 openreview.net/forum?id=ryGs6iA5Km](https://openreview.net/forum?id=ryGs6iA5Km).
- 666 Rongjian Yang, Zhijie Zhang, Weiguo Zheng, and Jeffrey Xu Yu. Fast continuous subgraph matching  
667 over streaming graphs via backtracking reduction. *Proceedings of the ACM on Management of*  
668 *Data*, 1(1):1–26, 2023.
- 669 Ruosong Ye, Caiqi Zhang, Runhui Wang, Shuyuan Xu, and Yongfeng Zhang. Language is all a  
670 graph needs. In Yvette Graham and Matthew Purver (eds.), *Findings of the Association for*  
671 *Computational Linguistics: EACL 2024*, pp. 1955–1973, St. Julian’s, Malta, March 2024a.  
672 Association for Computational Linguistics. URL [https://aclanthology.org/2024.  
673 findings-eacl.132/](https://aclanthology.org/2024.findings-eacl.132/).
- 674 Yutong Ye, Xiang Lian, and Mingsong Chen. Efficient exact subgraph matching via gnn-based path  
675 dominance embedding. *Proceedings of the VLDB Endowment*, 17(7):1628–1641, 2024b.
- 676 Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. Graph transformer  
677 networks. *Advances in neural information processing systems*, 32, 2019.
- 678 Stéphane Zampelli, Yves Deville, and Christine Solnon. Solving subgraph isomorphism problems  
679 with constraint programming. *Constraints*, 15:327–353, 2010.
- 680 Shijie Zhang, Shirong Li, and Jiong Yang. Gaddi: distance index based subgraph matching in  
681 biological networks. In *Proceedings of the 12th international conference on extending database*  
682 *technology: advances in database technology*, pp. 192–203, 2009.
- 683 Zhijie Zhang, Yujie Lu, Weiguo Zheng, and Xuemin Lin. A comprehensive survey and experimental  
684 study of subgraph matching: trends, unbiasedness, and interaction. *Proceedings of the ACM on*  
685 *Management of Data*, 2(1):1–29, 2024.
- 686 Wei Zhuo and Guang Tan. Efficient graph similarity computation with alignment regularization.  
687 *Advances in Neural Information Processing Systems*, 35:30181–30193, 2022.
- 688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701

## A TECHNICAL APPENDICES AND SUPPLEMENTARY MATERIAL

### A.1 PROOFS OF PROPERTIES OF SUBGRAPH RELATIONSHIP

We provide proofs for the properties related to the subgraph relationship below, including reflexivity, transitivity, and antisymmetry.

**Theorem A.1** (Reflexivity).  $G \sqsubseteq G$

*Proof.* Let  $G = (V, E)$ . Since  $V \subseteq V$ ,  $E \subseteq \{(u, v) \in E \mid u, v \in V\}$ , it follows from the definition of the subgraph relationship that  $G \sqsubseteq G$ .  $\square$

**Theorem A.2** (Reflexivity of Graph Isomorphism).  $G \cong G$

*Proof.* Consider a bijection  $f : V \xrightarrow{\sim} V$  defined by  $f(v) = v$  for all  $v \in V$ . This satisfies

$$\forall v_i, v_j \in V, (v_i, v_j) \in E \iff (f(v_i), f(v_j)) = (v_i, v_j) \in E.$$

By the definition of graph isomorphism,  $G$  and  $G$  are isomorphic, i.e.,  $G \cong G$ .  $\square$

**Theorem A.3** (Transitivity).  $(G_1 \sqsubseteq G_2) \wedge (G_2 \sqsubseteq G_3) \implies G_1 \sqsubseteq G_3$ .

*Proof.* Let  $G_1 = (V_1, E_1)$ ,  $G_2 = (V_2, E_2)$ , and  $G_3 = (V_3, E_3)$ . From  $G_1 \sqsubseteq G_2$ , we have:

$$V_1 \subseteq V_2 \quad \text{and} \quad E_1 \subseteq \{(u, v) \in E_2 \mid u, v \in V_1\}.$$

Similarly, from  $G_2 \sqsubseteq G_3$ , we have:

$$V_2 \subseteq V_3 \quad \text{and} \quad E_2 \subseteq \{(u, v) \in E_3 \mid u, v \in V_2\}.$$

Since  $V_1 \subseteq V_2$  and  $V_2 \subseteq V_3$ , it follows that:

$$V_1 \subseteq V_3.$$

For edges, from  $E_1 \subseteq \{(u, v) \in E_2 \mid u, v \in V_1\}$  and  $E_2 \subseteq \{(u, v) \in E_3 \mid u, v \in V_2\}$ , we deduce:

$$\forall (u, v) \in E_1, (u, v) \in E_2 \subseteq E_3 \quad \text{with} \quad u, v \in V_1 \subseteq V_3.$$

Therefore, we conclude:

$$E_1 \subseteq \{(u, v) \in E_3 \mid u, v \in V_1\}.$$

$\square$

**Theorem A.4** (Antisymmetry).  $(G_1 \sqsubseteq G_2) \wedge (G_2 \sqsubseteq G_1) \iff G_1 = G_2$ .

*Proof.* Let  $G_1 = (V_1, E_1)$  and  $G_2 = (V_2, E_2)$ . If  $G_1 \sqsubseteq G_2$  and  $G_2 \sqsubseteq G_1$ , we have:

$$V_1 \subseteq V_2 \quad \text{and} \quad E_1 \subseteq \{(u, v) \in E_2 \mid u, v \in V_1\},$$

and

$$V_2 \subseteq V_1 \quad \text{and} \quad E_2 \subseteq \{(u, v) \in E_1 \mid u, v \in V_2\}.$$

From  $V_1 \subseteq V_2$  and  $V_2 \subseteq V_1$ , we conclude  $V_1 = V_2$ . Similarly,  $E_1 \subseteq E_2$  and  $E_2 \subseteq E_1$  imply  $E_1 = E_2$ . Therefore, we have  $G_1 = G_2$ .

Conversely, if  $G_1 = G_2$ , then we have:

$$V_1 = V_2 \quad \text{and} \quad E_1 = \{(u, v) \in E_2 \mid u, v \in V_1\},$$

and

$$V_2 = V_1 \quad \text{and} \quad E_2 = \{(u, v) \in E_1 \mid u, v \in V_2\}.$$

Therefore,  $G_1 \sqsubseteq G_2$  and  $G_2 \sqsubseteq G_1$ .  $\square$

## A.2 COROLLARIES OF SUBGRAPH ISOMORPHISM RELATIONSHIP

**Corollary A.5** (Subgraph Isomorphism Partial Order). *Building upon the properties mentioned in Section 3.2, the subgraph isomorphism relation  $\preceq$ :*

$$G_1 \preceq G_2 \iff \exists G' = (V', E') \sqsubseteq G_2 \text{ s.t. } G_1 \cong G' \quad (13)$$

*satisfies the fundamental characteristics of a partial ordering on the graph space  $\mathcal{G}$ :*

- **Reflexivity:**

$$\forall G \in \mathcal{G}, G \preceq G$$

- **Transitivity:**

$$(G_1 \preceq G_2) \wedge (G_2 \preceq G_3) \implies G_1 \preceq G_3$$

- **Antisymmetry:**

$$(G_1 \preceq G_2) \wedge (G_2 \preceq G_1) \iff G_1 \cong G_2$$

We provide proofs for the properties related to the subgraph isomorphism below, including reflexivity, transitivity, and antisymmetry.

**Theorem A.6** (Reflexivity).  $\forall G \in \mathcal{G}, G \preceq G$

*Proof.* By Theorem A.1, we have  $G \sqsubseteq G$ ; by Theorem A.2, we have  $G \cong G$ , then we have  $G \preceq G$ .  $\square$

**Theorem A.7** (Transitivity).  $(G_1 \preceq G_2) \wedge (G_2 \preceq G_3) \implies G_1 \preceq G_3$

*Proof.* Let  $G_1 = (V_1, E_1)$ ,  $G_2 = (V_2, E_2)$ , and  $G_3 = (V_3, E_3)$ . Given  $G_1 \preceq G_2$ , there exists a mapping  $f_1 : V_1 \rightarrow V_2$  such that

$$\forall v_i, v_j \in V_1, (v_i, v_j) \in E_1 \implies (f_1(v_i), f_1(v_j)) \in E_2.$$

Similarly, from  $G_2 \preceq G_3$ , there exists a mapping  $f_2 : V_2 \rightarrow V_3$  satisfying

$$\forall v_i, v_j \in V_2, (v_i, v_j) \in E_2 \implies (f_2(v_i), f_2(v_j)) \in E_3.$$

Define a composite mapping  $g : V_1 \rightarrow V_3 = f_1 \circ f_2$  as  $g(v) = f_2(f_1(v))$ . Then:

$$\forall v_i, v_j \in V_1, (v_i, v_j) \in E_1 \implies (f_1(v_i), f_1(v_j)) \in E_2 \implies (g(v_i), g(v_j)) \in E_3.$$

Construct  $G'_3 = (\{g(v) \mid v \in V_1\}, \{(g(v_i), g(v_j)) \mid v_i, v_j \in V_1\})$ . This satisfies:

$$G'_3 \sqsubseteq G_3 \text{ with } G_1 \cong G'_3.$$

Therefore,  $G_1 \preceq G_3$  holds.  $\square$

**Theorem A.8** (Antisymmetry).  $(G_1 \preceq G_2) \wedge (G_2 \preceq G_1) \iff G_1 \cong G_2$

*Proof.* Let  $G_1 = (V_1, E_1)$ , and  $G_2 = (V_2, E_2)$ . Given  $G_1 \preceq G_2$ , there exists a mapping  $f_1 : V_1 \rightarrow V_2$  such that

$$\forall v_i, v_j \in V_1, (v_i, v_j) \in E_1 \implies (f_1(v_i), f_1(v_j)) \in E_2.$$

Similarly, from  $G_2 \preceq G_1$ , there exists a mapping  $f_2 : V_2 \rightarrow V_1$  satisfying

$$\forall v_i, v_j \in V_2, (v_i, v_j) \in E_2 \implies (f_2(v_i), f_2(v_j)) \in E_1.$$

The mappings  $f_1$  and  $f_2$  induce a bijective correspondence  $f' : V_1 \xrightarrow{\sim} V_2$  that preserves adjacency:

$$\forall v_i, v_j \in V_1, (v_i, v_j) \in E_1 \iff (f'(v_i), f'(v_j)) \in E_2 \quad (14)$$

Therefore,  $G_1 \cong G_2$  holds.

Conversely, if  $G_1 \cong G_2$ , there exists a bijective mapping  $f : V_1 \xrightarrow{\sim} V_2$  satisfying:

$$\forall v_i, v_j \in V_1, (v_i, v_j) \in E_1 \iff (f(v_i), f(v_j)) \in E_2 \quad (15)$$

The inverse mapping  $f^{-1} : V_2 \xrightarrow{\sim} V_1$  consequently preserves:

$$\forall v_i, v_j \in V_2, (v_i, v_j) \in E_2 \iff (f^{-1}(v_i), f^{-1}(v_j)) \in E_1$$

establishing  $G_2 \cong G_1$ .

Through constructive verification:

- 810 • Let  $G' = G_2$ . By definition:

811 
$$G' \sqsubseteq G_2 \quad \text{and} \quad G_1 \cong G' \implies G_1 \preceq G_2$$

812

- 813 • Let  $G'' = G_1$ . Similarly:

814 
$$G'' \sqsubseteq G_1 \quad \text{and} \quad G_2 \cong G'' \implies G_2 \preceq G_1$$

815

816 For isomorphic graphs  $G_1 \cong G_2$ , the partial order relation becomes symmetric:

817 
$$(G_1 \preceq G_2) \wedge (G_2 \preceq G_1) \iff G_1 \cong G_2$$

818 □

819

820 By Corollary A.5, the subgraph isomorphism relation satisfies the partial order properties. Conse-  
821 quently, the partial order-aware GNN proposed in this work demonstrates superior effectiveness in  
822 solving subgraph isomorphism problems.

### 824 A.3 PROOF OF PARTIAL ORDER-AWARE GNN

825

826 We demonstrate that the order constraints in Equation 8 are preserved under the composition of  
827 multiple message-passing layers in GNNs, particularly in simple models like the Graph Isomorphism  
828 Network (GIN) Xu et al. (2019).

829 **Proof Strategy:** By leveraging mathematical induction, we prove that GIN-style models inherently  
830 maintain these properties. Consider a  $k$ -layer GNN encoding nodes  $u$  (in search graphs) and  $v$  (in  
831 query graphs):

- 832
- 833 1. **Base Case:** For the trivial case of a single node  $v$  (i.e., 0-hop neighborhood where  
834  $\mathcal{N}_0(v) = \{v\}$ ), the partial ordering relation  $\preceq$  trivially satisfies:

835 
$$\forall v \in V, \mathbf{h}_v^{(0)} \preceq \mathbf{h}_v^{(0)}$$

836 This follows directly from the reflexive property of partial orders. The 0-hop neighborhood  
837 contains only the node itself, making the order embedding comparison degenerate to self-  
838 comparison.

- 839 2. **Inductive Step:** If the  $k$ -hop neighborhood of  $u$  forms a subgraph of  $v$ 's  $k$ -hop neighborhood  
840 ( $\mathcal{N}_k(u) \sqsubseteq \mathcal{N}_k(v)$ ), then:

841 
$$\forall s \in \mathcal{N}(v), \exists t \in \mathcal{N}(u) \text{ s.t. } \mathcal{N}_{k-1}(s) \sqsubseteq \mathcal{N}_{k-1}(t)$$

842

843 The  $(k-1)$ -hop neighborhoods of  $u$ 's neighbors are subgraphs of corresponding neigh-  
844 borhoods of  $v$ 's neighbors. This inductive process guarantees the preservation of order  
845 constraints through layer composition.

846 **Order Embedding Guarantee:** Suppose all GNN embeddings at layer  $(k-1)$  satisfy order  
847 constraints after transformation. Using sum-based neighborhood aggregation:

848 
$$\mathbf{h}_v^{(k)} = f_{\text{agg}} \left( \left\{ \mathbf{h}_u^{(k-1)} \mid u \in \mathcal{N}(v) \right\} \right) \quad (16)$$

849

850 where  $f_{\text{agg}}$  is order-preserving under summation. Then:

851 
$$\mathbf{h}_v^{(k)} \preceq \mathbf{h}_{v'}^{(k)} \text{ if } \forall u \in \mathcal{N}(v), \exists w \in \mathcal{N}(v') \text{ with } \mathbf{h}_u^{(k-1)} \preceq \mathbf{h}_w^{(k-1)} \quad (17)$$

852

853 **Subgraph Composition Property:** This corresponds to the fundamental property of subgraph  
854 composition:

855 **Theorem A.9** (Order Preservation). *Given GIN's update rule:*

856 
$$\mathbf{h}_v^{(l+1)} = \text{MLP}^{(l)} \left( (1 + \epsilon^{(l)}) \cdot \mathbf{h}_v^{(l)} + \sum_{u \in \mathcal{N}(v)} \mathbf{h}_u^{(l)} \right) \quad (18)$$

857

858

859 and the graph embedding by concatenating sum-pooled features across layers:

860 
$$h_G = \sum_{l=0}^L \gamma_l \cdot \sum_{v \in V} h_v^{(l)} \quad (19)$$

861

862

863

the GNN preserves partial order relationships between subgraphs.

864 *Proof.* The MLP’s universal approximation capability combined with sum aggregation maintains the  
865 partial order structure through Lipschitz continuity. The loss function explicitly enforces:  
866

$$867 \mathbf{h}_q \preceq \mathbf{h}_t \iff G_q \sqsubseteq G_t$$

868 making the embedding space order-isomorphic to the subgraph lattice.  $\square$   
869

870 **Definition A.10** (Max-Margin Order Loss). The structural constraint is enforced through our margin-  
871 based loss function:

$$872 \mathcal{L}(\mathbf{h}_q, \mathbf{h}_t) = \sum_{(\mathbf{h}_q, \mathbf{h}_t) \in P} E(\mathbf{h}_q, \mathbf{h}_t) \\ 873 + \sum_{(\mathbf{h}_q, \mathbf{h}_t) \in N} \max(0, \alpha - E(\mathbf{h}_q, \mathbf{h}_t)) \quad (20)$$

874 where

$$875 E(\mathbf{h}_q, \mathbf{h}_t) = \|\max(0, \mathbf{h}_q - \mathbf{h}_t)\|_2^2 \quad (21)$$

876 **Proposition A.11** (Learning Dynamics). *The loss design induces asymmetric gradient signals:*

- 877 • For **positive pairs**  $P = \{(q, t) \in \mathcal{MB} \mid \mathcal{N}(q) \sqsubseteq \mathcal{N}(t)\}$ , gradients dominate in dimensions  
878 where  $h_q^i > h_t^i$
- 879 • For **negative pairs**  $N = \mathcal{MB} \setminus P$ , gradients activate only when  $E(\mathbf{h}_q, \mathbf{h}_t) > \alpha$

880 *This creates faster convergence for order-satisfying pairs compared to violating ones.*

881 **Theorem A.12** (Order-Preserving Embedding). *After sufficient training iterations, the GIN embed-  
882 dings satisfy the vector partial order:*

$$883 \forall i \in \{1, \dots, D\}, \mathbf{h}_{G_1}^i \leq \mathbf{h}_{G_2}^i \iff G_1 \sqsubseteq G_2 \quad (22)$$

#### 884 A.4 PROBABILISTIC COMPLETENESS

885 We establish that SIREN guarantees to either find an exact solution or prove the non-existence of a  
886 solution given sufficient time. This probabilistic completeness relies on three fundamental properties:  
887 a finite search space, persistent exploration, and Markov chain ergodicity.

888 **1) Finite State Space:** The search space forms a finite state transition graph where each state  $s$   
889 represents a partial isomorphism mapping  $\mathcal{M} \subseteq V_Q \times V_T$  that satisfies the required topologi-  
890 cal constraints. For  $|V_Q| = n$  and  $|V_T| = m$ , the cardinality of the state space is bounded by  
891  $\sum_{k=0}^n \binom{n}{k} P(m, k) = O(m^n)$ , ensuring finiteness. Our state encoding scheme explicitly tracks  
892 partial mappings to preserve this property.

893 **2) Non-Zero Action Probabilities:** SIREN employs an  $\epsilon$ -greedy policy with a non-decaying  $\epsilon$ ,  
894 ensuring that the probability of selecting any valid branch remains strictly positive. The DQN is  
895 used solely to prioritize branch exploration (e.g., by favoring branches with higher Q-values) without  
896 performing irreversible pruning; even branches with low Q-values remain explorable in subsequent  
897 steps. Furthermore, the algorithm does not permanently exclude any branch due to state memorization.  
898 Consequently, the probability of selecting any feasible path branch remains non-zero.

899 **3) Ergodic Markov Chain:** The search process constitutes a finite Markov chain  $(S, P)$  with transi-  
900 tion probabilities  $P(s'|s) = \sum_{a: s \rightarrow s'} \pi_\theta(a|s)$ . This chain is irreducible (enabled by backtracking  
901 actions that permit state revisitation) and aperiodic (due to the existence of self-transitions). By the  
902 ergodic theorem Bellet (2006), the chain visits all states with probability 1 over infinite time.

903 **4) Robustness:** Pathological cases are mitigated through the following mechanisms:

- 904 • Permanent availability of all valid actions ( $\pi_\theta(a|s) > 0$  for any valid  $a$ );
- 905 • Explicit backtracking mechanisms;
- 906 • Absence of irreversible pruning, with the exception of Bidomain-based pruning that defini-  
907 tively precludes solution existence.

918 These features ensure exhaustive exploration even for deceptive graph structures.

919 In summary, given the finite state space, persistent exploration via strictly positive transition proba-  
920 bilities, ergodic properties, and robust backtracking mechanisms, SIREN guarantees probabilistic  
921 completeness.

#### 922 A.5 TRAINING PROCESS OF THE RL-ENHANCED SUBGRAPH ISOMORPHISM SOLVER

923 The training process consists of three consecutive phases: pretraining, imitation learning, and  
924 reinforcement learning. Each phase follows the same core workflow of experience collection followed  
925 by network updates.

926 **Pretraining Phase** (1250 iterations): Heuristic methods (specifically, the traditional subgraph  
927 matching algorithm VF3 Carletti et al. (2017)) are used to generate search trajectories and populate  
928 the experience replay buffer. This approach avoids the inefficiency of initial random exploration. The  
929 objective is to provide the network with high-quality initial samples of effective matching patterns,  
930 establishing a solid foundation for subsequent reinforcement learning.

931 **Imitation Learning Phase** (2500 iterations): The agent continues to use VF3 Carletti et al. (2017) to  
932 generate demonstration trajectories. However, the DQN now begins to learn by imitating the Q-value  
933 distribution of this heuristic policy. The goal is for the network to quickly converge to a performance  
934 level comparable to that of the traditional algorithm, thereby reducing the cost of exploration.

935 **Reinforcement Learning Phase** (6250 iterations): The agent relies entirely on the DQN’s  $\epsilon$ -greedy  
936 policy for exploration. The Q-values are optimized through environmental feedback (rewards),  
937 allowing the policy to be progressively refined. The final objective is to enable the network to surpass  
938 the performance of traditional heuristic methods by discovering more efficient matching paths.

---

#### 939 **Algorithm 1** Training Process of the RL-Enhanced Subgraph Isomorphism Solver

---

```

940 1: Initialize replay buffer  $\mathcal{D}$ , policy network  $\pi_\theta$ , target network  $\pi_{\theta^-}$ 
941 2: for episode = 1 to  $M$  do
942 3:   Sample graph pair  $(G_Q, G_T)$ , initialize state  $s_0$ 
943 4:   for  $t = 1$  to  $T_{\max}$  do
944 5:     Select action  $a_t$  via  $\epsilon$ -greedy:  $a_t \sim \pi_\theta(s_t)$ 
945 6:     Execute  $a_t$ , observe  $s_{t+1}$ , reward  $r_t$ 
946 7:     Store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{D}$ 
947 8:     Sample batch  $\mathcal{D} \sim \mathcal{D}$ , update  $\theta$  via DQN loss  $\mathcal{L}$ 
948 9:     Update target network:  $\theta^- \leftarrow \tau\theta + (1 - \tau)\theta^-$ 
949 10:  end for
950 11: end for

```

---

951 SIREN learns via deep Q-learning with prioritized experience replay Mnih (2013), as outlined in  
952 Algorithm 1. The framework maintains three core components: a replay buffer  $\mathcal{D}$  storing transition  
953 trajectories  $(s_t, a_t, r_t, s_{t+1})$ , a policy network  $\pi_\theta$  parameterized by  $\theta$  for action selection, and a target  
954 network  $\pi_{\theta^-}$  with delayed weight updates to stabilize training. The core training loop in each phase  
955 consists of two main stages: *experience collection* (populating the replay buffer) and *network update*  
956 (learning from the buffer).

957 The experience collection stage involves the following steps:

- 958 • Search initialization
- 959 • Action selection and search expansion
- 960 • Reward calculation and experience storage

961 The network update stage comprises the following steps:

- 962 • Experience sampling from the replay buffer
- 963 • Target Q-value computation
- 964 • Loss calculation and parameter optimization

- Periodic synchronization of the target network

At the start of each episode, a graph pair  $(G_Q, G_T)$  is sampled from the training distribution  $\mathbb{P}_{\text{train}}$ , and an initial state  $s_0 = \text{InitState}(G_Q, G_T)$  is constructed. The agent interacts with the environment for  $T_{\text{max}}$  timesteps using an  $\epsilon$ -greedy exploration strategy: with probability  $\epsilon$ , random actions are selected from the available action space  $\mathcal{A}_t$ , otherwise the optimal action  $a_t = \arg \max_a Q_\theta(s_t, a)$  is chosen by the policy network. Executing action  $a_t$  yields a new state  $s_{t+1}$  and immediate reward  $r_t$ , with the transition tuple  $(s_t, a_t, r_t, s_{t+1})$  stored in  $\mathcal{D}$  for subsequent learning.

Parameter updates are performed through minimization of the Huber loss  $\mathcal{L}(\theta)$  Huber (1992):

$$\mathbb{E}_{(s,a,r,s') \sim \mathcal{D}} \left[ \left( r + \gamma \max_{a'} Q_{\theta^-}(s', a') - Q_\theta(s, a) \right)^2 \right] \quad (23)$$

where  $\gamma$  denotes the discount factor, and  $r = +1$  for the immediate reward. The target network parameters  $\theta^-$  are softly updated using the rule:

$$\theta^- \leftarrow \tau \theta + (1 - \tau) \theta^- \quad \text{with } \tau \ll 1 \quad (24)$$

## A.6 TRAINING METHODOLOGY OF THE PARTIAL-ORDER-AWARE GNN

The training protocol for our partial-order-aware GNN comprises two coordinated phases: (1) *Training Data Generation* and (2) *Loss-Driven Optimization*.

### A.6.1 TRAINING DATA GENERATION

---

#### Algorithm 2 Contrastive Subgraph Sampling

---

- 1: Sample anchor node  $u \in V_T$  from target graph  $G_T$
  - 2: Generate  $G_u \subseteq G_T$  via randomized BFS with edge traversal probability  $p = 0.8$
  - 3: Construct query graph  $G_q$  by reapplying same BFS protocol on  $G_u$  anchored at  $u$
- Ensure:**  $G_q \subseteq G_u$  (preserves subgraph isomorphism via construction)
- 4: Generate negative pairs:
    - 5: **Type I:** Random anchor  $u' \neq u$  with BFS-generated  $G_{u'}$
    - 6: **Type II:** Perturb  $G_q$  via edge deletions/additions violating  $G_q \subseteq G_T$
- 

The data generation process ensures diverse yet controlled learning signals:

- **Positive Pairs:** For each target subgraph  $G_u \subseteq G_T$ , we systematically construct isomorphic queries  $G_q$  through duplicate randomized breadth-first search (BFS) traversals with edge sampling probability  $p = 0.8$ . This procedural generation guarantees  $G_q \subseteq G_u$  by design.
- **Negative Pairs:** We implement two challenging negative sampling strategies:
  1. *Non-anchored Negatives:* Random anchor selection with independent BFS generations that break subgraph relationships
  2. *Structurally Damaged Negatives:* Adversarial edge perturbations (15% edge flip probability) that invalidate subgraph isomorphism

### A.6.2 LOSS COMPUTATION AND OPTIMIZATION

Let  $\mathbf{h}_q, \mathbf{h}_u$  denote GNN embeddings of query  $G_q$  and target subgraph  $G_u$ . The contrastive loss from Definition A.10 is computed as:

$$\mathcal{L} = \sum_{(q,u) \in P} \|\max(0, \mathbf{h}_q - \mathbf{h}_u)\|^2 + \sum_{(q,u') \in N} \max(0, \alpha - \|\mathbf{h}_q - \mathbf{h}_{u'}\|^2) \quad (25)$$

Backpropagation updates both the GNN parameters  $\theta$  and the anchor-aware embedding space through:

$$\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L} \quad (26)$$

where  $\eta$  is the learning rate.

**Algorithm 3** SIREN’s RL-Enhanced Search Framework.

---

```

1026 1: Input: Query Graph  $G_Q$ , Target Graph  $G_T$ 
1027 2: Output: Solution Set  $\mathcal{S}_G = \{G' \sqsubseteq G_T \mid G_Q \cong G'\}$ 
1028 3:  $\mathcal{ST} \leftarrow \phi, \mathcal{S}_G \leftarrow \phi$ 
1029 4:  $\mathcal{ST}.\text{push}(\epsilon)$ 
1030 5: while  $\mathcal{ST} \neq \phi$  do
1031 6:    $s_t \leftarrow \mathcal{ST}.\text{pop}()$ 
1032 7:    $curG \leftarrow s_t.\text{getMCS}()$ 
1033 8:   if  $|V(curG)| = |V(G_Q)|$  then
1034 9:      $\mathcal{S}_G \leftarrow \mathcal{S}_G \cup curG$ 
1035 10:    continue
1036 11:   end if
1037 12:    $UB \leftarrow |curG| + \text{bound}(s_t)$ 
1038 13:   if  $|UB| \leq |V(G_Q)|$  or  $|s_t.\text{actions}| = 0$  then
1039 14:     continue
1040 15:   end if
1041 16:    $a_t^* \leftarrow \text{Policy}(s_t)$ 
1042 17:    $s_{t+1} \leftarrow \text{Action}(s_t, a_t^*)$ 
1043 18:    $s_t.\text{actions} \leftarrow s_t.\text{actions} \setminus \{a_t^*\}$ 
1044 19:    $\mathcal{ST}.\text{push}(s_t)$ 
1045 20:    $\mathcal{ST}.\text{push}(s_{t+1})$ 
1046 21: end while

```

---

## A.7 REINFORCEMENT LEARNING-ENHANCED HEURISTIC SEARCH

In SIREN, we established a reinforcement learning-enhanced heuristic search framework for subgraph isomorphism by augmenting a classical branch-and-bound maximum common subgraph search paradigm. As formalized in Algorithm 3, SIREN maintains a state heap  $\mathcal{ST}$  to track partial mapping states during subgraph isomorphism search. At each iteration, the algorithm retrieves the top state  $s_t = \mathcal{M}_t$ , where  $\mathcal{M}_t$  represents the current node mapping set between the pattern graph  $G_Q$  and the target graph  $G_T$ . The subgraph  $curG$  corresponding to  $\mathcal{M}_t$  is then extracted which satisfies both  $curG \sqsubseteq G_T$  and  $curG \sqsubseteq G_Q$ . A critical pruning decision is made based on a bidomain-estimated upper bound  $UB$ , as detailed in Section A.8. If this bound falls below the number of nodes of  $G_Q$ , i.e.  $|V(G_Q)|$ , the branch is pruned. When  $curG$  achieves full isomorphism with  $G_Q$  ( $curG \cong G_Q$ ), it is added to the solution set  $\mathcal{S}_G$ .

## A.8 ADDITIONAL EMBEDDINGS AND UPPERBOUND ESTIMATION VIA BIDOMAINS

In our methodology, inspired by GLSearch Bai et al. (2021), bidomains are introduced to provide more information for our partial-order aware GIN model, and facilitate pruning through estimating the upper bound  $UB$ . For a given state  $s_t = \mathcal{M}_t$ , the  $k$ -th bidomain  $\mathcal{B}_k$  is defined as:

$$\mathcal{B}_k = \langle V_{kQ}, V_{kT} \rangle \quad (27)$$

where  $V_{kQ} \subseteq V(G_Q)$  and  $V_{kT} \subseteq V(G_T)$  exhibit identical connectivity patterns with respect to the already matched node sets  $\mathcal{M}_t = \langle V_Q^s, V_T^s \rangle$ ,  $V_Q^s \subseteq V(G_Q)$  and  $V_T^s \subseteq V(G_T)$ . For the state  $s_t$  with  $n$  matched node pairs, i.e.,

$$\begin{cases} |V_Q^s| = |V_T^s| = n \\ V_Q^s = \{v_{Q1}^s, v_{Q2}^s, \dots, v_{Qn}^s\} \\ V_T^s = \{v_{T1}^s, v_{T2}^s, \dots, v_{Tn}^s\} \end{cases} \quad (28)$$

there are a total of  $2^n$  bidomains. The  $2^n$  bidomains are  $\mathcal{B}_{(d_1 \dots d_n)_2} = \langle V_{(d_1 \dots d_n)_2Q}, V_{(d_1 \dots d_n)_2T} \rangle$ , where  $d_1, \dots, d_n \in \{0, 1\}$ . For  $i \in \{1, 2, \dots, n\}$ , The nodes in  $\mathcal{B}_{(d_1 \dots d_n)_2}$  satisfy:

$$\begin{cases} d_i = 1 \iff \forall v \in V_{(d_1 \dots d_n)_2Q}, \exists e = (v, v_{Qi}^s) \in E(G_Q) \\ d_i = 0 \iff \forall v \in V_{(d_1 \dots d_n)_2Q}, \nexists e = (v, v_{Qi}^s) \in E(G_Q) \\ d_i = 1 \iff \forall v \in V_{(d_1 \dots d_n)_2T}, \exists e = (v, v_{Ti}^s) \in E(G_T) \\ d_i = 0 \iff \forall v \in V_{(d_1 \dots d_n)_2T}, \nexists e = (v, v_{Ti}^s) \in E(G_T) \end{cases} \quad (29)$$

1080  
1081  
1082  
1083  
1084  
1085  
1086  
1087  
1088  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

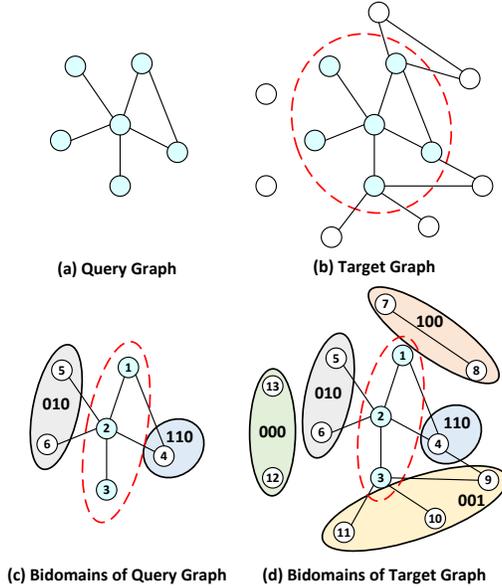


Figure 4: Examples of subgraph isomorphism and bidomains. (a) Query graph  $G_Q$ , (b) Target graph  $G_T$ , (c-d) Bidomains of  $G_Q$  and  $G_T$ .

Figure 4 presents an example of subgraph isomorphism and bidomains. Figure 4(a) shows the query graph  $G_Q$ , while Figure 4(b) displays the target graph  $G_T$ . Figures 4(c) and 4(d) illustrate the bidomains when the currently matched nodes are  $\langle v_{Q1}, v_{T1} \rangle$ ,  $\langle v_{Q2}, v_{T2} \rangle$ , and  $\langle v_{Q3}, v_{T3} \rangle$ . As shown in Figures 4(c), the nodes  $v_{Q5}$  and  $v_{Q6}$  are in the set  $V_{(010)_2Q}$ , and the node  $v_{Q4}$  is in the set  $V_{(110)_2Q}$ . As shown in Figures 4(d), the nodes  $v_{T5}$  and  $v_{T6}$  are in the set  $V_{(010)_2T}$ , the node  $v_{T4}$  is in the set  $V_{(110)_2T}$ .  $v_{T7}$  and  $v_{T8}$  are in  $V_{(100)_2T}$ .  $v_{T9}$ ,  $v_{T10}$  and  $v_{T11}$  are in  $V_{(001)_2T}$ .  $v_{T12}$  and  $v_{T13}$  are in  $V_{(000)_2T}$ .

To estimate the upperbound, note that each bidomain can contribute at most  $\min(|V_{kQ}|, |V_{kT}|)$  nodes to the future best solution. Therefore, the upperbound can be estimated as:

$$\sum_{B_k \in \mathcal{B}} \min(|V_{kQ}|, |V_{kT}|) \tag{30}$$

which is the `bound()` function in Algorithm 3.

We also incorporate the bidomain information as node labels into our partial-order-aware GNN. The resulting embeddings are used to enhance the input to our DQN model.

### A.9 EXAMPLE OF THE SUBGRAPH ISOMORPHISM

Below, we give an example of the subgraph isomorphism in real applications of loop mapping in coarse-grained reconfigurable architectures (CGRAs).

**Example 1 (Loop mapping):** In CGRA loop mapping, Figure 5 illustrates a  $2 \times 2$  processing element (PE) array case where subgraph isomorphism verifies if the compiler-generated data flow graph (DFG) (Figure 5(a)) embeds into the time-extended CGRA (TEC) model (Figure 5(b)) under a target initiation interval (II). This spatiotemporal mapping solution (Figure 5(c)) explicitly defines PE operations per cycle and inter-PE routing, systematically optimizing parallelism and resource utilization through isomorphic correspondence between DFG and TEC structures.

### A.10 STATISTICS OF REAL-WORLD GRAPH DATASETS

In Section 4.4, we selected the same datasets as GNN-PE Ye et al. (2024b) from TUDataset Morris et al. (2020) for testing, i.e., Yeast, Human, HPRD, WordNet, DBLP, Youtube, and US Patents. Statistics of these real graphs are summarized in Table 4.

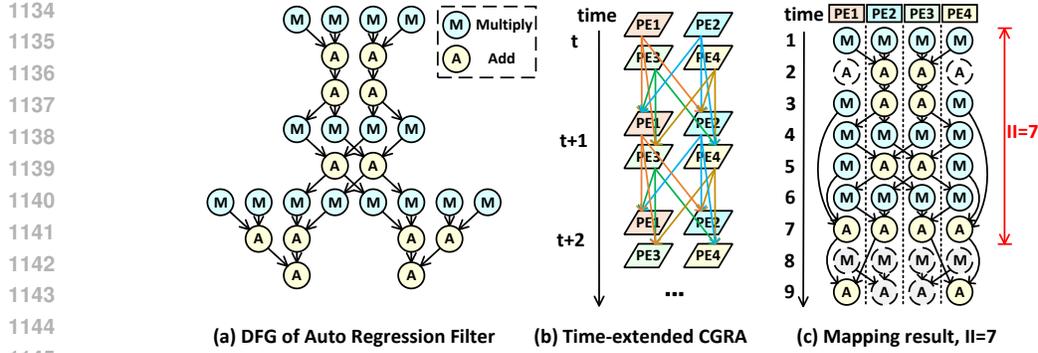


Figure 5: Example of loop mapping on CGRA with subgraph isomorphism. (a) DFG of Auto Regression Filter, (b) Time-extended CGRA of a  $2 \times 2$  PEA, (c) mapping result,  $II = 2$ .

Table 4: Statistics of Real-World Graph Datasets.

Datasets	$ V(G) $	$ E(G) $	$ \Sigma $	$\overline{deg}(G)$
Yeast	3112	12519	71	8.0
Human	4674	86282	44	36.9
HPRD	9460	34998	307	7.4
WordNet	76853	120399	5	3.1
DBLP	317080	1049866	15	6.6
Youtube	1134890	2987624	25	5.3
US Patents	3774768	16518947	20	8.8

#### A.11 RL FRAMEWORK ABLATION STUDY AND SCALABILITY BENCHMARKING

To analyze the impact of different reinforcement learning paradigms, we conducted comparisons between Deep Q-Networks (DQN) and Proximal Policy Optimization (PPO) Schulman et al. (2017). Our baseline DQN implementation follows the parametrization:

$$\begin{aligned}
 Q_{\theta}(s_t, a_t) &= \mathcal{F} \left( \underbrace{\text{GNN}_{\text{enc}}(G_Q, G_T, s_t)}_{\text{state embedding } \mathbf{h}_{s_t} \in \mathbb{R}^d} \oplus \underbrace{\mathbf{W}_a \phi(a_t)}_{\text{action embedding } \mathbf{h}_{a_t} \in \mathbb{R}^d} \right) \\
 &= \text{MLP}(\mathbf{h}_{s_t} \parallel \mathbf{h}_{G_Q} \parallel \mathbf{h}_{G_T} \parallel \mathbf{h}_{a_t})
 \end{aligned} \tag{31}$$

where  $\oplus$  denotes vector concatenation. For comparison, we adapt the PPO framework with dual-network architecture:

$$\begin{aligned}
 \pi_{\theta}(a_t | s_t) &= \text{Softmax} \left( \mathbf{W}_{\pi} \cdot \text{GNN}_{\text{enc}}(G_Q, G_T, s_t) \right) \\
 V_{\phi}(s_t) &= \mathbf{w}_v^{\top} \cdot \text{GNN}_{\text{enc}}(G_Q, G_T, s_t)
 \end{aligned} \tag{32}$$

*Remark A.13.* The key difference lies in DQN’s *action-conditioned* Q-function versus PPO’s *state-conditioned* policy distribution. Our GNN encoder maintains identical architecture across frameworks to isolate RL algorithm effects.

We validate scalability on a million-scale graph from the US Patent dataset ( $|V| = 3,774,768$ ). Through random bipartition, we generate query-target pairs with matched cardinality ( $|V_Q| = |V_T|$ ). The MCS identification time complexity is measured for four methods: GLSearch Bai et al. (2021), GNN-PE Ye et al. (2024b), SIREN and PPO-based SIREN (SIREN-PPO). The results are detailed in 5.

Despite being optimized for subgraph isomorphism, SIREN achieves  $3.83 \times$  speedup over GLSearch at 5K-node scale. The performance gap widens exponentially with graph size ( $R^2 = 0.98$  for

Table 5: Time cost (seconds) for MCS identification across graph sizes. Bold: best performance.

Method	Graph Size (Nodes)				
	1K	2K	3K	4K	5K
GLSEARCH	364.5	664.6	1820.2	2163.2	21507.2
GNN-PE	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$
SIREN	<b>165.4</b>	<b>357.3</b>	<b>639.8</b>	<b>939.8</b>	<b>5612.3</b>
SIREN-PPO	245.8	518.1	888.6	1119.6	7241.6

quadratic fit), demonstrating our method’s superior asymptotic complexity. The GNN-PE’s failure highlights the necessity of partial-order-preserving architectures for structural tasks.

**Assumption A.14** (RL Framework Efficacy). The superior empirical performance of our DQN-based framework compared to the PPO variant (Table 5) stems from two synergistic factors:

- **Partial-Order Preservation:** The  $\preceq$ -preserving GNN embeddings intrinsically capture subgraph containment relationships through the lens of order embeddings  $\mathbf{h}_G \in \mathbb{R}_+^D$  (Theorem A.12)
- **Compatibility with Value Iteration:** The subgraph isomorphism search dynamics are naturally expressible through Q-learning’s state-action value formulation:

$$Q(s, a) = \mathbb{E} \left[ R(s, a) + \gamma \max_{a'} \underbrace{\langle \mathbf{h}_{s'}, \mathbf{h}_{G_T} \rangle}_{\text{order alignment}} \right] \quad (33)$$

where  $R(s, a)$  encodes topological validity rewards

**Assumption A.15** (Policy Gradient Limitations). The relative underperformance of policy-based methods suggests:

$$\nabla_{\theta} \mathcal{J}(\pi_{\theta}) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[ \sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \hat{A}_t \right] \quad (34)$$

suffers from high variance in credit assignment for structural actions. Future work will develop SIREN-PPO+ with:

- Action space decomposition leveraging partial order constraints
- GNN-based advantage estimation  $\hat{A}_t = f_{\phi}(\mathbf{h}_{s_t}, \mathbf{h}_{G_T})$

## A.12 ADDITIONAL RESULTS ON EFFECTIVENESS

We have supplemented the experiments in Table 1 with more detailed results, including HITS@20 in Table 6, Mean Reciprocal Rank (MRR) in Table 7, and Precision@20 in Table 8, to strengthen the validity of our conclusions. Below, we elaborate on the experimental setup: for a given  $G_Q$ , a set of candidate  $G_T$  graphs is evaluated to determine whether  $G_Q$  is isomorphic to a subgraph of  $G_T$ . For certain distance-based approximate matching methods, we rank the results by ascending distance to provide the final judgment. The experimental results are presented as follows.

## A.13 ADDITIONAL RESULTS ON $|V_Q|$ -SCALING

Below, we present the results (average time:s) of the detailed  $|V_Q|$ -scaling experiments in Table 9.

## A.14 ABLATION STUDY ON GNN MODULE

We conducted experiments by replacing our GNN with SAGE and GCN modules, which do not preserve partial-order relations. The results in Table 10 show that while well-trained SAGE and GNN modules do assist DQN exploration (both improving MAP from 0.9 to 1.0), they are less effective than our partial-order-aware GNN in pruning the search process—that is, in optimizing

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

Table 6: Additional Results on HITS@20

Method	AIDS	MUTAG	PTC-FM	PTC-FR	PTC-MM	PTC-MR
SIMGNN	0.115	0.121	0.221	0.152	0.142	0.103
GRAPHSIM	0.048	0.071	0.077	0.058	0.053	0.083
GEDGNN	0.159	0.371	0.218	0.255	0.268	0.114
GOTSIM	0.137	0.168	0.238	0.143	0.210	0.224
ERIC	0.275	0.362	0.387	0.370	0.375	0.393
GMN-MATCH	0.392	0.465	0.455	0.477	0.422	0.485
NEUROMATCH	0.221	0.354	0.405	0.367	0.273	0.359
ISO NET	0.482	0.568	0.601	0.553	0.571	0.588
SUBMDSE-LATE	0.541	0.525	0.648	0.569	0.594	0.625
SUBMDSF-EARLY	0.677	0.726	0.785	0.730	0.716	0.752
SIREN	1.000	1.000	1.000	1.000	1.000	1.000

1256

1257

1258

1259

Table 7: Additional Results on Mean Reciprocal Rank (MRR)

Method	AIDS	MUTAG	PTC-FM	PTC-FR	PTC-MM	PTC-MR
SIMGNN	0.585	0.600	0.815	0.724	0.699	0.553
GRAPHSIM	0.345	0.401	0.492	0.357	0.402	0.480
GEDGNN	0.725	0.870	0.791	0.846	0.856	0.576
GOTSIM	0.672	0.726	0.811	0.699	0.792	0.796
ERIC	0.865	0.903	0.875	0.888	0.892	0.886
GMN-MATCH	0.892	0.917	0.917	0.951	0.913	0.945
NEUROMATCH	0.793	0.862	0.896	0.885	0.854	0.887
ISO NET	0.914	0.974	0.965	0.961	0.959	0.965
SUBMDSE-LATE	0.962	0.946	1.000	0.988	1.000	1.000
SUBMDSF-EARLY	1.000	1.000	0.998	0.998	0.997	0.999
SIREN	1.000	1.000	1.000	1.000	1.000	1.000

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

Table 8: Additional Results on Precision@20

Method	AIDS	MUTAG	PTC-FM	PTC-FR	PTC-MM	PTC-MR
SIMGNN	0.409	0.423	0.614	0.481	0.467	0.391
GRAPHSIM	0.193	0.233	0.279	0.204	0.221	0.276
GEDGNN	0.480	0.734	0.599	0.646	0.670	0.408
GOTSIM	0.450	0.527	0.616	0.472	0.597	0.604
ERIC	0.681	0.721	0.767	0.708	0.711	0.762
GMN-MATCH	0.765	0.801	0.799	0.828	0.800	0.826
NEUROMATCH	0.600	0.694	0.788	0.704	0.678	0.707
ISO NET	0.797	0.876	0.886	0.862	0.870	0.877
SUBMDSE-LATE	0.832	0.833	0.920	0.879	0.893	0.909
SUBMDSF-EARLY	0.916	0.941	0.931	0.919	0.911	0.929
SIREN	1.000	1.000	1.000	1.000	1.000	1.000

1290

Table 9: Additional Results on  $|V_Q|$ -scaling

$ V_Q $	RI	VF	GraphQL	QuickSI	GNN-PE	VF3	RM	VEQ	CaliG	DPiso	SIREN
4	2483.4	2795.5	14653.0	2030.7	87.3	828.5	101.2	111.8	609.6	131.1	50.7
6	4966.3	3678.2	14412.6	2108.7	125.0	1306.7	203.2	120.4	689.8	173.3	79.2
8	4672.3	3929.0	24677.7	4530.9	181.0	2163.8	303.0	126.2	1224.0	180.4	105.8
12	11231.3	9013.4	48940.1	5314.4	265.3	3847.0	438.9	216.3	1822.8	255.5	301.6
16	18212.3	19085.3	89025.2	15454.8	931.0	6791.3	1496.5	851.2	5897.7	1162.3	632.0
64	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	$> 10^5$	65258.2

search efficiency. Therefore, we conclude that both the partial-order constraint and the inherent representational capacity of GNNs play crucial roles in subgraph isomorphism tasks. While replacing our GNN module with other GNN variants is feasible, it does lead to a certain degree of performance degradation.

Table 10: Ablation Study on GNN Module

Model	Eval.	FR	FM	MM	MR	MUTAG	AIDS
GCN	MAP	1.000	1.000	1.000	1.000	1.000	1.000
	Avg. Runtime (ms)	190.9	114.9	138.9	191.5	111.4	181.6
SAGE	MAP	1.000	1.000	1.000	1.000	1.000	1.000
	Avg. Runtime (ms)	82.5	67.1	117.1	154.8	125.6	96.9
partial-order-aware GNN (SIREN)	MAP	1.000	1.000	1.000	1.000	1.000	1.000
	Avg. Runtime (ms)	50.2	44.8	43.0	56.4	49.7	53.0

## A.15 EXTENSIONS TO RELATED PROBLEMS

Below we outline how our method can be extended to address similar problems currently solved by existing approaches.

### A.15.1 GRAPH SIMILARITY

Regarding the Graph Similarity problem, if the goal is to quickly estimate similarity, we believe that methods based on spectral graph theory—such as computing the Laplacian eigenvalues of graphs—already offer a relatively efficient and well-established solution. Numerous existing works have demonstrated the effectiveness of this approach. However, it is challenging to extend such spectral methods into a complete subgraph isomorphism solution, particularly for graphs that contain only structural information without node labels.

As for our model, we consider that if the objective is solely to evaluate graph similarity, our partial-order preserved GNN model is capable of providing a meaningful assessment—which may not be perfectly precise but still offers valuable reference. In Table 11, we present a preliminary experimental evaluation applying our partial-order aware GNN to the graph similarity task.

Table 11: Precision on AIDS Dataset

Method	MAP	Precision@20
SimGNN	0.421	0.514
GraphSim	0.372	0.410
GOTSim	0.437	0.544
SIREN(GNN)	0.501	0.582

### A.15.2 GRAPH EDIT DISTANCE

We would like to highlight that our method can be naturally extended to solve the graph edit distance problem. The key idea is to adapt SIREN to find the maximum common subgraph (MCS) between  $G_Q$  and  $G_T$ . This can be achieved by simply recording the largest common subgraph identified during the subgraph isomorphism search process.

Once the MCS, denoted as  $G_C$ , is found, the graph edit distance between  $G_Q$  and  $G_T$  can be accurately computed by calculating the edit distance from  $G_C$  to both  $G_Q$  and  $G_T$  respectively. Since  $G_C$  is a subgraph of both  $G_Q$  and  $G_T$ , this computation is straightforward. Therefore, as long as our method correctly identifies the MCS, it guarantees an exact graph edit distance.

Currently, our approach can effectively solve the MCS problem. If an exhaustive search strategy is employed, it ensures the identification of the correct MCS. We acknowledge that, due to the inherent complexity of the search process, our method may be less efficient compared to algorithms specifically designed for graph edit distance. However, we emphasize that our approach offers high precision similar to traditional exhaustive search methods. While a fast estimation of graph edit distance could likely be achieved by adding an additional prediction layer to our partial-order-aware GNN, this

1350 direction falls outside the core focus of our paper. Therefore, we did not conduct experiments related  
1351 to this task.

1352

### 1353 A.15.3 MAXIMUM CLIQUE

1354

1355 We acknowledge that, as both are NP-complete problems, the maximum clique problem and the  
1356 subgraph isomorphism problem can be transformed into each other with strict logical equivalence due  
1357 to the nature of NP-completeness. A transformation from subgraph isomorphism to the maximum  
1358 clique problem was introduced by Levi (1973), though this is not the focus of our paper and is  
1359 mentioned here for reference. Below, we outline a method for reducing the maximum clique problem  
1360 to subgraph isomorphism, which can then be solved using our proposed approach.

1361 The key idea is that finding a maximum clique in a graph  $G$  is equivalent to identifying the largest  
1362 complete subgraph contained in  $G$ . This can be achieved by enumerating possible clique sizes—either  
1363 incrementally or via binary search—and checking whether a complete graph of that size is isomorphic  
1364 to a subgraph of  $G$ . Thus, our subgraph isomorphism method can be directly applied to solve the  
1365 maximum clique problem.

1366 It should be noted, however, that while our approach is applicable, it may not match the performance  
1367 of algorithms specifically optimized for the maximum clique problem. As stated in the paper, we  
1368 plan to explore further improvements in subsequent work.

1369

### 1370 A.16 COMPLEXITY ANALYSIS

1371

1372 The time complexity of our GNN-based query processing is approximately  $\mathcal{O}(k|V_T||V_Q|)$ , where  $k$  de-  
1373 notes the number of GNN layers. The per-iteration computational complexity remains  $\mathcal{O}(k|V_T||V_Q|)$ .  
1374 Although the theoretical worst-case time complexity suggests exponential scaling, experimental  
1375 results reveal a practical polynomial-time behavior within  $\mathcal{O}((|V_T||V_Q|)^4)$  to  $\mathcal{O}((|V_T||V_Q|)^5)$ .

1376

### 1377 A.17 LIMITATIONS

1378

1379 While SIREN performs well on moderately sized graphs, we acknowledge the scalability limitations  
1380 inherent in RL-driven search and GNN embedding generation when applied to massive graphs.  
1381 However, our systematic benchmarking reveals a critical industry-wide challenge: When confronted  
1382 with ultra-dense, large-scale graphs typical of real-world applications, neither traditional heuristic  
1383 approaches (VF3 Carletti et al. (2017), RI Bonnici et al. (2013), GraphQL He & Singh (2008), ...)  
1384 nor modern neural network-based solutions (GNN-PE Ye et al. (2024b), GLSearch Bai et al. (2021))  
1385 demonstrate viable computational tractability. Specifically, as benchmarked on graphs exceeding  $10^4$   
1386 nodes with density  $\rho > 0.3$ , all existing methods exhibit exponential time complexity growth beyond  
1387 practical feasibility thresholds.

### 1388 A.18 ETHICS STATEMENT

1389

1390 All authors have read and will adhere to the ICLR Code of Ethics throughout the submission process.  
1391 To the best of our knowledge, this work does not present any potential ethical concerns.

1392

### 1393 A.19 REPRODUCIBILITY STATEMENT

1394

1395 We are committed to ensuring the reproducibility of our work and welcome discussions with reviewers  
1396 regarding any reproducibility concerns during the review process. The supplementary material  
1397 includes the core implementation of SIREN. Upon acceptance of the paper, we will release the  
1398 full source code, trained models, datasets, checkpoints, and related resources to ensure complete  
1399 reproducibility.

1400

### 1401 A.20 THE USE OF LARGE LANGUAGE MODELS (LLMs)

1402

1403 The use of Large Language Models (LLMs) in this work was solely restricted to polishing and  
refining the linguistic expression of the manuscript. LLMs were not employed in any other aspect of

1404 the research, including but not limited to: related work survey, code development, model training and  
1405 testing, or conducting proofs.  
1406  
1407  
1408  
1409  
1410  
1411  
1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457