



Self-supervised Graph Representation Learning for Black Market Account Detection

Zequan Xu*
School of Informatics
Xiamen University
Xiamen, Fujian, China
xuzequan@stu.xmu.edu.cn

Lianyun Li*
School of Informatics
Xiamen University
Xiamen, Fujian, China
lilianyun@stu.xmu.edu.cn

Hui Li†
School of Informatics
Xiamen University
Xiamen, Fujian, China
hui@xmu.edu.cn

Qihang Sun
Tencent Inc.
Guangzhou, Guangdong, China
aaronqhsun@tencent.com

Shaofeng Hu
Tencent Inc.
Guangzhou, Guangdong, China
hugohu@tencent.com

Rongrong Ji
School of Informatics
Xiamen University
Xiamen, Fujian, China
rrj@xmu.edu.cn

ABSTRACT

Nowadays, Multi-purpose Messaging Mobile App (MMA) has become increasingly prevalent. MMAs attract fraudsters and some cybercriminals provide support for frauds via black market accounts (BMAs). Compared to fraudsters, BMAs are not directly involved in frauds and are more difficult to detect. This paper illustrates our BMA detection system SGRL (Self-supervised Graph Representation Learning) used in WeChat, a representative MMA with over a billion users. We tailor Graph Neural Network and Graph Self-supervised Learning in SGRL for BMA detection. The workflow of SGRL contains a pretraining phase that utilizes structural information, node attribute information and available human knowledge, and a lightweight detection phase. In offline experiments, SGRL outperforms state-of-the-art methods by 16.06%-58.17% on offline evaluation measures. We deploy SGRL in the online environment to detect BMAs on the billion-scale WeChat graph, and it exceeds the alternative by 7.27% on the online evaluation measure. In conclusion, SGRL can alleviate label reliance, generalize well to unseen data, and effectively detect BMAs in WeChat.

CCS CONCEPTS

• **Computing methodologies** → **Anomaly detection**; • **Security and privacy** → **Social network security and privacy**.

KEYWORDS

black market account detection, multi-purpose messaging mobile app, self-supervised learning, graph neural network

*The first two authors contributed equally. Work done when they were interns at Tencent Inc.

†Corresponding author.

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.

WSDM '23, February 27-March 3, 2023, Singapore, Singapore

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9407-9/23/02...\$15.00

<https://doi.org/10.1145/3539597.3570466>

ACM Reference Format:

Zequan Xu, Lianyun Li, Hui Li, Qihang Sun, Shaofeng Hu, and Rongrong Ji. 2023. Self-supervised Graph Representation Learning for Black Market Account Detection. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining (WSDM '23)*, February 27-March 3, 2023, Singapore, Singapore. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3539597.3570466>

1 INTRODUCTION

Nowadays, one single mobile app tends to integrate multiple functionalities. Such Multi-purpose Messaging Mobile Apps (MMAs) bring convenience to mobile users. WeChat¹ is a representative MMA with over a billion users. One can use WeChat for text/voice messaging and voice/video calls, login to mobile games, or use the digital payment service to directly transfer money to other users.

The great convenience brought by WeChat also attracts cybercriminals [23]. They deceive normal users through socializing over WeChat. After gaining trust, cybercriminals recommend stocks or illegal gambling so that normal users transfer money to cybercriminals' WeChat accounts or reveal their passwords to cybercriminals. Since malicious socialization and payment take place in one single app, the time left for normal users to realize the fraud is limited. The massive profit brought by MMA frauds expedites the development of the *black market* that supports frauds. As depicted in Fig. 1, certain cybercriminals create WeChat accounts to sell bulk quantities of WeChat accounts, bank accounts, phone numbers, national identification numbers, devices, or IP proxies that can be used to perpetrate frauds. In this paper, *black market accounts* (BMAs) refer to such WeChat accounts that are not directly involved in frauds but provide support services for frauds. BMAs are the basis of MMA frauds. Detecting and banning BMAs can significantly increase the difficulty of perpetrating MMA frauds.

This paper illustrates our designed system for the *BMA Detection* task (the *BAD* task) in WeChat. Due to the social nature of WeChat, we construct a user-user interaction graph and conduct the *BAD* task. This way, the *BAD* task is closely related to Graph-based Anomaly Detection (GBAD). However, *the BAD task in WeChat has unique properties and it is more challenging than other GBAD tasks*:

¹<https://www.wechat.com/en>




Figure 1: The role of black market accounts in WeChat.

P1: Huge and diverse data. Massive users with diverse behaviors are using WeChat. It is a time-consuming and human-intensive process to manually identify BMA patterns. Even if a few patterns are found, using them in rule-based detection may not provide promising results: the diversity of WeChat data (e.g., interactions with some patterns only occur between certain users) makes it unlikely that these patterns are ubiquitous in the large graph.

P2: Demand for generalization. New users join WeChat every day. Hence, the detection system should adapt properly to unseen accounts. On the other hand, the large WeChat data requires a strong generalization ability: training should be offline on *small* data and later the system can be used online for detecting BMAs over the *large* WeChat graph to reduce the cost at detection time.

P3: Lack of labels and great difficulty of manual labeling due to users' privacy. Unlike financial fraud/risk detection that uses financial records after the amortization period as labels [22] or MMMA fraudster detection that uses victim reports as labels, labels for the BAD task are mostly *unavailable*. Moreover, since BMAs do not directly get involved in frauds, manual labeling is *difficult*, let alone its high cost: (1) The behavior types (e.g., chat and transactions) and behavior frequency of some BMAs, are not much different compared to other accounts. (2) Private information like chat content may help increase the accuracy of BMA labeling, but it is inaccessible in order to protect user privacy.

To combat with BMAs in WeChat, we develop SGRL, a Self-supervised Graph Representation Learning based BMA detection system. We summarize our contributions as follows:

C1: An Inductive Graph Neural Network (GNN) encoder for BMA representation learning. To cope with P1 and P2, we adopt an Inductive GNN encoder (IG-Encoder) as the backbone of SGRL to automatically capture rich information in WeChat graph and reduce the burden of BMA pattern investigation.

C2: Novel structural/attribute encoders based on Graph Self-supervised Learning (GSSL). To overcome P3, we explore GSSL [12] to alleviate the reliance on BMA labels and design two types of encoders for modeling structures and node attributes, two essential parts in graph data, respectively:

- **C2.1: Contrastive structural encoders.** We observe that the surrounding structure (i.e., subgraph) of nodes is important in the BAD task. Thus, we design two contrastive structural encoder SSS-Encoder and SS-Encoder to encode the subgraph as a context into node representation.
- **C2.2: A self-supervised node attribute encoder.** Some node attributes are highly correlated with node labels. Hence, we design a node attribute encoder (SSA-Encoder) and use attributes as supervision. It captures attribute correlations through our

designed Graph-based Multi-label prediction task with Missing Labels (GMML), which benefits the downstream BAD task.

GSSL also helps SGRL handle P2: representation learning and detection are decoupled in SGRL, and offline trained SGRL on small data can be used for the online BAD task on the large graph.

C3: An industrial BMA detection system deployed in WeChat. We deploy SGRL in WeChat with over a billion users. Offline and online experiments based on the large-scale MMMA data of WeChat show that SGRL has better detection performance than existing methods and it can block the supports from BMAs for frauds.

2 OUR SYSTEM SGRL

2.1 Overview

We first construct a user-user interaction graph and extract node features from WeChat (Sec. 2.2). The graph is used to detect BMAs. As depicted in Fig. 2, the workflow of SGRL consists of two phases: pretraining and detection. The core part of SGRL contains four encoders: SSS-Encoder, SS-Encoder, SSA-Encoder and IG-Encoder_{detect}. They are constructed based on an inductive GNN encoder IG-Encoder (Sec. 2.3). SSS-Encoder and SS-Encoder are contrastive structural encoders (Sec. 2.4) that capture the structural information via maximizing the mutual information between a node and its surrounding subgraph. Through GMML, SSA-Encoder (Sec. 2.5) models attribute correlations that play a pivotal role in predicting BMAs. We train the four encoders simultaneously so that they can embed nodes into high-quality representations which are later used for detecting BMAs in the lightweight detection phase (Sec. 2.6).

2.2 Graph Construction and Feature Extraction

BMA detection should not violate users' privacy. Hence, SGRL cannot access private information like users' chat content (text, video or speech) in WeChat. *The data we used is chosen through a strict investigation process in order to protect users' privacy.*

We construct a user-user interaction graph where each node represents a WeChat account and each edge indicates that the connected nodes have some relations (e.g., they are in the same chat group, they are direct friends, or they have direct transactions). The graph is homogeneous and multiple relations between two nodes, if they exist, are merged as one edge. A more complex design like a heterogeneous graph [29] that distinguishes different node types and edge types can be used. However, using a homogeneous graph significantly reduces the number of nodes and edges, which is important for processing billion-scale accounts. In the meantime, in our experiments, we observe that using a homogeneous graph already provides satisfactory detection performance.

For each node, we extract a 7-dimensional binary vector where 0/1 in each dimension indicates the existence/non-existence of an attribute. Node attributes are designed by human experts and they may provide clues for detecting BMAs.

2.3 Inductive Graph Neural Network Encoder

We design an Inductive Graph Neural Network Encoder (IG-Encoder) as the backbone in SGRL. Note that the choice of the core GNN encoder is orthogonal to the design of SSS-Encoder, SS-Encoder and

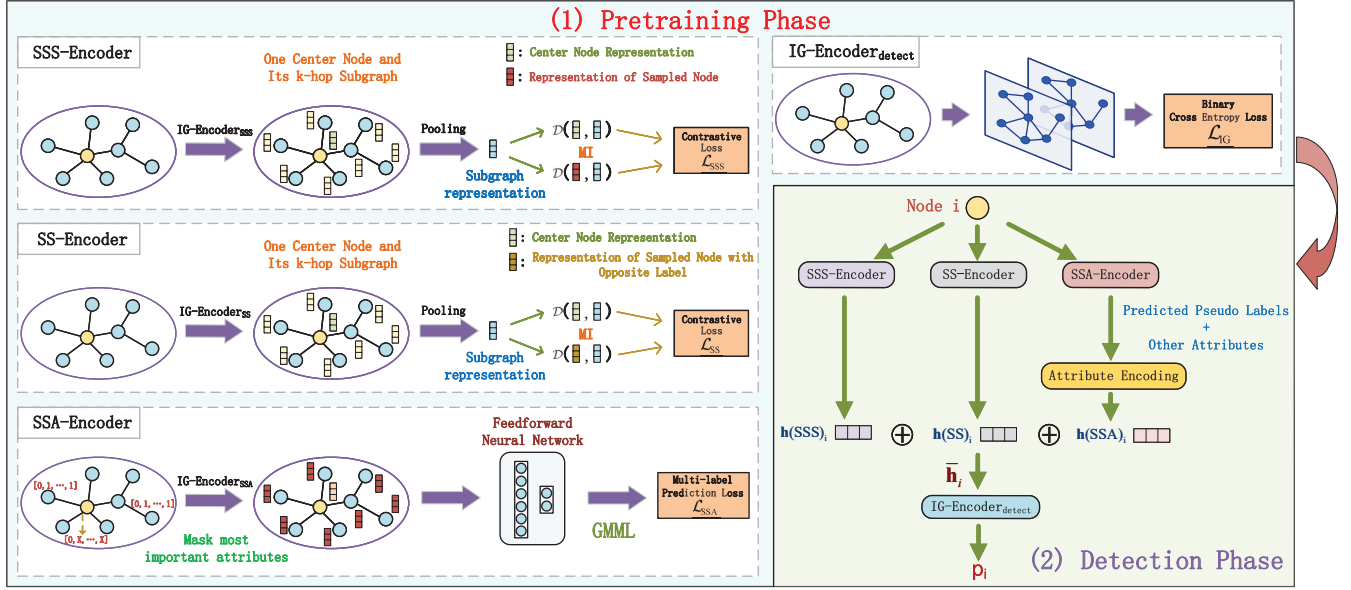


Figure 2: Overview of SGRL. It consists of a pretraining phase and a detection phase.

SSA-Encoder in SGRL, and other GNNs can be used to replace IG-Encoder. However, in experiments (Sec. 4.2.1), IG-Encoder shows much better performance than other GNNs for the BAD tasks.

Recall that each node i is associated with a binary attribute vector $\mathbf{a}_i \in \mathbb{R}^7$. We use an *attribute encoding mechanism* to encode node attributes. For each attribute in \mathbf{a}_i , IG-Encoder maintains two f -dimensional feature vectors initialized from $\mathcal{N}(0, 1)$ for its existence and non-existence. Thus, for each attribute vector \mathbf{a}_i , IG-Encoder has 7 feature vectors and we concatenate them to construct the initial representation vector $\mathbf{h}_i^{(0)} \in \mathbb{R}^{7f}$ of the node i .

We adopt the idea of spatial GNNs [21] for IG-Encoder. Firstly, information from neighboring nodes is aggregated:

$$\begin{aligned} \mathbf{x}_{N_i}^{(t+1)} &= \text{mean}(\mathbf{h}_{j_1}^{(t)}, \dots, \mathbf{h}_{j_s}^{(t)}) & \mathbf{h}_{N_i}^{(t+1)} &= \mathbf{x}_{N_i}^{(t+1)} \oplus \mathbf{y}_{N_i}^{(t+1)} \oplus \mathbf{z}_{N_i}^{(t+1)} \\ \mathbf{y}_{N_i}^{(t+1)} &= \max(\mathbf{h}_{j_1}^{(t)}, \dots, \mathbf{h}_{j_s}^{(t)}) & \mathbf{g}_{N_i}^{(t+1)} &= \mathbf{W}_g \mathbf{h}_{N_i}^{(t+1)} + \mathbf{b}_g \\ \mathbf{z}_{N_i}^{(t+1)} &= \text{sum}(\mathbf{h}_{j_1}^{(t)}, \dots, \mathbf{h}_{j_s}^{(t)}) \end{aligned} \quad (1)$$

where the superscript “ (t) ” indicates t -th iteration, \oplus is the concatenation operation, N_i denotes the neighbor set of node i and $j_s \in N_i$. $\text{mean}(\cdot)$, $\max(\cdot)$ and $\text{sum}(\cdot)$ are average pooling, max pooling and sum pooling, respectively. \mathbf{W}_g and $\mathbf{b}_g \in \mathbb{R}^f$ are learnable parameters. We use two different \mathbf{W}_g ($\mathbf{W}_g \in \mathbb{R}^{f_1 \times f_2}$ for $\mathbf{h}_{N_i}^{(1)}$ and $\mathbf{W}_g \in \mathbb{R}^{f_1 \times f_1}$ for $\mathbf{h}_{N_i}^{(t)}$ ($t > 1$)) and two different \mathbf{b}_g for different iterations in Eq. 1 since $\mathbf{h}_{N_i}^{(1)}$ and $\mathbf{h}_{N_i}^{(t)}$ ($t > 1$) have different shapes (See Appendix A in our online report²). We do not employ any attention mechanism that considers the different importance of neighbor nodes in neighbor aggregation since the number of nodes in WeChat graph is huge and attention mechanisms will incur additional overhead.

Then, IG-Encoder adds a self-connection to each node so that the information from the original node attributes will not vanish

during the message passing procedure:

$$\mathbf{s}_i^{(t+1)} = \mathbf{W}_s \mathbf{h}_i^{(0)} + \mathbf{b}_s, \quad \mathbf{e}_i^{(t+1)} = \text{RELU}(\mathbf{g}_{N_i}^{(t+1)} \oplus \mathbf{s}_i^{(t+1)}), \quad (2)$$

where $\mathbf{W}_s \in \mathbb{R}^{f_1 \times f_2}$ and $\mathbf{b}_s \in \mathbb{R}^f$ are learnable parameters, and $\text{RELU}(\cdot)$ is the Rectified Linear Unit. Although \mathbf{W}_s and \mathbf{b}_s are updated during training, IG-Encoder always uses the initial representation vector $\mathbf{h}_i^{(0)}$ in Eq. 2 to retain the impact of the original node attributes that come from the knowledge of human experts.

After that, IG-Encoder passes \mathbf{e}_i to a feedforward neural network followed by an L_2 normalization:

$$\mathbf{q}_i^{(t+1)} = \text{RELU}(\mathbf{W}_q \mathbf{e}_i^{(t+1)} + \mathbf{b}_q), \quad \mathbf{h}_i^{(t+1)} = \mathbf{q}_i^{(t+1)} / \|\mathbf{q}_i^{(t+1)}\|, \quad (3)$$

where $\mathbf{W}_q \in \mathbb{R}^{f_1 \times 2f_1}$ and $\mathbf{b}_q \in \mathbb{R}^{f_1}$ are learnable parameters.

IG-Encoder stacks two of the above GNN layers (i.e., Eqs. 1, 2 and 3) to extract the representation of node i . The output representation $\mathbf{h}_i^{(t+1)}$ is fed into a prediction component with a feedforward neural network to estimate the suspicious score of node i :

$$p_i^{(t+1)} = \sigma \mathbf{w}_{p(3)}^T (\mathbf{W}_{p(2)} (\mathbf{W}_{p(1)} \cdot \mathbf{h}_i^{(t+1)} + \mathbf{b}_{p(1)}) + \mathbf{b}_{p(2)}) + \mathbf{b}_{p(3)}, \quad (4)$$

where $\mathbf{w}_{p(3)} \in \mathbb{R}^{f_1}$, $\mathbf{W}_{p(2)} \in \mathbb{R}^{f_1 \times 2f_1}$, $\mathbf{W}_{p(1)} \in \mathbb{R}^{2f_1 \times f_1}$, $\mathbf{b}_{p(1)} \in \mathbb{R}^{2f_1}$, $\mathbf{b}_{p(2)} \in \mathbb{R}^{f_1}$ and $\mathbf{b}_{p(3)}$ are learnable parameters, and $\sigma(\cdot)$ is the sigmoid function. If p_i is larger than a pre-defined threshold ρ , node i will be labeled as a BMA.

IG-Encoder can be optimized with a standard binary cross entropy loss (denoted as \mathcal{L}_{IG} in Fig. 2) over *labeled* nodes (in the training data) for the BAD task. However, IG-Encoder does not learn node embeddings and use them for the BAD task. Instead, the global $7f$ -dimensional feature vectors for node attributes and the parameters of IG-Encoder are updated and later used to produce representations for new nodes during detection. Hence, graph representation learning based on IG-Encoder is indeed *inductive* (i.e., the trained model can be used for unseen nodes), which is essential

²Our online report is available at: <https://arxiv.org/abs/2212.02679>.




Figure 3: Two exemplifying surrounding structures: (1) Red node: BMA. (2) Blue node: Account for transactions. (3) Gray node: Account for delivering services. (4) Green node: Camouflage account. (5) Yellow node: Fraudster (buyer).

for the BAD task. SGRL, which uses IG-Encoder as its backbone for representation learning, is therefore also able to generate representations for unseen nodes.

2.4 Contrastive Structural Encoders

Through manual check, we find that surrounding structures of nodes are useful in the BAD task. Fig. 3 shows that, to avoid being easily detected, black market sellers manipulate several accounts for different purposes like camouflage, attracting fraudsters, transactions, delivering support services, etc. However, these accounts are controlled by the same person/gang, and their surroundings often display similar characteristics. For example, for the purpose of transactions, suspect accounts with similar attributes (blue nodes in Fig. 3) appear in the surrounding structures of various BMAs.

Although it is difficult to leverage manual rules of surrounding patterns, as discussed in P1 of Sec. 1, the surrounding structures (i.e., subgraphs) indeed have rich clues for detecting BMAs. Hence, we believe it is beneficial to maximize the correlation between a node and its subgraph and encode the subgraph as a context into node representation. Along this direction, we design contrastive structural encoders to maximize the *local* mutual information (MI) between the *subgraph-level* summary s_i and the *node-level* representation h_i of the center node i via graph contrastive learning [12].

We first use IG-Encoder to encode each node into node-level representation h . Then, for building contrastive structural encoders, two issues remain: (1) How do we define the surrounding structure of a node? (2) How to retrieve subgraph-level summary?

For the first issue, an elaborate way is to design representative subgraph schemas. For instance, human experts can design subgraph schemas according to structural features (e.g., node degree or graph density) and rule-based features (e.g., frequency of transactions). Apparently, this is costly for a large graph. Moreover, as new data comes into the graph and cybercriminals are trying different strategies to escape the detection, schemas may soon become invalid. Hence, we choose a simple yet effective design: use the complete k -hop neighbors of the center node as its subgraph. There are other sophisticated approaches to automatically retrieve subgraphs, e.g., use neighbors having largest personalized pagerank values w.r.t. the center node [7], or conduct random walks starting from the center node [14]. But they also introduce randomness: the subgraph for the same center node may vary at each iteration, making preprocessing subgraphs difficult. Differently, using complete k -hop neighbors can significantly reduce the time to retrieve a subgraph since they are fixed and can be preprocessed.

For the second issue, given a k -hop subgraph, we deploy a *read-out* function $\mathcal{R} : \mathbb{R}^{n \times f_1} \rightarrow \mathbb{R}^{f_1}$, where n is the number of nodes in the subgraph, to summarize f_1 -dimensional subgraph-level representations. Here, we apply average pooling as the readout function: $s_i = \frac{1}{|S_i|} \sum_{j \in S_i} h_j$, where S_i is the set of all nodes in the k -hop subgraph of node i (including i) and $|S_i|$ indicates the number of nodes in S_i . Note that, from now, we omit the superscript “(t)” for simplicity. Attention mechanisms can be employed in the readout function to distinguish different importance of nodes. But they typically require additional parameters and costly operations (e.g., the Softmax function when inputs contain many elements) that will incur high overhead on the large WeChat graph. Thus, we do not incorporate an attention mechanism in current SGRL.

Then, we adopt a self-supervised MI objective for maximizing MI between the subgraph-level summary representation and the center node representation:

$$\mathcal{L}_{MI} = \sum_{i=1}^{N_{pos}} (\mathbb{E}_{pos} [\log (\mathcal{D}(h_i, s_i))] + \mathbb{E}_{neg} [\log (1 - \mathcal{D}(h_{i'}, s_i))]), \quad (5)$$

where N_{pos} is the number of nodes in the positive node set, $h_{i'}$ is the representation of the negative node i' w.r.t. to i (we will discuss how to generate negative nodes later), and $\mathcal{D}(h, s)$ is a discriminator that estimates the divergence and assigns probability to the node-subgraph representation pair $\langle h, s \rangle$. We use a feedforward neural network as the discriminator but other designs can also be applied:

$$\mathcal{D}(h_i, s_i) = \sigma \left(\mathbf{w}_{\mathcal{D}}^{(2)T} (\mathbf{W}_{\mathcal{D}}^{(1)} \cdot (h_i \oplus s_i) + \mathbf{b}_{\mathcal{D}}^{(1)}) + b_{\mathcal{D}}^{(2)} \right), \quad (6)$$

where $\mathbf{w}_{\mathcal{D}}^{(2)} \in \mathbb{R}^{f_1}$, $\mathbf{W}_{\mathcal{D}}^{(1)} \in \mathbb{R}^{f_1 \times 2f_1}$, $\mathbf{b}_{\mathcal{D}}^{(1)} \in \mathbb{R}^{f_1}$ and $b_{\mathcal{D}}^{(2)}$ are learnable parameters. Eq. 5 is a noise-contrastive type objective with a standard binary cross-entropy loss. Optimizing Eq. 5 maximizes MI between h_i and s_i based on the Jensen-Shannon divergence between the joint distribution (positive samples) and the product of marginals (negative samples) [5, 18].

Depending on how positive and negative node-subgraph pairs are generated, two contrastive structural encoders are used in SGRL:

(1) SSS-Encoder, short for Self-supervised Contrastive Structural Encoder, is designed for the complete self-supervised setting without any labels. SSS-Encoder treats each node i and its subgraph S_i as a positive pair $\langle i, S_i \rangle$. For each $\langle i, S_i \rangle$, SSS-Encoder randomly selects another node i' , and use i' and S_i as the negative pair $\langle i', S_i \rangle$. For SSS-Encoder, each node i is a positive node. In the meantime, i can be randomly selected as a negative node when SSS-Encoder encodes other nodes in the graph.

(2) SS-Encoder, short for Supervised Contrastive Structural Encoder, is designed for the supervised setting, where a handful of labels are available. SS-Encoder uses each BMA i and its subgraph S_i as a positive node-subgraph pair $\langle i, S_i \rangle$. For each $\langle i, S_i \rangle$, SS-Encoder randomly selects another node i' from non-BMA set and use i' and S_i as the negative pair $\langle i', S_i \rangle$. BMAs and non-BMAs are labeled by human experts.

For both encoders, negative pairs are randomly regenerated at each iteration. IG-Encoder used in SSS-Encoder (SS-Encoder) for encoding nodes is denoted as IG-Encoder_{SSS} (IG-Encoder_{SS}).

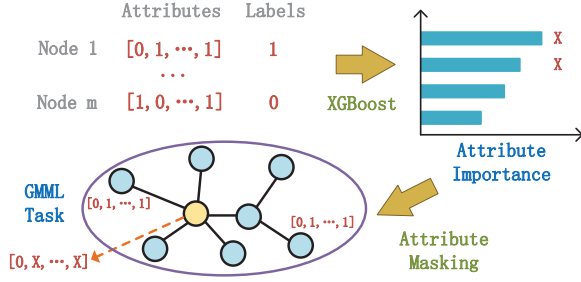


Figure 4: The workflow of SSA-Encoder.

in Fig. 2. SSS-Encoder (SS-Encoder) is optimized using the loss function in Eq. 5 and we denote it as \mathcal{L}_{SSS} (\mathcal{L}_{SS}) in Fig. 2. The output representation for node i from SSS-Encoder (SS-Encoder) is indicated as $\mathbf{h}(\text{SSS})_i$ ($\mathbf{h}(\text{SS})_i$) in Fig. 2. We design two related but different contrastive structural encoders so that human experts are allowed to get involved in the detection. Although it is difficult to manually label most nodes, a small number of labeled nodes can be used as supervision in SS-Encoder to improve representation learning. And we can observe from experiments (Sec. 4.2.2) that using two encoders together can achieve better performance than solely using either SSS-Encoder or SS-Encoder.

2.5 Self-supervised Node Attribute Encoder

We further design a Self-supervised Node Attribute Encoder (SSA-Encoder). Although most labels are unavailable, we can use node attributes as the attribute-level supervision to alleviate the reliance on the label-level supervision. The intuition is that improving the prediction accuracy of some node attributes (known), which are strongly correlated with node labels (unknown), can help the attribute encoder capture attribute correlations that benefit the downstream BAD task. Fig. 4 depicts the workflow of SSA-Encoder.

Specifically, we first adopt the non-GNN based method XGBoost [3], a gradient boosting algorithm that only relies on data attributes and shows promising results in various prediction tasks, and train it on a small number of labeled nodes to predict node labels. After training, XGBoost is able to estimate the importance of each attribute to label prediction. Then, we pick two most crucial attributes to the prediction and treat them as *pseudo labels* of each node. The reason for using more than one attribute as pseudo labels is that we find no attribute has a dominated importance value. On the other hand, including more attributes as pseudo labels may introduce noises since not all attributes have high importance values. Hence, we choose the two most crucial attributes as pseudo labels to balance both sides.

After that, we formulate a Graph-based Multi-label prediction task with Missing Labels (GMML for short) as a self-supervised learning task for SSA-Encoder:

DEFINITION 1 (GMML). Given partial attributes of the target node i , surrounding nodes in the k -hop subgraph of i and their complete node attributes, predict two pseudo labels of node i .

Multi-label learning with missing labels is well-studied for multi-label image classification [11], but it receives less attention in graph-relevant tasks. SSA-Encoder processes all nodes in the graph in each iteration. For each node i , when SSA-Encoder predicts its missing labels, we mask two dimensions of its attributes that correspond

to pseudo labels as missing labels. When SSA-Encoder predicts missing labels of other nodes that have k -hop subgraphs containing i , all node attributes of i are known to SSA-Encoder.

SSA-Encoder uses an IG-Encoder (i.e., IG-Encoder_{SSA}) without the prediction component (Eq. 4) to encode a node i into representation $\hat{\mathbf{h}}(\text{SSA})_i \in \mathbb{R}^{f_1}$ and then predicts pseudo labels of i :

$$\mathbf{r}_i = \sigma(\mathbf{W}_{\mathbf{r}(2)}(\mathbf{W}_{\mathbf{r}(1)} \cdot \hat{\mathbf{h}}(\text{SSA})_i + \mathbf{b}_{\mathbf{r}(1)}) + \mathbf{b}_{\mathbf{r}(2)}), \quad (7)$$

where $\mathbf{r}_i \in \mathbb{R}^2$ are the predicted probabilities of two pseudo labels for node i being 1. $\mathbf{W}_{\mathbf{r}(2)} \in \mathbb{R}^{2 \times \frac{f_1}{2}}$, $\mathbf{W}_{\mathbf{r}(1)} \in \mathbb{R}^{\frac{f_1}{2} \times f_1}$ and $\mathbf{b}_{\mathbf{r}(1)} \in \mathbb{R}^{\frac{f_1}{2}}$ and $\mathbf{b}_{\mathbf{r}(2)} \in \mathbb{R}^2$ are learnable weights.

We utilize a multi-label prediction loss for optimizing SSA-Encoder:

$$\mathcal{L}_{SSA} = -\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^2 [y_{ij} \log(r_{ij}) + (1 - y_{ij}) \log(1 - r_{ij})], \quad (8)$$

where N is the number of nodes in the graph, y_{ij} indicates the j -th pseudo label of node i , and r_{ij} is the predicted probability for the j -th pseudo label of node i being 1.

Compared to GSSL-based methods [6] that randomly mask node features and then recover, SSA-Encoder predicts attributes that are crucial to node label prediction. Hence, it “indirectly” links attribute prediction to node label prediction via attribute importance.

2.6 Putting All Together

Overall, detecting BMAs using SGRL involves two phases:

Pretraining Phase. We pretrain SSS-Encoder, SS-Encoder and SSA-Encoder independently on the training data with limited labels to capture the intrinsic properties of WeChat graph from different aspects. An IG-Encoder (i.e., IG-Encoder_{detect}) used later in detection is also trained over the limited training labels. All encoders can be optimized using gradient descent based optimization methods and we adopt Adam. Note that IG-Encoder_{SSS}, IG-Encoder_{SS}, IG-Encoder_{SSA} and IG-Encoder_{detect} do not share parameters.

Detection Phase. The four pretrained encoders are used for the *lightweight* detection on the test data. For a target node i (it may not exist in the training data), SSS-Encoder and SS-Encoder generate two node representations $\mathbf{h}(\text{SSS})_i \in \mathbb{R}^{f_1}$ and $\mathbf{h}(\text{SS})_i \in \mathbb{R}^{f_1}$. We then adopt SSA-Encoder to predict pseudo labels for i . The predictions are used to replace the two attributes in \mathbf{a}_i that correspond to pseudo labels to construct another attribute vector \mathbf{a}'_i for i . After that, we utilize the attribute encoding mechanism introduced in Sec. 2.3 with feature vectors in IG-Encoder_{detect} to encode \mathbf{a}'_i into $\mathbf{h}(\text{SSA})_i \in \mathbb{R}^{7f}$. Finally, the concatenation of $\mathbf{h}(\text{SSS})_i$ and $\mathbf{h}(\text{SS})_i$ and $\mathbf{h}(\text{SSA})_i$ is used as the initial representation for node i and fed into IG-Encoder_{detect} to estimate the suspicious score of node i :

$$\tilde{\mathbf{h}}_i^{(0)} = \mathbf{h}(\text{SSS})_i \oplus \mathbf{h}(\text{SS})_i \oplus \mathbf{h}(\text{SSA})_i, \quad (9)$$

where $\tilde{\mathbf{h}}_i^{(0)} \in \mathbb{R}^{2f_1+7f}$. In IG-Encoder_{detect}, shapes of \mathbf{W}_g and \mathbf{b}_g in Eq. 1 are modified according to the shape of $\tilde{\mathbf{h}}_i^{(0)}$ (See Appendix A in our online report).

The above method that replaces true attributes with predicted attributes from SSA-Encoder outperforms using original attributes (see Sec. 4.2.7). Recall that the *direct* optimization objective of SSA-Encoder is to capture the correlation between partial attributes

and pseudo labels. Since we choose two most crucial attributes to node label prediction as pseudo labels, predicting pseudo labels accurately can *indirectly* increase the prediction accuracy of node labels (the ultimate goal of SGRL). Nevertheless, for some nodes, the importance values of their pseudo labels do not dominate those of other attributes, i.e., the original attribute values for the two pseudo labels are not strongly connected to node labels and they are not “good” attributes for the BAD task. For such cases, SSA-Encoder can give predictions that are not the ground-truth pseudo labels, but they follow captured attribute correlations which can lead SSA-Encoder to better benefit BMA detection.

3 DISCUSSION

In this section, we provide some discussions about SGRL:

Model Size. We provide the detailed analysis in Appendix B in our online report. The model size S_{IG} of IG-Encoder is in the magnitude of $f_1 f_2$ where f_1 and f_2 are up to the order of tens (setting $f_1 = 32$ and $f_2 = 56$ in our experiments brings sufficiently good performance). Note that prevalent GNN-based methods involve at least one weight matrix in the neural network and their model size is at least $f_1 f_2$. The model size S_{SGRL} of the complete SGRL, without considering XGBoost, is four to five times larger than S_{IG} . Nevertheless, S_{SGRL} is still in the magnitude of $f_1 f_2$. In practice, our BMA detection system is trained offline for the pretraining phase on small, million-scale data and then its detection phase is performed online on the large WeChat graph. The additional parameter cost of SGRL, compared to using IG-Encoder only, is acceptable for the offline pretraining phase where training time and storage cost are not strictly limited. The online lightweight detection phase does not involve parameter update. Considering that SGRL achieves much better performance than other methods (Sec. 4), the model size of SGRL is acceptable. In our experiments (Sec. 4.3), we also show that SGRL can be deployed on billion-scale WeChat graph.

Relation to DGI and DCI. The idea of SGRL is closely related to Deep Graph Infomax (DGI) [18] and Deep Cluster Infomax (DCI) [19]. DGI contrasts the whole graph with a node (i.e., *global* MI) so that the global information can be embedded into node representation. When deploying on a large-scale graph like WeChat, summarizing the complete graph brings a global context with much noise. As discussed in Sec. 2.4, the local subgraph provides rich clues and SGRL contrasts the subgraph with the node in it. In addition, SGRL maximizes local MI instead of global MI, removing the obstacle to processing a billion-scale graph where summarizing the global graph-level representation has high cost. DCI can be viewed as a cluster-based DGI [19] as it contrasts a cluster obtained from K-Means with the node in it. Although DCI reduces the context level from the complete graph to a cluster, a cluster in a large-scale graph is still too large to reveal BMA patterns. Compared to DCI, SGRL not only uses local surrounding structures in contrastive learning (SSS-Encoder and SS-Encoder) but also captures attribute-level correlations to benefit downstream BAD task (SSA-Encoder).

Deployment challenge and design tradeoffs. Due to the large volume of MMMA data in WeChat, the greatest deployment challenge is how to achieve good detection performance and avoid high overhead. SGRL is currently retrained and updated on a daily

basis to allow sufficient training time. Other design tradeoffs for reducing cost, which have been discussed in Sec. 2, include using a homogeneous graph instead of a heterogeneous graph (Sec. 2.2), avoiding using attention mechanisms (Secs. 2.3 and 2.4), using complete k -hop neighbors instead of sophisticated subgraph schemas (Sec. 2.4) and using two pseudo labels (Sec. 2.5).

4 EXPERIMENTS

4.1 Experiment Setup

4.1.1 Data. In the *offline test*, we use a million-scale dataset sampled from WeChat. The graph constructed on this dataset has 4 million nodes and 13.4 million edges. 100,000 nodes are manually labeled by human experts: half of them are BMAs and the remaining nodes are non-BMAs. The labels for other 3.9 million nodes are unknown. Note that we also investigate the impact of the number of labels in our experiments (Sec. 4.2.3). We randomly divide labels into 80%/20% for training/test in the offline test. In the *online test* on the billion-scale WeChat graph, the SGRL pretrained on the training set of the offline data is directly deployed for detection.

4.1.2 Baselines. We compare SGRL and its backbone IG-Encoder with the following state-of-the-art methods (details are provided in Appendix C in our online report):

- (1) **Non-GNN classification methods:** XGBoost [3] and MLP.
- (2) **GNN-based representation learning:** GCN [9], GAT [17] and GeniePath [13].
- (3) **Self-supervised graph representation learning methods:** IG-Encoder_{MTL} [26], IG-Encoder_{M3S} [15] and Deep Graph Infomax (DGI) [18]. IG-Encoder_{MTL} enhances IG-Encoder using node clustering as additional self-supervised task and optimizes the encoder in a multi-task learning framework [26]. IG-Encoder_{M3S} leverages the self-training method [15] to improve IG-Encoder.
- (4) **Self-supervised graph anomaly detection method:** Deep Cluster Infomax (DCI) [19]. Similar to IG-Encoder_{MTL}, DCI adopts node clustering as a self-supervised task. It also decouples representation learning phase and detection phase as SGRL.

We modify MLP and GCN, which are originally transductive, with the attribute encoding mechanism used in IG-Encoder so that they can generalize to new data. In the reported results, we use abbreviations “XGB”, “IG”, “SSS”, “SS” and “SSA” to indicate XGBoost, IG-Encoder, SSS-Encoder, SS-Encoder and SSA-Encoder, respectively. Other abbreviations are explained when they are used.

4.1.3 Hyper-parameters and other settings. Details are provided in Appendix A in our online report. By default, we use 1-hop subgraph in SSS-Encoder and SS-Encoder. We test several values for the threshold ρ used for judging whether the probability p output in Eq. 4 or other prediction layers in baselines indicates a BMA or not, and the threshold r for predicting binary labels in SSA-Encoder according to r output in Eq. 7. Finally, we set $\rho = 0.5$ and $r = 0.5$ which are sufficient for all methods to work well. For a fair comparison, we set $f = 64$ for the attribute encoding in all methods in the detection phase. We use the Adam optimizer for all methods if applicable. *We tune hyper-parameters of all methods so that they achieve good results.*

Table 1: Offline performance. Results of SGRL and best baselines are shown in bold. Percentages indicate improvements of SGRL over best baselines.

Model	AUC	ACC	KS	Recall	Precision	F1-score
XGB	0.6843	0.6605	0.3253	0.6342	0.6749	0.6540
MLP	0.6851	0.6623	0.3260	0.6612	0.6678	0.6645
GCN	0.7197	0.6607	0.3229	0.6272	0.6778	0.6516
GAT	0.7514	0.6897	0.3802	0.6709	0.7022	0.6862
GeniePath	0.7738	0.7014	0.4100	0.7123	0.7018	0.7070
IG	0.7934	0.7223	0.4331	0.6665	0.7153	0.6901
IG _{MTL}	0.7962	0.7241	0.4570	0.6550	0.7429	0.6962
IG _{M3S}	0.7949	0.7243	0.4592	0.7154	0.7006	0.7080
DGI	0.7897	0.6900	0.3818	0.6940	0.7098	0.7018
DCI	0.8072	0.6987	0.3991	0.7138	0.7172	0.7155
SGRL	0.9402 ↑ 16.48%	0.8606 ↑ 18.82%	0.7263 ↑ 58.17%	0.8622 ↑ 20.52%	0.8622 ↑ 16.06%	0.8622 ↑ 20.50%

Table 2: Offline Performance of different variations of SGRL. Best performance is shown in bold.

Model	AUC	ACC	KS	Recall	Precision	F1-score
IG	0.7934	0.7223	0.4331	0.6665	0.7153	0.6901
SSS	0.8754	0.7892	0.5790	0.7673	0.8064	0.7864
SSS _{nn}	0.8070	0.7330	0.4680	0.6945	0.7575	0.7246
SS	0.8521	0.7695	0.5427	0.7254	0.8002	0.7610
SS ₍₂₎	0.7893	0.7194	0.4419	0.6570	0.7562	0.7031
SS _{sg}	0.7884	0.7153	0.4421	0.6175	0.7739	0.6869
SS _{ss}	0.8402	0.7586	0.5230	0.6786	0.8068	0.7371
SS _{nn}	0.7900	0.7194	0.4426	0.6350	0.7700	0.6960
SSA	0.8199	0.7412	0.4882	0.6421	0.8068	0.7151
SGRL _{SA}	0.9046	0.8143	0.6380	0.7833	0.8479	0.8143
SGRL _{S2}	0.9104	0.8268	0.6559	0.8117	0.8403	0.8258
SGRL	0.9402	0.8606	0.7263	0.8622	0.8622	0.8622

4.2 Offline Test

AUC, ACC (accuracy, i.e., the percentage of correctly predictions on test labels), KS (Kolmogorov-Smirnov statistic [8]), Recall, Precision and F1-score are used for offline evaluation.

4.2.1 Overall Performance. Tab. 1 reports results of different approaches in the offline test and we can observe that:

- (1) IG-Encoder significantly outperforms other GNN-based methods GCN, GAT and GeniePath. This observation has supported our decision of using IG-Encoder as the backbone of SGRL. SGRL exceeds IG-Encoder by a large margin. Thus the backbone IG-Encoder is not the only reason for SGRL’s remarkable performance.
- (2) Compared to IG-Encoder, IG-Encoder_{MTL} and IG-Encoder_{M3S} generally have better results, showing the effectiveness of recently proposed self-supervised reinforcements [15, 26] on the BAD task.
- (3) Compared to state-of-the-art self-supervised learning methods IG-Encoder_{MTL}, IG-Encoder_{M3S}, DGI and DCI, SGRL consistently shows noticeably superior performance.

4.2.2 Contributions of Each Encoder. We show the performance of different variations of SGRL in Tab. 2:

- (1) SSA in Tab. 2 indicates we uses SSA-Encoder and IG-Encoder_{detect} together for detection. We can see that all encoders (SSS, SS, SSA) in SGRL exceed original IG-Encoder. But using any single encoder does not perform well as using the complete SGRL.
- (2) In Tab. 2, SGRL_{S2} denotes that the output representations from original attribute encoding, SSS-Encoder and SS-Encoder are concatenated as node representations. We can find that using




Figure 5: F1-score of DGI, DCI and SGRL when reducing number of labels.




Figure 6: Representation visualization: (1) Red: BMAs. (2) Blue: non-BMAs.

- SSS-Encoder and SS-Encoder together (i.e., SGRL_{S2}) yields better results than only using either SSS-Encoder or SS-Encoder.
- (3) Incorporating SSA-Encoder into SGRL_{S2} (i.e., SGRL in Tab. 2) achieves better performance than solely using SGRL_{S2}.

Based on above observations, we can conclude that each encoder indeed contributes to the performance of SGRL.

4.2.3 Results Using Fewer or No Labels. Labels in the BAD task is scarce and the robustness of the detection methods, when facing fewer or no labels, is crucial to their practicability. We randomly mask labels in the offline dataset and derive two new datasets with same numbers of nodes and edges as the original dataset. One retains 60,000 node labels and the other retains 20,000 node labels. Labels are randomly divided into 80%/20% for training/test as the original dataset with 100,000 node labels. Fig. 5 reports F1-score of DGI, DCI and SGRL on three datasets. We can find that SGRL drops *insignificantly* as we reduce the number of labels and it consistently outperforms DGI and DCI. We also test SGRL when no labels are available. For such cases, SGRL degrades to SGRL_{SA} (i.e., the outputs from SSA-Encoder and SSS-Encoder are concatenated as node representations) reported in Tab. 2. We can see that SGRL_{SA} performs well without labels and it is even better than those baselines reported in Tab. 1 that uses labels. In summary, regardless of available labels, SGRL can consistently provide accurate predictions.

4.2.4 Visualization of Representations. We adopt t-SNE [16] to visualize node representations generated by SGRL for the test set. Fig. 6 depicts the result: most BMAs are distributed in the right while other non-BMAs are located in the left and middle. Representations of different accounts have a clear distinction, showing that SGRL produces high-quality node representations for the BAD task.

4.2.5 Choice of MI. We choose to maximize the local MI between a node and its subgraph in SSS-Encoder and SS-Encoder. There are other options for the MI. In the following, “ \leftrightarrow ” indicates maximizing MI between the connected two objects. In Tab. 2, we show performance of other possible choices for MI:

- **subgraph \leftrightarrow global graph (SS_{sg}):** We use the subgraph representation for a node i and the global graph as a positive pair, and the subgraph representation for a randomly chosen node i' with opposite label to i and the global graph as the negative pair in Eq. 5 for SS-Encoder. Apparently, the global graph contains too much noise and it is hard to tell whether the MI between a positive subgraph and the global graph should be maximized. Therefore, this option shows worst performance for SS-Encoder.
- **subgraph \leftrightarrow subgraph (SS_{ss}):** We maximize MI between subgraphs of nodes with same labels in SS-Encoder. According to

Tab. 2, this alternative design ranks the second among all SS-Encoder variations, showing the importance of subgraphs in the BAD task. For SSS-Encoder, we do not know a way to define the positive subgraph pairs since node labels are unknown. Hence, we do not experiment SSS_{SS}.

- **node↔node (SSS_{nn} and SS_{nn}):** We maximize MI between two randomly picked nodes (SSS-Encoder) or nodes with same labels (SS-Encoder). These variations show much worse performance compared to default SSS-Encoder and SS-Encoder. This observation, again, shows the importance of subgraphs.

As a summary, our design that maximizes the local MI between a node and its subgraph is most appropriate for SGRL.

4.2.6 Choice of k for Subgraph. By default, we use 1-hop subgraphs ($k = 1$). We also report the result of SS-Encoder for $k = 2$, i.e., SS₍₂₎ in Tab. 2: it noticeably degrades the performance of SS-Encoder. A possible reason is that larger subgraphs introduce noise and first-order neighbors are more important in subgraphs. Since we do not use any attention mechanism to distinguish different importance of nodes in subgraphs, using larger subgraphs degrades the performance. Besides, larger k means longer time and larger memory cost for the pooling operation over subgraphs. When setting $k = 2$, the cost is unaffordable and that is why we do not report SSS₍₂₎.

4.2.7 Inputs to Pretrained SSA-Encoder. By comparing SSA and IG (it is fed with original node attributes) in Tab. 2, we can see that SSA-Encoder show better performance. As mentioned in Sec. 2.6, SSA-Encoder can predict pseudo labels that may not be the same as the original attribute values of pseudo labels but follow the captured attribute correlations which help SSA predict node labels better.

4.3 Online Test

In the online test, we deploy IG-Encoder and SGRL that are pre-trained on the million-scale “small” data with 100,000 labels used in the offline test to detect BMAs over the large, billion-scale WeChat graph. In other words, we do not perform the pretraining phase.

In the online environment, we do not have test labels beforehand. In practice, BMAs and non-BMAs are severely imbalanced, and normal detection methods will label most accounts as non-BMAs. Thus, human experts do not check the large number of predicted non-BMAs as such manual checks are highly costly, and non-BMAs do not trigger warning in the detection system anyway. However, human experts will check all the predicted BMAs before they are banned by the system to avoid disturbing normal WeChat users. Considering the above online scenario where predicted BMAs are more important, we define an online evaluation measure called *Detection Success Rate (DSR)* as the evaluation measure. DSR is the percentage that predicted BMAs are BMAs in reality.

IG-Encoder and SGRL were deployed to detect BMAs over the WeChat graph once every day within a month. After the test period, we summarize each method’s one-month result and remove duplicated predicted BMAs (an account can be labeled as a BMA on different days). Then, we measure DSR on the summarized result for each method. The resulting DSR is 0.7495 and 0.8040 for IG-Encoder and SGRL, respectively. We can see that both methods achieve satisfactory performance in the online test. SGRL further

exceeds IG-Encoder by 7.27%, showing its superiority. Due to security reasons, we are not allowed to disclose the exact number of identified BMAs and detailed statistics of the WeChat graph. But we can roughly report that both methods correctly discover several hundred thousand BMAs over the large WeChat graph. This is an encouraging result as a few human experts may not correctly identify “unique” BMAs in the same magnitude within a month. Moreover, most of the identified BMAs are not even included in the “small” training set used in the offline pretraining phase, meaning that our system has a strong generalization ability and it can adapt properly to new, unseen data. In conclusion, the online test demonstrates that our system can effectively detect BMAs in WeChat.

5 RELATED WORK

Graph Self-Supervised Learning (GSSL). Previous works on graph representation learning mostly focus on (semi-)supervised learning, resulting in heavy label reliance and poor generalization [12]. Recently, GSSL [7, 15, 26], which extracts useful information from pretext tasks that do not rely on manual labels, has attracted great interest. Various pretext tasks are explored to alleviate label reliance and strengthen model generalization. Examples include masking features and predicting [6, 26], contrasting graph instances [14, 18, 25], node clustering [15, 26], and graph partitioning [26]. The above works have inspired the design of SGRL.

Graph-based Anomaly Detection (GBAD). GBAD has many applications [1] and various techniques have been used to improve GBAD: adopt Heterogeneous Information Networks [22, 24, 29], use GSSL to decouple learning and detection [19], capture the temporal information [27, 28], to name a few. Compared to these methods, SGRL encodes both local structure and attribute-level correlations into node representations and scales well to large-scale graphs.

Black Market Account Detection in Social Networks. Lee et al [10, 20] study real money trading detection in online games. But their methods are rule-based and inflexible. Supervised classification [4] and Weighted Generalized Canonical Correlation Analysis [2] are adopted to identify black market activities in Twitter. Compared to SGRL, these methods heavily rely on annotations that are almost unavailable in the BAD task in WeChat.

6 CONCLUSION

In this paper, we illustrate our BMA detection system SGRL used in WeChat. SGRL is designed based on GNN and GSSL. When detecting BMAs in WeChat, SGRL shows not only promising detection results but also a great generalization ability. Both offline and online experiments demonstrate that SGRL is able to detect BMAs over the large MMMA data in WeChat and block their supports for frauds. In the future, we plan to explore how to reduce the cost of advanced designs mentioned in Sec. 3 that are not used in current SGRL, and further improve the detection performance of SGRL.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (No. 62002303, 42171456), the Natural Science Foundation of Fujian Province of China (No. 2020J05001), and 2021 Tencent WeChat Rhino-Bird Focused Research Program.

REFERENCES

- [1] Leman Akoglu, Hanghang Tong, and Danai Koutra. 2015. Graph based anomaly detection and description: a survey. *Data Min. Knowl. Discov.* 29, 3 (2015), 626–688.
- [2] Udit Arora, Hridoy Sankar Dutta, Brihi Joshi, Aditya Chetan, and Tanmoy Chakraborty. 2020. Analyzing and Detecting Collusive Users Involved in Black-market Retweeting Activities. *ACM Trans. Intell. Syst. Technol.* 11, 3 (2020), 35:1–35:24.
- [3] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *KDD*. 785–794.
- [4] Hridoy Sankar Dutta, Aditya Chetan, Brihi Joshi, and Tanmoy Chakraborty. 2018. Retweet Us, We will Retweet You: Spotting Collusive Retweeters Involved in Blackmarket Services. In *ASONAM*. 242–249.
- [5] R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. Learning deep representations by mutual information estimation and maximization. In *ICLR*.
- [6] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay S. Pande, and Jure Leskovec. 2020. Strategies for Pre-training Graph Neural Networks. In *ICLR*.
- [7] Yizhu Jiao, Yun Xiong, Jiawei Zhang, Yao Zhang, Tianqi Zhang, and Yangyong Zhu. 2020. Sub-graph Contrast for Scalable Self-Supervised Graph Representation Learning. In *ICDM*. 222–231.
- [8] Frank J. Massey Jr. 1951. The Kolmogorov-Smirnov Test for Goodness of Fit. *Journal of the American statistical Association* 46, 253 (1951), 68–78.
- [9] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- [10] Eunjo Lee, Jiyoung Woo, Hyoungshick Kim, and Huy Kang Kim. 2018. No Silk Road for Online Gamers!: Using Social Network Analysis to Unveil Black Markets in Online Games. In *WWW*. 1825–1834.
- [11] Weiwei Liu, Haobo Wang, Xiaobo Shen, and Ivor W. Tsang. 2022. The Emerging Trends of Multi-Label Learning. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 11 (2022), 7955–7974.
- [12] Yixin Liu, Shirui Pan, Ming Jin, Chuan Zhou, Feng Xia, and Philip S. Yu. 2021. Graph Self-Supervised Learning: A Survey. *arXiv Preprint* (2021). <https://arxiv.org/abs/2103.00111>
- [13] Ziqi Liu, Chaochao Chen, Longfei Li, Jun Zhou, Xiaolong Li, Le Song, and Yuan Qi. 2019. GeniePath: Graph Neural Networks with Adaptive Receptive Paths. In *AAAI*. 4424–4431.
- [14] Jiezhong Qiu, Qibin Chen, Yuxiao Dong, Jing Zhang, Hongxia Yang, Ming Ding, Kuansan Wang, and Jie Tang. 2020. GCC: Graph Contrastive Coding for Graph Neural Network Pre-Training. In *KDD*. 1150–1160.
- [15] Ke Sun, Zhouchen Lin, and Zhanxing Zhu. 2020. Multi-Stage Self-Supervised Learning for Graph Convolutional Networks on Graphs with Few Labeled Nodes. In *AAAI*. 5892–5899.
- [16] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing Data using t-SNE. *J. Mach. Learn. Res.* 9 (2008), 2579–2605.
- [17] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR*.
- [18] Petar Velickovic, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R. Devon Hjelm. 2019. Deep Graph Infomax. In *ICLR*.
- [19] Yanling Wang, Jing Zhang, Shasha Guo, Hongzhi Yin, Cuiping Li, and Hong Chen. 2021. Decoupling Representation Learning and Classification for GNN-based Anomaly Detection. In *SIGIR*. 1239–1248.
- [20] Kyungmoon Woo, Hyukmin Kwon, Hyunuchul Kim, Chong-kwon Kim, and Huy Kang Kim. 2011. What can free money tell us on the virtual black market?. In *SIGCOMM*. 392–393.
- [21] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2021. A Comprehensive Survey on Graph Neural Networks. *IEEE Trans. Neural Networks Learn. Syst.* 32, 1 (2021), 4–24.
- [22] Bingbing Xu, Huawei Shen, Bing-Jie Sun, Rong An, Qi Cao, and Xueqi Cheng. 2021. Towards Consumer Loan Fraud Detection: Graph Neural Networks with Role-Constrained Conditional Random Field. In *AAAI*. 4537–4545.
- [23] Zequan Xu, Siqiang Luo, Jieming Shi, Hui Li, Chen Lin, Qihang Sun, and Shaofeng Hu. 2022. Efficiently Answering k-hop Reachability Queries in Large Dynamic Graphs for Fraud Feature Extraction. In *MDM*. 238–245.
- [24] Zequan Xu, Qihang Sun, Shaofeng Hu, Jiguang Qiu, Chen Lin, and Hui Li. 2022. Multi-view Heterogeneous Temporal Graph Neural Network for “Click Farming” Detection. In *PRICAI (1)*, Vol. 13629. 148–160.
- [25] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph Contrastive Learning with Augmentations. In *NeurIPS*. 5812–5823.
- [26] Yuning You, Tianlong Chen, Zhangyang Wang, and Yang Shen. 2020. When Does Self-Supervision Help Graph Convolutional Networks?. In *ICML*. 10871–10880.
- [27] Wenchao Yu, Wei Cheng, Charu C. Aggarwal, Kai Zhang, Haifeng Chen, and Wei Wang. 2018. NetWalk: A Flexible Deep Embedding Approach for Anomaly Detection in Dynamic Networks. In *KDD*. 2672–2681.
- [28] Li Zheng, Zhenpeng Li, Jian Li, Zhao Li, and Jun Gao. 2019. AddGraph: Anomaly Detection in Dynamic Graph Using Attention-based Temporal GCN. In *IJCAI*. 4419–4425.
- [29] 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*. 785–795.