# **Cluster Aware Graph Anomaly Detection**

Anonymous Author(s)

#### Abstract

1

2

3

4

5

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

56

Graph anomaly detection has gained significant attention across various domains, particularly in critical applications like fraud detection in e-commerce platforms and insider threat detection in cybersecurity. Usually, these data are composed of multiple types (e.g., user information and transaction records for financial data), thus exhibiting view heterogeneity. However, in the era of big data, the heterogeneity of views and the lack of label information pose substantial challenges to traditional approaches. Existing unsupervised graph anomaly detection methods often struggle with highdimensionality issues, rely on strong assumptions about graph structures or fail to handle complex multi-view graphs. To address these challenges, we propose a cluster aware multi-view graph anomaly detection method, called CARE. Our approach captures both local and global node affinities by augmenting the graph's adjacency matrix with the pseudo-label (i.e., soft membership assignments) without any strong assumption about the graph. To mitigate potential biases from the pseudo-label, we introduce a similarity-guided loss. Theoretically, we show that the proposed similarity-guided loss is a variant of contrastive learning loss, and we present how this loss alleviates the bias introduced by pseudo-label with the connection to graph spectral clustering. Experimental results on several datasets demonstrate the effectiveness and efficiency of our proposed framework. Specifically, CARE outperforms the second-best competitors by more than 39% on the Amazon dataset with respect to AUPRC and 18.7% on the YelpChi dataset with respect to AUROC. The code of our method is available at the anonymous GitHub link: https://anonymous.4open.science/r/CARE-demo-1C7F.

#### **1** Introduction

Graph-based anomaly detection has been an important research area across diverse domains for decades, particularly within highimpact applications, such as fraud detection within e-commerce platforms [50, 53, 61] and insider threat detection in the cybersecurity domain [9, 19, 28]. For instance, in the realm of e-commerce, leveraging a graph-based anomaly detection algorithm proves invaluable for identifying fraudulent sellers by analyzing the properties (i.e., attributes) and connections (i.e., structure) among users [30]. Similarly, in the context of insider threat detection, constructing a graph based on users' activities allows investigators to discern anomalous users in the organization by exploring the substructure of the graph [20].

Conference acronym 'XX, June 03-05, 2018, Woodstock, NY 55

57 https://doi.org/XXXXXXXXXXXXXXX

58

In the era of big data, the collected data often exhibit heterogeneous views (e.g., various data) and lack labeled data. For example, in the e-commerce platform, multiple types of data can be collected for heterogeneous graph construction, including the user's shopping history, search trends, and product ratings [29]; in credit card fraud detection, the data are composed of both cardholder information and transaction records (e.g., online purchase records) [73]. However, obtaining labels is often impractical due to the expense associated with labeling services and the demand for domain-specific expertise in discerning malicious patterns [70]. This highlights the need for innovative approaches that can deal with the intricacies of heterogeneous and unlabeled datasets.

Until then, many unsupervised anomaly detection methods [11, 23, 55] have been proposed and they can be categorized into two branches, including feature reconstruction methods and self-supervised learning methods. The feature reconstruction approaches focus on minimizing the reconstruction error of node attributes or structures [2, 13, 14]. However, these feature reconstruction based methods tend to suffer from the curse of high dimensionality [17], especially for citation network where the word occurrence is extracted as the node attributes. Another direction is the self-supervised learning methods [2, 27, 38, 49, 66, 72], which aim to design a proxy task related to anomaly detection, whereas these methods tend to have strong assumptions regarding the graph structure, thus only performing well on some graphs with certain structures. For instance, TAM [49] holds the one-class homophily assumption that normal nodes tend to have much stronger affinity with each other than with the abnormal nodes and the authors propose to maximize the local node similarity. However, TAM fails to consider a situation where the normal nodes and their normal neighbors might come from different classes, indicating the distinct features for these connected normal nodes. A general limitation for both branches is that many of these methods [14, 26, 72] are designed for single-view graphs, suffering from the presence of multiple views.

To address these limitations, we propose a unified Cluster AwaRE graph anomaly detection method to identify anomalous nodes in multi-view graphs, named CARE. In this work, we propose to capture both local and global node affinity and design an anomaly score function to assigning higher anomaly scores to nodes that are less similar to their neighbors based on both node attributes and structural similarity. Since the raw adjacency matrix in the graph only contains the local node affinity information, we measure the global node affinity by leveraging the pseudo-label (i.e., soft membership assignments) to augment the original adjacency matrix without any strong assumption about the graph. To reduce the potential bias introduced by the pseudo-label during the optimization, we propose a similarity-guided loss that utilizes the soft assignment to build a similarity map to help the model learn robust representations. Theoretically, we analyze that the proposed similarity-guided loss is a variant of the contrastive loss, and we present how the proposed regularization mitigates the potential bias by connecting it with

1

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97

98

99

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

<sup>© 2018</sup> Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-XXXX-X/18/06

graph spectral clustering. Our main contributions are summarized below:

- A novel self-supervised framework for detecting anomalies in multi-view graphs.
- A novel similarity-guided contrastive loss for learning graph contextual information and its theoretical analysis showing the connection to graph spectral clustering.
- Theoretical analysis showing the negative impact of other types of contrastive loss.
- Experimental results on six datasets demonstrating the effectiveness and efficiency of the proposed framework.

The rest of this paper is organized as follows. After briefly reviewing the related work in Section 2, we then introduce a multi-view graph anomaly detection framework in Section 3. Next, we conduct the systematic evaluation of the proposed framework on several datasets in Section 4 before we conclude the paper in Section 5.

#### 2 Related Work

117

118

119

120

121

123

124

125

126

127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145

146

147

148

149

150

151

153

154

155

156

157

158

159

160

174

In this section, we briefly review the related work on clustering and anomaly detection.

# 2.1 Clustering

In the past decades, clustering methods gradually evolved from traditional shallow methods, (e.g., Non-Negative Matrix Factorization (NMF) methods [4, 52, 68] and spectral clustering methods [5, 37, 43]), to deep learning-based methods (e.g., deep NMF [56, 69], graph neural network clustering methods [8, 39, 60], autoencoderbased clustering methods [15, 24]). For example, the authors of [52] propose an NMF-based clustering method that models the intrinsic geometrical structure of the data by assuming that several neighboring points should be close in the low dimensional subspace. The work in [4] presents a semi-NMF clustering method by taking advantage of the mutual reinforcement between data reduction and clustering tasks. Additionally, [69] extends shallow NMF to a deep Semi-NMF for multi-view clustering by learning the hierarchical structure of multi-view data and maximizing the mutual information of each pair of views. Different from these methods, this paper first follows the idea of the graph pooling method [67] to get the soft assignment, aiming to capture the global node affinity information and regularize the learned representations by the similarity-guided contrastive loss.

#### 2.2 Anomaly Detection

161 Anomaly detection has been studied for decades [22, 35, 42, 48]. The increasing demand in many domains, such as financial fraud detec-162 163 tion, anomaly detection in cybersecurity, etc., has attracted many 164 researchers' attention, and a variety of outstanding algorithms have 165 been proposed, ranging from shallow algorithms [1, 10, 40, 65] to deep models [3, 12, 34, 45, 46, 74]. To name a few, [65] presents an 166 algorithm named Outlier Pursuit, which projects the raw data to the 167 low-dimensional subspace and identifies the corrupted points with 168 PCA; [12] encodes observed co-occurrence in different events into 169 a hidden space and utilizes the weighted pairwise interactions of 170 different entity types to define the event probability. AnoGAN [3] 171 172 is a generative adversarial network (GAN) based anomaly detec-173 tion method, which utilizes the generator and the discriminator to capture the normal patterns while the anomalies are detected based on the residual score and the discrimination score. HCM-A [27] designs a self-supervised learning loss by forcing the prediction of the shortest path length between pairs of nodes. ComGA [41] proposes a community-aware attributed graph anomaly detection method to detect community structure of the graph. CONAD [66] integrates human knowledge of different anomaly types via data augmentation and introduces a contrastive learning-based method for graph anomaly detection. TAM [49] proposes a scoring measure by assigning large score to the nodes that are less affiliated with their neighbors and introduce truncated affinity maximization to reduce the bias during the optimization. In contrast, this paper proposes to capture both local and global node affinity information, and we propose similarity-guided contrastive learning loss to learn robust representations and to mitigate potential bias.

# 3 Proposed CARE Framework

In this section, we present our proposed framework, CARE, for multi-view graph anomaly detection. We begin by defining the notation and then introduce cluster-aware node affinity learning alongside the similarity-guided contrastive learning loss. Next, we outline the overall objective function and the inference process for detecting anomalies. Finally, we analyze the limitations of using weakly supervised contrastive loss in our method.

# 3.1 Notation

Throughout this paper, we use regular letters to denote scalars (e.g.,  $\alpha$ ), boldface lowercase letters to denote vectors (e.g., x), and boldface uppercase letters for matrices (e.g., A). Given an undirected graph  $\mathcal{G} = (V, E^1, ..., E^v, X^1, ..., X^v)$ , our objective is to identify anomalous nodes in the graph, where v represents the number of views, V consists of n vertices,  $E^v$  consists of  $m^v$  edges,  $X^v \in \mathbb{R}^{n \times d_v}$ denotes the feature matrix of the *v*-th view and  $d_v$  is the feature dimension. For clarity, we denote  $u_i$  as node  $i, \mathbf{x}_i^v \in \mathbb{R}^{d_v}$  as the node attributes of  $u_i$  for the *v*-th view,  $h_i \in \mathbb{R}^d$  as the embedding of node  $u_i$  by any type of GNNs and d is the feature dimensionality of the hidden representation.  $H^v \in \mathbb{R}^{n \times d^v}$  is the node embedding matrix. We let  $A^{\hat{v}} \in \mathbb{R}^{n \times n}$  denote the adjacency matrix of the *v*-th view where  $A_{ij}^v = 1$  iff node  $u_i$  and node  $u_j$  are connected,  $D^v \in \mathbb{R}^{n \times n}$ denotes the diagonal matrix of vertex degrees for the v-th view, and  $I \in \mathbb{R}^{n \times n}$  denotes the identity matrix. The symbols are summarized in Table 1.

# 3.2 Cluster-Aware Node Affinity Learning

Many self-supervised learning methods [2, 27, 38, 49, 66, 72] aim to design a proxy task relevant to the anomaly detection. One branch [16, 49] is to maximize the local node similarity for a multiview graph as follows:

$$\mathcal{L}_1 = \sum_{a=1}^{v} \sum_{u_i} \frac{1}{|\mathcal{N}^a(i)|} \sum_{u_j \in \mathcal{N}^a(i)} \operatorname{sim}(\boldsymbol{h}_i^a, \boldsymbol{h}_j^a)$$

$$=\sum_{a=1}\sum_{u_i}\sum_{u_j}\frac{1}{D_i^a}A_{ij}^a\cdot\sin(\boldsymbol{h}_i^a,\boldsymbol{h}_j^a) \tag{1}$$

175

176

177

178

179

180

181

182

183

184

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231



Figure 1: The overview of CARE. It first extracts the global node affinity based on the soft assignment by graph clustering method, and then combines the global node affinity and local node affinity together. Similarity-guided graph contrastive loss is then introduced to mitigate the potential bias.

Table 1	D	efinition	of	S٦	vmbols
---------	---	-----------	----	----	--------

Symbols	Definition
V	The set of vertices
$E^{a}$	The set of edges for the <i>a</i> -th view
$A^a$	The <i>a</i> -th view adjacency matrix
$X^a$	The <i>a</i> -th view node attribute matrix
v(n)	The number of views (nodes)
$D^a$	The degree matrix for the <i>a</i> -th view
$H^{a}$	The the <i>a</i> -th view node representations
Â	The adjacency matrix augmented by cluster similarity
$M^{a}$	The soft membership matrix for the <i>a</i> -th view
Ĥ	The average node representations
Ā	The average soft membership matrix
Ã	The normalized augmented adjacency matrix

where  $sim(\boldsymbol{h}_{i}^{a}, \boldsymbol{h}_{j}^{a}) = \frac{\boldsymbol{h}_{i}^{a}(\boldsymbol{h}_{j}^{a})^{T}}{|\boldsymbol{h}_{i}^{a}||\boldsymbol{h}_{j}^{a}|}$  measures the similarity of node embedding between the node  $u_{i}$  and its neighbor  $u_{j}$  based on the *a*-th view of the graph,  $\mathcal{N}^{a}(i)$  denotes the neighbors of the node  $u_{i}$  and  $D_{i}^{a}$  is the degree of the node  $u_{i}$  for the *a*-th view. However, these methods typically rely on *one-class homophily assumption*, which posits that normal nodes tend to exhibit strong affinities with each other, while the affinities among abnormal nodes are significantly weaker. This assumption is overly restrictive, as it focuses exclusively on extracting local node affinities while neglecting global node affinities.

A naive solution to relax this constraint is to incorporate high-order information by extending the local 1-hop neighbors to k-hop neighbors. However, this approach presents two main issues. First, including high-order neighbors inevitably introduces more abnor-mal nodes during the node affinity maximization process, resulting in a sub-optimal solution. Second, it overlooks a crucial scenario where normal anchor nodes and their neighbors may belong to different classes, indicating that these connected normal nodes may have distinct features. Since both normal neighbors from different classes and abnormal neighbors possess features that differ from those of normal anchor nodes, incorporating these distinct features in node affinity learning decreases the likelihood of detecting ab-normal nodes. Furthermore, adding high-order neighbors in node affinity learning also increases the probability of including neigh-bors from different classes, exacerbating the problem. To address this issue, we propose incorporating label information into the local

node affinity maximization as follows:

$$\mathcal{L}_{2} = \sum_{a=1}^{v} \sum_{u_{i}} \sum_{u_{j}} \sum_{u_{j}} \frac{1}{\sum_{u_{j}} (A_{ij}^{a} + S_{ij}^{a})} (A_{ij}^{a} + S_{ij}^{a}) \cdot \sin(h_{i}^{a}, h_{j}^{a})$$
(2)

where  $S_{ij} = 1$  if node  $u_i$  and node  $u_j$  belong to the same class and  $S_{ij} = 0$  otherwise. Compared to  $\mathcal{L}_1$ ,  $\mathcal{L}_2$  further encodes the label information in the node affinity learning. However, label information is often unavailable due to the high costs of labeling services and the rapid growth of new data. To address this issue, we propose replacing the unavailable label information with pseudo-labels derived from a graph clustering method. Following the idea of differential graph pooling [67], we employ a one-layer Graph Convolutional Network (GCN) [32] with a softmax activation function to model soft membership assignments as follows:

$$M^{a} = \text{GCN}^{a}(A^{a}, X^{a}, W^{a})$$
$$\bar{M} = \frac{1}{v} \sum_{a=1}^{v} M^{a}$$
(3)

where  $W^a \in \mathbb{R}^{d_a \times c}$  is the weight matrix of GCN for the *a*-th view and *c* is the number of clusters and we aggregate soft membership assignments from all views to obtain  $\overline{M}$ . We then augment the adjacency matrix by incorporating the clustering results as follows:

$$\hat{A} = (1 - \alpha)\bar{A} + \alpha\bar{M}\bar{M}^{T}$$
(4)

where  $\bar{A} = \frac{1}{v} \sum_{a=1}^{v} A^a$  and  $\alpha \in [0, 1]$  is a hyper-parameter balancing the importance between raw adjacency matrix and the similarity of the soft membership assignments. Our goal is to maximize the cluster-aware node affinity as follows:

$$\mathcal{L}_{A}(u_{i}) = \sum_{u_{j}} \frac{1}{\hat{D}_{i}} \hat{A}_{ij} \cdot \sin(\bar{h}_{i}, \bar{h}_{j}) + \sum_{a=1}^{v} ||h_{i}^{a} - \bar{h}_{i}||_{2}$$
(5)

where  $\hat{D}_i = \sum_{u_j} \hat{A}_{ij}$  measures the degree of node  $u_i$  in the new adjacency matrix and  $\bar{h}_i = \frac{1}{v} \sum_{a=1}^{v} h_i^a$  is the average node representation. The second term enforces consistency in node embeddings across all views. However, optimization using Eq. 5 may be significantly biased by low-quality soft membership assignments at the early stages. To address this issue, we introduce the similarity-guided graph contrastive regularization.

# 3.3 Similarity-guided Graph Contrastive Regularization

To mitigate potential bias introduced by low-quality soft membership assignments, we propose a similarity-guided graph contrastive loss that minimizes the difference between the similarity of the soft assignment and the similarity of node representations for any pair of nodes. This is formulated as follows:

$$\mathcal{L}_2 = \min_{\bar{H}} ||\bar{M}\bar{M}^T - \bar{H}\bar{H}^T||_F^2 \tag{6}$$

where  $\bar{M}$  is the soft assignment matrix computed in Eq. 3.  $\mathcal{L}_2$  aims to learn the hidden representations such that the representations of node  $u_i$  and node  $u_j$  are expected to be close in the latent space if their soft membership assignments are similar. Following the idea of many existing methods restoring graph structure along with the node attributes of the graph [2, 27, 49, 72], we propose to take the graph topological structure into consideration, reformulating the similarity-guided graph contrastive loss as:

$$\mathcal{L}_C = \min_{\tilde{H}} ||\tilde{A} - \bar{H}\bar{H}^T||_F^2 \tag{7}$$

where  $\tilde{A} = \tilde{D}^{-1/2} \hat{A} \tilde{D}^{-1/2}$  is the normalized augmented adjacency matrix, and  $\tilde{D} \in \mathbb{R}^{n \times n}$  is the diagonal matrix with  $\tilde{D}_{ii} = \sum_{u_j} \hat{A}_{ij}$ .

**Theoretical Justification.** We aim to provide a theoretical analysis of how the proposed similarity-guided graph contrastive regularization mitigates potential bias by connecting it with graph spectral clustering. Before delving into the direct analysis of bias mitigation, we first demonstrate that  $\mathcal{L}_C$  functions as a contrastive learning loss.

**Lemma 3.1.** (Similarity-guided Graph Contrastive Loss) Let  $\overline{\mathbf{M}}$  be the output of a one-layer graph neural network defined in Eq. 3. Then, we have

$$\mathcal{L}_C = \mathcal{L}_f + C \tag{8}$$

where  $\mathcal{L}_f = -\sum_{i=1}^n \sum_{j=1}^n \log \frac{\exp(2\tilde{A}_{ij}\tilde{h}_i\tilde{h}_j^T)}{\prod_{k=1}^n \exp((\tilde{h}_i\tilde{h}_k^T)^2)^{1/n}}$  is a graph contrastive loss and C is a constant.

#### See proof in Appendix A.1.

Remark: Compared to traditional contrastive learning losses, the denominator of  $\mathcal{L}_f$  in Lemma 3.1 is the product of the exponential similarity between two node embeddings rather than a summation. The weakly supervised contrastive methods [57, 71] impose a constraint that forces a given pair of nodes to form a positive/negative pair based on their likelihood of being assigned to the same cluster (further discussion can be found in Subsection 3.5). Unlike these "discrete" formulations regarding positive and negative pairs, our approach is in a continuous form. In  $\mathcal{L}_f$ ,  $S_{ij} = \bar{M}_i \bar{M}_i^T$ can be interpreted as the similarity measurement of a node pair  $(u_i, u_j)$  in terms of soft assignment, guiding the similarity of the representations of two nodes in the latent space. If we ignore the influence of the adjacency matrix in  $\tilde{A} = \tilde{D}^{-1/2} \hat{A} \tilde{D}^{-1/2} = \tilde{D}^{-1/2} (\alpha \bar{M} \bar{M}^T + (1-\alpha) \frac{1}{v} \sum_{a=1}^{v} A) \tilde{D}^{-1/2}$  by setting  $\alpha$  to be 1 (i. e.,  $\tilde{A} = \tilde{D}^{-1/2} \bar{M} \bar{M}^T \tilde{D}^{-1/2}$ ), it would be interesting to see that when two nodes are sampled from two distant clusters,  $\tilde{A}_{ij} \approx 0$  and  $\log \frac{\exp(2S_{ij}\bar{h}_i\bar{h}_j^T)}{\prod_{k=1}^n \exp((\bar{h}_i\bar{h}_k^T)^2)^{\frac{1}{n}}} = 0 \text{ for this pair of nodes. Notice that for}$ 

the node pair  $(u_i, u_j)$  with high confidence being assigned to the same cluster, the weight  $S_{ij}$  is larger than the weight  $S_{ik}$  for the uncertain node pair  $(u_i, u_k)$ , where the node  $u_k$  has low confidence to be assigned to the same cluster as  $u_i$ . By reducing the weight for these unreliable positive pairs, the negative impact of uncertain pseudo-labels can be alleviated.

Next, we aim to clarify our proposed contrastive loss by demonstrating the connection between  $\mathcal{L}_C$  and graph spectral clustering.

**Lemma 3.2.** (Graph Contrastive Spectral Clustering) Let  $\overline{\mathbf{M}}$  be the output of a one-layer graph neural network defined in Eq. 3 and  $\overline{\mathbf{h}}_i$  and  $\overline{\mathbf{h}}_j$  be unit vectors. Then, minimizing  $\mathcal{L}_C$  is equivalent to minimizing the following loss function:

$$\min \mathcal{L}_C = \min[2Tr(\bar{H}^T L\bar{H}) + R(\bar{H})]$$
(9)

where  $\mathbf{L} = \mathbf{I} - \tilde{A}$  can be considered as the normalized graph Laplacian,  $\mathbf{I}$  is the identity matrix and  $R(\bar{H}) = ||\bar{H}\bar{H}^T||_F^2$  is the regularization term.

#### See proof in Appendix A.2.

**Remark:** Based on Lemma 3.2,  $\mathcal{L}_C$  can be considered as the graph spectral clustering. Graph spectral clustering [59] aims to find clusters that minimize connections between different clusters while maximizing the connections within each cluster. Traditional graph spectral clustering [6] aims to find the embedding  $\bar{H}$  such that  $Tr(\bar{H}^T L'\bar{H})$  is minimized, where L' is the normalized graph Laplacian. The first term of Eq. 9 is similar to the objective function in traditional graph spectral clustering, but we enhance it by incorporating the similarity measurement  $S_{ij} = \bar{M}_i \bar{M}_i^T$  into the normalized graph Laplacian. By including the similarity measurement in the graph spectral clustering, we reinforce  $\mathcal{L}_C$  to mitigate the bias introduced by the clustering method defined in Eq. 3. It is important to note that in Lemma 3.2, the constraint that  $\bar{h}_i$  and  $\bar{h}_j$  are unit vectors is a common practice in many existing works [25, 62, 75]. This constraint can be easily implemented through normalization, i.e.,  $\bar{h}_i = \frac{\bar{h}_i}{||\bar{h}_i||_2}$ 

## 3.4 Objective Function and Inference

Now, we are ready to introduce the overall objective function:

$$\min J = -\sum_{u_i} \mathcal{L}_A(u_i) + \lambda \mathcal{L}_C \tag{10}$$

where  $\mathcal{L}_A$  is the cluster-aware node affinity loss and  $\mathcal{L}_C$  is the similarity-guided graph contrastive loss.  $\lambda$  is a constant parameter balancing these two terms. During the inference stage, we directly use the cluster-aware node affinity loss  $\mathcal{L}_A$  as the abnormal score:

$$score_i = -\mathcal{L}_A(u_i) \tag{11}$$

# 3.5 Why can't weakly supervised contrastive learning loss be used as a regularization?

One effective remedy to reduce uncertainty and learn high-quality representations in an unsupervised setting is through contrastive learning loss, which has demonstrated significant performance improvements in representation quality [25, 51, 58, 70]. However, an existing study [70] has theoretically proved that simply applying vanilla contrastive learning loss (i. e., InfoNCE [58]) can easily lead to the suboptimal solution. Similarly, according to [62], both normal Table 2: Statistics of the datasets, including the number of nodes, anomalies, and edges for two views.

Name	V	$ E^{1} $	$ E^{2} $	# Anomalies
BlogCatalog	5,196	171,743	-	298
Amazon	10,244	175,608	-	445
YelpChi	24,741	49,315	-	597
CERT	1,000	24,213	22,467	70
IMDB	4,780	1,811,262	419,883	334
DBLP	4,057	299,499	520,440	283

and abnormal nodes are uniformly distributed in the unit hypersphere in the latent space by minimizing the vanilla contrastive learning loss, which leads to worse performance for the anomaly detection task. To address this issue, many weakly supervised contrastive losses [57, 71] are proposed by incorporating the semantic information, such as the clustering results, into a contrastive regularization term as follows:

$$\mathcal{L}_{3} = -\sum_{i=1}^{n} \sum_{j \in C(i), j \neq i} \log \frac{\sin(\bar{\boldsymbol{h}}_{i}, \bar{\boldsymbol{h}}_{j})}{\sin(\bar{\boldsymbol{h}}_{i}, \bar{\boldsymbol{h}}_{j}) + \sum_{k \notin C(i)} \sin(\bar{\boldsymbol{h}}_{i}, \bar{\boldsymbol{h}}_{k})} \quad (12)$$

where  $\bar{\mathbf{h}}_i = \sum_{a=1}^{v} \mathbf{h}_i^a$  is the representation for node  $u_i$  aggregated over all views,  $\sin(\bar{h}_i, \bar{h}_j) = \exp(\bar{h}_i \bar{h}_j^T / \tau)$  and  $\tau$  is the temperature.  $j \in C(i)$  means that node  $u_i$  and node  $u_i$  are assigned into the same cluster or form a positive pair, while  $k \notin C(i)$  means that node  $u_k$  and node  $u_i$  are assigned into two different clusters, resulting in a negative pair. The intuition of the above equation is that if two nodes are from the same cluster, they should be close in the latent space by maximizing their similarity. However, we find out that the construction of the positive and negative pairs in Eq. 12 heavily relies on the quality of the soft assignment, while directly converting the soft assignment to the binary membership inevitably introduces bias/noise during the training phase. This bias will be amplified further in the node affinity learning as we mentioned in Subsection 3.2. (We also validate this in the ablation study in Subsection 4.3.) Here, we theoretically analyze that including this bias in the weakly supervised contrastive loss defined in Eq. 12 leads to suboptimal solution. Formally, we first define what is a true positive pair and a false positive pair respectively, and then introduce Theorem 3.4 to show the issue in Eq. 12.

**Definition 3.3.** Given a sample  $x_i$ , we say  $(x_i, x_j)$  is a true positive pair (or a false negative pair), if their optimal representations satisfy  $\exp(\bar{h}_i \bar{h}_j^T / \tau) > 1$  for a small positive value  $\tau$ . Similarly, we say  $(x_i, x_k)$  is a false positive pair (or a true negative pair), if their optimal representations satisfy  $\exp(\bar{h}_i \bar{h}_k^T / \tau) \approx 0$  for a small positive value  $\tau$ .

**Theorem 3.4.** Given the contrastive learning loss function  $\mathcal{L}_3$ , if there exists one false positive sample in the batch during training, the contrastive learning loss will lead to a sub-optimal solution.

See proof in Appendix A.3.

#### 4 Experimental Results

In this section, we demonstrate the performance of our proposed framework in terms of both effectiveness and efficiency by comparing it with state-of-the-art methods.

# 4.1 Experiment Setup

4.1.1 Datasets: We evaluate the performance of our proposed framework on six datasets for both single-view and multi-view graph anomaly detection scenarios, including the Insider Threat Test (CERT) [20], DBLP [63], IMDB [63], BlogCatalog [54], Amazon [18] and YelpChi [33] datasets. Among these datasets, CERT, IMDB, and DBLP are multi-view graphs, while BlogCatalog, Amazon, and YelpChi are single-view graphs. CERT, Amazon, and YelpChi are real-world datasets, whereas IMDB, DBLP, and BlogCatalog are semi-synthetic graphs. (See Appendix B.1 for the details of generating anomalous nodes.) Specifically, the CERT dataset is a collection of synthetic insider threat test datasets that provides both synthetic background data and data from synthetic malicious actors. This dataset does not include a feature matrix, so we use node2vec [21] to extract two feature matrices as two views. IMDB is a movie network, where each node corresponds to a movie, and two adjacency matrices indicate whether two movies share the same actor or director. DBLP is a citation network, where each node corresponds to an academic research paper, and two adjacency matrices indicate whether two papers share the same authors or if one paper cites another. Amazon is a review network, where each node represents a product in the musical instruments category, and its attributes are extracted from product reviews. Similarly, the Yelp dataset contains hotel and restaurant reviews, either filtered (spam) or recommended (legitimate) by Yelp [18]. The statistics of these graphs are summarized in Table 2.

4.1.2 Experiment Setting: The neural network structure of the proposed framework is GCN [32]. The hyper-parameters  $\alpha$  and  $\lambda$  for each dataset are specified in Appendix B.2. In all experiments, we set the initial learning rate to be 1e-5, the hidden feature dimension to be 128 and use Adam [31] as the optimizer. The similarity function sim(a, b) is defined as  $sim(a, b) = exp(\frac{a \cdot b^T}{|a||b|})$ . We use TAM as the backbone of our method to capture local node affinity. The number of GCN layers is set to 2. The experiments are performed on a Windows machine with a 24GB RTX 4090 GPU. The code of our framework can be found in the anonymous GitHub link \*.

4.1.3 *Evaluation Metrics:* Following [44, 49], all methods are evaluated based on Area Under the Receiver Operating Characteristic Curve (AUROC) and Area Under the Precision-Recall Curve (AUPRC). Higher AUROC/AUPRC indicates better performance. All of the experiments are repeated five times with different random seeds and the mean and standard deviation are reported.

*4.1.4 Baseline Methods:* In our experiments, we compare our proposed framework CARE with state-of-the-art methods in the following two settings.

For *single-view graphs*, we compare CARE with the following eight baseline methods: (1). **ANOMALOUS** [47]: a shallow method, jointly conducting attribute selection and anomaly detection as a whole based on CUR decomposition and residual analysis; (2). **Dominant** [14]: a graph auto-encoder-based deep neural network model for graph anomaly detection, which encodes both the topological structure and node attributes to node embedding; (3). **CoLA** [38]: a contrastive self-supervised graph anomaly detection method by

<sup>\*</sup>https://anonymous.4open.science/r/CARE-demo-1C7F

exploiting the local information; (4). SLGAD [72]: an unsupervised 581 framework for outlier detection based on unlabeled in-distribution 582 583 data, which uses contrastive learning loss as a regularization; (5). HCM-A [27]: a self-supervised learning by forcing the prediction of 584 the shortest path length between pairs of nodes; (6). ComGA [41]: a 585 community-aware attributed graph anomaly detection framework; 586 587 (7). CONAD [66]: a contrastive learning-based graph anomaly detection method; (8). TAM [49]: an unsupervised anomaly method, 588 589 proposing a scoring measure by assigning large score to the nodes 590 that are less affiliated with their neighbors.

For multi-view graphs, we compare CARE with the following 591 five baseline methods: (1). MLRA [36]: a multi-view non-negative 592 matrix factorization-based method for anomaly detection, which 593 performs cross-view low-rank analysis for revealing the intrinsic 594 structures of data; (2). NSNMF [2]: an NMF based method, in-595 corporating the neighborhood structural similarity information 596 into the NMF framework to improve the anomaly detection per-597 formance; (3). SRLSP [64]: a multi-view detection method based 598 599 on the local similarity relation and data reconstruction; (4). NC-600 MOD [13]: an auto-encoder-based multi-view anomaly detection method, which proposes neighborhood consensus networks to en-601 602 code graph neighborhood information; (5) We also report the per-603 formance of TAM [49] on these multi-view graphs by averaging the adjacency matrix or concatenating two feature matrices together. 604 We omit the comparison with other baseline methods since TAM 605 outperforms them. 606

### 4.2 Experimental Analysis

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

638

Multi-view Graph Scenario. In this subsection, we evalu-4.2.1 ate the performance of CARE by comparing it with five baseline methods on three multi-view graphs, i. e., CERT, IMDB and DBLP datasets. The experimental results with respect to AUROC and AUPRC are presented in Table 3. We observe the following: (i) MLRA and NSNMF exhibit the poor performance among all baseline methods across most datasets. This can be attributed to their design, which is optimized for independent and identically distributed data (e.g., image data). Thus, these two methods struggle to capture graph topological information, resulting in lower performance in both AUPRC and AUROC. (ii) In contrast, NCMOD incorporates graph topological information into the learned representation through its proposed neighborhood consensus networks, thus achieving the second-best performance on the IMDB and DBLP datasets. (iii) TAM suffers from the inability to handle multi-view graphs, performing worse than several self-supervised learning methods designed for multi-view graphs (e.g., NCMOD). (iv) Our proposed method outperforms all baselines in terms of AUROC and AUPRC across all datasets. Notably, graph-based anomaly detection methods designed for single views (e.g., TAM) falter in the presence of view heterogeneity as simply concatenating input features from multiple views distorts the feature space, resulting in significant performance challenges. In contrast, our method excels by encoding both local and global node affinities in the learned representation and mitigating biases through the proposed similarity-guided graph contrastive loss.

*4.2.2 Single-view Graph Scenario.* Next, we evaluate the performance of CARE by comparing it to eight baseline methods on three

639

640

641

642

643

644

645

646

647

648

649

650

651

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

single-view graphs, i. e., BlogCatalog, Amazon and YelpChi datasets. The experimental results with respect to AUROC and AUPRC are presented in Table 4. We have the following observations: (i) TAM outperforms most baseline methods across three single-view graphs. We attribute its great performance to the design of normal structurepreserved graph truncation to remove edges connecting normal and abnormal nodes. (ii) While TAM excels on BlogCatalog, CARE still ranks as a close second, showcasing its robustness in handling single-view graphs. (iii) CARE proves to be a highly competitive approach, demonstrating superior performance on the Amazon and YelpChi datasets. Notably, CARE surpasses TAM, the second-best method, by over 15.9% in AUROC and 39.3% in AUPRC on the Amazon dataset. Similarly, on the YelpChi dataset, CARE improves AUROC by over 18.7% compared to TAM. The ability of CARE to capture both local and global node affinity information in the learned representations enables it to perform effectively on the Amazon and YelpChi datasets. By leveraging its similarity-guided graph contrastive loss, CARE enhances representation learning, allowing it to better distinguish between normal and anomalous nodes. Overall, these results emphasize that while TAM leverages truncated affinity maximization techniques to tailor the raw adjacency matrix, CARE offers a more flexible and powerful approach by incorporating global and local affinities, making it a highly effective method in the single-view graph anomaly detection scenario. Its ability to significantly outperform other baselines on challenging datasets like Amazon and YelpChi demonstrates the effectiveness of the proposed method.

## 4.3 Ablation Study

In this subsection, we conduct an ablation study to demonstrate the necessity of each component of CARE and validate the effectiveness of similarity-guided graph contrastive loss over vanilla and weakly supervised contrastive loss. Specifically, CARE-A removes the similarity measurement of soft membership and replaces  $\mathcal{L}_A$  with  $\mathcal{L}_2$ . CARE-G refers to a variant of our proposed method by removing similarity-guided contrastive loss with the vanilla contrastive loss while CARE-WSC substitutes it with the weakly supervised loss as shown in Eq. 12. The experimental results with respect to AUROC on the BlogCatalog, Amazon, DBLP, and IMDB datasets are presented in Table 5. Our observations are as follows:

- Global Node Affinity: CARE shows slight improvements on Amazon and BlogCatalog compared to CARE-A. However, excluding the global node affinity matrix (i. e., the similarity map of the soft membership) in CARE-A results in significant performance drops of approximately 40% on DBLP and 56% on IMDB, highlighting the importance of capturing global affinities in certain datasets.
- Similarity-Guided Loss: Removing the similarity-guided graph contrastive loss (i.e., CARE-G) leads to an average AUROC performance drop of 15% across the four datasets, underscoring the critical role of this regularization in mitigating bias and improving performance.
- **Contrastive Loss Substitutions**: Replacing the similarityguided loss with either the vanilla contrastive loss (i.e., CARE-InfoNCE) or the weakly supervised contrastive loss

Table 3: Results on multi-view graphs (i. e., CERT, IMDB, DBLP) with respect to AUROC and AUPRC. We boldface the best performance and underline the second-best.

Method CE AUPRC	RT	IM	DB	DBLP		
	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC
MLRA	$0.0379 \pm 0.001$	$0.3829 \pm 0.003$	$0.2695 \pm 0.007$	$0.5926 \pm 0.005$	$0.2211 \pm 0.005$	$0.5568 \pm 0.005$
NSNMF	$0.0704 \pm 0.001$	$0.4578 \pm 0.001$	$0.0634 \pm 0.000$	$0.4969 \pm 0.001$	$0.1436 \pm 0.007$	$0.6418 \pm 0.001$
NCMOD	$0.0749 \pm 0.001$	$0.5133 \pm 0.001$	$0.6629 \pm 0.013$	$0.8030 \pm 0.007$	$0.4809 \pm 0.006$	$0.7271 \pm 0.004$
SRSLP	$0.0806 \pm 0.007$	$0.5405 \pm 0.003$	$0.5552 \pm 0.017$	$0.7343 \pm 0.003$	$0.0643 \pm 0.002$	$0.5228 \pm 0.001$
TAM	$0.0771 \pm 0.007$	$0.5400 \pm 0.005$	$0.6521 \pm 0.016$	$0.8233 \pm 0.013$	$0.3466 \pm 0.016$	$0.6690 \pm 0.005$
CARE	$0.1198 \pm 0.003$	$\textbf{0.6056} \pm \textbf{0.001}$	$0.8804 \pm 0.024$	$\textbf{0.8792} \pm \textbf{0.031}$	$\textbf{0.6380} \pm \textbf{0.027}$	$\textbf{0.8868} \pm \textbf{0.007}$

Table 4: Results on single-view datasets (i. e., BlogCatalog, Amazon, YelpChi) with respect to AUROC and AUPRC. We boldface the best performance and underline the second-best.

Method	BlogC	atalog	Ama	azon	YelpChi		
	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	
ANOMALOUS	$0.0652 \pm 0.005$	$0.5652 \pm 0.025$	$0.0558 \pm 0.001$	$0.4457 \pm 0.005$	$0.0519 \pm 0.002$	$0.4956 \pm 0.003$	
Dominant	$0.3102 \pm 0.011$	$0.7590 \pm 0.010$	$0.1424 \pm 0.002$	$0.5996 \pm 0.002$	$0.0395 \pm 0.020$	$0.4133 \pm 0.100$	
CoLA	$0.3270 \pm 0.000$	$0.7746 \pm 0.009$	$0.0677 \pm 0.001$	$0.5898 \pm 0.011$	$0.0448 \pm 0.002$	$0.4636 \pm 0.001$	
SLGAD	$0.3882 \pm 0.007$	$0.8123 \pm 0.002$	$0.0634 \pm 0.005$	$0.5937 \pm 0.005$	$0.0350 \pm 0.000$	$0.3312 \pm 0.035$	
HCM-A	$0.3139 \pm 0.001$	$0.7980 \pm 0.004$	$0.0527 \pm 0.015$	$0.3956 \pm 0.014$	$0.0287 \pm 0.012$	$0.4593 \pm 0.005$	
ComGA	$0.3293 \pm 0.028$	$0.7683 \pm 0.004$	$0.1153 \pm 0.005$	$0.5895 \pm 0.010$	$0.0423 \pm 0.000$	$0.4391 \pm 0.000$	
CONAD	$0.3284 \pm 0.004$	$0.7807 \pm 0.003$	$0.1372 \pm 0.009$	$0.6142 \pm 0.008$	$0.0405 \pm 0.002$	$0.4588 \pm 0.003$	
TAM	$\textbf{0.4182} \pm \textbf{0.225}$	$\textbf{0.8248} \pm \textbf{0.003}$	$0.2634 \pm 0.008$	$0.7064 \pm 0.008$	$0.0778 \pm 0.009$	$0.5643 \pm 0.007$	
CARE	$0.4043 \pm 0.010$	$0.8194 \pm 0.003$	$\textbf{0.6563} \pm \textbf{0.011}$	$\textbf{0.8656} \pm \textbf{0.002}$	$\textbf{0.1218} \pm \textbf{0.003}$	$\textbf{0.7516} \pm \textbf{0.003}$	

Table 5: Ablation study on BlogCatalog, Amazon, DBLP and IMDB datasets with respect to AUROC.

Model	BlogCatalog	Amazon	DBLP	IMDB	AVERAGE
CARE	$\textbf{0.8194} \pm \textbf{0.000}$	$\textbf{0.8656} \pm \textbf{0.002}$	$\textbf{0.8868} \pm \textbf{0.007}$	$\textbf{0.8792} \pm \textbf{0.031}$	0.8628
CARE-A	$0.8144 \pm 0.000$	$0.8514 \pm 0.003$	$0.4831 \pm 0.002$	$0.2138 \pm 0.013$	0.5907
CARE-G	$0.6313 \pm 0.002$	$0.8645 \pm 0.004$	$0.8639 \pm 0.003$	$0.4996 \pm 0.001$	0.7148
CARE-InfoNCE	$0.7685 \pm 0.001$	$0.7737 \pm 0.006$	$0.8560 \pm 0.005$	$0.8452 \pm 0.032$	0.8108
CARE-WSC	$0.7639 \pm 0.011$	$0.8019 \pm 0.004$	$0.8538 \pm 0.006$	$0.8581 \pm 0.005$	0.8194

(i.e., CARE-WSC) results in a 4% to 5% performance reduction on average. This aligns with the theoretical analysis in Theorem 3.4, suggesting that neither the vanilla nor weakly supervised contrastive losses are as effective in mitigating bias as the similarity-guided approach.

These results validate the necessity of the global affinity information and the similarity-guided contrastive loss, both of which are essential for the robustness and effectiveness of CARE.

#### 4.4 Parameter Analysis

In this subsection, we delve into the parameter sensitivity analysis of the CARE framework on the four datasets, specifically exploring the impact of  $\alpha$ ,  $\lambda$ , and the number of clusters. In the experiment, the mean and standard deviation of AUROC over five rounds are reported. We fix the number of clusters to be 10 and vary the values of  $\alpha$  and  $\lambda$  while recording the AUROC of CARE. The performance is visualized in Figure 2, where the x, y, and z axes represent  $\alpha$ , log( $\lambda$ ), and AUROC, respectively. From observation, CARE achieves the best performance when  $\alpha = 0.4$  on the Amazon and DBLP datasets, while it reaches its peak performance at  $\alpha = 0.01$  on the BlogCatalog dataset and  $\alpha = 1$  on the IMDB dataset. Upon investigation, we find 

that the number of edges in the first view on the IMDB dataset is unusually large compared to the number of edges in the second view as presented in Table 2. Due to this unusual pattern, we hypothesize that the raw adjacency matrix is not reliable and thus the global node affinity (i.e., similarity of soft assignment) plays a crucial role in detecting anomalous nodes. Thus, by setting  $\alpha = 1$ , we exclude the local node affinity in the loss function  $\mathcal{L}_A$  and therefore CARE achieves better performance solely relying on the global node affinity. In terms of the parameter analysis on  $\lambda$ , we find that CARE prefers a smaller value of  $\lambda$  (e. g.,  $\lambda = 0.1$ ). One explanation for this is that a large value of  $\lambda$  tends to overly dominate the optimization of the overall objective function as  $\lambda$  is used to balance the importance between the similarity-guided contrastive loss and the node affinity learning loss. In the subsequent experiment, we scrutinize the role of the number of clusters in shaping the anomaly detection criteria. Figure 3 illustrates the results, with the x and y axes representing the values of the number of clusters and AUROC, respectively. We observe that changing the number of the clusters does not greatly influence the performance of CARE across these four datasets. Notably, CARE achieves better performance when the number of clusters is 5 or 10 on most datasets, suggesting that the Conference acronym 'XX, June 03-05, 2018, Woodstock, NY







(c) DBLP dataset

(b) BlogCatalog dataset

proposed method prefers a small number of clusters. By reducing the number of clusters, our method tends to generalize better by capturing larger, more meaningful patterns in the data instead of overfitting to noise.

## 4.5 Efficiency Analysis

Time Complexity. In this subsection, we analyze the time com-plexity of our proposed CARE. Assume that the graph has *n* nodes, the input feature dimension is d, the hidden feature dimension is f, the number of nodes is n, the number of edges is |E|, and the number of clusters is k. For ease of explanation, we only consider the 1-layer case. Following [7], the time complexity of computing the soft mem-bership matrix using GCN is O(ndk + n|E|k + n). The complexity of the GCN to capture the hidden representation **h** is O(|E|d + ndf)with sparse computation. The complexity of the similarity-guided contrastive learning loss is  $O(n^2 f)$ . However, in the experiments, we can use the sampling strategy to sample  $p(p \ll n)$  nodes and thus the complexity can be reduced to  $O(p^2 f)$ . The total complexity of CARE is  $O(n(kd + fd + |E|k + 1) + |E|d + p^2 f)$ . 

Running Time Analysis. Next, we experiment on the YelpChi dataset to show the efficiency of our proposed CARE by changing the number of nodes and the number of layers. The reason why we only provide the efficiency analysis regarding the training/running time vs the number of nodes on the YelpChi dataset is that we can manually increase the number of nodes from 500 to 15,000 on this dataset. For the other datasets (e.g., IMDB, BlogCatalog, Amazon, DBLP), the number of nodes is less than 15,000 (see Table 2 for details) and we cannot get the running time if the number of nodes is set to be a value larger than 15,000. Thus, we do not provide the efficiency analysis on small datasets. In the first experiment, we fix the number of layers to be 1, the total number of iterations to be 10,000 and adjust the number of nodes by randomly selecting



(d) IMDB dataset

Figure 4: Efficiency analysis on the YelpChi dataset

k nodes, where  $k \in [500, 1000, 2000, 5000, 10000, 15000]$ . The experimental result is shown in Figure 4 (a). By observation, we find that the running time is quadratic to the number of nodes. In the second experiment, we fixed the number of nodes to be 2000, the total number of iterations to be 10,000, and adjusted the number of GCN layers from 1 layer to 4 layers. The experimental result is shown in Figure 4 (b). We observe that the running time is almost linear with respect to the number of layers.

# 5 Conclusion

In this paper, we proposed a novel self-supervised framework for anomaly detection in multi-view graphs, addressing key limitations of existing methods. By capturing both local and global node affinities and leveraging a similarity-guided contrastive loss, our approach effectively identifies anomalous nodes without relying on strong structural assumptions. The proposed method not only augments the graph structure with learned soft membership assignments but also mitigates the negative impact of low-quality assignments through a robust loss function, theoretically connected to graph spectral clustering. Our experimental results on six datasets demonstrate the superior performance of the proposed framework in detecting anomalies in complex, multi-view graphs.

Anon.

Cluster Aware Graph Anomaly Detection

Conference acronym 'XX, June 03-05, 2018, Woodstock, NY

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

#### 929 References

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

981

982

983

984

985

- Charu C. Aggarwal and Philip S. Yu. 2001. Outlier Detection for High Dimensional Data. In Proceedings of the 2001 ACM SIGMOD international conference on Management of data, Santa Barbara, CA, USA, May 21-24, 2001. ACM, 37-46.
- [2] Imtiaz Ahmed, Xia Ben Hu, Mithun P. Acharya, and Yu Ding. 2021. Neighborhood Structure Assisted Non-negative Matrix Factorization and Its Application in Unsupervised Point-wise Anomaly Detection. J. Mach. Learn. Res. 22 (2021), 34:1–34:32.
- [3] Samet Akcay, Amir Atapour Abarghouei, and Toby P. Breckon. 2018. GANomaly: Semi-supervised Anomaly Detection via Adversarial Training. In Computer Vision - ACCV 2018 - 14th Asian Conference on Computer Vision, Perth, Australia, December 2-6, 2018, Revised Selected Papers, Part III, Vol. 11363. Springer, 622–637.
- [4] Kais Allab, Lazhar Labiod, and Mohamed Nadif. 2017. A Semi-NMF-PCA Unified Framework for Data Clustering. IEEE Trans. Knowl. Data Eng. 29, 1 (2017), 2–16.
- [5] Francis R. Bach and Michael I. Jordan. 2003. Learning Spectral Clustering. In Advances in Neural Information Processing Systems 16 [Neural Information Processing Systems, NIPS 2003, December 8-13, 2003, Vancouver and Whistler, British Columbia, Canada]. MIT Press, 305–312.
- [6] Mikhail Belkin and Partha Niyogi. 2003. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural Comput.* 15, 6 (2003), 1373–1396.
- [7] Derrick Blakely, Jack Lanchantin, and Yanjun Qi. 2021. Time and space complexity of graph convolutional networks. Accessed on: Dec 31 (2021).
- [8] Deyu Bo, Xiao Wang, Chuan Shi, Meiqi Zhu, Emiao Lu, and Peng Cui. 2020. Structural Deep Clustering Network. In WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020. ACM / IW3C2, 1400-1410.
- [9] Oliver Brdiczka, Juan Liu, Bob Price, Jianqiang Shen, Akshay Patil, Richard Chow, Eugene Bart, and Nicolas Ducheneaut. 2012. Proactive Insider Threat Detection through Graph Learning and Psychological Context. In 2012 IEEE Symposium on Security and Privacy Workshops, San Francisco, CA, USA, May 24-25, 2012. IEEE Computer Society, 142–149.
- [10] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. 2011. Robust principal component analysis? Journal of the ACM (JACM) 58, 3 (2011), 1–37.
- [11] Bo Chen, Jing Zhang, Xiaokang Zhang, Yuxiao Dong, Jian Song, Peng Zhang, Kaibo Xu, Evgeny Kharlamov, and Jie Tang. 2023. GCCAD: Graph Contrastive Coding for Anomaly Detection. *IEEE Trans. Knowl. Data Eng.* 35, 8 (2023), 8037– 8051.
- [12] Ting Chen, Lu-An Tang, Yizhou Sun, Zhengzhang Chen, and Kai Zhang. 2016. Entity Embedding-Based Anomaly Detection for Heterogeneous Categorical Events. In Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016. IJCAI/AAAI Press, 1396–1403.
- [13] Li Cheng, Yijie Wang, and Xinwang Liu. 2021. Neighborhood Consensus Networks for Unsupervised Multi-view Outlier Detection. In Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Virtual Event, February 2-9, 2021. AAAI Press, 7099–7106.
- [14] Kaize Ding, Jundong Li, Rohit Bhanushali, and Huan Liu. 2019. Deep Anomaly Detection on Attributed Networks. In Proceedings of the 2019 SIAM International Conference on Data Mining, SDM 2019, Calgary, Alberta, Canada, May 2-4, 2019. SIAM, 594-602.
- [15] Kamran Ghasedi Dizaji, Amirhossein Herandi, Cheng Deng, Weidong Cai, and Heng Huang. 2017. Deep Clustering via Joint Convolutional Autoencoder Embedding and Relative Entropy Minimization. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017.* IEEE Computer Society, 5747–5756.
- [16] Xiangyu Dong, Xingyi Zhang, Yanni Sun, Lei Chen, Mingxuan Yuan, and Sibo Wang. 2024. SmoothGNN: Smoothing-based GNN for Unsupervised Node Anomaly Detection. *CoRR* abs/2405.17525 (2024).
- [17] David L Donoho et al. 2000. High-dimensional data analysis: The curses and blessings of dimensionality. AMS math challenges lecture 1, 2000 (2000), 32.
- [18] Yingtong Dou, Zhiwei Liu, Li Sun, Yutong Deng, Hao Peng, and Philip S. Yu. 2020. Enhancing Graph Neural Network-based Fraud Detectors against Camouflaged Fraudsters. In CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020, Mathieu d'Aquin, Stefan Dietze, Claudia Hauff, Edward Curry, and Philippe Cudré-Mauroux (Eds.). ACM, 315–324.
- [19] William Eberle, Jeffrey Graves, and Lawrence Holder. 2010. Insider threat detection using a graph-based approach. *Journal of Applied Security Research* 6, 1 (2010), 32–81.
- [20] Joshua Glasser and Brian Lindauer. 2013. Bridging the Gap: A Pragmatic Approach to Generating Insider Threat Data. In 2013 IEEE Symposium on Security and Privacy Workshops, San Francisco, CA, USA, May 23-24, 2013. IEEE Computer Society, 98–104.
- [21] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining. 855–864.
- [22] Frank E Grubbs. 1969. Procedures for detecting outlying observations in samples. Technometrics 11, 1 (1969), 1–21.

- [23] Xiaohong Guan, Wei Wang, and Xiangliang Zhang. 2009. Fast intrusion detection based on a non-negative matrix factorization model. J. Netw. Comput. Appl. 32, 1 (2009), 31–44.
- [24] Xifeng Guo, Xinwang Liu, En Zhu, and Jianping Yin. 2017. Deep Clustering with Convolutional Autoencoders. In Neural Information Processing - 24th International Conference, ICONIP 2017, Guangzhou, China, November 14-18, 2017, Proceedings, Part II, Vol. 10635. Springer, 373–382.
- [25] Jeff Z. HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. 2021. Provable Guarantees for Self-Supervised Deep Learning with Spectral Contrastive Loss. In Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual. 5000–5011.
- [26] Bryan Hooi, Hyun Ah Song, Alex Beutel, Neil Shah, Kijung Shin, and Christos Faloutsos. 2016. FRAUDAR: Bounding Graph Fraud in the Face of Camouflage. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016. ACM, 895–904.
- [27] Tianjin Huang, Yulong Pei, Vlado Menkovski, and Mykola Pechenizkiy. 2022. Hop-Count Based Self-supervised Anomaly Detection on Attributed Networks. In Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2022, Grenoble, France, September 19-23, 2022, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 13713), Massih-Reza Amini, Stéphane Canu, Asja Fischer, Tias Guns, Petra Kralj Novak, and Grigorios Tsoumakas (Eds.). Springer, 225–241.
- [28] Jianguo Jiang, Jiuming Chen, Tianbo Gu, Kim-Kwang Raymond Choo, Chao Liu, Min Yu, Weiqing Huang, and Prasant Mohapatra. 2019. Anomaly Detection with Graph Convolutional Networks for Insider Threat and Fraud Detection. In 2019 IEEE Military Communications Conference, MILCOM 2019, Norfolk, VA, USA, November 12-14, 2019. IEEE, 109–114.
- [29] Petra Jílková and Petra Králová. 2021. Digital consumer behaviour and ecommerce trends during the COVID-19 crisis. *International Advances in Economic Research* 27, 1 (2021), 83–85.
- [30] Ming Jin, Yixin Liu, Yu Zheng, Lianhua Chi, Yuan-Fang Li, and Shirui Pan. 2021. ANEMONE: Graph Anomaly Detection with Multi-Scale Contrastive Learning. In CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021, Gianluca Demartini, Guido Zuccon, J. Shane Culpepper, Zi Huang, and Hanghang Tong (Eds.). ACM, 3122–3126.
- [31] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings.
- [32] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net.
- [33] Srijan Kumar, Xikun Zhang, and Jure Leskovec. 2019. Predicting Dynamic Embedding Trajectory in Temporal Interaction Networks. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019, Ankur Teredesai, Vipin Kumar, Ying Li, Rómer Rosales, Evimaria Terzi, and George Karypis (Eds.). ACM, 1269–1278.
- [34] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. 2021. CutPaste: Self-Supervised Learning for Anomaly Detection and Localization. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, June 19-25, 2021.* Computer Vision Foundation / IEEE, 9664–9674.
- [35] Jianbo Li, Lecheng Zheng, Yada Zhu, and Jingrui He. 2021. Outlier Impact Characterization for Time Series Data. In Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021. AAAI Press, 11595–11603.
- [36] Sheng Li, Ming Shao, and Yun Fu. 2015. Multi-View Low-Rank Analysis for Outlier Detection. In Proceedings of the 2015 SIAM International Conference on Data Mining, Vancouver, BC, Canada, April 30 - May 2, 2015. SIAM, 748–756.
- [37] Zhenguo Li and Jianzhuang Liu. 2009. Constrained clustering by spectral kernel learning. In IEEE 12th International Conference on Computer Vision, ICCV 2009, Kyoto, Japan, September 27 - October 4, 2009. IEEE Computer Society, 421–427.
- [38] Yixin Liu, Zhao Li, Shirui Pan, Chen Gong, Chuan Zhou, and George Karypis. 2022. Anomaly Detection on Attributed Networks via Contrastive Self-Supervised Learning. *IEEE Trans. Neural Networks Learn. Syst.* 33, 6 (2022), 2378–2392.
- [39] Yue Liu, Wenxuan Tu, Sihang Zhou, Xinwang Liu, Linxuan Song, Xihong Yang, and En Zhu. 2022. Deep Graph Clustering via Dual Correlation Reduction. In Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022. AAAI Press, 7603–7611.
- [40] Chang-Tien Lu, Dechang Chen, and Yufeng Kou. 2003. Algorithms for Spatial Outlier Detection. In Proceedings of the 3rd IEEE International Conference on

1046

1049

1050

1081

1082

1083

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

Anon.

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

Data Mining (ICDM 2003), 19-22 December 2003, Melbourne, Florida, USA. IEEE Computer Society, 597–600.

- [41] Xuexiong Luo, Jia Wu, Amin Beheshti, Jian Yang, Xiankun Zhang, Yuan Wang, and Shan Xue. 2022. ComGA: Community-Aware Attributed Graph Anomaly Detection. In WSDM '22: The Fifteenth ACM International Conference on Web
  - Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 25, 2022, K. Selcuk Candan, Huan Liu, Leman Akoglu, Xin Luna Dong, and Jiliang Tang (Eds.). ACM, 657–665.
- [42] Xiaoxiao Ma, Jia Wu, Shan Xue, Jian Yang, Chuan Zhou, Quan Z Sheng, Hui Xiong, and Leman Akoglu. 2021. A comprehensive survey on graph anomaly detection with deep learning. *IEEE Transactions on Knowledge and Data Engineering* (2021).
- [43] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. 2001. On Spectral Clustering: Analysis and an algorithm. In Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001, December 3-8, 2001, Vancouver, British Columbia, Canada]. MIT Press, 849–856.
- [44] Guansong Pang, Anton van den Hengel, Chunhua Shen, and Longbing Cao. 2021.
   Toward Deep Supervised Anomaly Detection: Reinforcement Learning from Partially Labeled Anomaly Data. In *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14-18,* 2021, Feida Zhu, Beng Chin Ooi, and Chunyan Miao (Eds.). ACM, 1298–1308.
- [45] Guansong Pang, Cheng Yan, Chunhua Shen, Anton van den Hengel, and Xiao Bai. 2020. Self-Trained Deep Ordinal Regression for End-to-End Video Anomaly Detection. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020. Computer Vision Foundation / IEEE, 12170–12179.
- [46] Hyunjong Park, Jongyoun Noh, and Bumsub Ham. 2020. Learning Memory Guided Normality for Anomaly Detection. In 2020 IEEE/CVF Conference on Com puter Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19,
   2020. Computer Vision Foundation / IEEE, 14360–14369.
- [47] Zhen Peng, Minnan Luo, Jundong Li, Huan Liu, and Qinghua Zheng. 2018.
   ANOMALOUS: A Joint Modeling Approach for Anomaly Detection on Attributed
   Networks. In Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden, Jérôme Lang (Ed.). ijcai.org, 3513–3519.
- [48] Tahereh Pourhabibi, Kok-Leong Ong, Booi H Kam, and Yee Ling Boo. 2020. Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *Decision Support Systems* 133 (2020), 113303.
- [49] Hezhe Qiao and Guansong Pang. 2023. Truncated Affinity Maximization: Oneclass Homophily Modeling for Graph Anomaly Detection. In Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, New Orleans, LA, USA, December 10 - 16, 2023, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.).
- [50] Jagdish Ramakrishnan, Elham Shaabani, Chao Li, and Mátyás A. Sustik. 2019.
   Anomaly Detection for an E-commerce Pricing System. In *Proceedings of the* 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019, Ankur Teredesai, Vipin
   Kumar, Ying Li, Rómer Rosales, Evimaria Terzi, and George Karypis (Eds.). ACM, 1917–1926.
  - [51] Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. 2019. A Theoretical Analysis of Contrastive Unsupervised Representation Learning. In Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, Vol. 97. PMLR, 5628-5637.
- [52] Bin Shen and Luo Si. 2010. Non-Negative Matrix Factorization Clustering on Multiple Manifolds. In Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010. AAAI Press.
   [53] Payel Singenchuk and Anna Enishkina. 2019. Practical Liser and Entity Behavior.
  - [53] Pavel Slipenchuk and Anna Epishkina. 2019. Practical User and Entity Behavior Analytics Methods for Fraud Detection Systems in Online Banking: A Survey. 948 (2019), 83–93.
  - [54] Lei Tang and Huan Liu. 2009. Relational learning via latent social dimensions. In Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, June 28 - July 1, 2009. ACM, 817–826.
  - [55] Hanghang Tong and Ching-Yung Lin. 2011. Non-Negative Residual Matrix Factorization with Application to Graph Anomaly Detection. In Proceedings of the Eleventh SIAM International Conference on Data Mining, SDM 2011, April 28-30, 2011, Mesa, Arizona, USA. SIAM / Omnipress, 143–153.
  - [56] George Trigeorgis, Konstantinos Bousmalis, Stefanos Zafeiriou, and Björn W. Schuller. 2014. A Deep Semi-NMF Model for Learning Hidden Representations. In Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014, Vol. 32. JMLR.org, 1692–1700.

- [57] Yao-Hung Hubert Tsai, Tianqin Li, Weixin Liu, Peiyuan Liao, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2022. Learning Weakly-supervised Contrastive Representations. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net.
- [58] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation Learning with Contrastive Predictive Coding. *CoRR* abs/1807.03748 (2018).
  [59] Ulrike von Luxburg. 2007. A tutorial on spectral clustering. *Stat. Comput.* 17, 4
- (2007), 395–416. https://doi.org/10.1007/S11222-007-9033-Z
- [60] Chun Wang, Shirui Pan, Guodong Long, Xingquan Zhu, and Jing Jiang. 2017. MGAE: Marginalized Graph Autoencoder for Graph Clustering. In Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 - 10, 2017. ACM, 889–898.
- [61] Daixin Wang, Jianbin Lin, Peng Cui, Quanhui Jia, Zhen Wang, Yanming Fang, Quan Yu, Jun Zhou, Shuang Yang, and Yuan Qi. 2019. A semi-supervised graph attentive network for financial fraud detection. In 2019 IEEE International Conference on Data Mining (ICDM). IEEE, 598–607.
- [62] Tongzhou Wang and Phillip Isola. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, Vol. 119. PMLR, 9929-9939.
- [63] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S. Yu. 2019. Heterogeneous Graph Attention Network. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019.* ACM, 2022– 2032.
- [64] Yu Wang, Chuan Chen, Jinrong Lai, Lele Fu, Yuren Zhou, and Zibin Zheng. 2023. A Self-Representation Method with Local Similarity Preserving for Fast Multi-View Outlier Detection. ACM Trans. Knowl. Discov. Data 17, 1 (2023), 2:1–2:20.
- [65] Huan Xu, Constantine Caramanis, and Sujay Sanghavi. 2010. Robust PCA via Outlier Pursuit. In Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held 6-9 December 2010, Vancouver, British Columbia, Canada. Curran Associates, Inc., 2496–2504.
- [66] Zhiming Xu, Xiao Huang, Yue Zhao, Yushun Dong, and Jundong Li. 2022. Contrastive Attributed Network Anomaly Detection with Data Augmentation. In Advances in Knowledge Discovery and Data Mining - 26th Pacific-Asia Conference, PAKDD 2022, Chengdu, China, May 16-19, 2022, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 13281), João Gama, Tianrui Li, Yang Yu, Enhong Chen, Yu Zheng, and Fei Teng (Eds.). Springer, 444–457.
- [67] Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L. Hamilton, and Jure Leskovec. 2018. Hierarchical Graph Representation Learning with Differentiable Pooling. In Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada. 4805–4815.
- [68] Xianchao Zhang, Linlin Zong, Xinyue Liu, and Hong Yu. 2015. Constrained NMF-Based Multi-View Clustering on Unmapped Data. In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA. AAAI Press, 3174–3180.
- [69] Handong Zhao, Zhengming Ding, and Yun Fu. 2017. Multi-View Clustering via Deep Matrix Factorization. In Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA. AAAI Press, 2921–2927.
- [70] Lecheng Zheng, Jinjun Xiong, Yada Zhu, and Jingrui He. 2022. Contrastive Learning with Complex Heterogeneity. In KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022. ACM, 2594–2604.
- [71] Mingkai Zheng, Fei Wang, Shan You, Chen Qian, Changshui Zhang, Xiaogang Wang, and Chang Xu. 2021. Weakly Supervised Contrastive Learning. In 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021. IEEE, 10022–10031.
- [72] Yu Zheng, Ming Jin, Yixin Liu, Lianhua Chi, Khoa T Phan, and Yi-Ping Phoebe Chen. 2021. Generative and contrastive self-supervised learning for graph anomaly detection. *IEEE Transactions on Knowledge and Data Engineering* (2021).
- [73] Qiwei Zhong, Yang Liu, Xiang Ao, Binbin Hu, Jinghua Feng, Jiayu Tang, and Qing He. 2020. Financial Defaulter Detection on Online Credit Payment via Multi-view Attributed Heterogeneous Information Network. In WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020. ACM / IW3C2, 785–795.
- [74] Xiaokang Zhou, Yiyong Hu, Wei Liang, Jianhua Ma, and Qun Jin. 2021. Variational LSTM Enhanced Anomaly Detection for Industrial Big Data. *IEEE Trans. Ind. Informatics* 17, 5 (2021), 3469–3477.
- [75] Hao Zhu, Ke Sun, and Peter Koniusz. 2021. Contrastive Laplacian Eigenmaps. In Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual. 5682–5695.
- 1157 1158
- 1159 1160

Cluster Aware Graph Anomaly Detection

#### A Proof

#### A.1 Proof for Lemma 3.1

Lemma 3.1: (Similarity-guided Graph Contrastive Loss) Let  $\overline{M}$  be the output of a one-layer graph neural network defined in Eq. 3. Then,  $\mathcal{L}_C$  is equivalent to the following loss function:

$$\mathcal{L}_C = \mathcal{L}$$

where  $\mathcal{L}_f = -\sum_{i=1}^n \sum_{j=1}^n \log \frac{\exp(2\tilde{A}_{ij}\tilde{h}_i\tilde{h}_j^T)}{\prod_{k=1}^n \exp((\tilde{h}_i\tilde{h}_k^T)^2)^{1/n}}$  is a graph contrastive loss and C is a constant.

Proof.

$$\begin{split} \min_{\hat{H}} \mathcal{L}_{C} &= \min_{\hat{H}} ||\tilde{A} - \tilde{H}\tilde{H}^{T}||_{F}^{2} \\ &= \min_{\hat{H}} \sum_{i=1}^{n} \sum_{j=1}^{n} (\tilde{A}_{ij} - \bar{h}_{i}\bar{h}_{j}^{T})^{2} \\ &= \min_{\hat{H}} \sum_{i=1}^{n} \sum_{j=1}^{n} (\tilde{A}_{ij}^{2} - 2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T} + (\bar{h}_{i}\bar{h}_{j}^{T})^{2}) \end{split}$$

Notice that  $\overline{M}$  is independent of  $\overline{H}$ . When we fix the parameter  $\overline{M}$  to update  $\overline{H}$ , then  $\widetilde{A}_{ij}^2$  can be considered as a constant in this optimization problem. Thus, we have

$$\begin{split} \min_{\bar{H}} \mathcal{L}_{C} &= \min_{\bar{H}} \sum_{i=1}^{n} \sum_{j=1}^{n} (-2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T} + (\bar{h}_{i}\bar{h}_{j}^{T})^{2}) + C \\ &= \min_{\bar{H}} \sum_{i=1}^{n} \sum_{j=1}^{n} (-2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T} + \frac{1}{n} \sum_{k=1}^{n} (\bar{h}_{i}\bar{h}_{k}^{T})^{2}) + C \\ &= \min_{\bar{H}} - \sum_{i=1}^{n} \sum_{j=1}^{n} \log \frac{\exp(2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T})}{\prod_{k=1}^{n} \exp((\bar{h}_{i}\bar{h}_{k}^{T})^{2})^{\frac{1}{n}}} + C \\ &= \min_{\bar{H}} \mathcal{L}_{f} + C \end{split}$$

where  $\mathcal{L}_f = -\sum_{i=1}^n \sum_{j=1}^n \log \frac{\exp(2\tilde{A}_{ij}\tilde{h}_i\tilde{h}_j^T)}{\prod_{k=1}^n \exp((\tilde{h}_i\tilde{h}_k^T)^2)^{1/n}}$ . Thus, we have  $\mathcal{L}_C = \mathcal{L}_f + C$ , which completes the proof.

#### A.2 Proof for Lemma 3.2

**Lemma 3.2:** (Graph Contrastive Spectral Clustering) Let  $\overline{\mathbf{M}}$  be the output of a one-layer graph neural network defined in Eq. 3 and  $\overline{\mathbf{h}}_i$  and  $\overline{\mathbf{h}}_j$  be unit vectors. Then, minimizing  $\mathcal{L}_C$  is equivalent to minimizing the following loss function:

$$\min \mathcal{L}_C = \min[2Tr(\bar{H}^T L\bar{H}) + R(\bar{H})]$$

where  $\mathbf{L} = \mathbf{I} - \hat{\mathbf{A}}$  can be considered as the normalized graph Laplacian,  $\mathbf{I}$  is the identity matrix and  $R(\hat{\mathbf{H}}) = ||\hat{\mathbf{H}}\hat{\mathbf{H}}^T||_F^2$  is the regularization term.

PROOF. Based on the proof in Lemma 3.1, we have

$$\min \mathcal{L}_{C} = \sum_{i=1}^{n} \sum_{j=1}^{n} (-2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T} + (\bar{h}_{i}\bar{h}_{j}^{T})^{2})$$
$$= (\sum_{i=1}^{n} \sum_{j=1}^{n} (2\tilde{A}_{ij} - 2\tilde{A}_{ij} - 2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T} + (\bar{h}_{i}\bar{h}_{j}^{T})^{2})$$

Notice that  $\tilde{A} = \tilde{D}^{-1/2} (\alpha \bar{M} \bar{M}^T + (1 - \alpha) \frac{1}{v} \sum_{a=1}^{v} A^a) \tilde{D}^{-1/2}$  and  $\tilde{D}_{ii} = \sum_j (\alpha \bar{M}_i \bar{M}_j^T + (1 - \alpha) \frac{1}{v} \sum_{a=1}^{v} A_{ij}^a)$ . We have  $\sum_{j=1}^n \tilde{A}_{ij} = 1$  and thus  $\sum_{i=1}^n \sum_{j=1}^n 2\tilde{A}_{ij}$  is a constant, which can be ignored in this optimization problem. Since  $\bar{h}_i$  and  $\bar{h}_j$  are unit vectors, we have

$$\min \mathcal{L}_{C} = \sum_{i=1}^{n} \sum_{j=1}^{n} (2\tilde{A}_{ij} - 2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T} + (\bar{h}_{i}\bar{h}_{j}^{T})^{2})$$
$$= \sum_{i=1}^{n} \sum_{j=1}^{n} \tilde{A}_{ij} ||\bar{h}_{i}||_{2}^{2} + \sum_{i=1}^{n} \sum_{j=1}^{n} \tilde{A}_{ij} ||\bar{h}_{j}||_{2}^{2} - \sum_{i=1}^{n} \sum_{j=1}^{n} 2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T}$$

$$\sum_{i=1}^{n} \sum_{j=1}^{n} (\bar{\boldsymbol{h}}_i \bar{\boldsymbol{h}}_j^T)^2$$

$$=\sum_{i=1}^{n} I_{ii} ||\bar{h}_{i}||_{2}^{2} + \sum_{j=1}^{n} I_{jj} ||\bar{h}_{j}||_{2}^{2} - \sum_{i=1}^{n} \sum_{j=1}^{n} 2\tilde{A}_{ij}\bar{h}_{i}\bar{h}_{j}^{T}$$

$$+\sum_{i=1}\sum_{j=1}(\bar{\boldsymbol{h}}_i\bar{\boldsymbol{h}}_j^T)^2$$

$$= 2Tr(\bar{H}^T L\bar{H}) + R(\bar{H})$$

(13)

where  $L = I - \tilde{A}$  can be considered as the normalized graph Laplacian, I is the identity matrix and  $R(\tilde{H}) = ||\tilde{H}\tilde{H}^T||_F^2$ , which completes the proof.

#### A.3 **Proof for theorem 3.4**

**Definition 3.3** Given a sample  $\mathbf{x}_i$ , we say  $(\mathbf{x}_i, \mathbf{x}_j)$  is a false negative pair (or a true positive pair), if their optimal representations satisfy  $\exp(\bar{\mathbf{h}}_i \bar{\mathbf{h}}_j^T / \tau) > 1$  for a small positive value  $\tau$ . Similarly, we say  $(\mathbf{x}_i, \mathbf{x}_k)$  is a true negative pair (or a false positive pair), if their optimal representations satisfy  $\exp(\bar{\mathbf{h}}_i \bar{\mathbf{h}}_k^T / \tau) \approx 0$  for a small positive value  $\tau$ .

**Theorem 3.4** Given the contrastive learning loss function  $\mathcal{L}_3$ , if there exists one false positive sample in the batch during training, the contrastive learning loss will lead to a sub-optimal solution.

**PROOF.** We can rewrite  $\mathcal{L}_3$  as follows:

$$\begin{split} \mathcal{L}_{3} &= \sum_{i} \sum_{j \in C(i), j \neq i} [\log(\frac{\exp(\bar{h}_{i}\bar{h}_{j}^{T}/\tau) + \sum_{k \notin C(i)} \exp(\bar{h}_{i}\bar{h}_{k}^{T}/\tau)}{\exp(\bar{h}_{i}\bar{h}_{j}^{T}/\tau)})] \\ &= \sum_{i} \sum_{j \in C(i), j \neq i} [\log(\exp(\bar{h}_{i}\bar{h}_{j}^{T}/\tau) + \sum_{k \notin C(i)} \exp(\bar{h}_{i}\bar{h}_{k}^{T}/\tau)) - \bar{h}_{i}\bar{h}_{j}^{T}/\tau] \\ &= \sum_{j \in C(1), j \neq 1} [\log(\exp(\bar{h}_{1}\bar{h}_{j}^{T}/\tau) + \sum_{k \notin C(1)} \exp(\bar{h}_{1}\bar{h}_{k}^{T}/\tau)) - \bar{h}_{1}\bar{h}_{j}^{T}/\tau] \\ &+ \sum_{i} \sum_{j \in C(i), j \neq i} [\log(\exp(\bar{h}_{i}\bar{h}_{j}^{T}/\tau) + \sum_{k \notin C(i)} \exp(\bar{h}_{i}\bar{h}_{k}^{T}/\tau)) - \bar{h}_{i}\bar{h}_{j}^{T}/\tau] \end{split}$$

$$\sum_{i \neq 1} \sum_{j \in C(i), j \neq i} (\log(\exp(in_j/t) + \sum_{i \neq k \notin C(i)} \exp(in_k/t)) + \inf_{j \neq i} (14)$$

$$(14)$$

Here, we select  $x_1$  as the anchor node such that  $(x_i,x_1)$  is a true positive pair (i. e.,  $x_i$  and  $x_1$  are from the same cluster). Taking the derivative of  $\mathcal{L}_3$  with respect to  $\bar{h}_1$ , we have

$$\begin{split} \frac{\partial \mathcal{L}_{3}}{\partial \bar{h}_{1}} &= \frac{1}{\tau} \sum_{j \in C(1), j \neq 1} \left[ \frac{\exp(\bar{h}_{1}\bar{h}_{j}^{T}/\tau)\bar{h}_{j} + \sum_{k \notin C(1)} \exp(\bar{h}_{1}\bar{h}_{k}^{T}/\tau)\bar{h}_{k}}{\exp(\bar{h}_{1}\bar{h}_{j}^{T}/\tau) + \sum_{k \notin C(1)} \exp(\bar{h}_{1}\bar{h}_{k}^{T}/\tau)} - \bar{h}_{j} \right] \\ &+ \frac{1}{\tau} \sum_{i \neq 1, i \in C(1)} \left[ \frac{\exp(\bar{h}_{i}\bar{h}_{1}^{T}/\tau) + \sum_{k \notin C(i)} \exp(\bar{h}_{i}\bar{h}_{k}^{T}/\tau)}{\exp(\bar{h}_{1}\bar{h}_{1}^{T}/\tau) + \sum_{k \notin C(i)} \exp(\bar{h}_{i}\bar{h}_{k}^{T}/\tau)} - \bar{h}_{i} \right] \\ &= \frac{1}{\tau} \sum_{j \in C(1), j \neq 1} \left[ \frac{\sum_{k \notin C(1)} \exp(\bar{h}_{1}\bar{h}_{k}^{T}/\tau) + \sum_{k \notin C(1)} \exp(\bar{h}_{1}\bar{h}_{k}^{T}/\tau)}{\exp(\bar{h}_{1}\bar{h}_{j}^{T}/\tau) + \sum_{k \notin C(1)} \exp(\bar{h}_{1}\bar{h}_{k}^{T}/\tau)} \right] \\ &+ \frac{1}{\tau} \sum_{i \neq 1, i \in C(1)} \left[ \frac{-\sum_{k \notin C(i)} \exp(\bar{h}_{i}\bar{h}_{k}^{T}/\tau) + \sum_{k \notin C(i)} \exp(\bar{h}_{i}\bar{h}_{k}^{T}/\tau)}{\exp(\bar{h}_{i}\bar{h}_{k}^{T}/\tau)} \right] \end{split}$$

Setting the gradient to 0, we have

$$\sum_{j \in C(1), j \neq 1} \left[ \frac{\sum_{k \notin C(1)} \left[ \exp(\tilde{\boldsymbol{h}}_1 \boldsymbol{h}_k^T / \tau) \boldsymbol{h}_j - \exp(\tilde{\boldsymbol{h}}_1 \boldsymbol{h}_k^T / \tau) \boldsymbol{h}_k \right]}{\exp(\tilde{\boldsymbol{h}}_1 \tilde{\boldsymbol{h}}_j^T / \tau) + \sum_{k \notin C(1)} \exp(\tilde{\boldsymbol{h}}_1 \tilde{\boldsymbol{h}}_k^T / \tau)} \right]$$
  
+ 
$$\sum_{i \neq 1, i \in C(1)} \left[ \frac{\sum_{k \notin C(i)} \exp(\tilde{\boldsymbol{h}}_i \tilde{\boldsymbol{h}}_k^T / \tau) \tilde{\boldsymbol{h}}_i}{\exp(\tilde{\boldsymbol{h}}_i \tilde{\boldsymbol{h}}_1^T / \tau) + \sum_{k \notin C(i)} \exp(\tilde{\boldsymbol{h}}_i \tilde{\boldsymbol{h}}_k^T / \tau)} \right] = 0 \quad (15)$$

As  $(\mathbf{x}_i, \mathbf{x}_1)$  is a true positive pair, both  $(\mathbf{x}_i, \mathbf{x}_k)$  and  $(\mathbf{x}_1, \mathbf{x}_k)$  are true negative pairs. According to Definition 3.3,  $\exp(\bar{\mathbf{h}}_i \bar{\mathbf{h}}_k^T / \tau) \approx 0$  and  $\exp(\bar{\mathbf{h}}_i \bar{\mathbf{h}}_k^T / \tau) \approx 0$  for some positive small values  $\tau$  and both two terms of Eq. 15 is 0. Thus, Eq. 15 holds.

Next, we want to show that if there exists one false positive sample, we reach a contradiction. Assuming that  $x_1$  is one false positive sample of  $x_i$  in the batch, then both  $(x_1, x_j)$  and  $(x_i, x_1)$  are false positive pairs and  $(x_1, x_k)$  is a false negative pair for some k (e. g.,  $x_1$  and  $x_k$  are from the same cluster for some k). Similarly, as  $(x_i, x_k)$  is a true negative pair for all k,  $\exp(\bar{h}_i \bar{h}_k^T / \tau) \approx 0$  and the second term of Eq. 15 is approximately 0. Therefore, we have

$$\sum_{\substack{j \in C(1), j \neq 1}} \frac{\sum_{k \notin C(1)} [\exp(\bar{h}_1 \bar{h}_k^T / \tau) \bar{h}_j - \exp(\bar{h}_1 \bar{h}_k^T / \tau) \bar{h}_k]}{\exp(\bar{h}_1 \bar{h}_j^T / \tau) + \sum_{k \notin C(1)} \exp(\bar{h}_1 \bar{h}_k^T / \tau)} + 0 = 0$$

$$\sum_{\substack{j \in C(1), j \neq 1}} \sum_{k \notin C(1)} [\exp(\bar{h}_1 \bar{h}_k^T / \tau) \bar{h}_j - \exp(\bar{h}_1 \bar{h}_k^T / \tau) \bar{h}_k] = 0 \quad (16)$$

We multiply Eq. 16 by  $\bar{h}_i^T$ , where  $(x_i, x_j)$  is a true positive pair for any *j* (i. e., both  $x_i$  and  $x_j$  are from the same cluster) and  $(x_i, x_k)$  is a true negative pair for any *k*. Then, we have

$$\sum_{j \in C(1), j \neq 1} \sum_{k \notin C(1)} [\exp(\bar{\boldsymbol{h}}_1 \bar{\boldsymbol{h}}_k^T / \tau) \bar{\boldsymbol{h}}_j \bar{\boldsymbol{h}}_i^T + \exp(\bar{\boldsymbol{h}}_1 \bar{\boldsymbol{h}}_k^T / \tau) (-\bar{\boldsymbol{h}}_k \bar{\boldsymbol{h}}_i^T)] = 0$$
(17)

Since  $(\mathbf{x}_i, \mathbf{x}_j)$  is a true positive pair and  $(\mathbf{x}_i, \mathbf{x}_k)$  is a true negative pair, we have  $\exp(\bar{\mathbf{h}}_j \bar{\mathbf{h}}_i^T / \tau) > 1$  and  $\exp(\bar{\mathbf{h}}_k \bar{\mathbf{h}}_i^T / \tau) \approx 0$ , which means that  $\bar{h}_j \bar{h}_i^T > 0$  and  $\bar{h}_k \bar{h}_i^T < 0$ . Therefore, both two terms of Eq. 17 are non-negative and Eq. 17 holds if and only if  $\exp(\bar{h}_1 \bar{h}_k^T / \tau) \approx 0$  for any  $k \notin C(1)$  (i. e.,  $(x_1, x_k)$  is a true negative pair for any  $k \notin C(1)$ ).

If  $(x_1, x_k)$  is a true negative pair for any  $k \notin C(1)$ , then  $(x_1, x_i)$  has to be a true positive pair, and it contradicts our assumption that  $x_1$  is a false positive sample of  $x_i$ . Therefore, we reach a contradiction and we could not get the optimal solution for  $\bar{h}_1$ , which completes *j*the proof.

#### **B** Experiments

–́۱

In this section, we show the details of generating anomalous node for the semi-supervised datasets, including IMDB and DBLP. We also show the hyper-parameter specification for reproducing the experimental results.

#### **B.1** Anomalous node generation

We follow the published works [14] to generate anomalous nodes by perturbing the topological structure or node attributes of an attributed network. To perturb the topological structure of an attributed network, we adopt the method introduced by [14] to generate some small cliques as in many real-world scenarios. A small clique is a typical anomalous substructure due to larger node degrees than normal nodes. By [1], after we specify the clique size as m, we randomly select *m* nodes from the network and then make those nodes fully connected. Then all the *m* nodes in the clique are regarded as anomalies. In addition to the injection of structural anomalies, we adopt another attribute perturbation schema introduced by [14] to generate anomalies from an attribute perspective. For each selected node  $u_i$ , we randomly pick another k nodes and select node  $u_i$ whose attributes deviate the most from node  $u_i$  among the k nodes by maximizing the Euclidean distance  $||x_i - x_j||^2$ . Afterward, we then change the attributes  $x_i$  of node  $u_i$  to  $x_j$ .

#### **B.2** Reproducibility

All of the real-world data sets are publicly available. The experiments are performed on a Windows machine with a 24GB RTX 4090 GPU. We use TAM as the backbone of our method to capture the local node affinity. We set the number of clusters to be 10,  $\lambda = 0.1$  and  $\alpha = 0.8$  for the CERT dataset. For the IMDB dataset, we set the number of clusters to be 10, the value of  $\lambda = 1$  and  $\alpha = 0.8$ . For the DBLP dataset, we set the number of clusters to be 10, the value of  $\lambda = 0.01$  and  $\alpha = 0.8$ . For the BlogCatalog dataset, we set the number of clusters to be 5,  $\lambda = 0.01$  and  $\alpha = 0.01$ . For the Amazon dataset, we set the number of clusters to be 10,  $\lambda = 1$  and  $\alpha = 0.8$ . For the Yelp dataset, we set the number of clusters to be 10,  $\lambda = 1$  and  $\alpha = 0.8$ .