xFraud: Explainable Fraud Transaction Detection

Susie Xi Rao†, Shuai Zhang†, Zhichao Han⋆, Zitao Zhang⋆, Wei Min⋆, Zhiyao Chen⋆, Yinan Shan⋆, Yang Zhao⋆, Ce Zhang†

ETH Zurich† eBay China⋆

{raox,shuazhang,ce.zhang}@inf.ethz.ch

{zhihan,zitzhang,wmin,zhiyachen,yshan,yzhao5}@ebay.com

ABSTRACT

At online retail platforms, it is crucial to actively detect the risks of transactions to improve customer experience and minimize financial loss. In this work, we propose xFraud, an explainable fraud transaction prediction framework which is mainly composed of a detector and an explainer. The xFraud detector can effectively and efficiently predict the legitimacy of incoming transactions. Specifically, it utilizes a heterogeneous graph neural network to learn expressive representations from the informative heterogeneously typed entities in the transaction logs. The explainer in xFraud can generate meaningful and human-understandable explanations from graphs to facilitate further processes in the business unit. In our experiments with xFraud on real transaction networks with up to 1.1 billion nodes and 3.7 billion edges, xFraud is able to outperform various baseline models in many evaluation metrics while remaining scalable in distributed settings. In addition, we show that xFraud explainer can generate reasonable explanations to significantly assist the business analysis via both quantitative and qualitative evaluations.

PVLDB Reference Format:

Susie Xi Rao, Shuai Zhang, Zhichao Han, Zitao Zhang, Wei Min, Zhiyao Chen, Yinan Shan, Yang Zhao, Ce Zhang. xFraud: Explainable Fraud Transaction Detection. PVLDB, 15(3): 427 - 436, 2022.
doi:10.14778/3494124.3494128

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at https://github.com/eBay/xFraud.

1 INTRODUCTION

The online retail industry is reshaping our shopping behavior, and the resulting security risks are not negligible. Common threats in e-commerce include account acquisition, financial information theft, fake chargeback, money laundry, and many more. For instance, malicious attackers might try to steal customer's credit card information; the login credentials can also be acquired by hackers. These criminal activities can bring negative impacts on user experiences, cause financial losses, and seriously degrade the platform credibility. As such, it is critical to identify fraudulent behaviors and take every precaution to minimize risks.

Proceedings of the VLDB Endowment, Vol. 15, No. 3 ISSN 2150-8097.
doi:10.14778/3494124.3494128

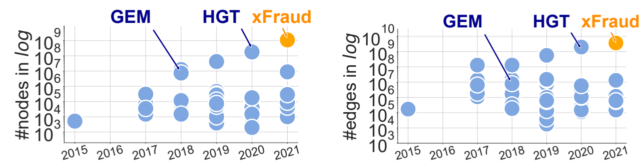

Figure 1: The node and edge numbers (*log*) of heterogeneous graph datasets in the literature. Full survey in Appendix A [34].

Fraud detection has been an emerging topic for e-commerce and social media companies. It is studied in various applications: malicious account detection (e.g., social networks [3], online payment systems [53], and online retailer platforms [4, 5, 28, 54]); anti-money laundry [10, 43]; spam reviews and news detection [38, 41]. Among these *transaction fraud detection* is an important topic [4, 5, 54]. In this work, we focus on automatic fraudulent transaction detection in a real-world e-commerce environment at eBay. For an incoming transaction, we aim to predict whether it is legitimate or not.

Challenges. Despite recent efforts in automatic fraudulent transaction detection [3, 8, 11, 17, 20, 22, 26, 28, 33, 38, 40, 44, 51, 53] with machine learning such as LSTM and graph neural networks (GNNs), we realize that three challenges still linger when it comes to our application scenario at eBay. ***(1. Information Heterogeneity)*** In our system, there are heterogeneous types of information concerning a transaction such as payment tokens, shipping addresses, email. Intuitively, such information is indicative in fraud detection. *How to effectively utilize such information in an end-to-end ML model?* ***(2. Scalability and Efficiency)*** Our platform can produce millions of transactions involving millions of users in a short span of time, which requires the detecting system to be *efficient and scalable for practical use.* Specifically, Figure 1 shows the landscape of heterogeneous graphs emerging in the last six years. In this paper, we are tackling a workload, to our best knowledge, that consists of one of the largest *heterogeneous* graphs for graph neural networks (see a more detailed survey in Appendix A [34]). This poses unique challenges in the system design and optimizations. ***(3. Explainability)*** Flagging a transaction to be fraudulent is not a trivial process. False decisions are likely to cause trouble to our customers and significantly degrade the platform's credibility. In general, this process requires extensive human efforts in cautiously reviewing the model's prediction, which is inefficient and costly. *How can we explain the outcome of an ML model, and more importantly, how close these explanations are to those developed by human experts in the business unit (BU)?*

Our Approach. To tackle the aforementioned challenges, we propose xFraud, an explainable fraud detection framework at eBay. xFraud is not only able to efficiently and effectively predict the legitimacy of a transaction but can also generate human readable explanations in to assist flagging transaction frauds. xFraud advances previous work in two ways.

First, xFraud builds upon a heterogeneous GNN to tackle *transaction fraud detection*. Specifically, xFraud consists of a detector and an explainer. In the detector, we tackle a transaction fraud detection task from the graph perspective. Different from the existing works [24, 27, 31], a heterogeneous graph of different node types (e.g., transactions, addresses, payment tokens) is constructed. To capture the heterogeneous relation patterns and learn more expressive node representations, a self-attentive heterogeneous graph neural network is adopted. The detector can automatically aggregate information from different types of nodes via disparate paths without manually predefined meta-paths. This is important because we do not need to predefine the meta-paths of risk propagation and preprocess the path representations, as required in [4, 17, 38, 53]. Additionally, we explore an efficient sampler in the message aggregation procedure, which empirically reduces the inference time significantly while achieves a competitive performance.

Second, to the best of our knowledge, we provide the first quantitative evaluation of input-level GNN explanation methods [45, 47] and its agreement with humans in a real-world application scenario. Unlike traditional classification tasks, the claim that a transaction is fraudulent should be made very cautiously to avoid harming customer experience and degrading the platform's credibility. As such, we integrate an explainer into our framework that can provide intuitive explanations for model predictions. Equipped with these explanations, our auditors, regulators, or decision makers know how a transaction is flagged by the detector, thus making more sensible decisions. One open question is *how well these explanations agree with the insights from human experts.* To this end, we conduct an extensive quantitative study to measure the agreement between human perception, GNNExplainer [47], and centrality measures, as well as provide case studies on how these explanations can help in practice. This study also reveals an interesting tradeoff between GNN-based explanations and traditional topological measures (e.g., centrality), which allows us to design a hybrid explainer that outperforms both strategies.

In summary, our technical contributions are as follows.

(1) We propose a heterogeneous GNN model (detector) to identify transaction frauds. Our model captures the heterogeneity in transaction graphs and applies to industrial-scale datasets. xFraud detector provides concrete analytical angles of fraudulent activities. Compared with HGT [18], in xFraud, we design a new sampling mechanism, inspired by the sparsity of our underlying graph; compared with GEM [28], xFraud uses a heterogeneous GNN architecture, which allows it to outperform a GEM-style model significantly.

(2) We add explainability into xFraud with a hybrid explainer. xFraud explainer computes the contributions of its neighboring node types and edges when predicting a node, and it also attends to global topological features learned from centrality measures. Thus, it enables explicit case studies of network predictions, which is beneficial for model trustworthiness. We perform the first quantitative evaluation between GNNExplainer and human judgments. We also compare edge weights computed via centrality measures with the weights learned by GNNExplainer, through which we identify a trade-off and propose a hybrid explainer in xFraud.

(3) We conduct careful system design and optimizations, which allow us to scale out, to our best knowledge, to one of the largest heterogeneous graphs being reported for ML workloads so far.

(4) We conduct experiments on real-world transaction networks to show the efficiency of xFraud in detecting transaction frauds and in facilitating the analysis of graph structural patterns.

2 RELATED WORK

xFraud builds upon recent successes of applying heterogeneous graphs to fraud detection (e.g., MAHINDER [53] and GEM [28] from Alibaba) and also recent efforts of GNN explainability [47]. However, it also makes significant improvements over these previous efforts, discussed as follows.

Fraud Detection. There are two lines of studies on fraud detection systems. One line of work leverages graph information with non-GNN methods (see anti-money laundry [10], spam detection [8, 20, 38], fraudster user detection [3, 11, 15, 17, 21, 23, 26, 33, 40, 53], fraud transaction detection [4, 5, 35, 54]). These models can be event/sequence based [3, 5, 22, 54], or meta-path based [4, 17, 38, 53]. Zhong et al. [53] propose MAHINDER which uses *heterogeneous graphs* in the context of defaulter detection by pre-defining a set of *meta-paths* in a heterogeneous graph of users and merchants. The preprocessed meta-path feature representations are trained with an attention mechanism and LSTM to measure the importance of nodes, links, and meta-paths at different timestamps. In xFraud, we focus on a different scenario, aiming at flagging each transaction of a user in various risk scenarios, as a legitimate user does not imply that all its transaction records are legitimate, e.g., once its payment token has been stolen. More importantly, we *focus on methods that do not need to define meta-paths a priori, instead are able to automatically learn these patterns using a GNN.*

The other line of work uses GNN methods [22, 24, 27, 28, 31, 41, 43, 44, 51]. Homogeneous graph has been widely applied in e-commerce applications (see spam review detection [41], anti-money laundry [43], risky/malicious account detection [24, 27, 31]). Recently, people start to solve real-world anomaly detection problems using heterogeneous graph (see spam review detection [22], suspicious user detection [28, 44, 51]) or to combine homogeneous and heterogeneous graphs [22], because it allows aggregating information propagation through various types of nodes/edges. GEM [28] by Liu et al. has utilized attention mechanisms in a device-account heterogeneous graph to capture user activity and device embeddings in each subgraph neighborhood. The heterogeneous graph in GEM is then fed into a GCN.

Explainability in GNN. Recently, how to interpret and explain GNN predictions has gained spotlight. There are two levels to explain: input level (GraphConsis[29], GNNExplainer [47], GraphLIME [19], [2, 45]), and model level (XGNN [48]). GNNExplainer [47], a GNN model agnostic explanation framework, proposes explaining the GNN predictions by maximizing the mutual information gain of the true node labels and the predicted labels using informative features. GNNExplainer enables a visualization of important subgraph patterns, which assists users to understand the feature contribution and node label propagation. GNN explainability in the financial domain has been addressed by Li et al. [45]. They have extended GNNExplainer techniques by (1) adding a regularization term that ensures at least one edge connected to each

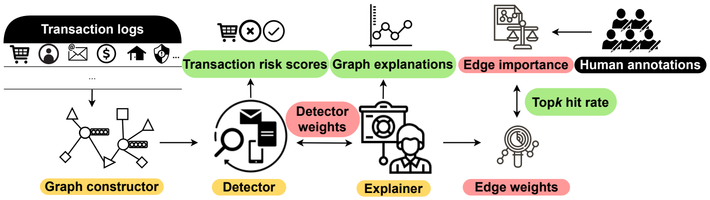

Figure 2: xFraud pipeline.

Transaction	Buyer	Email	Payment token	Shipping address	Transaction features
1	1	john_eth@ gmail.com	Credit card	Albert-Einstein-Strasse 1, Zurich	[..., ..., ...]
2	1	john_eth@ gmail.com	Payment slip	Hauptstrasse 1, Zurich	[..., ..., ...]

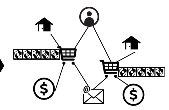

Figure 3: Transactions → a heterogeneous graph.

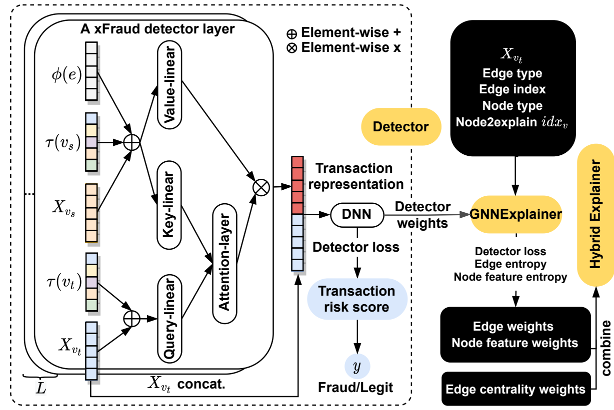

Figure 4: xFraud detector and explainer.

node is selected in the subgraph, (2) adding edge weighted graph attention to calculate the edge weights in the subgraph. Using GCN in a node classification task, they applied an explainer to identify informative graph patterns on financial transaction data like bitcoin over the counter (OTC) and account matching in bank transactions.

3 THE XFRAUD FRAMEWORK

Figure 2 illustrates the xFraud pipeline in a nutshell. First, we build a graph constructor to convert transaction logs into a graph abstraction (Sec. 3.1), which is then fed into a detector to generate a transaction risk score for each transaction record (Sec. 3.2). Then, to build a learnable hybrid explainer (Sec. 3.4), we combine the task-aware measures of predictions generated by the GNNExplainer, and the task-agnostic centrality measures.

One highlight of xFraud over previous efforts is its emphasis on explainability, especially its evaluation using insights from real-world experts in the business unit. We evaluate the efficacy of the explainer quantitatively (Sec. 5.1) and qualitatively (Sec. 5.2). Quantitatively, we calculate the agreement (top$k$ hit rate) between human annotations and explainer weights. Concretely, we first obtain human ground truth on edge importance in risk propagation. Then, we calculate edge importance scores from node importance scores with various aggregation methods. At last, we report the top$k$ hit rate between edge importance scores computed from human annotations and edge weights generated by the hybrid explainer. Qualitatively, we study in detail many cases where xFraud explainer assists the BU in better understanding complex fraudulent patterns.

3.1 Heterogeneous Graph Construction

Think of the critical entities involved under fraud scenarios. A credit card might be linked to both a legitimate user and a fraudulent user at different stages. The latter happens in a card stolen case. A common shipping address such as a warehouse is sometimes used in frauds. This linkage tends to be stable, compared with stolen financial instruments. If we formulate fraud detection as a semisupervised learning problem in an inductive setting [14] in a heterogeneous graph, we have the specification of the problem formulation as follows. In a heterogeneous transaction graph $\mathcal{G}$, $v \in \mathcal{V}$ has a type $\tau(v) \in \mathcal{A}$, where $\mathcal{A} := \{txn, pmt, email, addr, buyer\}$, referring to *transaction, payment token, email, shipping address, buyer*, respectively[1]. If a transaction has relation with another type of node in {*pmt, email, addr, buyer*}, we put an edge between those two nodes in the heterogeneous graph. Each *txn* node carries node attributes provided by a risk identification system. A transaction is represented as an ID in the transaction log. The item category in the purchase order relevant to one transaction (item-type info) is encoded in the transaction features. Each transaction is flagged legit or fraud. Figure 3 illustrates how to construct such a heterogeneous graph based on two transaction records sharing several entities.

3.2 The detector of xFraud

As we have seen in the literature, it is common to define meta-paths when analyzing graph structured data and then to extract corresponding features of nodes and edges on the meta-paths before feeding the features into a machine learning or deep learning model. However, in a fraud detection scenario, under many circumstances, it is by nature impossible to enumerate every possible scenario and their influential meta-paths. This is also one of the primary intuitions why a heterogeneous GNN is a desirable choice: it allows a network to learn the importance of meta-paths by itself based on the network structure and message passing.

3.2.1 ***xFraud detector***. We are inspired by Transformer [39] and HGT [18], when designing the xFraud detector incl. heterogeneous mutual attention and heterogeneous message passing with key, value, and query vector operations (self-attention mechanism). We do not allow target-specific aggregation on different node types, so that we reduce the cost in computing different weights for various node types. We see a better performance in our detector (see discussion in Sec. 4), when shared weights among different types of nodes are used. Moreover, we do not adopt relative temporal encoding in HGT when processing transactions with timestamps. Reasons are, we would like to keep track of all transactions a buyer executes, as well as the linking entities a transaction involves. We also model the relations between buyers and transactions. This makes our system adaptable to guest checkouts and their pertaining chargebacks, as those transactions could not be linked to any buyer accounts, but they could be linked to suspicious third-party

[1]For this study, we choose these attributes based on the homophilic tests [1]. It is shown that fraud exhibits homophilic effects [32], and entities with strong homophilic effects are considered in this work.

payment accounts or billing email addresses. In this manner, our system is able to capture disguised/missed fraud patterns of guest checkouts that could otherwise be neglected by (1) representing a transaction using buyers and timestamps as in HGT or by (2) representing the transactions in a homogeneous graph.

Comparison to GEM. Looking at the application scenarios, xFraud might seem similar to GEM (a malicious account detection system developed by Alibaba). However, xFraud differs from GEM [28] in that: (1) GEM is a system which directly applies a vanilla GCN to a heterogeneous graph, while our proposed xFraud considers the heterogeneous property of graphs in the underlying architecture (e.g., sampler, heterogeneous graph convolution). In this paper, we choose GEM as a representative of heterogeneous GCN in the evaluation. (2) We focus on a different application and have different node types. GEM focuses on fraudsters' detection, while we aim to find the anomaly transactions (a user may have both fraudulent and normal transactions due to account hacking). (3) The GEM model does not provide any explanation for its predictions. We have extensively discussed the GNN explainability and conducted experiments to understand the xFraud explainer.

In Figure 4, we show the detector architecture in detail (left).

(1) The detector takes a heterogeneous graph as input, incl. target and source node features X_{v_t}, X_{v_s}; target and source node types $\tau(v_t)$, $\tau(v_s)$; edge type $\phi(e)$. For the *txn* nodes, we have node features computed by a company risk identifier. For the other node types, the initial node features are empty and only get their inputs after the first convolution layer. The type features are in one-hot encoding of types. **(2)** *L heterogeneous convolution layers* process the graph with self-attention mechanism: the input layer $\mathrm{L}^{(0)}$ takes transaction features, node type embeddings (source and target), edge type embeddings as input, which are transformed into query, key, and value vectors. Attention scores are calculated for the source and target nodes and then layer-wise normalized, which are then fed into a *ReLU* activation function that emits input for the next convolution layer. In Sec. 3.2.2, we introduce with equations how heterogeneous mutual attention and message passing function in one heterogeneous convolution layer. **(3)** After L heterogeneous convolution layers, a *tanh* activation is applied to the transaction representations generated by GNN. Then these representations are concatenated with the original transaction features and fed into a feedforward connected network with two hidden layers. We then apply dropout, layer normalization, and *ReLU* transformation to calculate a predicted risk score and a label. **(4)** The loss function of xFraud detector is the cross entropy of the true label and the probability score calculated by *softmax* (see eq. 1 in Appendix D [34]).

3.2.2 *Heterogeneous Convolution Layer in xFraud Detector.*

We discuss the details of a xFraud detector layer shown in Figure 4. For a tuple $\langle \tau(v_s), \phi(e), \tau(v_t) \rangle$, where $e = (v_s, v_t)$, we initialize (1) the node type embeddings $\tau(v)^{emb}$ and the edge type embeddings $\phi(e)^{emb}$ with zero weights; (2) the attention weight matrices of source node $W^{att}_{\tau(v_s)}$ and of target node $W^{att}_{\tau(v_t)}$ with random weights subject to uniform distributions; and (3) the weight matrices for key, query, and value vectors denoted by W^K, W^Q, W^V, respectively, also with random values subject to uniform distributions. In a nutshell, a general attention-based heterogeneous convolution layer of the node v_t has three components, attention, message, and aggregate as shown in $H^l[v_t] \leftarrow Aggregate(Attention(v_s, v_t) \cdot Message(v_s))$. For each target node v_t, we create *query, key*, and *value* vector representations for self-attention mechanism with multiheads.

To construct the ith *query* vector for the target node v_t, we start with an input to the first layer by taking the transaction features of the target node $X^{txn}_{\tau(v_t)}$ and its node type embedding $\tau(v_t)^{emb}$ to calculate $Q^i(v_t) = \text{Q-Linear}^i_{\tau(v_t)}\left(X^{txn}_{\tau(v_t)} + \tau(v_t)^{emb}\right)$, Then, for $H^{(l-1)}$, where $l \neq 1$, we compute $Q^i(v_t) = \text{Q-Linear}^i_{\tau(v_t)}\left(H^{(l-1)}[v_t]\right)$, where $H^{(l-1)}[v_t]$ is the node representation of the node v_t on the $H^{(l-1)}$ layer.

To construct the ith *key* vector for the source node v_s, we start with an input to the first layer by taking the transaction features of the source node $X^{txn}_{\tau(v_s)}$, its node type embedding $\tau(v_s)^{emb}$ and the edge type embedding $\phi(e)^{emb}$ to calculate $K^i(v_s) = \text{K-Linear}^i_{\tau(v_s)}\left(X^{txn}_{\tau(v_s)} + \tau(v_s)^{emb} + \phi(e)^{emb}\right)$, Then, for $H^{(l-1)}$, where $l \neq 1$, we compute $K^i(v_s) = \text{K-Linear}^i_{\tau(v_s)}\left(H^{(l-1)}[v_s]\right)$, where $H^{(l-1)}[s]$ is the node representation of the node v_s on the $H^{(l-1)}$ layer.

To construct the ith *value* vector for source node v_s, we start with an input to the first layer by the taking transaction features of source node $X^{txn}_{\tau(v_s)}$, its node type embedding $\tau(v_s)^{emb}$ and the edge type embedding $\phi(e)^{emb}$ to calculate $V^i(v_s) = \text{V-Linear}^i_{\tau(v_s)}\left(X^{txn}_{\tau(v_s)} + \tau(v_s)^{emb} + \phi(e)^{emb}\right)$, Then, for $H^{(l-1)}$, where $l \neq 1$, we compute $V^i(v_s) = \text{V-Linear}^i_{\tau(v_s)}\left(H^{(l-1)}[v_s]\right)$.

We adopt the multiheaded attention to control the randomness of initial weights. First, we compute the attention output of one attention head, denoted by $\alpha\text{-head}^i(v_s, e, v_t)$, using this equation $\alpha\text{-head}^i(v_s, e, v_t) = \frac{\left(K^i(v_s)W^{att}_{\tau(v_s)} + Q^i(v_t)W^{att}_{\tau(v_t)}\right)}{\sqrt{d_k}}$, where $\sqrt{d_k}$ is the square root of the key vector's dimension.

The heterogeneous mutual attention of the target node query vector $Q^i(v_t)$ and the source node key vector $K^i(v_s)$ is then computed by $\alpha(v_s, e, v_t) = \underset{\forall v_s \in N(v_t)}{\text{softmax}}\left(\underset{i \in [1,h]}{\|} \alpha\text{-head}^i(v_s, e, v_t)\right)$, where $N(v_t)$ represents the neighbors of v_t, h the number of attention heads, $\|$ vector concatenation.

Finally, the message passing between $H^{(l)}$ and $H^{(l-1)}$ is given by $\text{msg}(v_s, e, v_t) = \underset{i \in [1,h]}{\|}\left(V^i(v_s) \cdot \text{dropout}\left(\alpha\text{-head}^i(v_s, e, v_t)\right)\right)$, where the right hand side is a concatenation of all msg-head^i, the message passing of one attention head at the ith query vector.

3.2.3 *xFraud Detector+: An Improvement over HGT.*

We implement xFraud detector and detector+, whose difference lies in the sampler. In detector, we use the original HGT implementation[2] and empirically show that HGSampling is computationally costly (see the inference time in Figure 10). Hence, we modify the sampling as in GraphSAGE and denote the efficient version of the xFraud detector as detector+. In detector+, the algorithm first samples

[2] https://github.com/acbull/pyHGT (last accessed: Oct 18, 2020).

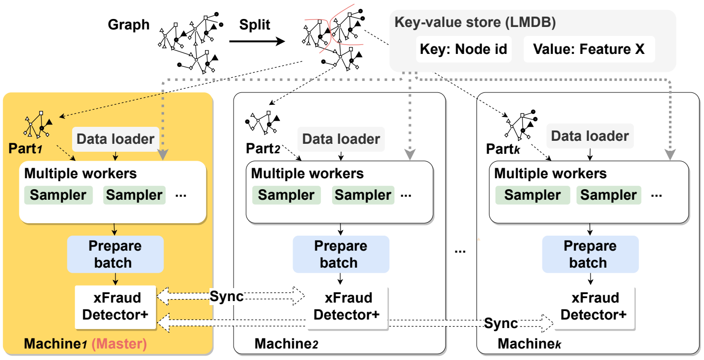

Figure 5: Architecture of Distributed xFraud Detector+.

k-hop neighborhood of a node and then aggregates feature information from neighbors and finally allows GNN to predict the label using aggregated information. In HGSampling, used by HGT, it tries to maintain a similar number of different $\tau(v)$ and $\phi(e)$ types after sampling and minimize the information loss and sample variance in the subgraph after sampling. However, in our datasets, the graph is much sparser (2.12 and 1.49 edges/node for *eBay-small* and *eBay-large*) compared with the Open Academic Graph (11.173 edges/node) used in HGT. Therefore, HGSampling is more costly than GraphSAGE because it requires all types of nodes and edges to be of similar size in the sampled subgraph. In the following sections, we consider xFraud detector equivalent to HGT and mainly focus on the evaluation of xFraud detector+.

3.3 Distributed xFraud Detector+

To make xFraud detector+ scalable to industrial-scale datasets, we designed a distributed learning architecture (see Figure 5). We briefly discuss its design and leave more details to Appendix C [34].

3.3.1 ***Graph Partitioning.*** We adopt the Power Iteration Clustering (PIC) algorithm [25] to partition the graph according to pairwise similarities of edge properties. PIC is effective for graph partition/clustering and well-suited to very large datasets due to its high efficiency. Each worker takes charge of a different partition during distributed training.

3.3.2 ***Distributed Learning.*** We utilize the DistributedDataParallel (DDP) tool provided by the flexible package, *PyTorch Ignite* [12], for distributed model learning. In terms of gradient synchronization, the gradients computed by different workers will be averaged following the default DDP gradient synchronization protocol. After that, parameters of the local model will be updated, and all models on different workers will be the same.

3.3.3 ***Data Loading.*** We use a lightweight KV-store to store all graph-related information. We choose to use Lightning Memory-Mapped Database (LMDB) as it allows us to have multiple data loaders simultaneously, where each worker has its own data loader. This alleviates the system bottleneck that we had when using LevelDB for the same purpose, which we found challenging to support multi-thread operations. This design decision turns out significant in reducing the training and inference time.

Table 1: Top$k$ hit rate computed on different explainability methods using various measures (on all 41 communities).

	Measures to calculate hit rate	H_{Top5}	H_{Top10}	H_{Top15}	H_{Top20}	H_{Top25}
1	edge betweenness	**0.469**	0.718	0.812	**0.903**	0.923
2	edge load	0.455	0.707	0.812	0.902	0.923
3	approximate current flow betweenness	0.450	0.690	**0.821**	0.899	0.923
4	betweenness	0.451	**0.724**	0.815	0.901	0.923
5	closeness	0.464	0.719	0.816	0.901	**0.924**
6	communicability betweenness	0.448	0.688	0.812	0.899	0.922
7	current flow betweenness	0.446	0.700	0.820	0.900	0.922
8	current flow closeness	0.441	0.691	0.815	0.900	**0.924**
9	degree	0.464	0.716	0.815	0.901	**0.924**
10	eigenvector	0.443	0.714	0.811	0.901	**0.924**
11	harmonic	0.464	0.719	0.816	0.901	**0.924**
12	load	0.452	**0.724**	0.815	0.901	0.923
13	subgraph	0.447	0.714	0.813	0.899	0.922
14	GNNExplainer weights	0.445	0.692	**0.821**	0.898	0.921
15	random weights	0.127	0.454	0.602	0.695	0.791

3.4 The Explainer of xFraud

We present a hybrid explainer (see Figure 4) in xFraud based on a trade-off between the *task-aware* GNNExplainer and the *task-agnostic* centrality measures.

3.4.1 ***Trade-off between GNNExplainer and edge centrality.*** GNNExplainer [47] is a model-agnostic explainer which assigns edge weights during node prediction. But, a fundamental question is — *Is GNNExplainer itself optimal for all scenarios? What if we replace the edge weights with other measures, such as random weights and edge centrality?* To answer this, we conduct a micro benchmark against other measures (see Table 1). For conducting a fair comparison, we design a metric called Top$k$ hit rate.

Metric: Top*k* hit rate. We create human annotations of edge importance and compare the explanation weights with the human annotations. The goal of this quantitative evaluation is to quantify the agreement between different edge importance measures and human annotations. Note that the average edge importance scores and edge weights are in different domains (see Figure 6 for an example): the former are discrete values $\in [0, 2]$, and the latter continuous numbers $\in [0, 1]$. We need a metric that reports the edge ranking of the most important ones in both domains. Hence, we compute the top$k$ hit rate H, defined as $\frac{\mathbb{N}^{+,k}_{human} \cap \mathbb{N}^{+,k}_{explainer}}{k}$, where $\mathbb{N}^{+,\mathbb{k}}_{-}$ denotes the set of edges ranked by human/explainer as top k. Concretely, we count the common edges in their top k selection and divide this count by k.

Trade-off and Intuition. Table 1 and Figure 7 illustrate a trade-off between GNNExplainer and edge centrality measures—they work well on different "communities" (test examples) and none of them dominates the other. GNNExplainer is developed to explain the predictions generated by a GNN network. GNNexplainer computes the importance scores of node features and assigns edge weights, with which we determine the most informative edges when a node prediction is made. On the contrary, edge centrality measures are popular methods to quantify the edge importance in a network, which is task-agnostic for node prediction. Intuitively, we should combine these two measures to generate an explainer which attends to both task-aware and task-agnostic measures.

3.4.2 ***xFraud Hybrid Explainer.*** Based on the above analysis, we propose a hybrid explainer and formulate a learning problem as follows. First, we learn two coefficients, namely, the centrality

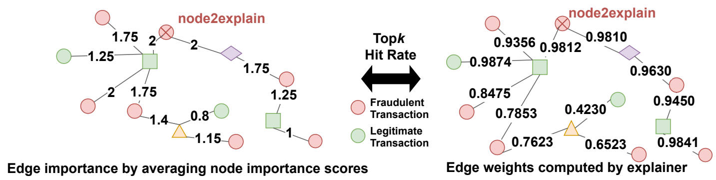

Figure 6: Edge importance scores (left) by human and edge weights by the explainer (right).

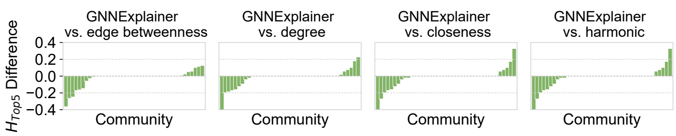

Figure 7: GNNExplainer and centrality measures work well for different communities, forming a trade-off.

coefficient A and explainer coefficient B to combine the weights from different explanation mechanisms: $A * w(c) + B * w(e)$. We can learn these two coefficients by either Ridge regression or directly maximizing the hit rate on the training set.

4 EXPERIMENTS OF XFRAUD DETECTOR+

We conduct extensive experiments on real-world transaction datasets sampled from the eBay commerce platform to verify the efficacy and efficiency of xFraud detector+. The statistics of the datasets are summarized in Table 2. The details on the graph construction process are in Appendix B [34]. We run end-to-end experiments on *eBay-xlarge* as it is a superset of *eBay-large* and *eBay-small*. Specifically, we run the distributed version of xFraud detector+ since *eBay-xlarge* is too large to be fit into a single machine. In the graph partition process, we first split the whole graph into 128 subgraphs using PIC. We then organize these 128 subgraphs into κ groups, where κ is the number of workers[3]. Different groups are handled by different workers. After the end-to-end evaluation (Sec. 4.1), we conducted an ablation study (Sec. 4.2) on *eBay-large* and *eBay-small* to study the trade-off between xFraud detector (i.e., HGT) and xFraud detector+.

4.1 End-to-end Experiments

We report the end-to-end results on *eBay-xlarge* in Table 3. In addition, we show precision-recall curve and ROC curve in Figure 8 and Figure 9 to further study prediction performance as *eBay-xlarge* is an extremely imbalanced graph dataset.

End-to-end Results. From Table 3, our detector+, achieves the best AUC (averaged across seeds) using 8 machines w.r.t. GEM and GAT. In terms of training efficiency, xFraud detector+ takes only slightly longer time per epoch compared to GEM in an 8-machine setting. If we increase the number of machines to 16, the training time is reduced by 1.89×, 1.84× and 1.82× for GAT, GEM, and our detector+, respectively. Despite roughly linear speedups, we observe lowered AUC compared with 8 machines. In our implementation,

[3]We first order the 128 subgraphs according to the total number of nodes in ascending order. Then, we put the first few subgraphs that cumulatively have $\lceil \frac{|\mathcal{V}|}{\kappa} \rceil$ nodes into the same group. We repeat this process until we get κ groups. In this way, we ensure that each machine receives a graph partition of similar total number of nodes.

Table 2: Dataset summary ("B":billion;"M":million;"K":thousand)
***The ratio of frauds is only reported on the sampled datasets.**

Dataset	Features	Graph type	#Nodes	#Edges	Fraud%*
eBay-xlarge	480	hetero	1.1B	3.7B	4.33%
eBay-small	114	hetero	289K	613K	4.30%
eBay-large	480	hetero	8.9M	13.2M	3.57%

Table 3: End-to-end performance on the dataset *eBay-xlarge* (epochs: 128). We report the average scores over two different seeds (A and B).

# machines	Model	AUC	Training time (s/epoch)	Inference time (s/batch)
8	GAT	0.8879	62.74	0.0557 ± 0.1966
	GEM	0.8961	**61.77**	**0.0167** ± **0.0054**
	xFraud detector+	**0.9074**	70.47	0.0799 ± 0.1868
16	GAT	0.8866	**33.11** (1.89×)	0.0557 ± 0.1966
	GEM	**0.8938**	33.56 (1.84×)	**0.0167** ± **0.0054**
	xFraud detector+	0.8892	38.72 (1.82×)	0.0799 ± 0.1868

each machine only has a subgraph; therefore, all three methods obtain a suboptimal model due to a restrained field of neighbors and edges. This phenomenon reveals a trade-off about our current way of handling large-scale graphs — one can use more resources to accelerate the model training but might have to compromise the model performance. It is an interesting future work to understand how to develop better distributed algorithms for training heterogeneous graph models. For more details on the system implementation and results, see Appendix C [34]. GEM (8 machines) takes the shortest time to do inference over a batch of 640 nodes due to the simplicity of its convolution layers. GAT and xFraud have longer inference time because their implementations of attention mechanisms. xFraud takes slightly longer than GAT due to its attention on heterogeneous types of nodes and edges. Since all methods take less than 0.1 second for a batch, all of them are practical to be deployed in production. Overall, xFraud is appealing in fraud detection, as it achieves the best model quality with a reasonably fast inference speed. Since xFraud detector+ is scalable to 16 machines, an online production scenario using xFraud can leverage historical and up-to-date transaction records to incrementally train a detector (Appendix H [34]).

Precision-recall Curve (P/R Curve). Figure 8 illustrates the P/R Curve using different settings. The trade-off of precision and recall is an ever-lasting goal for machine learning models; and xFraud detector+ achieves a better balance between precision and recall compared to GAT and GEM, which means that our model can return more accurate results (higher precision) as well as most of the true fraudulent transactions being found (higher recall).

ROC Curve. In *eBay-xlarge*, the majority of transactions are benign and the ratio of fraud transactions is very low. Besides, from the application perspective, the task of fraud detection is vulnerable to false positive cases, when benign transactions are flagged as fraud. This would cause an overwhelming human verification and significantly worsen user experience. Therefore, it is important to study the imbalance-aware metrics like true positive rate (FPR) and false positive rate (TPR). Specifically, if we restrict the FPR being lower than 0.1 as in Figure 9 (even this small ratio could

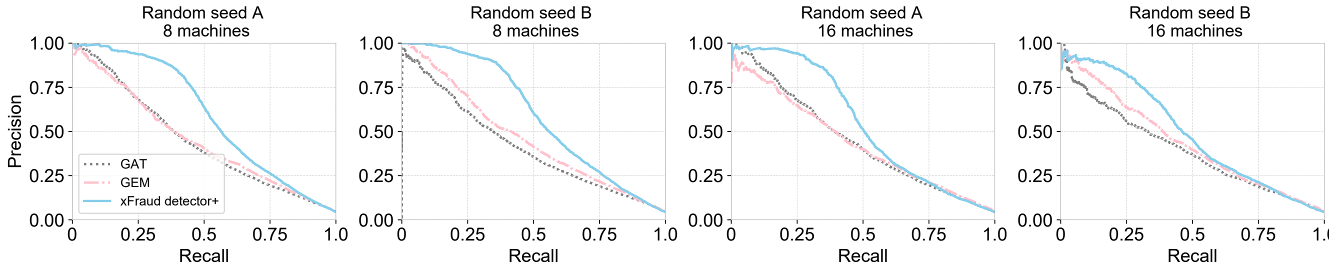

Figure 8: Precision-recall curves using different settings (seeds and # machines) on *eBay-xlarge*.

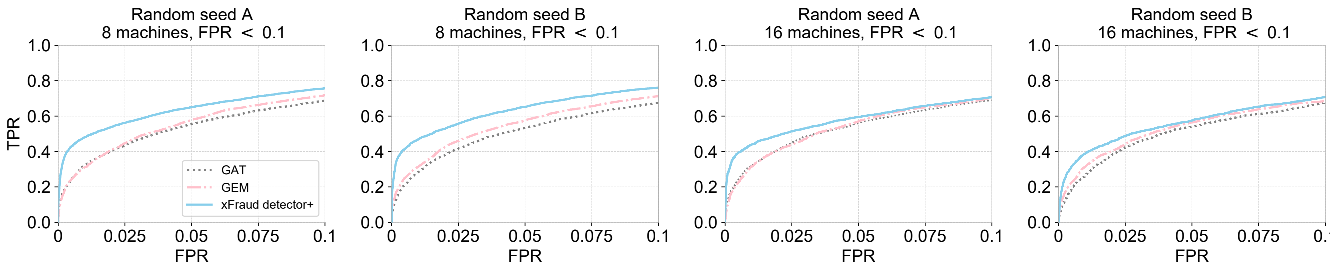

Figure 9: ROC curves using different settings (seeds and # machines) on *eBay-xlarge* for FPR < 0.1.

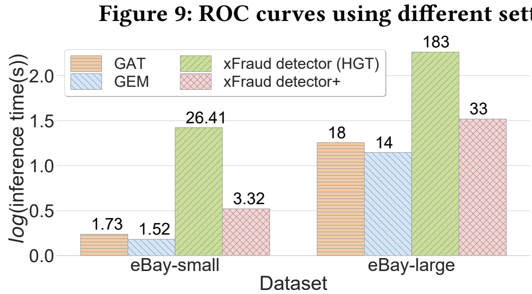

Figure 10: Total inference time (in *log*) on the test set of *eBay-small* and *eBay-large* (actual time in parentheses).

involve 85.7M transactions in *eBay-xlarge*), xFraud significantly outperforms GAT and GEM when only a small FPR is allowed. We plot the full range of FPR in Figure 15 of Appendix H [34], where the three models have a similar area under ROC curve (i.e., AUC-ROC), xFraud's ROC curve is consistently beyond GAT and GEM.

Discussion. All results we report are on a dataset after pre-filtering the fraudulent/benign transactions with rule/ML-based filters and down-sampling the benign transactions. In Appendix H [34] we discuss the implication of this pre-filtering step and the production scenario of xFraud. Even without downsampling benign transactions, xFraud achieves a reasonable precision and recall on industrial-scale data: from 3 fraud candidates investigated by the business unit, 1 will be a real fraud, with 0.1 of recall.

4.2 Ablation Study: xFraud detector+ vs. xFraud detector (i.e., HGT)

Here, we conduct an ablation study of xFraud detector (i.e., HGT) and xFraud detector+ to demonstrate the efficiency of the sampler. We run experiments on *eBay-small* and *eBay-large*, which are subsets of *eBay-xlarge* because our *eBay-xlarge* is too large such that xFraud detector (i.e., HGT) can no longer handle it. We report the inference time and AUC using a single machine in Figure 10. Concretely, xFraud detector+ achieves a 5 × speedups in terms of inference time (during testing) on *eBay-large* compared with xFraud detector (i.e., HGT), and the speedup on *eBay-small* is even larger, i.e., up to 7 ×. Meanwhile, using a simplified yet efficient sampler (GraphSAGE) will not sacrifice the model AUC (*eBay-small* vs. *eBay-large*): 0.7248 vs. 0.8683 for HGT, and 0.7262 vs. 0.8690 for xFraud detector+. Interestingly, we observe that the xFraud detector+ can even slightly outperform xFraud detector (i.e., HGT) on both datasets in terms of AUC.

5 EXPERIMENTS OF XFRAUD EXPLAINER

In this section, we discuss how we build an xFraud explainer on top of the detector. The main contribution of the explainer is to compute node features and edge masks of important features and nodes. As we see in the quantitative analysis (Sec. 5.1) and the case studies (Sec. 5.2), xFraud explainer ranks important edges with high agreement with expert human annotators and provides interesting insights for risk experts when analyzing risk propagation in a local heterogeneous graph. The output of the explainer carries the following meaning: node feature masks give high weights to the node feature dimensions influential in prediction; edge masks are the weights of edges in the subgraph, which indicate the strength of connectedness between pairs of nodes when flagging fraud. We visualize the subgraph structure with these outputs.

5.1 Quantitative Analysis of xFraud Explainer: Top*k* Hit Rate

First, we want to quantify the efficacy of xFraud explainer by studying the agreement between human perception and explainer output. Our approach is to compute the agreement between the edge importance scores based on human annotations and the edge weights generated by our hybrid explainer.

Sample. In total, we randomly select 41 communities from our test set, among which 18 communities has a fraudulent transaction

Table 4: Top$k$ hit rate in the test communities.

H(_)	Edge betweenness $H(c)$	GNNExplainer $H(e)$	Hybrid (ridge) $H(h)$	Hybrid (grid) $H(h)$
Top5	0.45540	0.44800	0.44890	**0.45550**
Top10	0.78175	0.77580	**0.81115**	0.78700
Top15	0.87763	0.88473	0.88963	**0.89410**
Top20	0.96205	0.95840	0.96198	**0.96275**
Top25	**0.96616**	0.95954	0.96614	0.96614

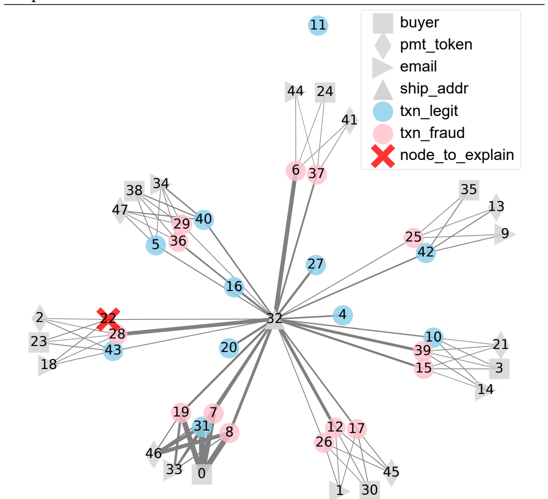

Figure 11: TP: xFraud helps to catch potential frauds.

as seed (label 1), 23 a legitimate transaction as seed (label 0). The AUC score of this test sample is 81.88%. A community is formed around a transaction seed node, where all connected nodes and edges are taken. In total, we have 1,591 nodes of five types (buyer, transaction, shipping address, email, and payment token), and 3,344 edges.[4] On average, there are 81.56 edges per community.

Human annotations and edge importance score. We have created the human evaluation of edge importance scores of all edges in these 41 communities. The annotation protocol and score calculations are listed in Appendix E [34]. Out of 41 communities, we take the first 21 communities as the training set, the last 20 as the test set. We trained two versions of the hybrid explainer: (1) via ridge regression on the human annotations on the training set, and (2) via directly optimizing the average hit rate on the training set.

Results. Table 4 illustrates the result. We see that the hybrid explainer consistently outperforms both GNNExplainer and centrality measures. This is not surprising, as shown in our previous discussion of the tradeoff. It is an exciting future direction to come up with better ways to combine these different metrics together to form an even better explainer for graphs.

5.2 Qualitative Analysis of xFraud Explainer: Case Studies

To visualize the subgraph for a certain node, we use the node index, edge indices and their masks, true labels of nodes as input. The thicker an edge is, the stronger the connection is. We visualize the connections with nondirectional edges and use the ground truth labels for transactions. **True positive (TP): flagging frauds.** In Figure 11, we see a generic shipping address (node 32, a warehouse) connected to both fraudulent/benign transactions related to various buyers using various payment tokens/emails. According to BU, one explanation for this pattern is there is often a lag between user chargebacks and when the frauds have taken place, not to mention it is possible that a card stolen claim might never be forwarded to eBay from some banks[5]. As a result, we cannot fully trust the positive labels in such a case, where it is clearly unusual for such a community to have an extremely mixed benign/transactions across buyers. And it could also be a case where defaulters disguise their true purposes by "cultivating" some legitimate accounts to execute a few legit transactions. For the 2^{nd} assumption, the BU needs extra evidence to further examine. Either way, xFraud is able to flag the node-to-predict as fraudulent by learning from the important edges (the thicker ones), and to inform the BU that this set of buyers are highly suspicious and should be under more detailed examinations. This shows the importance in detecting frauds on the transaction level as we propose in xFraud, instead of just on the account level as in GEM [28]. Currently, the BU is only using a rule based system[6] to filter the suspicious transactions stored in the tabular format. xFraud explainer is innovative to a traditional BU annotation routine, because it allows experts to combine graph level and feature level information. For extensive case studies on **false positive (FP): benign → fraud** and **false negative (FN): fraud → benign**, we discuss in Appendix G [34], where we also discuss system limitations and potential solutions to improve xFraud.

6 CONCLUSION

In this paper, we propose xFraud, a system for detecting fraud transaction and explaining model prediction. Specifically, a heterogeneous graph is constructed and a self-attentive heterogeneous graph neural network is leveraged for risky transaction scoring. We further design a learnable hybrid explainer that leverages both GNNExplainer and centrality measures to learn node- and edge-level explanations simultaneously. Through extensive experiments on real-world datasets, we show the proposed xFraud detector+ can efficiently process billion-scale heterogeneous graphs and outperform the competitive baselines. More importantly, xFraud is the first work that quantifies a strong agreement between human perception and explainer outputs. Real-world case studies illustrate that with the hybrid xFraud explainer, we can generate convincing explanations to assist further decision-making of business units.

ACKNOWLEDGMENTS

CZ and the DS3Lab gratefully acknowledge the support from Swiss National Science Foundation (Project Number 200021_184628 and 197485), Innosuisse/SNF BRIDGE Discovery (Project Number 40B2-0_187132), European Union Horizon 2020 Research and Innovation Programme (DAPHNE, 957407), Botnar Research Centre for Child Health, Swiss Data Science Center, Alibaba, Cisco, eBay, Google Focused Research Awards, Kuaishou Inc., Oracle Labs, Zurich Insurance, and the Department of Computer Science at ETH Zurich.

[4]Note that explainer assumes directions and assigns two weights to bidirectional edges connecting a pair of nodes. Since human annotation is on the node level, and it is generally hard for annotators to consider directions, we remove directions in the explainer weights by taking the larger weight.

[5]This is out of control of eBay and is due to the inconsistency of reporting systems at some banks.

[6]*skope-rules* which perform rule mining on tabular data, https://github.com/scikit-learn-contrib/skope-rules (last accessed: Oct 18, 2020).

REFERENCES

[1] Bart Baesens, Veronique Van Vlasselaer, and Wouter Verbeke. 2015. *Fraud analytics using descriptive, predictive, and social network techniques: a guide to data science for fraud detection.* John Wiley & Sons.

[2] Federico Baldassarre and Hossein Azizpour. 2019. Explainability techniques for graph convolutional networks. *arXiv preprint arXiv:1905.13686* (2019).

[3] Adam Breuer, Roee Eilat, and Udi Weinsberg. 2020. Friend or Faux: Graph-Based Early Detection of Fake Accounts on Social Networks. In *Proceedings of The Web Conference 2020.* 1287–1297.

[4] Bokai Cao, Mia Mao, Siim Viidu, and S Yu Philip. 2017. HitFraud: a broad learning approach for collective fraud detection in heterogeneous information networks. In *2017 IEEE international conference on data mining (ICDM).* IEEE, 769–774.

[5] Shaosheng Cao, XinXing Yang, Cen Chen, Jun Zhou, Xiaolong Li, and Yuan Qi. 2019. TitAnt: online real-time transaction fraud detection in Ant Financial. *Proceedings of the VLDB Endowment* (2019).

[6] Yukuo Cen, Xu Zou, Jianwei Zhang, Hongxia Yang, Jingren Zhou, and Jie Tang. 2019. Representation learning for attributed multiplex heterogeneous network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining.* 1358–1368.

[7] Shiyu Chang, Wei Han, Jiliang Tang, Guo-Jun Qi, Charu C Aggarwal, and Thomas S Huang. 2015. Heterogeneous network embedding via deep architectures. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining.* 119–128.

[8] Sarthika Dhawan, Siva Charan Reddy Gangireddy, Shiv Kumar, and Tanmoy Chakraborty. 2019. Spotting collective behaviour of online frauds in customer reviews. *arXiv preprint arXiv:1905.13649* (2019).

[9] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining.* 135–144.

[10] Paul Emmerich, Maximilian Pudelko, Sebastian Gallenmüller, and Georg Carle. 2017. FlowScope: Efficient packet capture and storage in 100 Gbit/s networks. In *2017 IFIP Networking Conference (IFIP Networking) and Workshops.* IEEE, 1–9.

[11] Dhivya Eswaran, Stephan Günnemann, Christos Faloutsos, Disha Makhija, and Mohit Kumar. 2017. Zoobp: Belief propagation for heterogeneous networks. *Proceedings of the VLDB Endowment* 10, 5 (2017), 625–636.

[12] V. Fomin, J. Anmol, S. Desroziers, J. Kriss, and A. Tejani. 2020. High-level library to help with training neural networks in PyTorch. https://github.com/pytorch/ignite.

[13] Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. 2020. Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding. In *Proceedings of The Web Conference 2020.* 2331–2341.

[14] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in neural information processing systems.* 1024–1034.

[15] Bryan Hooi, Hyun Ah Song, Alex Beutel, Neil Shah, Kijung Shin, and Christos Faloutsos. 2016. Fraudar: Bounding graph fraud in the face of camouflage. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* 895–904.

[16] Binbin Hu, Yuan Fang, and Chuan Shi. 2019. Adversarial learning on heterogeneous information networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining.* 120–129.

[17] Binbin Hu, Zhiqiang Zhang, Chuan Shi, Jun Zhou, Xiaolong Li, and Yuan Qi. 2019. Cash-out user detection based on attributed heterogeneous information network with a hierarchical attention mechanism. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 946–953.

[18] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous graph transformer. In *Proceedings of The Web Conference 2020.* 2704–2710.

[19] Qiang Huang, Makoto Yamada, Yuan Tian, Dinesh Singh, Dawei Yin, and Yi Chang. 2020. GraphLIME: Local interpretable model explanations for graph neural networks. *arXiv preprint arXiv:2001.06216* (2020).

[20] Parisa Kaghazgaran, James Caverlee, and Anna Squicciarini. 2018. Combating crowdsourced review manipulators: A neighborhood-based approach. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining.* 306–314.

[21] Srijan Kumar, Bryan Hooi, Disha Makhija, Mohit Kumar, Christos Faloutsos, and VS Subrahmanian. 2018. Rev2: Fraudulent user prediction in rating platforms. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining.* 333–341.

[22] Ao Li, Zhou Qin, Runshi Liu, Yiqun Yang, and Dong Li. 2019. Spam review detection with graph convolutional networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management.* 2703–2711.

[23] Xiang Li, Wen Zhang, Jiuzhou Xi, and Hao Zhu. 2018. HGsuspector: Scalable Collective Fraud Detection in Heterogeneous Graphs. (2018).

[24] Chen Liang, Ziqi Liu, Bin Liu, Jun Zhou, Xiaolong Li, Shuang Yang, and Yuan Qi. 2019. Uncovering Insurance Fraud Conspiracy with Network Learning. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval.* 1181–1184.

[25] Frank Lin and William W Cohen. 2010. Power iteration clustering. In *ICML*.

[26] Shenghua Liu, Bryan Hooi, and Christos Faloutsos. 2017. Holoscope: Topology-and-spike aware fraud detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management.* 1539–1548.

[27] 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 *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 4424–4431.

[28] Ziqi Liu, Chaochao Chen, Xinxing Yang, Jun Zhou, Xiaolong Li, and Le Song. 2018. Heterogeneous graph neural networks for malicious account detection. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management.* 2077–2085.

[29] Zhiwei Liu, Yingtong Dou, Philip S Yu, Yutong Deng, and Hao Peng. 2020. Alleviating the Inconsistency Problem of Applying Graph Neural Network to Fraud Detection. *arXiv preprint arXiv:2005.00625* (2020).

[30] Qingsong Lv, Ming Ding, Qiang Liu, Yuxiang Chen, Wenzheng Feng, Siming He, Chang Zhou, Jianguo Jiang, Yuxiao Dong, and Jie Tang. 2021. Are we really making much progress? Revisiting, benchmarking, and refining heterogeneous graph neural networks. (2021).

[31] Jun Ma, Danqing Zhang, Yun Wang, Yan Zhang, and Alexey Pozdnoukhov. 2018. GraphRAD: A Graph-based Risky Account Detection System. (2018).

[32] Wei Min, Zhengyang Tang, Min Zhu, Yuxi Dai, Yan Wei, and Ruinan Zhang. 2018. Behavior language processing with graph based feature generation for fraud detection in online lending. In *Proceedings of Workshop on Misinformation and Misbehavior Mining on the Web.*

[33] Hamed Nilforoshan and Neil Shah. 2019. SliceNDice: Mining Suspicious Multi-Attribute Entity Groups with Multi-View Graphs. In *2019 IEEE International Conference on Data Science and Advanced Analytics (DSAA).* IEEE, 351–363.

[34] Susie Xi Rao, Shuai Zhang, Zhichao Han, Zitao Zhang, Wei Min, Zhiyao Chen, Yinan Shan, Yang Zhao, and Ce Zhang. 2021. Appendix for xFraud: Explainable Fraud Transaction Detection. https://github.com/eBay/xFraud/blob/master/documents/Appendix_XFraud_VLDB.pdf.

[35] Yuxiang Ren, Hao Zhu, Jiawei ZHang, Peng Dai, and Liefeng Bo. 2019. EnsemFDet: An Ensemble Approach to Fraud Detection based on Bipartite Graph. *arXiv preprint arXiv:1912.11113* (2019).

[36] Chuan Shi, Binbin Hu, Wayne Xin Zhao, and S Yu Philip. 2018. Heterogeneous information network embedding for recommendation. *IEEE Transactions on Knowledge and Data Engineering* 31, 2 (2018), 357–370.

[37] Yu Shi, Fangqiu Han, Xinwei He, Xinran He, Carl Yang, Jie Luo, and Jiawei Han. 2018. mvn2vec: Preservation and collaboration in multi-view network embedding. *arXiv preprint arXiv:1801.06597* (2018).

[38] Kai Shu, Deepak Mahudeswaran, Suhang Wang, and Huan Liu. 2020. Hierarchical propagation networks for fake news detection: Investigation and exploitation. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 626–637.

[39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems.* 5998–6008.

[40] Haibo Wang, Chuan Zhou, Jia Wu, Weizhen Dang, Xingquan Zhu, and Jilong Wang. 2018. Deep structure learning for fraud detection. In *2018 IEEE International Conference on Data Mining (ICDM).* IEEE, 567–576.

[41] Jianyu Wang, Rui Wen, Chunming Wu, Yu Huang, and Jian Xion. 2019. Fdgars: Fraudster detection via graph convolutional networks in online app review system. In *Companion Proceedings of The 2019 World Wide Web Conference.* 310–316.

[42] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. 2019. Heterogeneous graph attention network. In *The World Wide Web Conference.* 2022–2032.

[43] Mark Weber, Giacomo Domeniconi, Jie Chen, Daniel Karl I Weidele, Claudio Bellei, Tom Robinson, and Charles E Leiserson. 2019. Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics. *arXiv preprint arXiv:1908.02591* (2019).

[44] Rui Wen, Jianyu Wang, Chunming Wu, and Jian Xiong. 2020. ASA: Adversary Situation Awareness via Heterogeneous Graph Convolutional Networks. In *Companion Proceedings of the Web Conference 2020.* 674–678.

[45] P. Reddy X. Li, J. Saude and M. Veloso. 2020. Classifying and understanding financial data using graph neural network. In *AAAI*.

[46] Carl Yang, Yuxin Xiao, Yu Zhang, Yizhou Sun, and Jiawei Han. 2020. Heterogeneous network representation learning: A unified framework with survey and benchmark. *IEEE Transactions on Knowledge and Data Engineering* (2020).

[47] Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. 2019. Gnnexplainer: Generating explanations for graph neural networks. In *Advances in neural information processing systems.* 9244–9255.

[48] Hao Yuan, Jiliang Tang, Xia Hu, and Shuiwang Ji. 2020. XGNN: Towards Model-Level Explanations of Graph Neural Networks. *arXiv preprint arXiv:2006.02587* (2020).

[49] Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. 2019. Graph transformer networks. *Advances in Neural Information Processing Systems* 32 (2019), 11983–11993.

[50] Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V Chawla. 2019. Heterogeneous graph neural network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 793–803.

[51] Yiming Zhang, Yujie Fan, Yanfang Ye, Liang Zhao, and Chuan Shi. 2019. Key Player Identification in Underground Forums over Attributed Heterogeneous Information Network Embedding Framework. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 549–558.

[52] Kai Zhao, Ting Bai, Bin Wu, Bai Wang, Youjie Zhang, Yuanyu Yang, and Jian-Yun Nie. 2020. Deep adversarial completion for sparse heterogeneous information network embedding. In *Proceedings of the Web Conference 2020*. 508–518.

[53] 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 *Proceedings of The Web Conference 2020*. 785–795.

[54] Yongchun Zhu, Dongbo Xi, Bowen Song, Fuzhen Zhuang, Shuai Chen, Xi Gu, and Qing He. 2020. Modeling Users' Behavior Sequences with Hierarchical Explainable Network for Cross-domain Fraud Detection. In *Proceedings of The Web Conference 2020*. 928–938.