
Relational Curriculum Learning for Graph Neural Network

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Graph neural networks (GNNs) have achieved great success in representing data
2 with dependencies by recursively propagating and aggregating messages along
3 the edges. However, edges in real-world graphs often have varying degrees of
4 difficulty, and some edges may even be noisy to the downstream tasks. Therefore,
5 existing GNNs may lead to suboptimal learned representations because they usually
6 treat every edge in the graph equally. On the other hand, curriculum learning
7 (CL), which mimics the human learning principle of learning data samples in a
8 meaningful order, has been shown to be effective in improving the generalization
9 ability and robustness of representation learners by gradually proceeding from easy
10 to more difficult samples during training. Unfortunately, existing CL strategies
11 are designed for independent data samples and cannot be trivially generalized to
12 handle data dependencies. To address these issues, we propose a novel CL method
13 to gradually incorporate more edges into training according to their difficulty from
14 easy to hard, where the degree of difficulty is measured by how well the edges
15 are expected given the model training status. We demonstrate the strength of our
16 proposed method in improving the generalization ability and robustness of learned
17 representations through extensive experiments on nine synthetic datasets and nine
18 real-world datasets.

19 1 Introduction

20 Inspired by cognitive science studies [7, 28] that humans can benefit from the sequence of learning
21 basic (easy) concepts first and advanced (hard) concepts later, curriculum learning (CL) [2] suggests
22 training a machine learning model with easy data samples first and then gradually introducing more
23 hard samples into the model according to a designed pace, where the difficulty of samples can usually
24 be measured by their training loss [22]. Many previous studies have shown that this easy-to-hard
25 learning strategy can effectively improve the generalization ability of the model [2, 17, 13, 10, 30, 37],
26 and some studies [17, 13, 10] have shown that CL strategies can also increase the robustness of the
27 learned model against noisy training samples. An intuitive explanation is that in CL settings noisy
28 data samples correspond to harder samples and CL learner spends less time with the harder (noisy)
29 samples to achieve better generalization performance and robustness.

30 Although CL strategies have achieved great success in many fields such as computer vision and
31 natural language processing, existing methods are designed for independent data (such as images)
32 while designing effective CL methods for data with dependencies has been largely underexplored.
33 For example, in a citation network, two researchers with highly related research topics (e.g. machine
34 learning and data mining) are more likely to collaborate with each other, while the reason behind
35 a collaboration of two researchers with less related research topics (e.g. computer architecture and
36 social science) might be more difficult to understand. Prediction on one sample impacts that of
37 another, forming a graph structure that encompasses all samples connected by their dependencies.

38 There are many machine learning techniques for such graph-structured data, ranging from traditional
39 models like conditional random field [31], graph kernels [32], to modern deep models like GNNs [40].
40 However, traditional CL strategies are insufficient for them, which require not only considering the
41 difficulty in individual samples but also the difficulty of their dependencies to determine how to
42 gradually composite correlated samples for learning.

43 As previous CL strategies indicated that an easy-to-hard learning sequence on data samples can
44 improve the generalization and robustness performance, an intuitive question is whether a similar
45 strategy on data dependencies that iteratively involves easy-to-hard edges in learning can also benefit.
46 Unfortunately, there exists no trivial way to directly generalize existing CL strategies on independent
47 data to handle data dependencies due to several unique challenges: (1) **Difficulty in quantifying**
48 **edge selection criteria.** Existing CL studies on independent data often use supervised computable
49 metrics (e.g. training loss) to quantify sample difficulty, but how to quantify the difficulties of
50 understanding the dependencies between data samples which has no supervision is challenging.
51 (2) **Difficulty in designing an appropriate curriculum to gradually involve edges.** Existing CL
52 methods usually design a fixed pacing function to include a certain ratio of samples during training.
53 Similar to the human learning process, the model should ideally retain a certain degree of freedom to
54 adjust the pacing of including edges according to its own learning status. (3) **Difficulty in ensuring**
55 **convergence and a numerical steady process for CL in graphs.** Discrete changes in the number of
56 edges can cause drift in the optimal model parameters between training iterations. How to guarantee
57 a numerically stable learning process for CL on edges is challenging.

58 In order to address the aforementioned challenges, in this paper, we propose a novel CL algorithm
59 named **Relational Curriculum Learning (RCL)** to improve the generalization ability and robustness
60 of representation learners on data with dependencies. To address the first challenge, we propose an
61 approach to select the edges by quantifying their corresponding difficulties in a self-supervised learn-
62 ing manner. Specifically, for each training iteration, we choose K easiest edges whose corresponding
63 relations are most well-expected by the current model. Second, to design an appropriate learning
64 pace for gradually involving more edges in training, we present the learning process as a concise
65 optimization model, which automatically lets the model gradually increase the number K to involve
66 more edges in training according to its own status. Third, to ensure convergence of optimizing the
67 model, we propose a proximal optimization algorithm with a theoretical convergence guarantee and
68 an edge reweighting scheme to smooth the graph structure transition. Finally, we demonstrate the
69 superior performance of RCL compared to state-of-the-art comparison methods through extensive
70 experiments on both synthetic and real-world datasets.

71 2 Related work

72 **Curriculum learning (CL).** [2] first proposed the idea of CL in the context of machine learning,
73 aiming to improve model performance by gradually including easy to hard samples in training the
74 model. Self-paced learning [22] measures the difficulty of samples by their training loss, which
75 addressed the issue in previous works that difficulties of samples are generated by prior heuristic
76 rules. Therefore, the model can adjust the curriculum of samples according to its own training
77 status. Following works [16, 15, 44] further proposed many supervised measurement metrics for
78 determining curriculums, for example, the diversity of samples [15] or the consistency of model
79 predictions [44]. Meanwhile, many empirical and theoretical studies were proposed to explain why
80 CL could lead to generalization improvement from different perspectives. For example, studies such
81 as MentorNet [17] and Co-teaching [13] empirically found that utilizing CL strategy can achieve
82 better generalization performance when the given training data are noisy. [10] provided theoretical
83 explanations on the denoising mechanism that CL learners waste less time with the noisy samples as
84 they are considered harder samples. Some studies [2, 30, 37, 11, 21] also realized that CL can help
85 accelerate the optimization process of non-convex objectives and improve the speed of convergence
86 in the early stages of training.

87 Despite great success, most of the existing designed CL strategies are for independent data such as
88 images, and there is little work on generalizing CL strategies to handle samples with dependencies.
89 Few existing attempts on graph-structured data [23], such as [35, 5, 36, 25], simply treat nodes as
90 independent samples and then apply CL strategies on independent data, which ignore the fundamental
91 and unique dependency information that carried by the structure in data, thus can not well handle the
92 correlation between data samples. Furthermore, these models are mostly based on heuristic-based
93 sample selection strategies [5, 36, 25], which largely limit the generalizability of these methods.

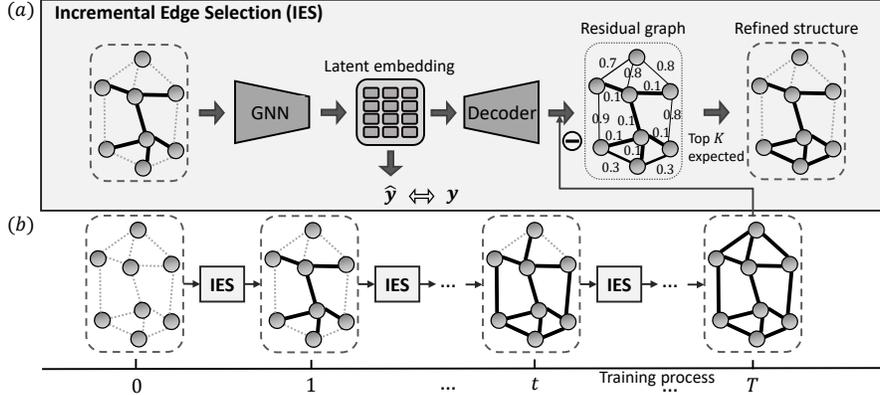


Figure 1: The overall framework of RCL. (a) The *Incremental Edge Selection* module first extracts the latent node embedding by the GNN model given the current training structure, then jointly learn the node prediction label \mathbf{y} and reconstructs the input structure by a decoder. A small residual error on an edge indicates the corresponding dependency is well expected and thus can be added to the refined structure for the next iteration. (b) The iterative learning process of RCL. The model starts with an empty structure and gradually includes more edges until the training structure converges to the input structure.

94 **Graph structure learning.** Another stream of existing studies that are related to our work is
 95 *graph structure learning*. Recent studies have shown that GNN models are vulnerable to adversarial
 96 attacks on graph structure [6, 39]. In order to address this issue, studies in *graph structure learning*
 97 usually aim to jointly learn an optimized graph structure and corresponding graph representations.
 98 Existing works [8, 4, 18, 43, 26] typically consider the hypothesis that the intrinsic graph structure
 99 should be sparse or low rank from the original input graph by pruning “irrelevant” edges. Thus, they
 100 typically use pre-deterministic methods [6, 45, 8] to preprocess graph structure such as singular value
 101 decomposition (SVD), or dynamically remove “redundant” edges according to the downstream task
 102 performance on the current sparsified structure [4, 18, 26]. However, modifying the graph topology
 103 will inevitably lose potential useful information lying in the removed edges. More importantly, the
 104 modified graph structure is usually optimized for maximizing the performance on the training set,
 105 which can easily lead to overfitting issues.

106 3 Preliminaries

107 **Graph Neural Networks** Graph neural networks (GNNs) are a class of methods that have shown
 108 promising progress in representing structured data in which data samples are correlated with each
 109 other. Typically, the data samples are treated as nodes while their dependencies are treated as edges
 110 in the constructed graph. Formally, we denote a graph as $G = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$
 111 is a set of nodes that $N = |\mathcal{V}|$ denotes the number of nodes in the graph and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set
 112 of edges. We also let $\mathbf{X} \in \mathbb{R}^{N \times b}$ denote the node attribute matrix and $\mathbf{A} \in \mathbb{R}^{N \times N}$ represents the
 113 adjacency matrix. Specifically, $A_{ij} = 1$ denotes there is an edge connecting nodes v_i and $v_j \in \mathcal{V}$,
 114 otherwise $A_{ij} = 0$. A GNN model f maps node feature matrix \mathbf{X} associated with the adjacency
 115 matrix \mathbf{A} to the model predictions $\hat{\mathbf{y}} = f(\mathbf{X}, \mathbf{A})$, and get the loss $L_{\text{GNN}} = L(\hat{\mathbf{y}}, \mathbf{y})$, where L is the
 116 objective function and \mathbf{y} is the ground-truth label of nodes. The loss on one node v_i is denoted as
 117 $l_i = L(\hat{y}_i, y_i)$.

118 **Curriculum Learning** In order to leverage the information carried by the various difficulties of
 119 data samples into the training process, Curriculum Learning (CL) [2, 22], which is inspired by the
 120 cognitive process of human learning principles that learning concepts in a meaningful order [7], is a
 121 popular training strategy that can improve the generalization ability and robustness of representation
 122 learners. Specifically, instead of randomly presenting all training samples to the model as in traditional
 123 machine learning algorithms, CL learners start with easy samples and gradually include harder ones
 124 during the training process, where the difficulty of samples can be measured by a predetermined
 125 policy or a supervised computable metric (e.g. training loss).

126 4 Methodology

127 As previous CL methods have shown that an easy-to-hard learning sequence of independent data
 128 samples can improve the generalization ability and robustness of the representation learner, the goal

129 of this paper is to develop an effective CL method on data with dependencies, which is extremely
 130 difficult due to several unique challenges: (1) Difficulty in designing a feasible principle to select
 131 edges by properly quantifying their difficulties. (2) Difficulty in designing an appropriate pace of
 132 curriculum to gradually involve more edges in training based on model status. (3) Difficulty in
 133 ensuring convergence and a numerical steady process for optimizing the CL model.

134 In order to address the above challenges, we propose a novel CL method named **Relational Curriculum**
 135 **Learning (RCL)**. The sequence, which gradually includes edges from easy to hard, is called *curricu-*
 136 *lum* and learned in different grown-up stages of training. In order to address the first challenge, we
 137 propose a self-supervised module *Incremental Edge Selection (IES)*, which is shown in Figure 1(a), to
 138 select the K easiest edges at each training iteration that are mostly expected by the current model. The
 139 details are elaborated in Section 4.1. To address the second challenge, we present a joint optimization
 140 framework to automatically increase the number of selected edges K given its own training status.
 141 The framework is elaborated in Figure 1(b) and details can be found in Section 4.2. Finally, to ensure
 142 convergence of optimization and steady the numerical process, we propose a proximal optimization
 143 algorithm with theoretical convergence guarantee in Section 4.2 Algorithm 1 and an edge reweighting
 144 scheme to smooth the discrete edge incrementing process in Section 4.3.

145 4.1 Incremental edge selection by quantifying difficulties of sample dependencies

146 Here we propose a novel way to select edges by first quantifying their difficulty levels. Existing works
 147 on independent data typically use supervised metrics such as training loss of samples to quantify their
 148 difficulty level, but there exists no supervised metrics on edges. To address this issue, we propose a
 149 self-supervised module *Incremental Edge Selection (IES)*. We first quantify the difficulty of edges by
 150 measuring how well the edges are expected from the currently learned embeddings of their connected
 151 nodes. Then the most well-expected edges are selected as the easiest edges for the next iteration of
 152 training. As shown in Figure 1(a), given the currently selected edges at iteration t , we first feed them
 153 to the GNN model to extract the latent node embeddings. Then we restore the latent node embeddings
 154 to the original graph structure through a decoder, which is called the reconstruction of the original
 155 graph structure. The residual graph \mathbf{R} , which is defined as the degree of mismatch between the
 156 original adjacency matrix \mathbf{A} and the reconstructed adjacency matrix $\tilde{\mathbf{A}}^{(t)}$, can be considered a strong
 157 indicator for describing how well the edges are expected by the current model. Specifically, a smaller
 158 residual error indicates a higher probability of being a well-expected edge.

159 With the developed self-supervised method to measure the difficulties of edges, here we formulate the
 160 key learning paradigm of selecting the top K easiest edges. To obtain the training adjacency matrix
 161 $\mathbf{A}^{(t)}$ that will be fed into the GNN model $f^{(t)}$, we introduce a learnable binary mask matrix \mathbf{S} with
 162 each element $\mathbf{S}_{ij} \in \{0, 1\}$. Thus, the training adjacency matrix at iteration t can be represented as
 163 $\mathbf{A}^{(t)} = \mathbf{S}^{(t)} \odot \mathbf{A}$. To filter out the edges with K smallest residual error, we penalize the summarized
 164 residual errors over the selected edges, which can be represented as $\sum_{i,j} \mathbf{S}_{ij} \mathbf{R}_{ij}$. Therefore, the
 165 learning objective can be presented as follows:

$$\begin{aligned} \min_{\mathbf{w}} L_{\text{GNN}} + \beta \sum_{i,j} \mathbf{S}_{ij} \mathbf{R}_{ij}, \\ \text{s.t. } \|\mathbf{S}\|_1 \geq K, \end{aligned} \quad (1)$$

166 where the first term $L_{\text{GNN}} = L(f(\mathbf{X}, \mathbf{A}^{(t)}; \mathbf{w}), \mathbf{y})$ is the node-level predictive loss, e.g. cross-entropy
 167 loss for node classification task. The second term $\sum_{i,j} \mathbf{S}_{ij} \mathbf{R}_{ij}$ aims at penalizing the residual errors
 168 over the edges selected by the mask matrix \mathbf{S} . β is a hyperparameter to tune the balance between
 169 terms. The constraint is to guarantee only the most K well-expected edges are selected.

171 More concretely, the value of a residual edge $\tilde{\mathbf{A}}_{ij}^{(t)} \in [0, 1]$ can be computed by a non-parametric
 172 kernel function $\kappa(\mathbf{z}_i^{(t)}, \mathbf{z}_j^{(t)})$, e.g. the inner product kernel. Then the residual error \mathbf{R}_{ij} between
 173 the input structure and the reconstructed structure can be defined as $\|\tilde{\mathbf{A}}_{ij}^{(t)} - \mathbf{A}_{ij}\|$, where $\|\cdot\|$ is
 174 commonly chosen to be the squared ℓ_2 -norm.

175 4.2 Automatically control the pace of increasing edges

176 In order to dynamically include more edges into training, an intuitive way is to iteratively increase
 177 the value of K in Equation 1 to allow more edges to be selected. However, it is difficult to determine
 178 an appropriate value of K respect to the training status of the model. Besides, directly solving

Algorithm 1 Proximal Alternating Minimization Algorithm for Optimizing Equation 2

Input: Node features \mathbf{X} , adjacency matrix \mathbf{A} , a stepsize μ and hyperparameter γ

Output: The learnable parameter \mathbf{w} of GNN model f

```
1: Initialize  $\mathbf{w}^{(0)}, \mathbf{S}^{(0)}, \lambda$ 
2: while Not converged do
3:    $\mathbf{w}^{(t)} = \arg \min_{\mathbf{w}} L(f(\mathbf{X}, \mathbf{A}^{(t-1)}; \mathbf{w}), \mathbf{y}) + \beta \sum_{i,j} \mathbf{S}_{ij} \left\| \tilde{\mathbf{A}}_{ij}^{(t-1)} - \mathbf{A}_{ij} \right\| + \frac{\gamma}{2} \|\mathbf{w} - \mathbf{w}^{(t-1)}\|$ 
4:   Given  $\mathbf{w}^{(t)}$ , extract latent nodes embedding  $\mathbf{Z}^{(t)}$  from GNN model  $f$ 
5:   Calculate reconstructed structure  $\tilde{\mathbf{A}}_{ij}^{(t)} = \kappa(\mathbf{z}_i^{(t)}, \mathbf{z}_j^{(t)})$  for all pairs of  $i, j$ 
6:    $\mathbf{S}^{(t)} = \arg \min_{\mathbf{S}} \beta \sum_{i,j} \mathbf{S}_{ij} \left\| \mathbf{A}_{ij} - \tilde{\mathbf{A}}_{ij}^{(t)} \right\| + g(\mathbf{S}; \lambda) + \frac{\gamma}{2} \|\mathbf{S} - \mathbf{S}^{(t-1)}\|$ 
7:   Compute  $\mathbf{A}^{(t)} = \mathbf{S}^{(t)} \odot \mathbf{A}$ 
8:   if  $\mathbf{A}^{(t)} \neq \mathbf{A}$  then
9:     Increase  $\lambda$  by stepsize  $\mu$ 
10:  end if
11: end while
```

179 Equation 1 is difficult since \mathbf{S} is a binary matrix where each element $\mathbf{S}_{ij} \in \{0, 1\}$, optimizing \mathbf{S}
180 would require solving a discrete constraint program at each iteration. To address this issue, we first
181 relax the problem into continuous optimization so that each \mathbf{S}_{ij} can be allowed to take any value
182 in the interval $[0, 1]$. Then we treat the constraint as a Lagrange multiplier and solve an equivalent
183 problem by substituting the constraint to a regularization term $g(\mathbf{S}; \lambda)$, thus, our overall loss function
184 can be rewritten as:

$$\min_{\mathbf{w}, \mathbf{S}} L_{\text{GNN}} + \beta \sum_{i,j} \mathbf{S}_{ij} \mathbf{R}_{ij} + g(\mathbf{S}; \lambda), \quad (2)$$

185 where $g(\mathbf{S}; \lambda) = \lambda \|\mathbf{S} - \mathbf{A}\|$ and $\|\cdot\|$ is commonly chosen to be the squared ℓ_2 -norm. Since the
186 training adjacency matrix $\mathbf{A}^{(t)} = \mathbf{S}^{(t)} \odot \mathbf{A}$, as $\lambda \rightarrow \infty$, more edges in the input structure are included
187 until the training adjacency matrix $\mathbf{A}^{(t)}$ converges to the input adjacency matrix \mathbf{A} . Specifically, the
188 regularization term $g(\mathbf{S}; \lambda)$ controls the learning scheme by the *age parameter* λ , where $\lambda = \lambda(t)$
189 grows with the number of iterations. By monotonously increasing the value of λ , the regularization
190 term $g(\mathbf{S}; \lambda)$ will push the mask matrix gradually converge to the input adjacency matrix \mathbf{A} , resulting
191 in more edges automatically involved in the training structure.

192 **Optimization of learning objective.** It is worth noting that optimizing our objective function in
193 Equation 2 requires jointly optimizing parameter \mathbf{w} of GNN model f and the mask matrix \mathbf{S} . In
194 order to address this problem, we propose a proximal alternating optimization schema to iteratively
195 update \mathbf{w} and \mathbf{S} in sequence. The full algorithm is presented in Algorithm 1. As we can see, our
196 algorithm takes the input of node feature matrix \mathbf{X} and original adjacency matrix \mathbf{A} , a stepsize μ
197 to control the increasing pace of age parameter λ , and a hyperparameter γ to tune the proximal
198 terms. After initializing the parameters \mathbf{w} and \mathbf{S} , it alternates between two updating steps until it
199 finally converges: (1) Step 3 first learns the optimal parameter of GNN model f with the current
200 training adjacency matrix; (2) Step 4 & 5 extracts the latent node embedding by fixing the GNN
201 model parameter and build the reconstructed adjacency matrix by the kernel function; (3) Step 6
202 learns the optimal mask matrix \mathbf{S} with the reconstructed adjacency matrix and regularization term;
203 (4) Step 7 refines the training adjacency matrix with respect to the updated mask matrix; (5) The
204 age parameter λ is increased when the training adjacency matrix $\mathbf{A}^{(t)}$ is still different from the input
205 adjacency matrix \mathbf{A} , thus more edges will be included in the next iteration of the training.

206 **Theorem 4.1.** *We have the following convergence guarantees for Algorithm 1:*

- 207 • **Avoidance of Saddle Points.** *If the second derivatives of $L(f(\mathbf{X}, \mathbf{A}^{(t)}; \mathbf{w}), \mathbf{y})$ and $g(\mathbf{S}; \lambda)$ are*
208 *continuous, then for sufficiently large γ , any bounded sequence $(\mathbf{w}^{(t)}, \mathbf{S}^{(t)})$ generated by Algorithm*
209 *1 with random initializations will not converge to a strict saddle point of F almost surely.*
210 • **Second Order Convergence.** *If the second derivatives of $L(f(\mathbf{X}, \mathbf{A}^{(t)}; \mathbf{w}), \mathbf{y})$ and $g(\mathbf{S}; \lambda)$ are*
211 *continuous, and $L(f(\mathbf{X}, \mathbf{A}^{(t)}; \mathbf{w}), \mathbf{y})$ and $g(\mathbf{S}; \lambda)$ satisfy the Kurdyka-Łojasiewicz (KL) property*
212 *[33], then for sufficiently large γ , any bounded sequence $(\mathbf{w}^{(t)}, \mathbf{S}^{(t)})$ generated by Algorithm 1 with*
213 *random initialization will almost surely converge to a second-order stationary point of F .*

214 The detailed proof can be found in Appendix B.

215 **4.3 Smooth structure transition by edge reweighting**

216 Note that in the Algorithm 1, the optimization process requires iteratively updating the parameters
 217 \mathbf{w} of the GNN model f and training adjacency matrix $\mathbf{A}^{(t)}$, where $\mathbf{A}^{(t)}$ varies discretely between
 218 iterations. However, GNN models mostly work in a message-passing fashion, which computes
 219 node representations by recursively aggregating information along edges from neighboring nodes.
 220 Discretely modifying the number of edges will result in a great drift of the optimal model parameters
 221 between iterations. Therefore, it can increase the difficulty of finding optimal parameters and even
 222 hurt the generalization ability of the model in some cases. Besides the numerical problem caused
 223 by discretely increasing the number of edges, another issue raised by the CL strategy in Section 4.1
 224 is the trustworthiness of the estimated edge difficulty, which is inferred by the residual error on the
 225 edges. Although the residual error can reflect how well edges are expected in the ideal case, the
 226 quality of the learned latent node embeddings may affect the validity of this metric and compromise
 227 the quality of the designed curriculum by the CL strategy.

228 To address both issues, we propose a novel edge reweighting scheme to (1) smooth the transition
 229 of the training structure between iterations, and (2) reduce the weight of edges that connect nodes
 230 with low-confidence latent embeddings. Formally, we use a smoothed version of structure $\bar{\mathbf{A}}^{(t)}$ to
 231 substitute $\mathbf{A}^{(t)}$ for training the GNN model f in step 3 of Algorithm 1, where the mapping from $\mathbf{A}^{(t)}$
 232 to $\bar{\mathbf{A}}^{(t)}$ can be represented as:

$$\bar{\mathbf{A}}_{ij}^{(t)} = \pi_{ij}^{(t)} \mathbf{A}_{ij}^{(t)}, \quad (3)$$

233 where $\pi_{ij}^{(t)}$ is the weight imposed to edge e_{ij} at iteration t . $\pi_{ij}^{(t)}$ is calculated by considering the
 234 counted occurrences of edge e_{ij} until the iteration t and the confidence of the latent embedding for
 235 the connected pair of nodes v_i and v_j :

$$\pi_{ij}^{(t)} = \psi(e_{ij})\rho(v_i)\rho(v_j), \quad (4)$$

236 where ψ is a function that reflects the number of edge occurrences and ρ is a function to reflect the
 237 degree of confidence for the learned latent node embedding. The details of these two functions are
 238 described as follow.

239 **Smooth the transition of the training structure between iterations.** In order to obtain a smooth
 240 transition of the training structure between iterations, we take the learned curriculum of selected edges
 241 into consideration. Formally, we model ψ by a smooth function of the edge selected occurrences
 242 compared to the model iteration occurrences before the current iteration:

$$\psi(e_{ij}) = t(e_{ij})/t, \quad (5)$$

243 where t is the number of current iterations and $t(e_{ij})$ represents the counting number of selecting
 244 edge e_{ij} . Therefore, we transform the original discretely changing training structure into a smoothly
 245 changing one by taking the historical edge selection curriculum into consideration.

246 **Reduce the influence of nodes with low confidence latent embeddings.** As introduced in our
 247 Algorithm 1 line 6, the estimated structure $\bar{\mathbf{A}}$ is inferred from the latent embedding \mathbf{Z} , which is
 248 extracted from the trained GNN model f . Such estimated latent embedding may possibly shift from
 249 the true underlying embedding, which results in the inaccurately reconstructed structure around the
 250 node. In order to alleviate this issue, we model the function ρ by the training loss on nodes, which
 251 indicates the confidence of their learned latent embeddings. This idea is similar to previous CL
 252 strategies on inferring the difficulty of data samples by their supervised training loss. Specifically, a
 253 larger training loss indicates a low confident latent node embedding. Mathematically, the weights
 254 $\rho(v_i)$ on node v_i can be represented as a distribution of their training loss:

$$\rho(v_i) \sim e^{-l_i} \quad (6)$$

255 where l_i is the training loss on node v_i . Therefore, a node with a larger training loss will result in a
 256 smaller value of $\rho(v_i)$, which reduces the weight of its connecting edges.

257 **5 Experiments**

258 In this section, the experimental settings are introduced first in Section 5.1, then the performance
 259 of the proposed method on both synthetic and real-world datasets are presented in Section 5.2. We
 260 further present the robustness test on our CL method against topological structure adversarial attack in
 261 Section 5.3. Intuitive visualizations of the edge selection curriculum are shown in Section 5.4. In addi-
 262 tion, we verify the effectiveness of framework components through ablation studies in Appendix A.2
 263 and measure the parameter sensitivity in Appendix A.2 due to the space limit.

Homo ratio	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
GCN	50.84±1.03	56.50±0.50	65.17±0.48	77.94±0.54	87.15±0.44	93.27±0.24	97.48±0.25	99.10±0.17	99.93±0.03
GNNSVD	<u>54.96±0.76</u>	<u>58.45±0.56</u>	63.06±0.63	70.23±0.61	80.51±0.41	85.02±0.46	90.31±0.27	94.23±0.22	96.74±0.23
ProGNN	47.87±0.87	<u>54.59±0.55</u>	65.39±0.44	76.96±0.49	87.76±0.51	93.16±0.34	97.60±0.31	99.04±0.19	99.94±0.03
NeuralSparse	51.42±1.35	57.99±0.69	65.10±0.43	75.37±0.34	87.40±0.29	93.54±0.28	97.16±0.15	99.01±0.22	99.83±0.07
PTDNet	48.21±1.98	55.52±2.82	65.82±0.94	79.37±0.45	89.17±0.39	94.19±0.18	98.61±0.12	99.51±0.09	99.81±0.05
CLNodes	50.37±0.73	56.64±0.56	65.04±0.66	77.52±0.48	86.85±0.44	93.10±0.47	97.34±0.25	99.02±0.18	99.88±0.04
RCL	57.57±0.43	62.06±0.28	73.98±0.55	84.54±0.75	92.69±0.09	97.42±0.17	99.62±0.05	99.89±0.02	99.93±0.06
GIN	48.33±1.89	53.62±1.39	64.08±0.99	77.55±1.10	85.31±0.75	90.57±0.36	97.82±0.18	99.59±0.11	<u>99.91±0.02</u>
GNNSVD	43.21±1.60	45.68±1.66	54.90±1.16	68.29±0.79	79.76±0.52	85.63±0.44	93.65±0.39	97.22±0.17	98.94±0.17
ProGNN	45.76±1.40	52.96±1.01	64.12±1.07	76.95±0.87	85.13±0.71	89.96±0.55	96.54±0.48	99.51±0.12	99.78±0.05
NeuralSparse	50.23±2.05	54.12±1.52	62.81±0.75	76.98±1.17	85.14±0.94	92.57±0.44	98.02±0.20	99.61±0.12	99.91±0.05
PTDNet	<u>53.23±2.76</u>	<u>56.12±2.03</u>	65.81±1.38	<u>77.81±1.02</u>	86.14±0.65	93.21±0.74	97.08±0.41	99.51±0.18	99.91±0.03
CLNodes	45.36±1.42	51.10±1.15	62.53±0.88	<u>75.83±1.07</u>	87.76±0.90	94.25±0.44	98.30±0.26	99.60±0.09	99.92±0.03
RCL	57.63±0.66	62.08±1.17	71.02±0.61	80.61±0.69	88.62±0.43	94.88±0.36	98.19±0.19	99.32±0.08	99.89±0.04
GraphSage	62.57±0.55	67.33±0.64	71.06±0.74	80.88±0.54	85.88±0.51	91.42±0.37	95.26±0.33	97.78±0.16	99.52±0.13
GNNSVD	64.42±0.80	65.71±0.39	67.12±0.58	68.47±0.50	77.70±0.65	82.86±0.50	87.81±0.71	91.61±0.55	95.01±0.50
ProGNN	58.57±2.09	66.75±0.91	72.14±0.64	81.27±0.44	<u>86.89±0.47</u>	<u>92.10±0.39</u>	95.21±0.30	97.51±0.23	99.50±0.11
NeuralSparse	61.70±0.77	66.65±0.66	70.60±0.79	79.65±0.45	84.19±0.91	91.31±0.54	94.86±0.53	97.16±0.23	99.55±0.19
PTDNet	65.72±1.08	69.25±0.92	72.60±0.77	79.65±0.45	86.54±0.56	91.79±0.53	96.10±0.58	97.98±0.13	99.78±0.08
CLNodes	69.41±0.66	<u>70.83±0.58</u>	<u>75.51±0.36</u>	<u>82.65±0.43</u>	87.08±0.56	91.58±0.41	95.91±0.38	<u>98.33±0.26</u>	99.57±0.14
RCL	68.03±0.37	71.39±0.51	76.99±0.99	83.76±0.55	88.24±0.30	93.34±0.56	97.66±0.52	98.86±0.28	99.64±0.08

Table 1: Node classification accuracy on synthetic datasets (%). The best-performing method on each backbone GNN model is highlighted in bold, while the second-best method is underlined. In situations where RCL’s performance is not strictly the best among all methods, we can see that almost all methods can achieve a near-perfect performance and RCL is still close to the best methods.

264 5.1 Experimental settings

265 **Synthetic datasets.** To evaluate the effectiveness of our proposed method on datasets with ground-
266 truth difficulty labels on edges, we follow previous studies [19, 1] to generate a set of synthetic
267 datasets, where the formation probability of an edge is designed to reflect its likelihood to positively
268 contribute to the node classification job, which indicates its ground-truth difficulty level. Specifically,
269 the nodes in a generated graph are divided into 10 equally sized node classes 1, 2, . . . , 10, and the
270 node features are sampled from overlapping multi-Gaussian distributions. Each generated graph is
271 associated with a *homophily coefficient (homo)* which indicates the probability of a node forming
272 an edge to another node with the same label. For the rest edges that are formed between nodes with
273 different labels, the probability of forming an edge is inversely proportional to the distances between
274 their labels. Nodes with close classes are more likely to be connected since the formation probability
275 decreases with the distance of the node label, and connections from nodes with close classes can
276 increase the likelihood of accurately classifying a node due to the homophily property of the designed
277 node classification task. Therefore, an edge with a high formation probability indicates a higher
278 chance to positively contribute to the node classification task because it connects a node with a close
279 class, and thus can be considered an easy edge. We vary the value of *homo* to generate nine graphs in
280 total. More details and visualization about the synthetic dataset can be found in Appendix A.1.

281 **Real-world datasets.** To further evaluate the performance of our proposed method in real-world
282 scenarios, nine benchmark real-world attributed network datasets, including four citation network
283 datasets Cora, Citeseer, Pubmed [42] and ogbn-arxiv [14], two coauthor network datasets CS and
284 Physics [29], two Amazon co-purchase network datasets Photo and Computers [29], and one protein
285 interaction network ogbn-proteins [14]. We follow the data splits from [3] on citation networks and
286 use a 5-fold cross-validation setting on coauthor and Amazon co-purchase networks. All datasets
287 are publicly available from Pytorch-geometric library [9] and Open Graph Benchmark (OGB) [14],
288 where basic statistics are reported in Table 2.

289 **Comparison methods.** We incorporate three commonly used GNN models, including GCN [20],
290 GraphSAGE [12], and GIN [41], as the baseline model and also the backbone model for RCL. In
291 addition to evaluating our proposed method against the baseline GNNs, we further leverage two
292 categories of state-of-the-art comparison methods in the experiments: (1) We incorporate four graph
293 structure learning methods GNNSVD [8], ProGNN [18], NeuralSparse [43], and PTDNet [26]; (2)
294 We further compare with a curriculum learning method named CLNode [36] which gradually select
295 nodes in the order of the difficulties defined by a heuristic-based strategy. More details about the
296 comparison methods can be found in Appendix A.1.

297 **Initializing graph structure by a pre-trained model.** It is worth noting that the model needs
298 an initial training graph structure $\mathbf{A}^{(0)}$ in the initial stage of training. An intuitive way is that we
299 can initialize the model to work in a purely data-driven scenario that starts only with isolated nodes
300 where no edges exist. However, an instructive initial structure can greatly reduce the search cost and
301 computational burden. Inspired by many previous CL works [37, 11, 17, 44] that incorporate prior
302 knowledge of a pre-trained model into designing curriculum for the current model, we initialize the

	Cora	Citeseer	Pubmed	CS	Physics	Photo	Computers	ogbn-arxiv	ogbn-proteins
# nodes	2,708	3,327	19,717	18,333	34,493	7,650	13,752	169,343	132,534
# edges	10,556	9,104	88,648	163,788	495,924	238,162	491,722	1,166,243	39,561,252
# features	1,433	3,703	500	6,805	8,415	745	767	100	8
GCN	85.74±0.42	78.93±0.32	87.91±0.09	93.03±0.32	96.55±0.15	93.25±0.70	88.09±0.40	71.74±0.29	72.51±0.35
GNNsVD	83.24±1.03	74.80±0.87	88.81±0.38	93.79±0.11	96.11±0.13	89.63±0.73	86.49±0.77	67.44±0.51	66.92±0.64
ProGNN	85.66±0.61	74.78±0.55	87.22±0.33	<u>94.04±0.19</u>	<u>96.75±0.26</u>	92.07±0.67	88.72±0.59	-	-
NeuralSparse	85.95±0.98	76.24±0.48	86.83±0.40	92.31±0.47	95.56±0.30	90.57±0.90	88.62±0.83	-	-
PTDNet	83.84±0.95	77.54±0.42	87.89±0.08	93.60±0.43	96.56±0.09	88.92±0.87	87.52±0.70	-	-
CLNode	85.67±0.33	78.99±0.57	89.50±0.28	93.83±0.24	95.76±0.16	93.39±0.83	89.28±0.38	70.95±0.18	71.40±0.32
RCL	87.15±0.44	79.79±0.55	89.79±0.12	94.66±0.32	97.02±0.23	94.41±0.76	90.23±0.23	74.08±0.33	75.19±0.26
GIN	84.43±0.65	74.87±0.20	85.72±0.40	91.48±0.36	95.62±0.30	93.02±0.91	86.94±1.58	69.26±0.34	74.51±0.32
GNNsVD	82.23±0.65	72.11±0.70	88.31±0.15	91.40±0.87	95.30±0.29	89.49±1.11	82.66±2.26	67.79±0.41	70.65±0.53
ProGNN	85.02±0.41	78.12±0.93	87.82±0.51	-	-	92.23±0.67	83.54±1.48	-	-
NeuralSparse	84.92±0.58	75.44±0.87	86.11±0.49	89.66±0.82	95.05±0.57	<u>93.28±0.83</u>	<u>87.22±0.54</u>	-	-
PTDNet	83.02±1.01	75.00±0.74	88.04±0.29	91.01±0.21	95.57±0.40	90.70±0.76	87.08±0.65	-	-
CLNode	83.52±0.77	75.82±0.58	86.92±0.61	91.71±0.41	95.75±0.46	92.78±0.90	85.93±1.53	70.58±0.17	73.97±0.31
RCL	86.64±0.39	77.60±0.18	89.17±0.29	93.92±0.27	96.75±0.17	93.88±0.51	89.76±0.19	72.55±0.15	78.76±0.22
GraphSage	86.22±0.27	77.27±0.23	88.50±0.16	94.22±0.18	96.26±0.34	93.82±0.51	88.62±0.21	71.49±0.27	77.68±0.20
GNNsVD	83.11±0.82	73.19±0.49	88.42±0.38	93.86±0.36	95.96±0.12	89.31±0.53	81.46±1.15	69.82±0.34	71.82±0.39
ProGNN	86.23±0.42	74.45±0.83	88.52±0.45	-	-	90.89±0.69	89.34±0.54	-	-
NeuralSparse	84.60±0.52	76.32±0.55	89.02±0.39	93.89±0.58	96.67±0.20	90.78±1.06	88.37±0.37	-	-
PTDNet	86.03±0.60	76.07±0.58	86.78±0.45	93.78±0.43	95.32±0.31	92.96±0.87	84.89±1.47	-	-
CLNode	86.60±0.64	77.23±0.54	88.76±0.57	94.13±0.34	96.87±0.45	93.90±0.42	89.57±0.62	71.54±0.20	78.40±0.41
RCL	86.90±0.39	78.95±0.18	90.14±0.43	95.05±0.23	96.88±0.19	95.06±0.52	90.47±0.38	73.13±0.14	79.89±0.35

Table 2: Node classification results on real-world datasets (%). The best-performing method on each backbone is highlighted in bold and second-best is underlined. (-) denotes an out-of-memory issue.

303 training structure $\mathbf{A}^{(0)}$ by a pre-trained vanilla GNN model f^* . Specifically, we follow the same
304 steps from line 4 to line 7 in the algorithm 1 to obtain the initial training structure $\mathbf{A}^{(0)}$ but the latent
305 node embedding is extracted from the pre-trained model f^* .

306 **Implementation Details** We use the baseline model (GCN, GIN, GraphSage) as the backbone
307 model for both our RCL method and all comparison methods. For a fair comparison, we require all
308 models follow the same GNN architecture with two convolution layers. For each split, we run each
309 model 10 times to reduce the variance in particular data splits. Test results are according to the best
310 validation results. General training hyperparameters (such as learning rate or the number of training
311 epochs) are equal for all models.

312 5.2 Effectiveness results

313 Table 1 presents the node classification results of the synthetic datasets. We report the average
314 accuracy and standard deviation for each model against the *homo* of generated graphs. From the
315 table, we observe that our proposed method RCL consistently achieves the best or most competitive
316 performance to all the comparison methods over three backbone GNN architectures. Specifically,
317 RCL outperforms the second best method on average by 4.17%, 2.60%, and 1.06% on GCN, GIN,
318 and GraphSage backbones, respectively. More importantly, the proposed RCL method performs
319 significantly better than the second best model when the *homo* of generated graphs is low (≤ 0.5), on
320 average by 6.55% on GCN, 4.17% on GIN, and 2.93% on GraphSage backbones. These demonstrate
321 that our proposed RCL method significantly improves the model’s capability of learning an effective
322 representation to downstream tasks especially when the edge difficulties vary largely in the data.

323 We report the experimental results of the real-world datasets in Table 2. The results demonstrate the
324 strength of our proposed method by consistently achieving the best results in all 9 datasets by GCN
325 backbone architecture, all 9 datasets by GraphSage backbone architecture, and 8 out of 9 datasets by
326 GIN backbone architecture. Specifically, our proposed method improved the performance of baseline
327 models on average by 1.86%, 2.83%, and 1.62% over GCN, GIN, and GraphSage, and outperformed
328 the second best models model on average by 1.37%, 2.49%, and 1.22% over the three backbone
329 models, respectively. The results demonstrate that the proposed RCL method consistently improves
330 the performance of GNN models in real-world scenarios.

331 Our experimental results are statically sound. In 43 out of 48 tasks our method outperforms the
332 second-best performing model with strong statistical significance. Specifically, we have in 30 out of
333 43 cases with a significance $p < 0.001$, in 8 out of 43 cases with a significance $p < 0.01$, and in 5
334 out of 43 cases with a significance $p < 0.05$. Such statistical significance results can demonstrate
335 that our proposed method can consistently perform better than the baseline models in both scenarios.

336 5.3 Robustness analysis against adversarial topological structure attack

337 To further examine the robustness of the RCL method on extracting powerful representation from
338 correlated data samples, we follow previous works [18, 26] to randomly inject fake edges into
339 real-world graphs. This adversarial attack can be viewed as adding random noise to the topological
340 structure of graphs. Specifically, we randomly connect M pairs of previously unlinked nodes in the

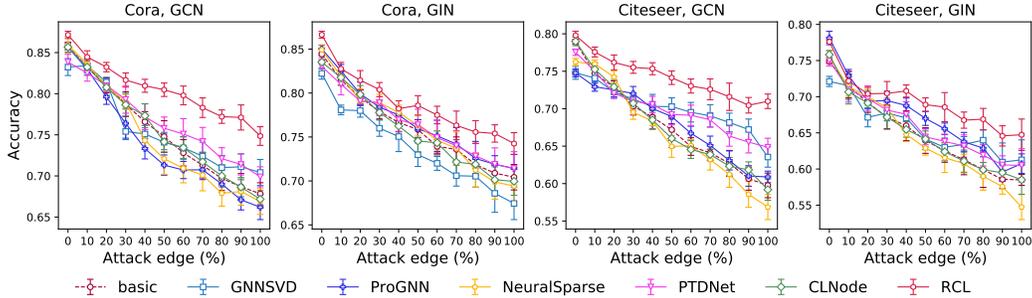


Figure 2: Node classification accuracy (%) on Cora and Citeseer under random structure attack. The attack edge ratio is computed versus the original number of edges, where 100% means that the number of inserted edges is equal to the number of original edges.

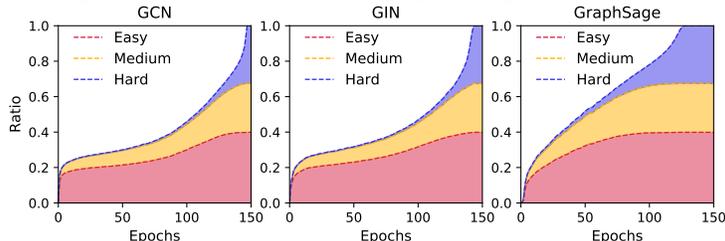


Figure 3: Visualization of edge selection process during training.

341 real-world datasets, where the value of M varies from 10% to 100% of the original edges. We then
 342 train RCL and all the comparison methods on the attacked graph and evaluate the node classification
 343 performance. The results are shown in Figure 2, we can observe that RCL shows strong robustness to
 344 adversarial structural attacks by consistently outperforming all compared methods on all datasets.
 345 Especially, when the proportion of added noisy edges is large ($> 50\%$), the improvement becomes
 346 more significant. For instance, under the extremely noisy ratio at 100%, RCL outperforms the second
 347 best model by 4.43% and 2.83% on Cora dataset, and by 6.13%, 3.47% on Citeseer dataset, with
 348 GCN and GIN backbone models, respectively.

349 5.4 Visualization of learned edge selection curriculum

350 Besides the effectiveness and robustness of the RCL method on downstream classification results,
 351 it is also interesting to verify whether the learned edge selection curriculum satisfies the rule from
 352 easy to hard. Since real-world datasets do not have ground-truth labels of difficulty on edges, we
 353 conduct visualization experiments on synthetic datasets, where the difficulty of each edge can be
 354 indicated by its formation probability. Specifically, we classify edges into three balanced categories
 355 according to their difficulty: easy, medium, and hard. Here, we define all homogenous edges that
 356 connect nodes with the same class as easy, edges connecting nodes with adjacent classes as medium,
 357 and the remaining edges connecting nodes with far away classes as hard. We report the proportion of
 358 edges selected for each category during training in Figure 3. We can observe that RCL can effectively
 359 select most of the easy edges at the early stage of training, then more easy edges and most medium
 360 edges are gradually included during training, and most hard edges are left unselected until the end
 361 stage of training. Such edge selection behavior is highly consistent with the core idea of designing
 362 a curriculum for edge selection, which verifies that our proposed method can effectively design
 363 curriculums to select edges according to their difficulty from easy to hard.

364 6 Conclusion

365 This paper focuses on developing a novel CL method to improve the generalization ability and
 366 robustness of GNN models on learning representations of data samples with dependencies. The
 367 proposed method **Relational Curriculum Learning (RCL)** effectively addresses the unique challenges
 368 in designing CL strategy for handling dependencies. First, a self-supervised learning module is
 369 developed to select appropriate edges that are expected by the model. Then an optimization model is
 370 presented to iteratively increment the edges according to the model training status and a theoretical
 371 guarantee of the convergence on the optimization algorithm is given. Finally, an edge reweighting
 372 scheme is proposed to steady the numerical process by smoothing the training structure transition.
 373 Extensive experiments on synthetic and real-world datasets demonstrate the strength of RCL in
 374 improving the generalization ability and robustness.

References

- 375
- 376 [1] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr
377 Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional
378 architectures via sparsified neighborhood mixing. In *international conference on machine*
379 *learning*, pages 21–29. PMLR, 2019.
- 380 [2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning.
381 In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48,
382 2009.
- 383 [3] Jie Chen, Tengfei Ma, and Cao Xiao. Fastgcn: Fast learning with graph convolutional networks
384 via importance sampling. In *International Conference on Learning Representations*, 2018.
- 385 [4] Yu Chen, Lingfei Wu, and Mohammed Zaki. Iterative deep graph learning for graph neural
386 networks: Better and robust node embeddings. *Advances in neural information processing*
387 *systems*, 33:19314–19326, 2020.
- 388 [5] Guanyi Chu, Xiao Wang, Chuan Shi, and Xunqiang Jiang. Cuco: Graph representation with
389 curriculum contrastive learning. In *IJCAI*, pages 2300–2306, 2021.
- 390 [6] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. Adversarial attack
391 on graph structured data. In *International conference on machine learning*, pages 1115–1124.
392 PMLR, 2018.
- 393 [7] Jeffrey L Elman. Learning and development in neural networks: The importance of starting
394 small. *Cognition*, 48(1):71–99, 1993.
- 395 [8] Negin Entezari, Saba A Al-Sayouri, Amirali Darvishzadeh, and Evangelos E Papalexakis. All
396 you need is low (rank) defending against adversarial attacks on graphs. In *Proceedings of the*
397 *13th International Conference on Web Search and Data Mining*, pages 169–177, 2020.
- 398 [9] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric.
399 *arXiv preprint arXiv:1903.02428*, 2019.
- 400 [10] Tieliang Gong, Qian Zhao, Deyu Meng, and Zongben Xu. Why curriculum learning & self-
401 paced learning work in big/noisy data: A theoretical perspective. *Big Data & Information*
402 *Analytics*, 1(1):111, 2016.
- 403 [11] Guy Hacoheh and Daphna Weinshall. On the power of curriculum learning in training deep
404 networks. In *International Conference on Machine Learning*, pages 2535–2544. PMLR, 2019.
- 405 [12] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large
406 graphs. *Advances in neural information processing systems*, 30, 2017.
- 407 [13] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi
408 Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels.
409 *Advances in neural information processing systems*, 31, 2018.
- 410 [14] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele
411 Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs.
412 *arXiv preprint arXiv:2005.00687*, 2020.
- 413 [15] Lu Jiang, Deyu Meng, Shoou-I Yu, Zhenzhong Lan, Shiguang Shan, and Alexander Hauptmann.
414 Self-paced learning with diversity. *Advances in neural information processing systems*, 27,
415 2014.
- 416 [16] Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander G Hauptmann. Self-paced
417 curriculum learning. In *Twenty-ninth AAAI conference on artificial intelligence*, 2015.
- 418 [17] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning
419 data-driven curriculum for very deep neural networks on corrupted labels. In *International*
420 *conference on machine learning*, pages 2304–2313. PMLR, 2018.

- 421 [18] Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. Graph
422 structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD*
423 *international conference on knowledge discovery & data mining*, pages 66–74, 2020.
- 424 [19] Fariba Karimi, Mathieu Génois, Claudia Wagner, Philipp Singer, and Markus Strohmaier.
425 Homophily influences ranking of minorities in social networks. *Scientific reports*, 8(1):1–12,
426 2018.
- 427 [20] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional
428 Networks. In *Proceedings of the 5th International Conference on Learning Representations*,
429 2017.
- 430 [21] Yajing Kong, Liu Liu, Jun Wang, and Dacheng Tao. Adaptive curriculum learning. In *Pro-*
431 *ceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5067–5076,
432 2021.
- 433 [22] M Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models.
434 *Advances in neural information processing systems*, 23, 2010.
- 435 [23] Haoyang Li, Xin Wang, and Wenwu Zhu. Curriculum graph machine learning: A survey. *arXiv*
436 *preprint arXiv:2302.02926*, 2023.
- 437 [24] Qiuwei Li, Zhihui Zhu, and Gongguo Tang. Alternating minimizations converge to second-order
438 optimal solutions. In *International Conference on Machine Learning*, pages 3935–3943. PMLR,
439 2019.
- 440 [25] Xiaohe Li, Lijie Wen, Yawen Deng, Fuli Feng, Xuming Hu, Lei Wang, and Zide Fan. Graph
441 neural network with curriculum learning for imbalanced node classification. *arXiv preprint*
442 *arXiv:2202.02529*, 2022.
- 443 [26] Dongsheng Luo, Wei Cheng, Wenchao Yu, Bo Zong, Jingchao Ni, Haifeng Chen, and Xiang
444 Zhang. Learning to drop: Robust graph neural network via topological denoising. In *Proceedings*
445 *of the 14th ACM international conference on web search and data mining*, pages 779–787,
446 2021.
- 447 [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,
448 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative
449 style, high-performance deep learning library. *Advances in neural information processing*
450 *systems*, 32, 2019.
- 451 [28] Douglas LT Rohde and David C Plaut. Language acquisition in the absence of explicit negative
452 evidence: How important is starting small? *Cognition*, 72(1):67–109, 1999.
- 453 [29] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann.
454 Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.
- 455 [30] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors
456 with online hard example mining. In *Proceedings of the IEEE conference on computer vision*
457 *and pattern recognition*, pages 761–769, 2016.
- 458 [31] Charles Sutton, Andrew McCallum, et al. An introduction to conditional random fields.
459 *Foundations and Trends® in Machine Learning*, 4(4):267–373, 2012.
- 460 [32] S Vichy N Vishwanathan, Nicol N Schraudolph, Risi Kondor, and Karsten M Borgwardt. Graph
461 kernels. *Journal of Machine Learning Research*, 11:1201–1242, 2010.
- 462 [33] Junxiang Wang, Hongyi Li, and Liang Zhao. Accelerated gradient-free neural network training
463 by multi-convex alternating optimization. *Neurocomputing*, 487:130–143, 2022.
- 464 [34] Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *IEEE Transactions*
465 *on Pattern Analysis and Machine Intelligence*, 44(9):4555–4576, 2021.
- 466 [35] Yiwei Wang, Wei Wang, Yuxuan Liang, Yujun Cai, and Bryan Hooi. Curgraph: Curriculum
467 learning for graph classification. In *Proceedings of the Web Conference 2021*, pages 1238–1248,
468 2021.

- 469 [36] Xiaowen Wei, Weiwei Liu, Yibing Zhan, Du Bo, and Wenbin Hu. Clnode: Curriculum learning
470 for node classification. *arXiv preprint arXiv:2206.07258*, 2022.
- 471 [37] Daphna Weinshall, Gad Cohen, and Dan Amir. Curriculum learning by transfer learning: Theory
472 and experiments with deep networks. In *International Conference on Machine Learning*, pages
473 5238–5246. PMLR, 2018.
- 474 [38] Richard Lee Wheeden and Antoni Zygmund. *Measure and integral*, volume 26. Dekker New
475 York, 1977.
- 476 [39] Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, and Liming Zhu. Ad-
477 versarial examples for graph data: deep insights into attack and defense. In *Proceedings of the*
478 *28th International Joint Conference on Artificial Intelligence*, pages 4816–4823, 2019.
- 479 [40] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A
480 comprehensive survey on graph neural networks. *IEEE transactions on neural networks and*
481 *learning systems*, 32(1):4–24, 2020.
- 482 [41] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
483 networks? In *International Conference on Learning Representations*, 2018.
- 484 [42] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning
485 with graph embeddings. In *International conference on machine learning*, pages 40–48. PMLR,
486 2016.
- 487 [43] Cheng Zheng, Bo Zong, Wei Cheng, Dongjin Song, Jingchao Ni, Wenchao Yu, Haifeng Chen,
488 and Wei Wang. Robust graph representation learning via neural sparsification. In *International*
489 *Conference on Machine Learning*, pages 11458–11468. PMLR, 2020.
- 490 [44] Tianyi Zhou, Shengjie Wang, and Jeff Bilmes. Robust curriculum learning: from clean label de-
491 tection to noisy label self-correction. In *International Conference on Learning Representations*,
492 2020.
- 493 [45] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. Adversarial attacks on neural
494 networks for graph data. In *Proceedings of the 24th ACM SIGKDD international conference on*
495 *knowledge discovery & data mining*, pages 2847–2856, 2018.

496 **A Additional experimental settings and results**

497 **A.1 Additional experimental settings**

498 **Synthetic datasets** To evaluate the effectiveness of our proposed method on datasets with ground-
499 truth difficulty labels on structure, we first follow previous studies [19, 1] to generate a set of synthetic
500 datasets, where the difficulty of edges in generated graphs are indicated by their formation probability.
501 Specifically, as shown in Figure 4, each generated graph is with 5,000 nodes, which are divided into
502 10 equally sized node classes 1, 2, . . . , 10. The node features are sampled from overlapping multi-
503 Gaussian distributions. Each generated graph is associated with a *homophily coefficient* (*homo*) which
504 indicates the likelihood of a node forming a connection to another node with the same label (same
505 color in Figure 4). For example, a generated graph with $homo = 0.5$ will have on average half of the
506 edges formed between nodes with the same label. For the rest edges that are formed between nodes
507 with different labels (different colors in Figure 4), the probability of forming an edge is inversely
508 proportional to the distances between their labels. Mathematically, the probability of forming an
509 edge between node u and node v follows $p_{u \rightarrow v} \propto e^{-|c_u - c_v|}$, where the distances between labels
510 $|c_u - c_v|$ means shortest distance of two classes on a circle. Therefore, the probability of forming
511 an edge in the synthetic graph can reflect how well this edge is expected. Specifically, edges with
512 a higher formation probability, e.g. connecting nodes with the same label or close labels, meaning
513 that there is a higher chance that this connection will positively contribute to the prediction (less
514 chance to be a noisy edge). Conversely, edges with a lower formation probability, e.g., connecting
515 nodes with faraway labels, mean that there is a higher chance that this connection will negatively
516 contribute to the prediction (higher chance to be a noisy edge). We vary the value of *homo* from
517 0.1, 0.2, . . . , 0.9 to generate nine graphs in total. Similar to previous works [19, 1], we randomly
518 partition each synthetic graph into equal-sized train, validation, and test node splits.

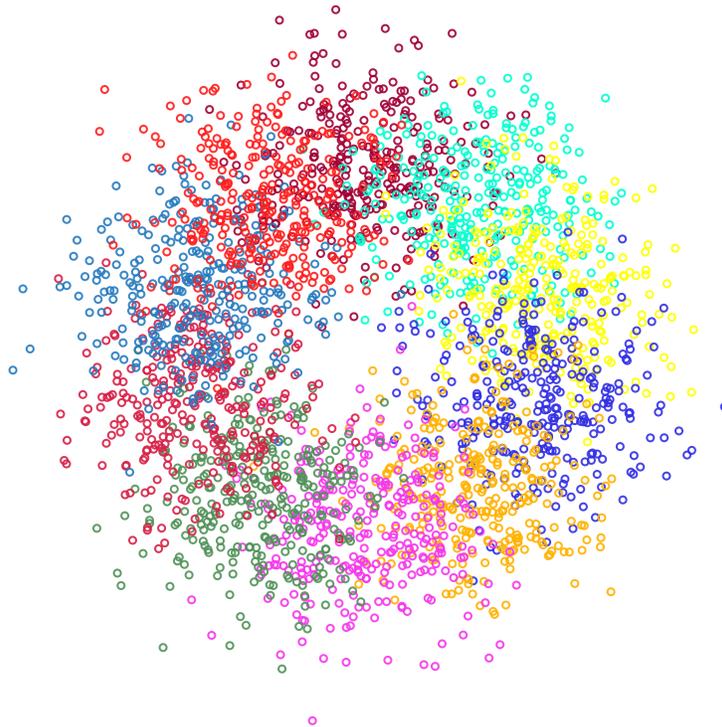


Figure 4: Visualization of synthetic datasets. Each color represents a class of nodes. Node attributes are sampled from overlapping multi-Gaussian distributions, where the attributes of nodes with close labels are likely to have short distances. Homogeneous edges represent edges that connect nodes of the same class (with the same color). The probability of connecting two nodes of different classes decreases with the distance between the center points of their class distribution. Therefore, the formation probability of a node denotes the edge difficulty, since edges between nodes with close classes are more likely to positively contribute to the prediction under the homogeneous assumption.

519 **Initializing graph structure by a pre-trained model.** It is worth noting that the model needs
520 an initial training graph structure $\mathbf{A}^{(0)}$ in the initial stage of training. An intuitive way is that we
521 can initialize the model to work in a purely data-driven scenario that starts only with isolated nodes
522 where no edges exist. However, an instructive initial structure can greatly reduce the search cost and
523 computational burden. Inspired by many previous CL works [37, 11, 17, 44] that incorporate prior
524 knowledge of a pre-trained model into designing curriculum for the current model, we initialize the
525 training structure $\mathbf{A}^{(0)}$ by a pre-trained vanilla GNN model f^* . Specifically, we follow the same
526 steps from line 4 to line 7 in the algorithm 1 to obtain the initial training structure $\mathbf{A}^{(0)}$ but the latent
527 node embedding is extracted from the pre-trained model f^* .

528 **Implementation Details** We use the baseline model (GCN, GIN, GraphSage) as the backbone
529 model for both our RCL method and all comparison methods. For a fair comparison, we require all
530 models follow the same GNN architecture with two convolution layers. For each split, we run each
531 model 10 times to reduce the variance in particular data splits. Test results are according to the best
532 validation results. General training hyperparameters (such as learning rate or the number of training
533 epochs) are equal for all models. For the pre-trained model to initialize the training structure, we
534 utilize the same model as the backbone model utilized by our method. For example, if we use GCN
535 as the backbone model for RCL, the pre-trained model to initialize is also GCN. All experiments are
536 conducted on a 64-bit machine with four NVIDIA Quadro RTX 8000 GPUs. The proposed method is
537 implemented with Pytorch deep learning framework [27].

538 The following describes the details of our comparison models.

539 **Graph Neural Networks (GNNs).** We first introduce three baseline GNN models as follows.

540 **(i) GCN.** Graph Convolutional Networks (GCN) [20] is a commonly used GNN, which introduces a
541 first-order approximation architecture of the Chebyshev spectral convolution operator;

542 **(ii) GIN.** Graph Isomorphism Networks (GIN) [41] is a variant of GNN, which has provably powerful
543 discriminating power among the class of 1-order GNNs;

544 **(iii) GraphSage.** GraphSage [12] is a GNN method that computes the hidden representation of the
545 root node by aggregating the hidden node representations hierarchically from bottom to top.

546 **Graph structure learning.** We then introduce four state-of-the-art methods for jointly learning the
547 optimal graph structure and downstream tasks.

548 **(i) GNNSVD.** GNNSVD [8] first apply singular value decomposition (SVD) on the graph adjacency
549 matrix to obtain a low-rank graph structure and apply GNN on the obtained low-rank structure;

550 **(ii) ProGNN.** ProGNN [18] is a method to defend against graph adversarial attacks by obtaining a
551 sparse and low-rank graph structure from the input structure;

552 **(iii) NeuralSparse.** NeuralSparse [43] is a method to learn robust graph representations by iteratively
553 sampling k -neighbor subgraphs for each node and sparsifying the graph according to the performance
554 on the node classification;

555 **(iv) PTDNet.** PTDNet [26] learns a sparsified graph by pruning task-irrelevant edges, where sparsity
556 is controlled by regulating the number of edges.

557 **Curriculum learning on graph data.** We introduce a recent curriculum learning work on node
558 classification as follows.

559 **(i) CLNode.** CLNode [36] regards nodes as data samples and gradually incorporates more nodes into
560 training according to their difficulty. They apply a heuristic-based strategy to measure the difficulty of
561 nodes, where the nodes that connect neighboring nodes with different classes are considered difficult.

562 **Searching space for hyperparameters.**

563 Number of epochs trained: $\{150, 500\}$;

564 Learning rate for model: $\{1e-2, 5e-3, 1e-3\}$;

565 Number of GNN layers: $\{2\}$;

566 Dimension of hidden state: $\{64\}$;

567 Age parameter $\lambda : \{1, 2, 3, 4, 5\}$ (A larger value indicates faster pacing for adding edges, where 1
568 denotes the training structure will converge to the input structure at the final iteration).

	Synthetic	Citeseer	Computers	ogbn-arxiv	ogbn-proteins
Vanilla	7.32s	3.90s	16.88s	55.22s	1438.23s
GNNsVD	11.49s	3.82s	35.96s	135.72s	2632.42s
CLNode	6.29s	3.96s	17.02s	58.53s	1545.53s
ProGNN	220.25s	72.42s	1953.23s	(-)	(-)
NeuralSparse	310.02s	88.91s	6553.34s	(-)	(-)
PTDNet	153.43s	48.42s	2942.02s	(-)	(-)
Ours	4.07s	2.42s	14.62s	71.49s	2239.05s

Table 3: Running time of our method and comparison methods. Here (-) denotes an out-of-memory error and Vanilla denotes the standard GNN model.

	Synthetic1	Synthetic2	Citeseer	CS	Computers
Full	73.98±0.55	97.42±0.17	79.79±0.55	94.66±0.22	90.23±0.23
Curriculum-linear	70.93±0.54	95.19±0.19	79.04±0.38	94.14±0.26	89.28±0.21
Curriculum-root	70.13±0.72	95.50±0.18	78.27±0.54	94.47±0.34	89.27±0.15
Random-linear	58.76±0.46	89.78±0.11	77.43±0.49	92.76±0.14	88.76±0.18
Random-root	61.04±0.20	91.04±0.09	76.81±0.35	92.92±0.15	88.81±0.28
w/o edge appearance	70.70±0.43	95.77±0.16	77.77±0.65	94.39±0.21	89.56±0.30
w/o node confidence	72.38±0.41	96.86±0.17	78.72±0.72	94.34±0.13	90.03±0.62
w/o pre-trained model	72.56±0.69	93.89±0.14	78.28±0.77	94.50±0.14	89.80±0.55

Table 4: Ablation study. Here “Full” represents the original method without removing any component. The best-performing method on each dataset is highlighted in bold.

569 A.2 Additional experiments

570 **Time complexity analysis** Here we consider GCN as the backbone. First, the time complexity of
571 an L -layer GCN is $O(L|\mathcal{E}|b + L|\mathcal{V}|b^2)$, where b is the number of node attributes. Second, the time
572 complexity of measuring the difficulty levels of edges by reconstruction is $O(|\mathcal{E}|d)$ where d is the
573 number of latent embedding dimensions. Third, the time complexity of selecting the edges to add is
574 $O(|\mathcal{E}|)$. Therefore, the total time complexity of our algorithm is $O(|\mathcal{E}|(Lb + d) + L|\mathcal{V}|b^2)$.

575
576 In addition, we compare the total running time of our method and all comparison methods in the
577 Table 3. We can observe that the running time of our proposed method is comparable to that of
578 standard GNN models in all datasets. Notably, our method is even faster than standard GNN models
579 in some datasets. One possible reason is that at the beginning of training, the graphs in our model
580 have much fewer edges than those in standard GNN models. Therefore, the computational cost of the
581 GNN model is also reduced.

582 **Ablation study** To investigate the effectiveness of our proposed model with some simpler heuristics,
583 we deploy a series of synthetic analysis. We first train the model with node classification task purely
584 and select the top K expected edges as suggested by the reviewer. Specifically, we follow previous
585 works [34, 36] using two classical selection pacing functions as follows:

$$\text{Linear: } K_{\text{linear}}(t) = \frac{t}{T}|\mathcal{E}|;$$

$$\text{Root: } K_{\text{root}}(t) = \sqrt{\frac{t}{T}}|\mathcal{E}|,$$

586 where t is the number of current iterations and T is the number of total iterations, and $|\mathcal{E}|$ is
587 the number of total edges. We name these two variants Curriculum-linear and Curriculum-root,
588 respectively. In addition, we also remove the edge difficulty measurement module and use random
589 selection instead. Specifically, we gradually incorporate more edges into training in random order
590 to verify the effectiveness of the learned curriculum. We name two variants as Random-linear and
591 Random-root with the above two mentioned pacing functions, respectively.

592 In order to further investigate the impact of the proposed components of RCL. We also first consider
593 variants of removing the edge smoothing components mentioned in Section 4.3. Specifically, we
594 consider two variants *w/o EC* and *w/o NC*, which remove the smoothing function of the edge
595 occurrence ratio and the component to reflect the degree of confidence for the latent node embedding
596 in RCL, respectively. In addition to examining the effectiveness of edge smoothing components, we
597 further consider a variant *w/o pre-trained model* that avoids using a pre-trained model to initialize

598 model, which is mentioned in Appendix A.1, to initialize the training structure by a pre-trained model
 599 and instead starts with inferred structure from isolated nodes with no connections.

600 We present the results of two synthetic datasets (*homophily coefficient*= 0.3, 0.6) and three real-world
 601 datasets in Table 4. We summarize our findings from the above table as below: (i) Our full model
 602 consistently outperforms the two variants Curriculum-linear and Curriculum-root by an average of
 603 1.59% on all datasets, suggesting that our pacing module can benefit model training. It is worth
 604 noting that these two variants also outperform the baseline vanilla GNN model Vanilla by an average
 605 of 1.92%, which supports the assumption that even a simple curriculum learning strategy can still
 606 improve model performance. (ii) We observe that the performance of the two variants Random-linear
 607 and Random-root on all datasets drops by 3.86% on average compared to the variants Curriculum-
 608 linear and Curriculum-root. Such behavior demonstrates the effectiveness of our proposed edge
 609 difficulty quantification module by showing that randomly involving edges into training cannot benefit
 610 model performance. (iii) We can observe a significant performance drop consistently for all variants
 611 that remove the structural smoothing techniques and initialization components. The results validate
 612 that all structural smoothing and initialization components can benefit the performance of RCL on
 613 the downstream tasks.

614 **Parameter sensitivity analysis** Recall that RCL learns a curriculum to gradually add edges in
 615 a given input graph structure to the training process until all edges are included. An interesting
 616 question is how the speed of adding edges will affect the performance of the model. Here we conduct
 617 experiments to explore the impact of age parameter λ which controls the speed of adding edges to
 618 the model performance. Here a larger value of λ means that the training structure will converge to
 619 the input structure earlier. For example, $\lambda = 1$ means that the training structure will probably not
 620 converge to the input structure until the last iteration, and $\lambda = 5$ means that the training structure will
 621 converge to the input structure around half of the iterations are complete, and then the model will
 622 be trained with the full input structure for the remaining iterations. We present the results on two
 623 synthetic datasets (*homophily coefficient*= 0.3, 0.6) and two real-world datasets in Figure 5. As can
 624 be seen from the figure, the classification results are steady that the average standard deviation is only
 625 0.41%. It is also worth noting that the peak values for all datasets consistently appear around $\lambda = 3$,
 626 which indicates that the best performance is when the training structure converges to the full input
 structure around two-thirds of the iterations are completed.

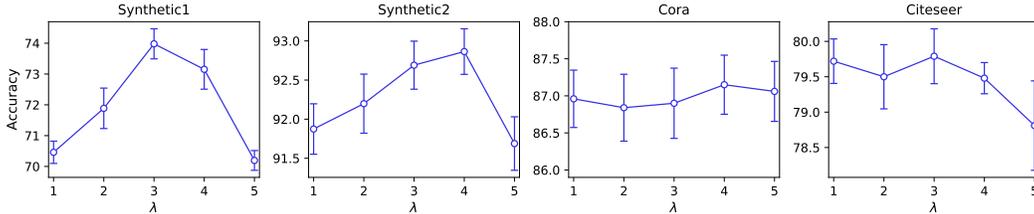


Figure 5: Parameter sensitivity analysis on four datasets. Here a larger value of λ means the training structure will converge to the original structure at an earlier training stage.

627

628 B Mathematical proof

629 **Theorem 1.** We have the following convergence guarantees for Algorithm 1:

- 630 • **Avoidance of Saddle Points** If the second derivatives of $L(f(\mathbf{X}, \mathbf{A}^{(t)}; \mathbf{w}), \mathbf{y})$ and $g(\mathbf{S}; \lambda)$ are
 631 continuous, then for sufficiently large γ , any bounded sequence $(\mathbf{w}^{(t)}, \mathbf{S}^{(t)})$ generated by Algorithm
 632 1 with random initializations will not converge to a strict saddle point of F almost surely.
- 633 • **Second Order Convergence** If the second derivatives of $L(f(\mathbf{X}, \mathbf{A}^{(t)}; \mathbf{w}), \mathbf{y})$ and $g(\mathbf{S}; \lambda)$ are
 634 continuous, and $L(f(\mathbf{X}, \mathbf{A}^{(t)}; \mathbf{w}), \mathbf{y})$ and $g(\mathbf{S}; \lambda)$ satisfy the Kurdyka-Łojasiewicz (KL) property
 635 [33], then for sufficiently large γ , any bounded sequence $(\mathbf{w}^{(t)}, \mathbf{S}^{(t)})$ generated by Algorithm 1 with
 636 random initialization will almost surely converges to a second-order stationary point of F .

637 *Proof.* We prove this theorem by Theorem 10 and Corollary 3 from [24].

638 **[Avoidance of Saddle Points]** Because the sequence $(\mathbf{w}^{(t)}, \mathbf{S}^{(t)})$ is bounded, and the second
 639 derivatives of L and g are continuous, then they are bounded. In other words, we have

640 $\max\{\|\nabla_{\mathbf{w}}^2 L(f(\mathbf{X}, \mathbf{A}^{(t)}; \mathbf{w}^{(t)}), \mathbf{y})\|, \|\nabla_{\mathbf{S}}^2 g(S^{(t)}; \lambda)\|\} \leq p$, where $p > 0$ is a constant. Simi-
641 larly, it is easy to check that the second derivative of the term $\sum_{i,j} \mathbf{S}_{ij} \|\tilde{\mathbf{A}}_{ij}^{(t)} - \mathbf{A}_{ij}\|_2^2$ is bounded,
642 i.e., $\max\{\|\nabla_{\mathbf{w}}^2 \sum_{i,j} \mathbf{S}_{ij} \|\tilde{\mathbf{A}}_{ij}^{(t)} - \mathbf{A}_{ij}\|_2^2\|, \|\nabla_{\mathbf{S}}^2 \sum_{i,j} \mathbf{S}_{ij} \|\tilde{\mathbf{A}}_{ij}^{(t)} - \mathbf{A}_{ij}\|_2^2\|\} \leq q$, where $q > 0$ is
643 constant and $\tilde{\mathbf{A}}$ is a function of \mathbf{w} . Therefore, it means that the objective F is bi-smooth, i.e.
644 $\max\{\|\nabla_{\mathbf{w}}^2 F\|, \|\nabla_{\mathbf{S}}^2 F\|\} \leq p + q$. In other words, F satisfies Assumption 4 from [24]. Moreover,
645 the second derivative of F is continuous. For any $\gamma > p + q$, any bounded sequence $(\mathbf{w}^{(t)}, \mathbf{S}^{(t)})$
646 generated by Algorithm 1 will not converge to a strict saddle of F almost surely by Theorem 10 from
647 [24].
648 **[Second Order Convergence]** From the above proof of avoidance of saddle points, we know that F
649 satisfies Assumption 4 from [24]. Moreover, because L and g satisfy the KL property, and the term
650 $\sum_{i,j} \mathbf{S}_{ij} \|\tilde{\mathbf{A}}_{ij}^{(t)} - \mathbf{A}_{ij}\|_2^2$ satisfies the KL property, we conclude that F satisfy the KL property as
651 well. From the proof above, we also know that the second derivative of F is continuous. Because
652 continuous differentiability implies Lipschitz continuity [38], it infers that the first derivative of
653 F is Lipschitz continuous. As a result, F satisfies Assumption 1 from [24]. Because F satisfies
654 Assumptions 1 and 4, then for any $\gamma > p + q$, any bounded sequence $(\mathbf{w}^{(t)}, \mathbf{S}^{(t)})$ generated by
655 Algorithm 1 will almost surely converges to a second-order stationary point of F by Corollary 3 from
656 [24]. \square

657 While the convergence of Algorithm 1 entails the second-order optimality conditions of f and g ,
658 some commonly used f such as the GNN with sigmoid or tanh activations and some commonly used
659 g such as the squared ℓ_2 norm satisfy the KL property, and Algorithm 1 is guaranteed to avoid a strict
660 saddle point and converges to a second-order stationary point.
661