

---

# Maximum Independent Set: Self-Training through Dynamic Programming

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 This work presents a novel graph neural network (GNN) framework for solving  
2 the maximum independent set (MIS) inspired by dynamic programming (DP).  
3 Specifically, given a graph, we propose a DP-like recursive algorithm based on  
4 GNNs that firstly constructs two smaller sub-graphs, predicts the one with the  
5 larger MIS, and then uses it in the next recursive call. To train our algorithm,  
6 we require annotated comparisons of different graphs concerning their MIS size.  
7 Annotating the comparisons with the output of our algorithm leads to a self-training  
8 process that results in more accurate self-annotation of the comparisons and vice  
9 versa. We provide numerical evidence showing the superiority of our method vs  
10 prior methods in multiple synthetic and real-world datasets.

## 11 1 Introduction

12 Deep neural networks (DNNs) have achieved unprecedented success in extracting intricate patterns  
13 directly from data without the need for handcrafted rules, while still generalizing well to new and  
14 previously unseen instances [He et al., 2016, Vaswani et al., 2017]. Among other applications,  
15 this success has led to the development of frameworks that utilize DNNs to solve combinatorial  
16 optimization (CO) problems, such as the Traveling Salesman Problem [Xing and Tu, 2020, Hu et al.,  
17 2020, Prates et al., 2019], the Job-Shop Scheduling Problem [Zhang et al., 2020, Park et al., 2021],  
18 and the Quadratic Assignment Problem [Nowak et al., 2017].

19 A core challenge for deep learning approaches on CO is the lack of training data. Annotating  
20 such data requires the solution of a huge number of instances of the CO, hence such supervised  
21 learning approaches are computationally infeasible for NP-hard problems [Yehuda et al., 2020].  
22 Circumventing this difficulty is key to unlocking the full potential of otherwise broadly applicable  
23 DNNs for CO.

24 Our work demonstrates how classical ideas in CO together with DNNs can lead to a scalable self-  
25 supervised learning approach, mitigating the lack of training data. Concretely, we focus on the  
26 Maximum Independent Set (MIS) problem: Given a graph  $G(V, E)$ , MIS asks for a set of nodes of  
27 maximum cardinality such that no two nodes in the selected set are connected with an edge. MIS  
28 is an NP-hard problem with several hand-crafted heuristics (e.g., *greedy heuristic*, *local search*).  
29 More recently, several deep learning approaches have been proposed [Karalias and Loukas, 2020,  
30 Toenshoff et al., 2019, Schuetz et al., 2022a]

31 Our approach involves the following steps to determine an MIS in a graph. We use graph neural  
32 networks (GNNs) [Wu et al., 2020] to enable a model to generate approximate maximum independent  
33 sets after training on data that was annotated by the model itself. For this purpose, we draw inspiration  
34 from dynamic programming (DP) and employ a DP-like recursive algorithm. Initially, we are given a  
35 graph. At each recursive step, we select a random vertex from that graph and create two sub-graphs:  
36 one by removing the selected vertex and another by removing all its neighboring vertices. We then

37 make a comparison between these sub-graphs to determine which sub-graph is likely to have a larger  
38 independent set, and we use the sub-graph with the highest estimated IS for the next recursive call.  
39 We repeat this process until we reach a graph consisting only of isolated vertices, which signifies the  
40 discovery of an independent set for the original graph.

41 Dynamic programming guarantees that if our predictions are accurate (i.e., we select the sub-graph  
42 with the largest MIS value), our recursive algorithm will always result in a maximum independent set.  
43 To make accurate predictions, we introduce “graph comparing functions,” which take two graphs as  
44 input and output a winner. We implement such graph-comparing functions with GNNs.

45 We adopt a self-training approach to train our graph-comparing function and optimize the parameters  
46 of the GNN. In each epoch, we update the graph-comparing function parameters to ensure it accurately  
47 fits the data it has seen so far. The data comprises pairs of graphs  $(G, G')$  along with a label  
48  $\text{Label}(G, G') \in \{0, 1\}$ . For annotating the labels, we utilize the output of the recursive algorithm  
49 that leverages the graph-comparing function. Supported by theoretical and experimental evidence,  
50 we demonstrate how the self-annotation process improves parameter selection.

51 We conduct a thorough validation of our self-training approach in three real-world graph distribution  
52 datasets. Our algorithm surpasses the performance of previous deep learning methods [Karalias and  
53 Loukas, 2020, Toenshoff et al., 2019, Ahn et al., 2020] in the context of the MIS problem. To further  
54 validate the efficacy of our method, we explore its robustness on out-of-distribution data. Notably,  
55 our results demonstrate that the induced algorithm achieves competitive performance, showcasing the  
56 generalization capability of the learned comparator across different graph structures and distributions.  
57 In addition, we extend the evaluation of our DP-based self-training approach to tackle the Minimum  
58 Vertex Cover (MVC) problem in Appendix E. Encouragingly, similar to the MIS case, our induced  
59 GNN-based algorithms for MVC admit competitive performance with respect to other deep-learning  
60 approaches.

## 61 2 Related work

62 Our work lies in the intersection of various domains, i.e., combinatorial optimization, Dynamic  
63 Programming, and (graph) neural networks. We review the most critical ideas in each domain here  
64 and defer a more detailed discussion in Appendix A.

65 **Graph Neural Networks (GNNs)** have gained widespread popularity due to their ability to learn  
66 representations of graph-structured data [Xiao et al., 2022, Zhang and Chen, 2018, Zhu et al., 2021,  
67 Errica et al., 2019] invariant to the size of the graph. More complex architectural blocks, such as  
68 the Graph Convolutional Network (GCN) [Kipf and Welling, 2017, Zhang et al., 2019], the Graph  
69 Attention Network (GAT) [Veličković et al., 2017], and the Graph Isomorphism Network [Xu et al.,  
70 2018] have become influential instances of GNNs. In our work, we utilize a simple GNN architecture  
71 to showcase the effectiveness of our proposed framework. While our choice of architecture is  
72 intentionally simple, we emphasize its modular nature, which enables us to incorporate more complex  
73 GNNs with ease.

74 **Combinatorial optimization:** Supervised learning approaches have been used for tackling CO tasks,  
75 such as the Traveling Salesman Problem (TSP) [Vinyals et al., 2015], the Vehicle Routing Problem  
76 (VRP) [Shalaby et al., 2021], and Graph Coloring [Lemos et al., 2019]. Due to the graph structure of  
77 the problems, GNNs are often used for tackling those tasks [Prates et al., 2019, Nazari et al., 2018,  
78 Schuetz et al., 2022b]. However, owing to the computational overhead of obtaining supervised labels,  
79 such supervised approaches often do not scale well. Instead, unsupervised approaches have been  
80 deployed recently [Wang and Li, 2023]. A popular approach relies on a continuous relaxation of the  
81 loss function [Karalias and Loukas, 2020, Wang et al., 2022, Wang and Li, 2023]. In contrast to the  
82 previous unsupervised works, we adopt Dynamic Programming techniques to diminish the overall  
83 time complexity of the algorithm. Another approach uses reinforcement learning (RL) methods to  
84 address CO tasks, such as in Covering Salesman Problem [Li et al., 2021], the TSP [Zhang et al.,  
85 2022], the VRP [James et al., 2019], and the Minimum Vertex Cover (MVC) [Tian and Li, 2021].  
86 However, applying RL to CO problems can be challenging because of the long learning time required  
87 and the non-differentiable nature of the loss function.

88 **Dynamic Programming** has been a powerful problem-solving technique since at least the 50s [Bell-  
89 man, 1954]. In recent years, researchers have explored the use of deep neural networks (DNNs) to

90 replace the function responsible for dividing a problem into subproblems and estimating the optimal  
 91 decision at each step [Yang et al., 2018]. Despite the progress, implementing CO tasks with Dynamic  
 92 Programming suffers from significant computational overheads, since the size of the search space  
 93 grows exponentially with the problem size [Xu et al., 2020]. Our approach overcomes this issue by  
 94 utilizing a model that approximates the standard lookup table from Dynamic Programming, meaning  
 95 that we avoid the exponential search space typically associated with DP.

### 96 3 An optimal solution to Maximum Independent Set (MIS)

97 Let us first introduce MIS and its relationship with Dynamic Programming.

98 **Notation:**  $G(V, E)$  denotes an undirected graph where  $V$  stands for vertices and  $E$  for the edges.  
 99  $\mathcal{N}(v)$  denotes the neighbors of vertex  $v \in V$ ,  $\mathcal{N}(v) = \{u \in V : (u, v) \in E\}$ . The *degree* of vertex  
 100  $v \in V$  is  $d(v) := |\mathcal{N}(v)|$ . Given a set of vertices  $S \subseteq V$ ,  $G/S$  denotes the remaining graph of  $G$   
 101 after removing all nodes  $v \in S$ .

102 **Definition 1** (Maximum Independent Set). *Given an undirected graph  $G(V, E)$ , find a maximum set  
 103 of nodes  $S \subseteq V$  such that  $(u, v) \notin E$  for all vertices  $u, v \in S$ . We denote with  $|\text{MIS}(G)|$  the size of  
 104 the maximum independent set of graph  $G$ .*

105 Dynamic Programming is a powerful technique for algorithmic design in which the optimal solution  
 106 of the instance of interest is constructed by combining the optimal solution of smaller sub-instances.  
 107 The combination step is governed by local optimality conditions which, in the context of MIS, take  
 108 the form of Theorem 1. Theorem 1 establishes that the decision to remove a node  $v \in V$  or its  
 109 neighbors  $\mathcal{N}(v)$  during the recursive process depends on whether  $|\text{MIS}(G/\mathcal{N}(v))| \geq |\text{MIS}(G/v)|$ .  
 110 This decision continues until a graph with no edges is reached. According to Theorem 1, if at each  
 111 step of the recursion, the choice is made based on whether  $|\text{MIS}(G/\mathcal{N}(v))| \geq |\text{MIS}(G/v)|$  or not,  
 112 then the resulting empty graph is guaranteed to be an optimal solution. The proof of this theorem can  
 113 be found in Appendix I.

114 **Theorem 1.** *Let a graph  $G(V, E) \in \mathcal{G}$ . Then for any vertex  $v \in V$  with  $d(v) \geq 1$ ,*

$$|\text{MIS}(G)| = \max(|\text{MIS}(G/\mathcal{N}(v))|, |\text{MIS}(G/\{v\})|)^1.$$

### 115 4 Graph Neural Network-based Algorithm for MIS

116 In this section, we present our approach for developing algorithms for MIS parameterized by  
 117 parameters  $\theta \in \Theta$ . Initially in Sec. 4.1 we present how any *graph-comparing function* taking as  
 118 input two different graphs and outputting a  $\{0, 1\}$  value can be used in the construction of an  
 119 algorithm computing an independent set (not necessarily optimal). In Sec. 4.2 we present how Graph  
 120 Neural Networks can be used in the construction of such graph-comparing functions. Finally, in  
 121 Sec. 4.3, we present our *inference algorithm* that computes an independent set of any graph  $G \in \mathcal{G}$ .

#### 122 4.1 MIS Algorithms induced by Graph Comparing Functions

123 Consider a function  $\text{CMP} : \mathcal{G} \times \mathcal{G} \mapsto \{0, 1\}$  that compares two graphs  $G, G'$  based on the size of their  
 124 MIS. Namely, if  $|\text{MIS}(G)| \geq |\text{MIS}(G')|$  then  $\text{CMP}(G, G') = 0$  and  $\text{CMP}(G, G') = 1$  otherwise.

125 Theorem 1 ensures that if we have access in such a *graph-comparing function* then we can compute  
 126 an independent set of maximum size for any graph  $G$ . From a starting node  $v$ , with an initial  
 127 graph from  $G$ , and by recursively selecting either  $G/\{v\}$  or  $G/\mathcal{N}(v)$  based on  $|\text{MIS}(G/\{v\})| \geq$   
 128  $|\text{MIS}(G/\mathcal{N}(v))|$ , we are ensured to end in an independent set of maximum size. The decision of  
 129 whether  $|\text{MIS}(G/\{v\})| \geq |\text{MIS}(G/\mathcal{N}(v))|$  at each recursive call can be made according to the  
 130 output of  $\text{CMP}(G/\{v\}, G/\mathcal{N}(v))$ .

131 The cornerstone idea of our approach is that *any graph-comparing function*  $\text{CMP}$  induces such  
 132 a recursive algorithm for a MIS. Recursively selecting  $G/\{v\}$  or  $G/\mathcal{N}(v)$  based on the output  
 133 of a graph generating function  $\text{CMP}(G/\{v\}, G/\mathcal{N}(v)) \in \{0, 1\}$  always guarantees to reach an

<sup>1</sup>Note: In the trivial case where  $G$  is an empty graph (i.e., it has no edges), the size of the maximum independent set is  $|V|$ .

---

**Algorithm 1** Comparator-Induced Algorithm

---

```
1: function  $\mathcal{A}^{\text{CMP}}(G(V, E))$  ▷ Algorithm  $\mathcal{A}^{\text{CMP}}(G)$  takes a graph  $G$  as input
2:   if  $|E| = 0$  then return  $V$ 
3:   end if
4:   pick a vertex  $v \in V$  with  $d(v) > 0$  uniformly at random.
5:    $G_0 \leftarrow G \setminus \{v\}$  and  $G_1 \leftarrow G \setminus \mathcal{N}(v)$ 
6:   if  $\text{CMP}(G_0, G_1) = 0$  then
7:      $G \leftarrow G_0$  ▷ Remove vertex  $v$ 
8:   else
9:      $G \leftarrow G_1$  ▷ Remove the neighbors of  $v$ 
10:  end if
11:  return  $\mathcal{A}^{\text{CMP}}(G)$ 
12: end function
```

---

134 independent set of the original graph. In case  $\text{CMP}(G, G') \neq \mathbb{I}[|\text{MIS}(G)| < |\text{MIS}(G')|]$ , where  $\mathbb{I}$  is  
135 the indicator function, it is not guaranteed that the computed independent set is of the maximum size.  
136 However, there might exist a reasonable graph comparing functions that *i*) are efficiently computable  
137 *ii*) lead to near-optimal solutions.

138 In Definition 2 and Algorithm 1 we formalize the idea above.

139 **Definition 2.** A comparator  $\text{CMP} : \mathcal{G} \times \mathcal{G} \mapsto \{0, 1\}$  is a function taking as input two graphs  $G, G'$   
140 and outputting a  $\{0, 1\}$  value.

141 **Proposition 1.** Any comparator  $\text{CMP} : \mathcal{G} \times \mathcal{G} \mapsto \{0, 1\}$  induces a randomized algorithm  $\mathcal{A}^{\text{CMP}}$   
142 (Algorithm 1).

143 **Remark 1.** Given a graph-comparing function  $\text{CMP} : \mathcal{G} \times \mathcal{G} \mapsto \{0, 1\}$ , the induced algorithm is  
144 randomized, since at Step 4 of Algorithm 1, a vertex  $v$  is randomly selected. Notice that Algorithm 1  
145 recursively proceeds until a subgraph with 0 edges is reached (see Step 2).

146 **Remark 2.** Two different comparators  $\text{CMP}$  and  $\text{CMP}'$  induce two difference algorithms  $\mathcal{A}^{\text{CMP}}$   
147 and  $\mathcal{A}^{\text{CMP}'}$  for estimating the maximum independent set.

## 148 4.2 Comparators through Graph Neural Networks

149 In this section, we discuss the architecture of a model  $M_\theta : \mathcal{G} \mapsto \mathbb{R}$ , parameterized by  $\theta \in \Theta$ , that is  
150 used for the construction of a comparator function

$$\text{CMP}_\theta(G, G') = \mathbb{I}[M_\theta(G) < M_\theta(G')] .$$

151 In order to embed graph-level information, we introduce a new GNN module, which we refer to as  
152 the Graph Embedding Module (GEM). Unlike standard GNN modules, this module captures different  
153 semantic meanings of differing embeddings of a node, its neighbors, and anti-neighbors.

### 154 Graph Embedding Module (GEM):

155 The GEM operates using the following recursive formula:

$$\mu_v^{k+1} = \text{LN} \left( \text{GELU} \left( \left[ \theta_0^k \mu_v^k \parallel \theta_1^k \sum_{u \in \mathcal{N}(v)} \mu_u^k \parallel \theta_2^k \sum_{u \notin \mathcal{N}(v)} \mu_u^k \right] \right) \right) . \quad (1)$$

156 Initially, all nodes in this graph have zeros embeddings  $\mu_v^0 = \vec{0} \in \mathbb{R}^{3p}$ . Here,  $\mu_v^0$  denotes the initial  
157 embedding vector of node  $v$ . In Eq. (1), for all iterations  $k \in [0, \dots, K - 1]$ , the embeddings of a  
158 node denoted by  $\mu_v^k \in \mathbb{R}^{3p}$ , its neighbors, and its anti-neighbors  $v$  are put through their own linear  
159 layers, denoted by  $\theta_0^k, \theta_1^k, \theta_2^k \in \mathbb{R}^{p \times 3p}$ , which are the parameters of the module. The bias term  
160 is omitted in the equation for readability purposes. We incorporate anti-neighbors in the GEM to  
161 capture complementary relationships between nodes. By using separate linear layers for different




Figure 1: Architecture of model  $M_\theta(G)$ . From left to right: initially, an input graph  $G$  is passed into the model with zeros as node embeddings, which are displayed as white in the figure. The striped green edges connect the anti-neighbors, which are also used in the GEM. After  $K$  iterations of the GEM module, the final node embeddings are obtained. These are then averaged to obtain a graph embedding  $\mu_G$ . Finally, the graph embedding is put through multiple fully-connected layers to obtain a final logit value for the input graph.

162 features, we emphasize the contrasting semantic meaning between neighbors and anti-neighbors,  
 163 representing negative and positive relationships in the graph. Then, the individual feature embeddings  
 164 are concatenated, which is denoted by  $[\dots \parallel \dots]$ , followed by a GELU activation function [Hendrycks  
 165 and Gimpel, 2016] and layer normalization [Ba et al., 2016a].

166 The complete model architecture is depicted in Fig. 1. At a high level, the architecture uses a Graph  
 167 Embedding Module to extract a global graph embedding from the input graph, which is then passed  
 168 through a set of fully connected layers to output a logit for that graph. During the training process of  
 169 the comparator function, we utilize  $\text{CMP}_\theta(G, G') = \text{softmax}([M_\theta(G) \parallel M_\theta(G')])$ , which forms a  
 170 differentiable loss function for classification.

### 171 4.3 Inference Algorithm

172 In the previous section, we discussed how a parameterization  $\theta \in \Theta$  defines the graph-comparing  
 173 function  $\text{CMP}_\theta(G, G') = \mathbb{I}[M_\theta(G) < M_\theta(G')]$ . As a result, the same parameterization  $\theta \in \Theta$   
 174 defines an algorithm  $\mathcal{A}^{\text{CMP}_\theta}$ , where at Step 6 of Algorithm 1, the comparing function  $\text{CMP}_\theta$  is used.

## 175 5 Self-Supervised Training through the Consistency Property

176 In this section, we present our methodology for selecting the parameters  $\theta \in \Theta$  so that the resulting  
 177 inference algorithm  $\mathcal{A}^{\text{CMP}_\theta}(\cdot)$  computes independent sets with (close to) the maximum value.

178 The most straightforward approach is to select the parameters  $\theta \in \Theta$  such that  $\text{CMP}_\theta(G, G') \simeq$   
 179  $\mathbb{I}[|\text{MIS}(G)| < |\text{MIS}(G')|]$  using labeled data. The problem with this approach is that a huge amount  
 180 of annotated data of the form  $\{(G, G'), \mathbb{I}[|\text{MIS}(G)| < |\text{MIS}(G')|]\}$  are required. Since finding the  
 181 MIS is an NP-Hard problem, annotating such data comes with an insurmountable computational  
 182 burden.

183 The **key idea** to overcome the latter limitation is to annotate the data of the form  $\{(G, G')\}$  by using  
 184 the algorithm  $\mathcal{A}^{\text{CMP}_\theta}(\cdot)$  that runs in polynomial time with respect to the size of the graph. Intuitively,  
 185 our proposed framework entails the optimization of the parameterized comparator function  $\text{CMP}_\theta$   
 186 on data generated using algorithm  $\mathcal{A}^{\text{CMP}_\theta}$ . A better comparator function leads to a better algorithm,  
 187 which leads to better data, and vice versa. This mutually reinforcing relationship between the two  
 188 components of our framework is theoretically indicated by Theorem 2 that we present in Section 5.1.  
 189 The exact steps are detailed below.

### 190 5.1 Consistent Graph Comparing Functions

191 In this section, we introduce the notion of a *consistent* graph-comparing function (Definition 3) that  
 192 plays a critical role in our self-supervised learning approach. Kindly take note that  $\mathcal{A}^{\text{CMP}}$  utilizes  
 193 the unparameterized variant of a comparator function, whereas  $\mathcal{A}^{\text{CMP}_\theta}$  utilizes its parameterized  
 194 counterpart.

---

**Algorithm 2** Basic Pipeline of our Training Approach

---

- 1: **Input:** A distribution  $\mathcal{D}$  over graphs.
  - 2: Initialize parameters  $\theta_0 \in \Theta$ .
  - 3: Initialize a *graph-buffer*  $\mathcal{B} \leftarrow \emptyset$ .
  - 4: **for** each epoch  $t = 0, \dots, T - 1$  **do**
  - 5:     Sample a graph  $G_{\text{init}} \sim \mathcal{D}$ .
  - 6:     Run  $\mathcal{A}^{\text{CMP}_{\theta_t}}(G)$  and store in  $\mathcal{B}$  all graphs produced during each recursive call of Algorithm 1.
  - 7:     Update the parameters  $\theta_{t+1} \in \Theta$  such that
 
$$\theta^{t+1} := \operatorname{argmin}_{\theta \in \Theta} \mathbb{E}_{(G, G') \sim \mathcal{B}} [\ell(\text{CMP}_{\theta}(G, G'), \mathbb{I}[\mathbb{E}[\mathcal{A}^{\text{CMP}_{\theta_t}}(G)] < \mathbb{E}[\mathcal{A}^{\text{CMP}_{\theta_t}}(G')]])]$$
  - 8: **end for**
- 

195 **Definition 3** (Consistency). A graph-comparing function  $\text{CMP} : \mathcal{G} \times \mathcal{G} \mapsto \{0, 1\}$  is called consistent  
 196 if and only if for any pair of graphs  $G, G' \in \mathcal{G}$ ,

$$\text{CMP}(G, G') = 0 \text{ if and only if } \mathbb{E}[\mathcal{A}^{\text{CMP}}(G)] \geq \mathbb{E}[\mathcal{A}^{\text{CMP}}(G')].$$

198 **Remark 3.** In Definition 3 we use  $\mathbb{E}[\mathcal{A}^{\text{CMP}}(G)]$ ,  $\mathbb{E}[\mathcal{A}^{\text{CMP}}(G')]$  since, as we have already  
 199 discussed, a comparator  $\text{CMP}$  induces a randomized algorithm  $\mathcal{A}^{\text{CMP}}$ .

200 In Theorem 2, we formally establish that any consistent graph-comparing function  $\text{CMP}$  induces an  
 201 optimal algorithm for the MIS.

202 **Theorem 2.** Let a consistent comparator  $\text{CMP} : \mathcal{G} \times \mathcal{G} \mapsto \{0, 1\}$ . Then the algorithm  $\mathcal{A}^{\text{CMP}}(\cdot)$   
 203 always computes a Maximum Independent Set,  $\mathbb{E}[\mathcal{A}^{\text{CMP}}(G)] = |\text{MIS}(G)|$  for all  $G \in \mathcal{G}$ .

204 Theorem 2 guarantees that any consistent graph comparing function  $\text{CMP}$  induces an optimal  
 205 algorithm  $\mathcal{A}^{\text{CMP}}$  for MIS. The proof for this theorem can be found in Appendix I. Hence, the  
 206 selection of parameters  $\theta^* \in \Theta$  should be selected such that  $\text{CMP}_{\theta^*}$  is consistent. More precisely:

207 **Goal of Training:** Find parameters  $\theta^* \in \Theta$  such that for all  $G, G' \in \mathcal{G}$ :

$$\text{CMP}_{\theta^*}(G, G') = 0 \text{ if and only if } \mathbb{E}[\mathcal{A}^{\text{CMP}_{\theta^*}}(G)] \geq \mathbb{E}[\mathcal{A}^{\text{CMP}_{\theta^*}}(G')].$$

## 208 5.2 Training a Consistent Comparator

209 The cornerstone idea of our self-supervised learning approach is to make the comparator more and  
 210 more consistent over time. Namely, the idea is to update the parameters as follows:

$$\theta^{t+1} := \operatorname{argmin}_{\theta \in \Theta} \mathbb{E}_{G, G'} [\ell(\text{CMP}_{\theta}(G, G'), \mathbb{I}[\mathbb{E}[\mathcal{A}^{\text{CMP}_{\theta_t}}(G)] < \mathbb{E}[\mathcal{A}^{\text{CMP}_{\theta_t}}(G')]])] \quad (2)$$

211 where  $\ell(\cdot, \cdot)$  is a binary classification loss. In Eq. (2),  $\theta_t$  are the fixed parameters of the previous  
 212 epoch. Thus, in the next few paragraphs, we only use the notation  $\mathcal{A}^{\text{CMP}_{\theta_t}}$  to denote the fixed  
 213 parameters. Gradient updates are only computed over  $\theta$ .

214 **Remark 4.** We remark that neither solving the non-convex minimization problem of Eq. (2) nor the  
 215 existence of parameters  $\theta^* \in \Theta$  such that  $\text{CMP}_{\theta^*}$  can be guaranteed. However, using a first-order  
 216 method for Eq. (2) and a large enough parameterization can lead to an approximately consistent  
 217 comparator with approximately optimal performance.

218 In Algorithm 2, we present the basic pipeline of the self-training approach that selects the parameters  
 219  $\theta \in \Theta$  such that the inference algorithm  $\mathcal{A}^{\text{CMP}_{\theta_t}}$  admits a competitive performance given as input  
 220 graphs  $G$  following a graph-distribution  $\mathcal{D} \subseteq \mathcal{G}$ . However, while the basic pipeline of our self-training  
 221 approach follows Algorithm 2, there are several differences and tweaks that we incorporate into our  
 222 training process.

223 **Creating the graph buffer  $\mathcal{B}$ :** We are given a shuffled dataset of graphs  $\mathcal{D}$ , which represents the  
 224 training data for the model. The core difference between the pipeline and the training process

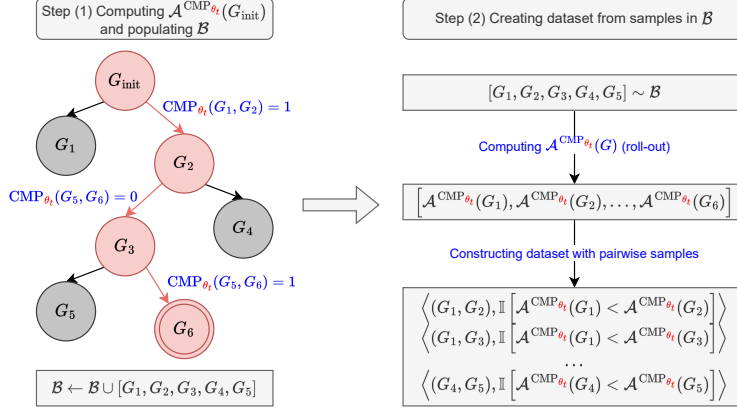


Figure 2: An example of data generation for the training. (Left) At the beginning of each training epoch (Step 5), Algorithm 2 samples  $G_{\text{init}} \sim \mathcal{D}$  and computes an independent set by using the comparator  $\text{CMP}_{\theta_t}$  and by following the branches in the *recursion tree* that are marked red, with the doubly circled one being the produced independent set. The generated graphs from this procedure are added to the buffer  $\mathcal{B}$ . (Right) Then, a dataset is created by sampling graphs from the buffer and then computing an estimate of their MIS size (based on  $\mathcal{A}^{\text{CMP}_{\theta_t}}$ ). Based on this estimate, a dataset is created with graph pairs  $(G, G')$  and their corresponding binary labels denoting which MIS estimates are larger.

225 comes from the graph buffer  $\mathcal{B}$ . In Algorithm 2, this buffer stores any graph  $G$  that is found  
 226 during the recursive call of  $\mathcal{A}^{\text{CMP}_{\theta_t}}(G_{\text{init}})$  on  $G_{\text{init}} \sim \mathcal{D}$  (Step 6 of Algorithm 2). However, in  
 227 the implementation of the graph buffer, it stores pairs of graphs  $(G, G')$  that were generated by  
 228  $\mathcal{A}^{\text{CMP}_{\theta_t}}(G_{\text{init}})$ , alongside a binary label that indicates which of the two graphs has a larger estimated  
 229 MIS size. How this estimate is generated, will be explained further down this section.

230 **The training process:** Prior to starting training, we first set two hyperparameters: one that specifies  
 231 the number of graphs used to populate the buffer before training the model, and another that determines  
 232 the number of graph pairs generated from  $\mathcal{A}^{\text{CMP}_{\theta_t}}(G)$  per graph  $G \sim \mathcal{D}$ . Then, a dataset is created  
 233 by generating these pairs for the set number of graphs. The dataset is then added to the graph buffer,  
 234 replacing steps 4, 5, and 6 in Algorithm 2, which only does this with one graph per epoch. Next,  
 235 training starts, and after completing a set number of epochs, a new dataset is created using the updated  
 236 model, and the process is repeated iteratively.

237 **Estimating the MIS:** Finally, the loss function in Step 7 of Algorithm 2, also operates slightly  
 238 differently. The main difference arises from  $\mathbb{I}[\mathbb{E}[|\mathcal{A}^{\text{CMP}_{\theta_t}}(G)|] < \mathbb{E}[|\mathcal{A}^{\text{CMP}_{\theta_t}}(G')|]]$ , since the  
 239 estimates  $|\text{MIS}(G)| \approx \mathbb{E}[|\mathcal{A}^{\text{CMP}_{\theta_t}}(G)|]$  and  $|\text{MIS}(G')| \approx \mathbb{E}[|\mathcal{A}^{\text{CMP}_{\theta_t}}(G')|]$  are not directly  
 240 utilized. Instead, we propose two other approaches, which are better approximations than the  
 241 expectations used in Algorithm 2, since they use a maximizing operator.

242 The first approach involves performing so-called "roll-outs" on the graph pairs generated  $G$  and  
 243  $G'$  by  $\mathcal{A}^{\text{CMP}_{\theta_t}}$ , in order to estimate their MIS sizes. To perform the roll-outs, we simply run  
 244  $\mathcal{A}^{\text{CMP}_{\theta_t}}$  on graphs  $G$  and  $G'$   $m$  times and use the maximum size of the found independent  
 245 sets as an estimate of their MIS. Formally, in a roll-out on a graph  $G$ , we sample the inde-  
 246 pendent sets  $\text{ISS}_1, \text{ISS}_2, \dots, \text{ISS}_m \sim \mathcal{A}^{\text{CMP}_{\theta_t}}(G)$ . Then, the estimate of the MIS size of  $G$  is  
 247  $\max(|\text{ISS}_1|, |\text{ISS}_2|, \dots, |\text{ISS}_m|)$ .

248 An example of the entire process of generating the dataset using roll-outs can be found in Fig. 2.

249 **Mixed roll-out variant:** We introduce a variant of the aforementioned method, which utilizes the  
 250 deterministic greedy algorithm. This greedy algorithm iteratively creates an independent set by  
 251 removing the node with the lowest degree and adding it to the independent set. This algorithm is often  
 252 an efficient approximation to the optimum solution. Our variant is constructed as follows: we compute  
 253 the maximum between the roll-outs of the model and the result of the greedy algorithm, which creates  
 254 a dataset with more accurate self-supervised approximations of the MIS values. This, in turn,  
 255 generates binary targets for the buffer that are more likely to be accurate. Thus, for this second variant,  
 256 the estimate of the MIS size of a graph  $G$  would be  $\max(|\text{Greedy}(G)|, |\text{ISS}_1|, |\text{ISS}_2|, \dots, |\text{ISS}_m|)$ .

## 257 6 Experiments

258 In this section, we conduct an evaluation of the proposed method for the MIS problem. Let us first  
259 describe the training setup, the baselines, and the datasets. Additional details and experiments on  
260 MIS are displayed in the Appendices C and D. Our method also generalizes well in MVC, as the  
261 results in Appendix E illustrate.

### 262 6.1 Training setup

263 **Our model:** We implement two comparator models: one using just roll-outs with the model, and  
264 another using the roll-outs together with greedy, called "mixed roll-out". We train each model  
265 using a graph embedding module with  $K = 3$  iterations, which takes in 32-dimensional initial node  
266 embeddings.

267 **Baselines:** We compare against the neural approaches *Erdos GNN* [Karalias and Loukas, 2020], *RUN-*  
268 *CSP* from Toenshoff et al. [2019], and a method specifically for the MIS problem: *LwDMIS* [Ahn  
269 et al., 2020]. Since we observe unexpected performance from *RUN-CSP* on the *COLLAB* and *RB*  
270 datasets, we have omitted those results from the table. We train every model for 300 epochs. Each  
271 experiment is performed on a single GPU with 6GB RAM.

272 Besides neural approaches, we use traditional baselines, such as the *Greedy MIS* [Wormald, 1995],  
273 *Simple Local Search* [Feo et al., 1994] and a *Random Comparator* as a sanity check. Furthermore,  
274 we implement two mixed-integer linear programming solvers: *SCIP 8.0.3* and the highly optimized  
275 commercial solver *Gurobi 10.0*.

276 **Datasets:** We evaluate our model on three standard datasets, following Karalias and Loukas [2020]:  
277 *COLLAB* [Yanardag and Vishwanathan, 2015], *TWITTER* [Leskovec and Krevl, 2014] and *RB* [Xu  
278 et al., 2007, Toenshoff et al., 2019]. In addition, we introduce the *SPECIAL* dataset that includes  
279 challenging graphs for handcrafted approaches as we detail in Appendix C.

### 280 6.2 Results

281 Table 1 reports the average approximation ratios on the test instances of the various datasets. The  
282 approximation ratio is computed by dividing a solution’s independent set size by the optimum solution,  
283 which is computed using the Gurobi solver with a time limit of 1 hour per graph.

284 The results indicate that the greedy algorithm performs strongly in three of the four datasets, which is  
285 consistent with the observation of Angelini and Ricci-Tersenghi [2022]. However, notice that our pro-  
286 posed approach outperforms the greedy in both the *Twitter* and the *SPECIAL* datasets, which validates  
287 that the greedy heuristic is not optimal in every case and is prone to failing in few cases. Importantly,  
288 among the neural approaches that are the main compared methods, our proposed method performs  
289 favorably in all datasets. The performance of our method indicates that the proposed self-training  
290 scheme is able to learn from diverse data distributions and generalize reasonably well in the test sets  
291 of the respective dataset. In addition, the proposed method is faster than the rest neural approaches.

292 The mixed roll-out model in Table 1 outperforms the normal roll-out model in almost all datasets,  
293 indicating the effectiveness of the greedy heuristic in roll-outs. This is particularly evident in the *RB*  
294 dataset. However, for *SPECIAL* instances, the normal model performs marginally better, possibly  
295 due to the unsuitability of the greedy guiding heuristic as a baseline for this dataset.

296 **Out of distribution** : We examine the performance of the learned comparator through its general-  
297 ization to new graph distributions. Concretely, we conduct an out-of-distribution analysis as follows:  
298 each model is trained in one graph distribution, indicated by the rows of Table 2. Then, the model  
299 is evaluated on different graph distributions, indicated by the columns of Table 2. The analysis is  
300 conducted on both our model and the approach of *Erdos GNN* [Karalias and Loukas, 2020].

301 Surprisingly, our model trained over *COLLAB* displays good generalization skills across different  
302 datasets, even outperforming the *RB*-trained model on the *RB* dataset. Conversely, *Erdos GNN* trained  
303 over *RB* performs poorly over the *COLLAB* dataset. Both models trained over the *RB* dataset perform  
304 more poorly in general, likely due to the highly specific graph distribution of the *RB* dataset. Moreover,  
305 our model, on the whole, exhibits good generalization skills over different graph distributions.



Table 1: Test set approximation ratios (higher is better; the best performance in bold) on four datasets. We report the average approximation ratios (along with std and time budget) on MIS. Notice that the proposed method outperforms all the deep-learning-based approaches across datasets.

Method ( $\downarrow$ ) Dataset ( $\rightarrow$ )	RB	COLLAB	TWITTER	SPECIAL
CMP (Normal Roll-outs)	$0.770 \pm 0.107$ (0.43 s/g)	$0.990 \pm 0.051$ (0.17 s/g)	$0.967 \pm 0.083$ (0.35 s/g)	<b><math>0.996 \pm 0.029</math></b> (0.04 s/g)
CMP (Mixed Roll-outs)	<b><math>0.836 \pm 0.083</math></b> (0.36 s/g)	<b><math>0.990 \pm 0.049</math></b> (0.21 s/g)	<b><math>0.977 \pm 0.031</math></b> (0.21 s/g)	$0.994 \pm 0.035$ (0.05 s/g)
Erdos' GNN	$0.813 \pm 0.107$ (1.39 s/g)	$0.952 \pm 0.142$ (0.60 s/g)	$0.935 \pm 0.078$ (1.37 s/g)	$0.921 \pm 0.218$ (1.03 s/g)
LwDMIS	$0.804 \pm 0.089$ (0.42 s/g)	$0.978 \pm 0.031$ (0.17 s/g)	$0.972 \pm 0.032$ (0.19 s/g)	$0.828 \pm 0.304$ (0.32 s/g)
RUN-CSP (Accurate)	–	–	$0.875 \pm 0.053$ (0.57 s/g)	$0.946 \pm 0.059$ (0.51 s/g)
Greedy MIS	$0.925 \pm 0.053$ (0.01 s/g)	$0.998 \pm 0.023$ (0.02 s/g)	$0.964 \pm 0.048$ (0.04 s/g)	$0.131 \pm 0.055$ (0.03 s/g)
Random CMP	$0.615 \pm 0.155$ (0.42 s/g)	$0.817 \pm 0.211$ (0.30 s/g)	$0.634 \pm 0.182$ (0.36 s/g)	$0.225 \pm 0.279$ (0.41 s/g)
Simple Local Search (10s)	$0.565 \pm 0.237$	$0.860 \pm 0.213$	$0.644 \pm 0.218$	$0.188 \pm 0.340$
SCIP 8.0.3 (1s)	$0.741 \pm 0.351$	$0.999 \pm 0.016$	$0.959 \pm 0.024$	1.000
SCIP 8.0.3 (5s)	$0.937 \pm 0.118$	1.000	$0.999 \pm 0.024$	1.000
Gurobi 10.0 (0.5s)	$0.969 \pm 0.070$	$0.981 \pm 0.068$	$0.985 \pm 0.085$	$0.940 \pm 0.237$
Gurobi 10.0 (1s)	$0.983 \pm 0.051$	1.000	1.000	1.000
Gurobi 10.0 (5s)	$0.999 \pm 0.008$	1.000	1.000	1.000

Table 2: Out-of-distribution approximation ratios during inference (higher is better). Every row denotes a model trained on a specific dataset. Every column considers a different test dataset. The CMP is trained using mixed roll-outs. Notice that the proposed method generalizes well in out-of-distribution structures. This is indicative of the learned comparator extracting robust patterns.

Model ( $\downarrow$ ) Dataset ( $\rightarrow$ )	RB	COLLAB	TWITTER
CMP RB	–	$0.903 \pm 0.186$	$0.668 \pm 0.187$
CMP COLLAB	$0.856 \pm 0.080$	–	$0.906 \pm 0.094$
CMP TWITTER	$0.773 \pm 0.101$	$0.927 \pm 0.148$	–
Erdos' GNN RB	–	$0.361 \pm 0.334$	$0.752 \pm 0.188$
Erdos' GNN COLLAB	$0.680 \pm 0.071$	–	$0.592 \pm 0.186$
Erdos' GNN TWITTER	$0.746 \pm 0.092$	$0.666 \pm 0.385$	–

## 306 7 Conclusion

307 Motivated by the principles of Dynamic Programming, we develop a self-training approach for  
 308 important CO problems, such as the Maximum Independent Set and the Minimum Vertex Cover.  
 309 Our approach embraces the power of self-training, offering the dual benefits of data self-annotation  
 310 and data generation. These inherent attributes are instrumental in providing an unlimited source of  
 311 data indicating that the performance of the induced algorithms can be significantly improved with  
 312 sufficient scaling on the computational resources. We firmly believe that a thorough investigation  
 313 into the interplay between Dynamic Programming and self-training techniques can pave the way for  
 314 new deep-learning-oriented approaches for demanding CO problems.

315 **Limitations:** Our current empirical approach lacks theoretical guarantees on the convergence or the  
 316 approximate optimality of the obtained algorithm. Additionally, the implemented GNN is using core  
 317 modules, while more complex modules could result in further empirical improvements, which can be  
 318 the next step in this direction.

## 319 References

320 Sungsoo Ahn, Younggyo Seo, and Jinwoo Shin. Learning what to defer for maximum independent  
 321 sets. *International Conference on Machine Learning (ICML)*, 2020.

- 322 Maria Chiara Angelini and Federico Ricci-Tersenghi. Cracking nuts with a sledgehammer: when  
323 modern graph neural networks do worse than classical greedy algorithms. *arXiv preprint*  
324 *arXiv:2206.13211*, 2022.
- 325 Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016a.
- 326 Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint*  
327 *arXiv:1607.06450*, 2016b.
- 328 Richard Bellman. The theory of dynamic programming. *BULL*, 60(6):503–515, 1954.
- 329 Richard Bellman. Dynamic programming treatment of the travelling salesman problem. *JACM*, 9(1):  
330 61–63, 1962.
- 331 John R Current and David A Schilling. The covering salesman problem. *Transportation science*, 23  
332 (3), 1989.
- 333 Federico Errica, Marco Podda, Davide Bacciu, and Alessio Micheli. A fair comparison of graph  
334 neural networks for graph classification. *arXiv preprint arXiv:1912.09893*, 2019.
- 335 Thomas A Feo, Mauricio GC Resende, and Stuart H Smith. A greedy randomized adaptive search  
336 procedure for maximum independent set. *Journal of Global Optimization*, 42(5):860–878, 1994.
- 337 Maxime Gasse, Didier Chételat, Nicola Ferroni, Laurent Charlin, and Andrea Lodi. Exact combina-  
338 torial optimization with graph convolutional neural networks. *NIPS*, 32, 2019.
- 339 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image  
340 recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778,  
341 2016.
- 342 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint*  
343 *arXiv:1606.08415*, 2016.
- 344 Yujiao Hu, Yuan Yao, and Wee Sun Lee. A reinforcement learning approach for optimizing multiple  
345 traveling salesman problems over graphs. *Knowledge-Based Systems*, 204:106244, 2020.
- 346 JQ James, Wen Yu, and Jiatao Gu. Online vehicle routing with neural combinatorial optimization  
347 and deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 20  
348 (10):3806–3817, 2019.
- 349 Steven James, George Konidaris, and Benjamin Rosman. An analysis of monte carlo tree search.  
350 *AAAI Conference on Artificial Intelligence*, 31(1), Feb. 2017. URL [https://ojs.aaai.org/  
351 index.php/AAAI/article/view/11028](https://ojs.aaai.org/index.php/AAAI/article/view/11028).
- 352 Nikolaos Karalias and Andreas Loukas. Erdos goes neural: an unsupervised learning framework  
353 for combinatorial optimization on graphs. *Advances in neural information processing systems*  
354 (*NeurIPS*), 33:6659–6672, 2020.
- 355 Maryam Karimi-Mamaghan, Mehrdad Mohammadi, Patrick Meyer, Amir Mohammad Karimi-  
356 Mamaghan, and El-Ghazali Talbi. Machine learning at the service of meta-heuristics for solving  
357 combinatorial optimization problems: A state-of-the-art. *EJOR*, 296(2):393–422, 2022.
- 358 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*  
359 *arXiv:1412.6980*, 2014.
- 360 Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks.  
361 *International Conference on Learning Representations (ICLR)*, 2017.
- 362 Henrique Lemos, Marcelo Prates, Pedro Avelar, and Luis Lamb. Graph colouring meets deep learning:  
363 Effective graph neural network models for combinatorial problems. In *International Conference*  
364 *on Tools for Artificial Intelligence (ICTAI)*, pages 879–885. IEEE, 2019.
- 365 Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. [http:  
366 //snap.stanford.edu/data](http://snap.stanford.edu/data), June 2014.

- 367 Kaiwen Li, Tao Zhang, Rui Wang, Yuheng Wang, Yi Han, and Ling Wang. Deep reinforcement  
368 learning for combinatorial optimization: Covering salesman problems. *IEEE Transactions on*  
369 *Cybernetics*, 52(12):13142–13155, 2021.
- 370 Zhuwen Li, Qifeng Chen, and Vladlen Koltun. Combinatorial optimization with graph convolutional  
371 networks and guided tree search. *Advances in neural information processing systems (NeurIPS)*,  
372 31, 2018.
- 373 Nina Mazyavkina, Sergey Sviridov, Sergei Ivanov, and Evgeny Burnaev. Reinforcement learning for  
374 combinatorial optimization: A survey. *COR*, 134:105400, 2021.
- 375 Christopher Morris, Nils M. Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion  
376 Neumann. TUDataset: A collection of benchmark datasets for learning with graphs. In *International*  
377 *Conference on Machine Learning (ICML)*, 2020. URL [www.graphlearning.io](http://www.graphlearning.io).
- 378 Mohammadreza Nazari, Afshin Oroojlooy, Lawrence Snyder, and Martin Takáč. Reinforcement  
379 learning for solving the vehicle routing problem. *Advances in neural information processing*  
380 *systems (NeurIPS)*, 31, 2018.
- 381 Alex W. Nowak, Soledad Villar, Afonso S. Bandeira, and Joan Bruna. A note on learning algorithms  
382 for quadratic assignment with graph neural networks. *ArXiv*, abs/1706.07450, 2017.
- 383 Junyoung Park, Sanjar Bakhtiyar, and Jinkyoo Park. Schedulenet: Learn to solve multi-agent  
384 scheduling problems with reinforcement learning, 2021.
- 385 Marcelo Prates, Pedro HC Avelar, Henrique Lemos, Luis C Lamb, and Moshe Y Vardi. Learning to  
386 solve np-complete problems: A graph neural network for decision tsp. In *AAAI Conference on*  
387 *Artificial Intelligence*, volume 33, pages 4731–4738, 2019.
- 388 Martin JA Schuetz, J Kyle Brubaker, and Helmut G Katzgraber. Combinatorial optimization with  
389 physics-inspired graph neural networks. *Nature Machine Intelligence*, 4(4):367–377, 2022a.
- 390 Martin JA Schuetz, J Kyle Brubaker, Zhihuai Zhu, and Helmut G Katzgraber. Graph coloring with  
391 physics-inspired graph neural networks. *Physical Review Research*, 4(4):043131, 2022b.
- 392 Ulrich Schwalbe and Paul Walker. Zermelo and the early history of game theory. *Games and*  
393 *economic behavior*, 34(1):123–137, 2001.
- 394 Mohamed A Wahby Shalaby, Ayman R Mohammed, and Sally S Kassem. Supervised fuzzy c-means  
395 techniques to solve the capacitated vehicle routing problem. *IAJIT*, 18(3A):452–463, 2021.
- 396 Hong Tian and Dazi Li. Graph convolutional neural networks with am-actor-critic for minimum  
397 vertex cover problem. In *2021 33rd Chinese Control and Decision Conference (CCDC)*, pages  
398 3841–3846. IEEE, 2021.
- 399 Jan Toenshoff, Martin Ritzert, Hinrikus Wolf, and Martin Grohe. Run-csp: unsupervised learning of  
400 message passing networks for binary constraint satisfaction problems. *CoRR*, 2019.
- 401 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz  
402 Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information*  
403 *processing systems (NeurIPS)*, pages 5998–6008, 2017.
- 404 Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua  
405 Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- 406 Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. *NeurIPS*, 28, 2015.
- 407 Haoyu Wang and Pan Li. Unsupervised learning for combinatorial optimization needs meta-learning.  
408 *arXiv preprint arXiv:2301.03116*, 2023.
- 409 Haoyu Peter Wang, Nan Wu, Hang Yang, Cong Hao, and Pan Li. Unsupervised learning for  
410 combinatorial optimization with principled objective relaxation. In *NIPS*, 2022.
- 411 Nicholas C Wormald. Differential equations for random processes and random graphs. *AAP*, pages  
412 1217–1235, 1995.

- 413 Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A  
414 comprehensive survey on graph neural networks. *IEEE transactions on neural networks and*  
415 *learning systems*, 32(1):4–24, 2020.
- 416 Shunxin Xiao, Shiping Wang, Yuanfei Dai, and Wenzhong Guo. Graph neural networks in node  
417 classification: survey and evaluation. *Mach. Vis. Appl.*, 33:1–19, 2022.
- 418 Zhihao Xing and Shikui Tu. A graph neural network assisted monte carlo tree search approach to  
419 traveling salesman problem. *IEEE Access*, 8:108418–108428, 2020.
- 420 Zhihao Xing, Shikui Tu, and Lei Xu. Solve traveling salesman problem by monte carlo tree search  
421 and deep neural network. *arXiv preprint arXiv:2005.06879*, 2020.
- 422 Ke Xu, Frédéric Boussemart, Fred Hemery, and Christophe Lecoutre. Random constraint satisfaction:  
423 Easy generation of hard (satisfiable) instances. *Artificial intelligence*, 171(8-9):514–534, 2007.
- 424 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural  
425 networks? *arXiv preprint arXiv:1810.00826*, 2018.
- 426 Ruiyang Xu and Karl Lieberherr. Learning self-play agents for combinatorial optimization problems.  
427 *KER*, 35:e11, 2020.
- 428 Shenghe Xu, Shivendra S Panwar, Murali Kodialam, and TV Lakshman. Deep neural network  
429 approximated dynamic programming for combinatorial optimization. In *Proceedings of the AAAI*  
430 *Conference on Artificial Intelligence*, volume 34, pages 1684–1691, 2020.
- 431 Pinar Yanardag and S. V. N. Vishwanathan. Deep graph kernels. *SIGKDD*, 2015.
- 432 Feidiao Yang, Tiancheng Jin, Tie-Yan Liu, Xiaoming Sun, and Jialin Zhang. Boosting dynamic  
433 programming with neural networks for solving np-hard problems. In *ACML*, pages 726–739.  
434 PMLR, 2018.
- 435 Gal Yehuda, Moshe Gabel, and Assaf Schuster. It’s not what machines can learn, it’s what we cannot  
436 teach. In *ICML*, pages 10831–10841. PMLR, 2020.
- 437 Cong Zhang, Wen Song, Zhiguang Cao, Jie Zhang, Puay Siew Tan, and Xu Chi. Learning to  
438 dispatch for job shop scheduling via deep reinforcement learning. *Advances in neural information*  
439 *processing systems (NeurIPS)*, 33:1621–1632, 2020.
- 440 Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. *NeurIPS*, 31, 2018.
- 441 Rongkai Zhang, Cong Zhang, Zhiguang Cao, Wen Song, Puay Siew Tan, Jie Zhang, Bihan Wen,  
442 and Justin Dauwels. Learning to solve multiple-tsp with time window and rejections via deep  
443 reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- 444 Si Zhang, Hanghang Tong, Jiejun Xu, and Ross Maciejewski. Graph convolutional networks: a  
445 comprehensive review. *Computational Social Networks*, 6(1):1–23, 2019.
- 446 Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal Xhonneux, and Jian Tang. Neural bellman-ford  
447 networks: A general graph neural network framework for link prediction. *NeurIPS*, 34:29476–  
448 29490, 2021.