Don't Ignore Alienation and Marginalization: Correlating Fraud Detection

Yilong Zang[1] , **Ruimin Hu**[1*] , **Zheng Wang**[1] , **Danni Xu**[2] ,
Jia Wu[3] , **Dengshi Li**[4] , **Junhang Wu**[1] , **Lingfei Ren**[1]
[1]National Engineering Research Center for Multimedia Software,
School of Computer Science, Wuhan University
[2]School of Computing, National University of Singapore
[3]School of Computing, Macquarie University
[4]School of Artificial Intelligence, Jianghan University
{zangyl, hrm, wangzwhu, wjh920925, renlingfei}@whu.edu.cn, dannixu.gracie@foxmail.com,
jia.wu@mq.edu.au, reallds@jhun.edu.cn

Abstract

The anonymity of online networks makes tackling fraud increasingly costly. Thanks to the superiority of graph representation learning, graph-based fraud detection has made significant progress in recent years. However, upgrading fraudulent strategies produces more advanced and difficult scams. One common strategy is synergistic camouflage —— combining multiple means to deceive others. Existing methods mostly investigate the differences between relations on individual frauds, that neglect the correlation among multi-relation fraudulent behaviors. In this paper, we design several statistics to validate the existence of synergistic camouflage of fraudsters by exploring the correlation among multi-relation interactions. From the perspective of multi-relation, we find two distinctive features of fraudulent behaviors, *i.e.*, alienation and marginalization. Based on the finding, we propose COFRAUD, a correlation-aware fraud detection model, which innovatively incorporates synergistic camouflage into fraud detection. It captures the correlation among multi-relation fraudulent behaviors. Experimental results on two public datasets demonstrate that COFRAUD achieves significant improvements over state-of-the-art methods.

1 Introduction

The vast amount of networks - social media, blogs, trading networks, and communication networks, provide a considerable space for fraudulent activity [Velampalli and Eberle, 2017]. The anonymity of online networks and the lack of supervision by administrators make fraud cases often more destructive and difficult to address [Hooi *et al.*, 2017]. Thanks to the superiority of graph representation learning, complex user interactions can be efficiently exploited to identify fraudsters. Graph-based fraud detection has become a research problem with full of opportunities and challenges in recent years. It focuses on utilizing graphs to establish the interactions with multiple types of relations (multi-relation) between nodes and design models to identify the fraudsters [Pourhabibi *et al.*, 2020].

Existing approaches have achieved some success in exploring fraudulent behavior. [Liu *et al.*, 2020; Dou *et al.*, 2020] proposed to reconstruct graph structure to tackle the inconsistency problem resulting from fraudsters' camouflage. [Liu *et al.*, 2021] sampled the neighbor information to alleviate the label imbalance of fraudsters. [Zhang *et al.*, 2021] refined the graph inconsistency problem from three aspects and proposed a unified GNN-based model to solve them. [Tang *et al.*, 2022] first explored the relevance between fraud and frequency spectral energy, and had notable success. These approaches focus on single-strategy fraud and therefore are more inclined to address initial simple scams.

However, upgrading fraudulent strategies produce more advanced and difficult scams. One common strategy is *synergistic camouflage* — combining multiple means to deceive others. Fraudsters tend to employ various techniques to prevent themselves from exposure. *E.g.*, telecom scammers use multiple communication methods, such as phone calls and text messages, together to covertly spread their words. It means that interactions of multiple relation types (multi-relation) are correlated to achieve the ultimate deceiving goal. Unfortunately, to our knowledge, synergistic camouflages have not been discussed by the current methods which usually consider user interactions in each relation type separately. These methods mostly investigate the differences between relations (described as relation inconsistency [Liu *et al.*, 2020; Zhang *et al.*, 2021]) on individual frauds, that neglect the correlation among multi-relation fraudulent behaviors.

Hence, to solve the above issues, we intend to answer the following two questions in this paper. 1) *Is there any synergistic camouflage phenomenon on the fraudsters' multi-relation network*? 2) *How can the correlation among multi-relation interactions be explored to solve the synergistic camouflage problem?*

To answer question 1, we plan to explore the correlation among individuals' multi-relation fraudulent behaviors.

*Corresponding Author

Specifically, we design several statistics to analyze the user interaction patterns on all relation-specific subgraphs. The details of our statistics are illustrated in Section 2.2. In conclusion, we have the following two observations: 1) **Alienation.** we calculate the cross-relation neighbor overlap of all users, and find that the neighbor overlap of fraudsters is lower than normal users. This illustrates that fraudsters' multi-relation interactions are alienated. They often do not establish regular social connections with their neighbors to avoid being identified by other users. 2) **Marginalization.** We then measure the Pagerank [Page *et al.*, 1999] of all users on the subgraphs and find that the fraudsters have lower Pagerank values than the normal users on most of the subgraphs. Understandably, in order to hide their activities, fraudsters tend to place themselves in an undetectable marginal position. The above two statistical findings demonstrate the existence of synergistic camouflage of fraudsters. They also demonstrate that the multi-relations in different fraudulent behavior are correlated. Furthermore, we gain the motivation that the correlation between multi-relational interactions can effectively help distinguish fraudsters from normal users.

Thus, as for question 2, we propose a correlation-aware fraud detection model, called COFRAUD, which captures the correlation among multi-relation interactions to solve the synergistic camouflage problem. Specifically, COFRAUD contains the following three modules: 1) in the intra-relation exploration module, we utilize the user frequency information in each relation-specific subgraph to obtain user representations and relation representations; 2) in the inter-relation exploration module, we pass information between different relation-specific node representations of the same user to represent the correlation among user multiple interactions; 3) in the relation fusion modu le, we not only fuse the information from intra-relation exploration and inter-relation exploration, but also aggregate the node and relation representations, to form the final user behavior representation. We conduct extensive experiments on two public datasets to evaluate the performance of our COFRAUD. Experimental results show that COFRAUD outperforms state-of-the-art methods for the fraud detection task.

The contributions of our paper are summarized as follows:

- We propose and prove the existence of synergistic camouflage by exploring the correlation among multi-relation interactions through statistical analysis. From the perspective of multi-relation, we found that fraudsters have alienation and marginalization which are apparently distinguished from normal users.

- We develop a novel correlation-aware fraud detection model, called COFRAUD, which captures multi-relation correlation information to solve the synergistic camouflage problem.

- We conduct adequate evaluations on two real-world datasets to validate the performance of COFRAUD. The results show that COFRAUD achieves significant improvements over state-of-the-art methods.

2 Preliminary

In this section, we give some necessary mathematical definitions and formulate the problem to be solved. Then we present our preliminary to describe the motivation.

2.1 Definition and Problem Formulation

Multi-relation graph. The graph is a heterogeneous graph constructed by users and their multi-relation types of interactions. We denote the $\mathcal{G} = (\mathcal{V}, \mathcal{X}, \mathcal{E}, \mathcal{R}, \mathcal{Y})$ as a multi-relations graph, where the $v_i \in \mathcal{V}$ means the node, the $r \in \mathcal{R}$ denotes the type of relations, the $e_{ij,r} \in \mathcal{E}$ represents the edge between node v_i and v_j under the r-relation, and the $x_i \in \mathcal{X}$ is the feature of node v_i. Node v_i is associated with the label $y_i \in 0, 1$, where 1 denotes the node is a fraudster.

Relation-specific subgraph. We define relation-specific subgraphs which are split by the multi-relation graph according to relation types. Each subgraph is homogeneous, *i.e.*, there is only one type of node and relation.

Problem formulation. The graph-based fraud detection is regarded as a semi-supervised classification task. Given a multi-relation graph $\mathcal{G}$ and partial nodes with labels y, we intend to infer the labels of another part of the nodes.

2.2 Motivation

To validate the existence of the synergistic camouflage phenomenon, we conduct two statistics on a real-world public dataset Amazon [McAuley and Leskovec, 2013] which has three types of relations: U-P-U, U-S-U, and U-V-U (referring Sec4.1). In order to analyze the behavior correlation across relations, we should leverage the relation-specific interaction information. Our intuition is that the variety of user interactions leads to the structural diversity of the relation-specific subgraphs. So we intend to study the behavior patterns of fraudsters from the subgraph structure. Then we conduct the following two statistics.

Alienation

First of all, we mine clues on the local structure of the subgraph, *i.e.*, the user neighbor information. Referring to Adamic–Adar index [Adamic and Adar, 2003], we creatively design the following formula to measure the cross-relation neighbor overlap of individuals:

$$\mathcal{O}(v_{r_1}, v_{r_2}) = \sum_{\mathcal{N}(v_{r_1}) \cap \mathcal{N}(v_{r_2})} \frac{1}{\log |\mathcal{N}(v_{r_1})| \cdot \log |\mathcal{N}(v_{r_2})|},$$

where $\mathcal{N}(v_{r_1})$ denotes the neighbor set of node v in relation r_1. Then we normalize the value by dividing by the degrees of the node on the two relation-specific subgraphs. This metric is designed to represent whether the user has more and a larger ratio of common neighbors between the two relations.

As shown in Figure 1, we calculate the distribution of $\mathcal{O}(\cdot, \cdot)$ of all users between every two relation-specific subgraphs. It is easy to find that the cross-relation neighbor overlap of fraudsters is closer to 0 in all the statistical results that clearly distinguish them from normal users. This indicates that the multi-relation interactions of fraudsters are alienated. normal users tend to build more frequent interactions while

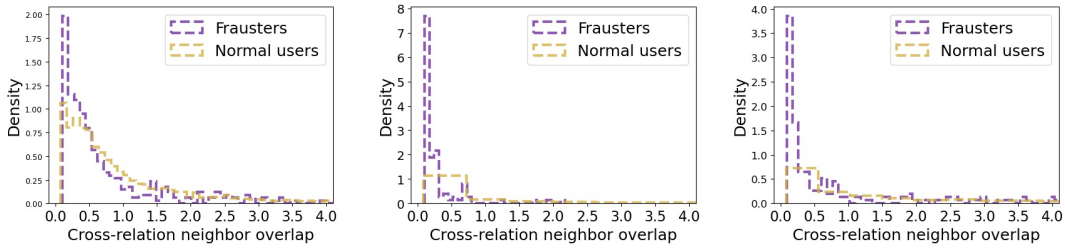

(a) U-P-U *vs.* U-S-U (b) U-P-U *vs.* U-V-U (c) U-S-U *vs.* U-V-U

Figure 1: The distribution of cross-relation neighbor overlap of all users between every two relation-specific subgraphs. We divide the users into two groups, fraudsters and normal users. The metric is designed to represent whether the user has more and a larger ratio of common neighbors between the two relations. A larger value means that the individual has more common neighbors in the two relation-specific subgraphs and the interactions in the two subgraphs of the user are more similar. For the intuitiveness of the results, we omitted users with values above 4, whose density is negligibly small.

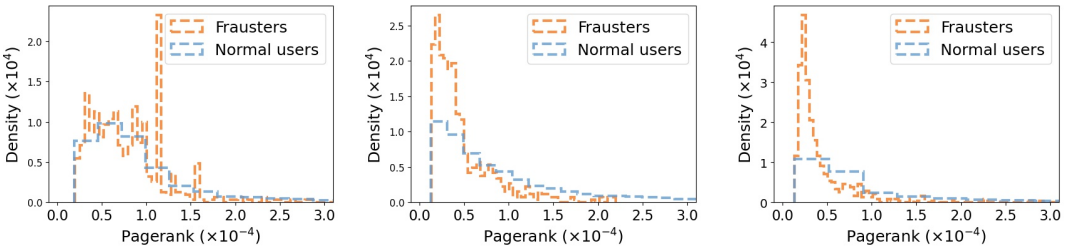

(a) U-P-U (b) U-S-U (c) U-V-U

Figure 2: The distribution of Pagerank of all users on the three relation-specific subgraphs. This metric measures the frequency of user interaction and social status in the subgraph structure. The higher the Pagerank value, the more important the node is in the graph and the more frequent the user's interactions are. We ignore the very few users with Pagerank values above 0.0003 but with very low density.

fraudsters often avoid establishing regular interactions with others. However, if the interactions of each relation type independently are observed independently, it is difficult to detect the fraudsters.

Marginalization

After that, we analyze the user's interaction information on the global structure of the relation-specific subgraphs. We utilize a metric that measures the importance of nodes, Pagerank [Page *et al.*, 1999], which has the following formula:

$$\mathcal{PR}(v_i) = \frac{1-\mathrm{d}}{|\mathcal{V}|} + \mathrm{d} \sum_{v_j \in \mathcal{N}(v_i)} \frac{PR(v_j)}{|\mathcal{N}(v_i)|},$$

where d is the damping factor. This metric measures the frequency of user interaction and social status on the subgraph.

As shown in Figure 2, we calculate the distribution of Pagerank of all users on the three relation-specific subgraphs. It can be observed that the fraudsters have their Pagerank value concentrated in the lower part in most relations, compared to the normal users. This implies that fraudster nodes tend to distribute at the margins of all the relation-specific subgraphs and are not notable. It is easy to understand that fraudsters usually hide their behavior, keeping themselves in an undetectable marginal position.

3 Method

Based on the above two observations, we gain a motivation that fraudulent behavior patterns are obviously different from the norm after correlating interactions of multiple relation types. Hence, we realize that passing the interaction information between different relation types for the same user seems to be meaningful for fraud detection. As we know, existing fraud detection methods focus on passing user representation information within each relation type and aggregating them [Zhang *et al.*, 2021; Tang *et al.*, 2022]. This motivates us to design a distinct information passing module that explores the correlation among multi-relation interactions. Then we introduce our model, COFRAUD, in this section.

3.1 Framework Overview

Figure 3 shows the pipeline of COFRAUD. our framework is comprised of three different modules, respectively the intra-relation exploration module, the inter-relation exploration module, and the relation fusion module. In the intra-relation exploration module, we capture node frequency information and represent node and relation representations. In the inter-relation exploration module, we intend to pass information between different relation-specific node representations of the same user to explore the correlation among user multiple interactions. In the final module, we not only fuse the infor-

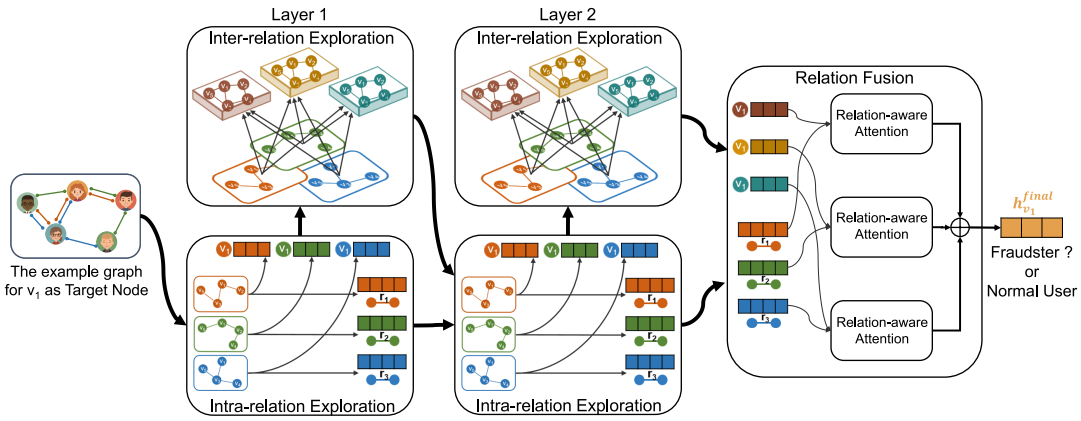

Figure 3: The framework of correlation-aware fraud detection model, COFRAUD. It consists of the intra-relation exploration module for capturing node frequency information and representing relation-specific node and relation representations. It also contains the inter-relation exploration module for exploring the correlation among multi-relation interactions. Finally, the relation fusion module is leveraged for not only fusing the information from intra-relation exploration and inter-relation exploration but also aggregating the node and relation representations.

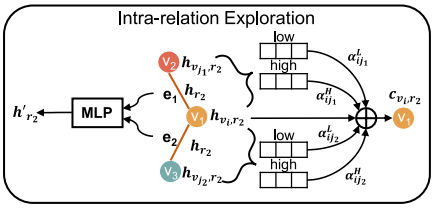

Figure 4: Details of the intra-relation exploration module.

mation from intra-relation exploration and inter-relation exploration but also aggregate the node and relation representations, to get the final user fraud behavior representation.

3.2 Intra-relation Exploration

In the intra-relation exploration module, the node representation and relation representation are processed on each relation-specific subgraph independently. In each subgraph, we leverage the frequency information of the user to suit the fraud detection task. The key to the GNN model's success is considered a special form of low-pass filter [Wu *et al.*, 2019]. However, fraud in social networks tends to focus on high-frequency information [Tang *et al.*, 2022]. This high-frequency information indicating the difference of nodes contributes more to fraud detection than low-frequency information indicating the similarity. Hence, we intend to distinctively represent relation-specific node frequency information to represent the differences between different user features.

As shown in Figure 4, we take the target node v_1 and the relation r_2 as the example. After dividing the original graph G into relation-specific subgraphs, we input node features h_{v,r_2} and relation features h_{r_2}. Motivated by the success of the frequency signal representation on GNN [Bo *et al.*, 2021], we use the attention mechanism to learn the high-frequency and low-frequency signals, α^H and α^L. The values are set as $\alpha^H + \alpha^L = 1$, and $\alpha = \alpha^L - \alpha^H$. Then we consider the features of both neighbor nodes and relations to learn the coefficient α:

$$z_{v_i,r_2} = W_v h_{v_i,r_2}, \tag{1}$$

$$z_{v_j,r_2} = W_v h_{v_j,r_2}, \tag{2}$$

$$z_{r_2} = W_{r_2} h_{r_2}, \tag{3}$$

$$\begin{aligned}\alpha_{ij,r_2} &= \alpha^L_{ij,r_2} - \alpha^H_{ij,r_2} \\ &= tanh\Big(z^T_{r_2}(z_{v_i,r_2} \parallel z_{v_j,r_2})\Big),\end{aligned} \tag{4}$$

where the weight matrix W_v and W_{r_2} are used to transform the node features and relation features into latent spaces. Then function $tanh(\cdot)$ limits the value of α in [-1,1] that represents not only the proportion of low-frequency and high-frequency signals but also the coefficients of neighbors in aggregation.

After calculating α_{ij}, the neighbor representations are aggregated as follow:

$$c_{v_i,r_2} = (1-\lambda) W_{align} h_{v_i,r_2} + \lambda \sum_{j \in \mathcal{N}_{v_i}} \frac{\alpha_{ij,r_2}}{\sqrt{d_{v_i,r_2} d_{v_j,r_2}}} z_{v_j,r_2}, \tag{5}$$

where we utilized the weighted residual connection to aggregate the neighbor information. The λ is a trainable parameter and W_{align} is used to align the dimensions of h_{v_i,r_2}. And the d_{v_i,r_2} and d_{v_j,r_2} are the degrees of the node v_i and v_j in the subgraph of relation r_2, which are used to normalize the coefficients.

In addition, we leverage the MLP model to represent the semantic information of the relations. The formula is as follows:

$$h'_{r_2} = W_{r_2,mlp} h_{r_2} + b_{r_2,mlp}, \tag{6}$$

where the $W_{r_2,mlp}$ and $b_{r_2,mlp}$ are the trainable parameters to update the relation features.

3.3 Inter-relation Exploration

As mentioned previously, the correlation among different types of fraudulent behavior has been verified and distinguished from normal behavior. This drives us to pass different relation-specific node information to each other before aggregating different representations.

Through the above relation-specific frequency adaptation representation, We obtain the frequency-aware representation of the target node which is specific to the relation types. Taking the target node v_1 in relation r_2 as an example, we correlate the multi-relation information as follows:

$$h'_{v_1,r_2} = \sum_{r_i \in \mathcal{R}} \beta_{r_2,r_i} \cdot c_{v_1,r_i}, \tag{7}$$

where β_{r_2,r_i} is the coefficient between relation r_i and relation r_2. we calculate it as follows:

$$\beta_{r_2,r_i} = \frac{exp\Big(LeakyReLU({q_{r_2}}^T c_{v_1,r_i})\Big)}{\sum_{r_j \in \mathcal{R}} exp\Big(LeakyReLU({q_{r_2}}^T c_{v_1,r_j})\Big)}, \tag{8}$$

where q_{r_2} is the trainable attention vector to control the information flow from all relations.

It is worth mentioning that we design two layers of intra- and inter-relation exploration modules. In the input of the first layer, the node representations are the original features x and the relation representations are one-hot encoding, where the value of 1 corresponds to the relation type. And the input in the second layer, h_{v_1,r_2} and h_{r_2}, are the output in the first layer, h'_{v_1,r_2} and h'_{r_2}.

3.4 Relation Fusion

After exploring intra- and inter-relation information, we need to design a suitable aggregator to converge them and leverage both node and relation representations. Aggregators of existing fraud detection methods focus only on the differences among relations, and simply summarize the node representations with weights. This cannot accurately simulate the complicated activities of real-world fraudsters. To better reflect the diversity and complexity of social interactions, we aggregate the information passed across different relation types, while fusing relation semantic information.

Specifically, the attention mechanism is used to aggregate the above information as follows (taking the target node v_1 and relation r_2 as an example):

$$\gamma_{v_1,r_2} = \frac{exp\Big(LeakyReLU\big((V_{r_2} h'_{v_1,r_2})^\top U_{r_2} h'_{r_2}\big)\Big)}{\sum_{r_i \in \mathcal{R}} exp\Big(LeakyReLU\big((V_{r_i} h'_{v_1,r_i})^\top U_{r_i} h'_{r_i}\big)\Big)}, \tag{9}$$

$$h_{v_1}^{final} = \sum_{r_i \in \mathcal{R}} \gamma_{v_1,r_i} \cdot V_{r_i} h'_{v_1,r_i}, \tag{10}$$

where V_r and U_r are the transformation matrix respectively for node representation and relation representation. Remarkably, we consider the relation semantic information while aggregating node representations containing multi-relation information. The significance of this is to increase the distinctiveness of node representations in each relation type.

Dataset	Users	Fraudsters	Relation	Edges
Amazon	11944	9.5%	U-P-U	175,608
			U-S-U	3,566,479
			U-V-U	1,036,737
Yelp	45954	14.5%	R-U-R	49,315
			R-T-R	573,616
			R-S-R	3,402,743

Table 1: Statistical details of two datasets.

Finally, we adjust the dimension of the final node representations $h_{v_1}^{final}$ through the fully connected layer and minimize the cross entropy loss to train the fraud detection task.

4 Experiments

4.1 Experimental Setup

Dataset

Our experiments are conducted on two real-world datasets, Amazon [McAuley and Leskovec, 2013] and Yelp [Rayana and Akoglu, 2015]. **Amazon.** It is extracted from the musical instrument comments on Amazon.com. There are three relations: U-P-U (users reviewing at least one same product), U-S-U (users having at least one same star rating within one week), and U-V-U (users with top-5% mutual review similarities). **Yelp.** It is comprised of spam reviews on restaurants and hotels. There are also three relations: R-U-R (the reviews posted by the same user), R-S-R (the reviews under the same product posted in the same star rating), R-T-R (the reviews under the same product posted in the same month). The statistical details of the two datasets are shown in Table 1.

Baselines

First, we chose some Normal GNN models as the baseline: GCN [Kipf and Welling, 2017], GAT [Veličković *et al.*, 2018], GraphSAGE [Hamilton *et al.*, 2017]. Then, some state of art GNN-based fraud detection methods were used to compare with our approach as follows: GraphConsis [Liu *et al.*, 2020], CARE-GNN [Dou *et al.*, 2020], PC-GNN [Liu *et al.*, 2021], FRAUDRE [Zhang *et al.*, 2021], and BWGNN [Tang *et al.*, 2022]. In the classical GNN model, we merge the multi-relation graph into a homogeneous graph. For fairness, we replaced the binary classification threshold search in BWGNN with the argmax operation which is the same as the other methods in the validation set. In addition to that, we follow the papers to choose the parameters.

Evaluation Metrics

We adopt three widely used metrics to measure the performance of all the methods, respectively AUC, Rec, and F1. The AUC is the area under the ROC Curve that can evaluate the performance of classification by eliminating the influence of imbalanced classes. The Rec and F1 both are the macro-average of the recall and F1-score of the two classes.

Implementation Details

We set the training ratio as 40% and 10% to compare the performance of different methods. For the remaining part of the samples, we divide the validation set and test set in the ratio of 1:2. In the validation set, we take the maximum value of F1 to determine the performance of the methods in the test

Category	Methods	Amazon						Yelp					
		10%			40%			10%			40%		
		AUC	Rec	F1	AUC	Rec	F1	AUC	Rec	F1	AUC	Rec	F1
Normal GNN	GCN	77.26	50.00	47.51	77.94	50.00	47.51	52.12	50.00	46.08	53.12	50.00	46.08
	GAT	76.96	50.00	47.50	77.35	50.00	47.50	50.14	50.00	46.08	49.67	50.00	46.08
	GraphSAGE	69.87	50.00	47.50	71.49	50.00	47.50	52.94	50.00	46.08	56.45	50.00	46.41
Graph-based fraud detection	GraphConsis	82.67	82.63	75.97	85.15	85.10	77.98	64.12	64.72	61.30	61.02	61.67	63.03
	CARE-GNN	88.16	88.19	88.21	87.36	83.90	88.36	69.73	65.68	52.86	70.99	66.80	56.47
	PC-GNN	93.31	**88.47**	81.77	94.64	87.37	88.21	79.04	71.63	66.90	83.05	72.56	70.90
	FRAUDRE	91.34	87.94	87.35	94.27	87.50	90.09	71.47	64.49	59.58	73.93	66.22	61.65
	BWGNN	93.73	78.38	83.72	96.69	84.85	87.73	84.08	62.99	71.55	89.97	70.04	76.21
	COFRAUD	**94.35**	87.64	**90.39**	**97.21**	**89.08**	**91.53**	**87.49**	**73.21**	**74.62**	**91.52**	**79.70**	**79.71**

Table 2: Performance on the two datasets under different percentages of the training data. (%)

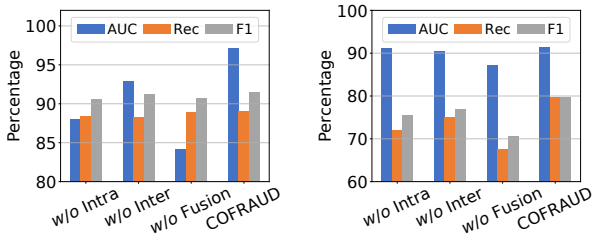

(a) Amazon (b) Yelp

Figure 5: The ablation analysis on Amazon and Yelp. (%)

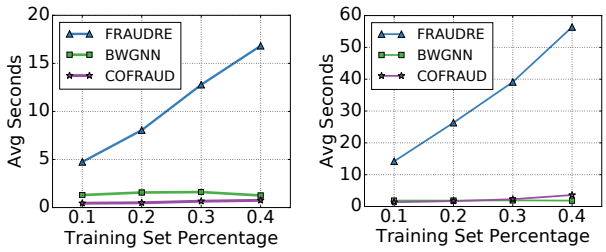

(a) Amazon (b) Yelp

Figure 6: The time efficiency on Amazon and Yelp.

Units	Layers	Amazon			Yelp		
		AUC	Rec	F1	AUC	Rec	F1
16	1	91.74	88.80	91.47	91.09	73.64	76.66
	2	93.73	87.86	89.82	91.14	75.49	77.45
32	1	91.74	88.80	91.47	91.22	76.88	78.36
	2	**97.21**	**89.08**	**91.53**	**91.52**	**79.70**	**79.71**
64	1	89.64	87.68	90.40	90.88	76.67	77.69
	2	93.01	88.39	89.73	91.05	77.63	78.37
128	1	91.48	79.21	79.18	91.47	77.47	78.37
	2	92.37	88.31	90.05	91.01	75.12	77.50

Table 3: The parameter analysis on Amazon and Yelp.(%)

set. We implement our method by Pytorch and DGL [Wang, 2019]. And we conduct all the experiments on the GPU resource of Google Colaboratory [Carneiro *et al.*, 2018].

4.2 Performance

Table 2 presents the results of comparing COFRAUD with other baselines under training ratio 10% and 40%. All the methods are evaluated by the AUC, Rec, and F1. From the table, we can observe that:

1) Importance of multi-relation information. Fraud detection methods based on multi-relation graphs generally have a higher performance compared to normal GNN models which utilize homogeneous graphs. This indicates that multi-relation information is essential and helpful for fraud detection.

2) Importance of frequency information. BWGNN achieves better results compared to other fraud detection methods. It indicates that frequency information effectively represents the differences between user features, which is the key to fraud detection. It also motivates us to utilize the node frequency information in the Intra-relation exploration module.

3) Superiority of COFRAUD. Our model achieves significant improvements in most evaluation metrics which proves the effectiveness of capturing the correlation of the multi-relation fraudulent behavior. In particular, COFRAUD achieves the best performance in all metrics on the Yelp dataset and improves significantly for Rec and F1 on the two datasets. Specifically, it improves 2.41% on the Amazon dataset and improves 3.07% on the Yelp dataset compared with BWGNN at F1 under 10% training ratio, as well as improves 3.82% on the Amazon dataset and 3.50% on the Yelp dataset compared with BWGNN at F1 under 10% training ratio.

4) Limitation of COFRAUD. Overall, all methods perform better on Amazon dataset than in the Yelp dataset. However, COFRAUD has a relatively unsatisfactory performance in the Amazon dataset. This may be due to the label imbalance caused by the small percentage of fraudsters in the Amazon dataset, which is a problem to be solved in COFRAUD.

4.3 Ablation Analysis

In this literature, we design a novel model to capture not only the frequency signal information by intra-relation but also the correlation of multi-relation fraud behavior by inter-relation. Then we aggregate the above two modules to obtain the final representations. To validate the effectiveness of these modules, we design the following three ablation variants: 1) *w/o*

intra. We replaced the node representation part of this module with the GAT, because completely removing the inter-relation exploration module would invalidate the other modules. This allows us to determine whether the frequency signal information can benefit our model. 2) *w/o* inter. We remove the inter-relation exploration module and let the node representations in intra-relation exploration directly input the relation fusion. 3) *w/o* fusion. We replace the relation fusion module with the weighted summation. The training ratio is set as 0.4.

In Figure 5, COFRAUD achieves the best results among all metrics compared to the other variants, which shows that each module in the model is meaningful. In both two datasets, *w/o* relation had the worst results in all variants, which illustrates the importance of effectively fusing information from all modules. The relation representation input in the relation fusion module can introduce relation semantic information to identify the differences among relation-specific node representations.

4.4 Parameter Analysis

In this set of experiments, we explore the effect of the number of feature dimensions of hidden units and the number of module layers of Intra-relation exploration and Inter-relation exploration on the performance of the model. We set the dimensions of all hidden units to 16, 32, 64, and 128, and module layers to 1 and 2, to record the performance of COFRAUD. As illustrated in Table 3, there is no doubt that setting the units to 32 and the layers to 2 is the best parameter configuration. Analyzing the overall trend, the performance of COFRAUD does not drop significantly with parameter changes, indicating the stability of the model.

4.5 Time Efficiency

To demonstrate the time efficiency of COFRAUD, we compare COFRAUD with two current states of art fraud detection methods, FRAUDRE and BWGNN. We record the average training time per epoch with the training ratio varying from 10% to 40%. We set the hidden units to 32 and the batch size to 1024 for all algorithms on the two datasets. As shown in Figure 6, compared to FRAUDRE, COFRAUD runs much faster on both two datasets. Compared to the well-performing BWGNN, COFRAUD has higher running efficiency on the Amazon dataset and is comparable to it on the Yelp dataset.

4.6 Case Study

To clearly illustrate the performance of COFRAUD and how COFRAUD tackles the synergistic camouflage, we show two cases of fraud detection results of the two models, FRAUDRE and COFRAUD. For both the two methods, we extract the test set results for a training set ratio of 0.4 on amazon.

Figure 7 depicts the distribution of prediction results for both two methods in terms of cross-relation neighbor overlap value. It's obvious that compared to FRAUDRE, COFRAUD has a significantly higher detection rate for both fraudsters and normal users, especially near the lower left side of the figure, *i.e.*, the overlap values are low. This shows that COFRAUD can overcome the fraudsters' alienation.

Figure 8 shows the distribution of prediction results in terms of Pagerank value. We can find that FRAUDRE seems

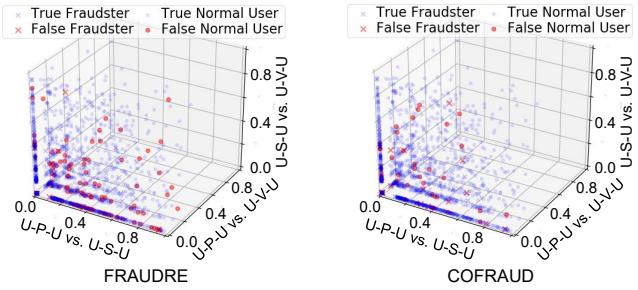

Figure 7: The case 1 of the prediction results of FRAUDRE and COFRAUD. We plot the 3d scatter figure to show the distribution of cross-relation neighbor overlap of users and mark the label and prediction outcome of them. The three axes represent the two by two overlap values among the three relation types. The users have two types of labels, normal users as circles, and fraudsters as crosses. The prediction results are represented in two colors, blue as correct, and red as false.

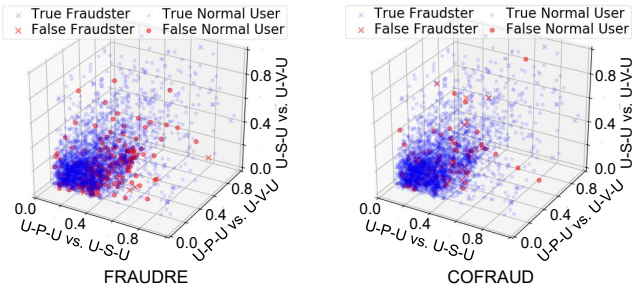

Figure 8: The case 2 of the prediction results of FRAUDRE and COFRAUD. Except for using the Pagerank values calculated in the three relation types as axes, the settings are the same as in Figure 7.

to be helpless for a large number of users with low Pagerank values that cluster in the lower left of the figure. Whereas COFRAUD exploits the marginalization of fraudsters and thus successfully predicts the results.

5 Conclusion

In this paper, we first validate the existence of synergistic camouflage by exploring the correlation among multi-relation interactions. Based on this, we propose a novel correlation-aware fraud detection model (COFRAUD) to solve the synergistic camouflage problem.

Border impact. This paper opens up a new perspective on graph-based anomaly user detection. In the past, these methods focused on analyzing variations in behavior patterns among users on the graph. Our findings suggest that correlating multiple types of behaviors of users themselves is also of great value.

Limitations. 1) Our model needs in-depth exploration of the label imbalance problem. So we plan to apply a superior loss function to incorporate anti-imbalance into COFRAUD. 2) From a practical perspective, COFRAUD should validate its performance in more application scenarios. We intend to apply the model to other fraud datasets such as telecom fraud and insurance fraud.

Acknowledgements

This work was supported by the National Nature Science Foundation of China (No.U22A2035, U1803262, U1736206).

References

[Adamic and Adar, 2003] Lada A Adamic and Eytan Adar. Friends and neighbors on the web. *SOCIAL NETWORKS*, 2003.

[Bo *et al.*, 2021] Deyu Bo, Xiao Wang, Chuan Shi, and Huawei Shen. Beyond low-frequency information in graph convolutional networks. In *AAAI*, 2021.

[Carneiro *et al.*, 2018] Tiago Carneiro, Raul Victor Medeiros Da Nóbrega, Thiago Nepomuceno, Gui-Bin Bian, Victor Hugo C De Albuquerque, and Pedro Pedrosa Reboucas Filho. Performance analysis of google colaboratory as a tool for accelerating deep learning applications. *IEEE Access*, 6:61677–61685, 2018.

[Dou *et al.*, 2020] Yingtong Dou, Zhiwei Liu, Li Sun, Yutong Deng, Hao Peng, and Philip S Yu. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *CIKM*, 2020.

[Hamilton *et al.*, 2017] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *NeurIPS*, 2017.

[Hooi *et al.*, 2017] Bryan Hooi, Kijung Shin, Hyun Ah Song, Alex Beutel, Neil Shah, and Christos Faloutsos. Graph-based fraud detection in the face of camouflage. *TKDD*, 2017.

[Kipf and Welling, 2017] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.

[Liu *et al.*, 2020] Zhiwei Liu, Yingtong Dou, Philip S Yu, Yutong Deng, and Hao Peng. Alleviating the inconsistency problem of applying graph neural network to fraud detection. In *SIGIR*, 2020.

[Liu *et al.*, 2021] Yang Liu, Xiang Ao, Zidi Qin, Jianfeng Chi, Jinghua Feng, Hao Yang, and Qing He. Pick and choose: a gnn-based imbalanced learning approach for fraud detection. In *WWW*, 2021.

[McAuley and Leskovec, 2013] Julian John McAuley and Jure Leskovec. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *WWW*, 2013.

[Page *et al.*, 1999] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.

[Pourhabibi *et al.*, 2020] Tahereh Pourhabibi, Kok-Leong Ong, Booi H Kam, and Yee Ling Boo. Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *DECIS SUPPORT SYST*, 2020.

[Rayana and Akoglu, 2015] Shebuti Rayana and Leman Akoglu. Collective opinion spam detection: Bridging review networks and metadata. In *ACM SIGKDD*, 2015.

[Tang *et al.*, 2022] Jianheng Tang, Jiajin Li, Ziqi Gao, and Jia Li. Rethinking graph neural networks for anomaly detection. In *ICML*, 2022.

[Velampalli and Eberle, 2017] Sirisha Velampalli and William Eberle. Novel graph based anomaly detection using background knowledge. In *The Thirtieth International Flairs Conference*, 2017.

[Veličković *et al.*, 2018] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018.

[Wang, 2019] Minjie Yu Wang. Deep graph library: Towards efficient and scalable deep learning on graphs. In *ICLR workshop*, 2019.

[Wu *et al.*, 2019] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *ICML*, 2019.

[Zhang *et al.*, 2021] Ge Zhang, Jia Wu, Jian Yang, Amin Beheshti, Shan Xue, Chuan Zhou, and Quan Z Sheng. Fraudre: fraud detection dual-resistant to graph inconsistency and imbalance. In *ICDM*, 2021.