

HitFraud: A Broad Learning Approach for Collective Fraud Detection in Heterogeneous Information Networks

Bokai Cao*, Mia Mao†, Siim Viidu† and Philip S. Yu*‡

*Department of Computer Science, University of Illinois at Chicago, IL, USA; {caobokai, psyu}@uic.edu

†Electronic Arts, Redwood City, CA, USA; {mmao, sviidu}@ea.com

‡Institute for Data Science, Tsinghua University, Beijing, China

Abstract—On electronic game platforms, different payment transactions have different levels of risk. Risk is generally higher for digital goods in e-commerce. However, it differs based on product and its popularity, the offer type (packaged game, virtual currency to a game or subscription service), storefront and geography. Existing fraud policies and models make decisions independently for each transaction based on transaction attributes, payment velocities, user characteristics, and other relevant information. However, suspicious transactions may still evade detection and hence we propose a broad learning approach leveraging a graph based perspective to uncover relationships among suspicious transactions, *i.e.*, inter-transaction dependency. Our focus is to detect suspicious transactions by capturing common fraudulent behaviors that would not be considered suspicious when being considered in isolation. In this paper, we present HITFRAUD that leverages heterogeneous information networks for collective fraud detection by exploring correlated and fast evolving fraudulent behaviors. First, a heterogeneous information network is designed to link entities of interest in the transaction database via different semantics. Then, graph based features are efficiently discovered from the network exploiting the concept of meta-paths, and decisions on frauds are made collectively on test instances. Experiments on real-world payment transaction data from ELECTRONIC ARTS demonstrate that the prediction performance is effectively boosted by HITFRAUD with fast convergence.

Index Terms—collective fraud detection, inter-transaction dependency, heterogeneous information network.

I. INTRODUCTION

Fraud detection has attracted significant research efforts in recent years for various tasks including finance, security and web services. In this work, we investigate the fraud detection problem on electronic game platforms where it is desirable to identify suspicious payment transactions in an early stage in order to avoid chargebacks and to enhance normal users' experience. Current fraud models make independent decisions for each transaction, and detection becomes harder when intelligent adversaries are used, *e.g.*, proxy IP addresses. Graph based methods can detect frauds by leveraging the linkage information between entities of interest [3, 4, 6]. Such methods are relatively harder to evade because making a fraud payment transaction unavoidably generates links in the graph which reveals *inter-transaction dependency*.

This work was partially done while the author was an intern at Electronic Arts.

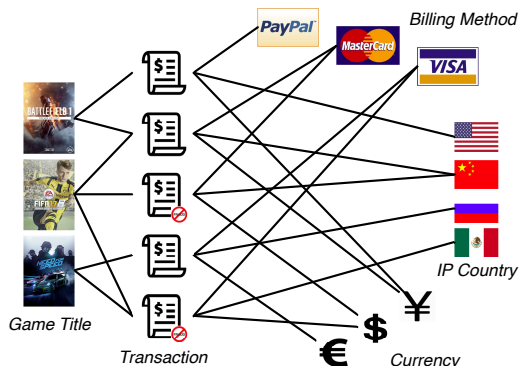


Fig. 1. A simplified example of the constructed heterogeneous information network.

Existing graph based fraud detection approaches heavily focus on homogeneous information networks and bipartite graphs. Heterogeneous information networks (HINs) [9] are a special type of information networks that involve multiple types of nodes or multiple types of links, as shown in Figure 1. In a HIN, different types of nodes and links have different semantic meanings. Such complex and semantically enriched networks possess great potential for knowledge discovery [5, 7, 10]. Its applications to fraud detection, however, are largely unexplored. Machine learning approaches that leverage a variety of data sources or fuse heterogeneous information can be referred to as *broad learning*.

Therefore, we are motivated to investigate how to leverage HINs to facilitate the fraud detection task. Most importantly, we seek to capture the inter-transaction dependency. It is critical to explore such relationships among suspicious transactions because fraudulent behaviors are often *correlated* and *fast evolving*. (1) HINs provide us with an effective and compact representation of linked transactions in various semantics, *e.g.*, the same currency, the same IP address, and the same game titles. The statistics of the label information (*i.e.*, fraud or normal) of these linked transactions can be aggregated, and thereby add a new dimension of measurements to distinguish suspicious transactions from normal ones based on the correlated fraudulent behaviors. (2) In order to tackle the problem of fast evolving fraudulent behaviors, we should

not only consider the inter-transaction dependency across training transactions and test transactions, but also include the dependency among test transactions. Hence, suspicious transactions are identified in a semi-supervised manner by iteratively obtaining the predicted labels of test transactions and updating the statistics of linked transactions in alternation. Such a collective prediction procedure has the potential to detect a suspicious transaction even if it appears to be normal by itself but its linked transactions (other test transactions in a batch sharing categorical variables) are identified as very suspicious, and thereby improve the recall metric. The main contributions of this work are fourfold:

- As fraud payment transactions generally do not occur in isolation, *i.e.*, fraudulent behaviors are often correlated and fast evolving, we formulate the fraud detection task as a collective prediction problem in a HIN to capture relationships among fraud payment transactions. (Section II)
- To address the daunting challenge on huge feature space, we design a HIN that can effectively capture the various relationships among transactions. Here meta-paths (a sequence of link types) are explored to identify the relevant inter-transaction dependency features. (Section III)
- We propose an effective and efficient algorithm to compute meta-path based features in the framework of collective fraud detection. (Section IV)
- We validate that the observation that fraudulent behaviors are often correlated and fast evolving can indeed be explored to more effectively capture fraud payment transactions. We evaluate the proposed framework on real-world payment transaction data from ELECTRONIC ARTS payment system, and results show that recall and F-score can significantly be improved by exploring inter-transaction dependency via the proposed collective fraud detection method based on HINs. (Section V)

A full version of this paper is available in [2].

II. PROBLEM DEFINITION

Heterogeneous Information Network (HIN). A heterogeneous information network is a type of information network with multiple types of nodes or multiple types of links [9, 10]. It can be represented as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. $\mathcal{V} = \mathcal{V}^1 \cup \dots \cup \mathcal{V}^m$ denotes the set of nodes involving m node types: $\mathcal{V}^1 = \{v_1^1, \dots, v_{n_1}^1\}, \dots, \mathcal{V}^m = \{v_1^m, \dots, v_{n_m}^m\}$ where v_p^i represents the p -th node of type i . $\mathcal{E} = \mathcal{E}^1 \cup \dots \cup \mathcal{E}^r \subseteq \mathcal{V} \times \mathcal{V}$ denotes the set of links between nodes in \mathcal{V} involving r link types. Mathematically, a link type k starting from source nodes of type i and ending at target nodes of type j is described by an adjacency matrix $\mathbf{A}^k \in \mathbb{R}^{n_i \times n_j}$ where $\mathbf{A}^k[p, q] = 1$ if there exists a link in \mathcal{E}^k between v_p^i and v_q^j , otherwise $\mathbf{A}^k[p, q] = 0$. We can write this link type as “ $\mathcal{V}^i \xrightarrow{\mathcal{E}^k} \mathcal{V}^j$ ”.

Collective Fraud Detection. We may assume without loss of generality that nodes in \mathcal{V}^1 are the target entities, *i.e.*, transactions, in the case of fraud detection. The number of target entities is denoted as $n = n_1$. Formally, we are given a data matrix $\mathbf{X} = [\mathbf{x}_1^T; \dots; \mathbf{x}_n^T]$ where $\mathbf{x}_i \in \mathbb{R}^d$ is the

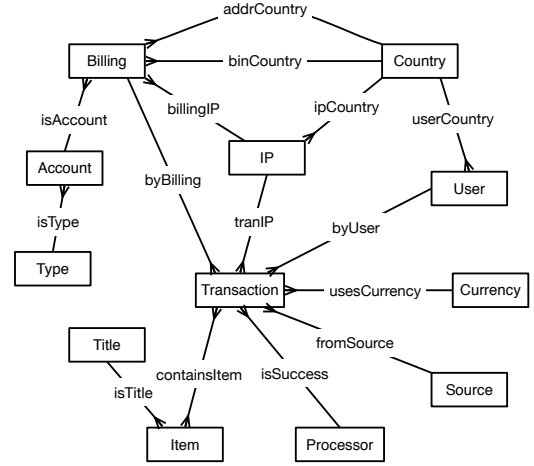


Fig. 2. The network schema of EA payment transaction data. Each rectangle represents a node type, and each line represents a link type.

feature vector of the i -th transaction which typically involves extensive feature engineering and domain knowledge. Labels are denoted as $\mathbf{y} = [y_1, \dots, y_n]$ where $y_i = 1$ if the i -th transaction is a fraud, otherwise $y_i = 0$. The target is to identify the label y_u of an unseen transaction \mathbf{x}_u (typically, with different items, user accounts, billing accounts, and IP addresses, *etc.*).

In this paper, we investigate how to leverage a HIN to facilitate the fraud detection task. The problem of collective fraud detection in a HIN reduces to learning a predictive function $f : (\mathcal{V}, \mathcal{E}, \mathbf{X}) \rightarrow \mathbf{y}$. It is a nontrivial problem due to threefold challenges:

- We need to design a network that can capture various relationships among transactions.
- We need to discover relevant network features efficiently, whose time complexity should ideally be linear to the number of transactions and insensitive to the type of relationships.
- We need to effectively explore dependency among suspicious transactions, especially among the test ones, in order to detect the correlated and fast evolving fraudulent behaviors.

III. DATASET

In this work, we use ELECTRONIC ARTS (EA) payment transaction data as an example to do the study. All transactions on EA digital platform go through a set of policies, rules and models to determine their levels of risk, and a subset of them are sent for additional manual review. An experienced team reviews those transactions and decides if they should be rejected or approved, and this review decision is used as the ground-truth for training and evaluating fraud detection algorithms. We collected manual review data for $n = 130K$ transactions during a recent period. Each transaction is associated with a $d = 2K$ dimensional feature vector, including transaction attributes, payment velocities, user characteristics, and other relevant information.

TABLE I
TOP-10 DISCRIMINATIVE META-PATHS AND THEIR SEMANTICS.

ID	Meta-path	Semantics
#1	transaction $\xrightarrow{byBilling}$ billing $\xrightarrow{billingIP}$ IP $\xrightarrow{billingIP^{-1}}$ billing $\xrightarrow{byBilling^{-1}}$ transaction	from the same billing IP address
#2	transaction \xrightarrow{tranIP} IP $\xrightarrow{tranIP^{-1}}$ transaction	from the same transaction IP address
#3	transaction $\xrightarrow{byBilling}$ billing $\xrightarrow{byBilling^{-1}}$ transaction	using the same billing account
#4	transaction \xrightarrow{byUser} user $\xrightarrow{byUser^{-1}}$ transaction	using the same user account
#5	transaction \xrightarrow{tranIP} IP $\xrightarrow{billingIP^{-1}}$ billing $\xrightarrow{byBilling^{-1}}$ transaction	transaction IP and billing IP are the same
#6	transaction $\xrightarrow{byBilling}$ billing $\xrightarrow{billingIP}$ IP $\xrightarrow{tranIP^{-1}}$ transaction	billing IP and transaction IP are the same
#7	transaction $\xrightarrow{containsItem}$ item $\xrightarrow{containsItem^{-1}}$ transaction	containing the same item
#8	transaction $\xrightarrow{containsItem}$ item $\xrightarrow{isTitle}$ title $\xrightarrow{isTitle^{-1}}$ item $\xrightarrow{containsItem^{-1}}$ transaction	containing the same game title
#9	transaction $\xrightarrow{fromSource}$ source $\xrightarrow{fromSource^{-1}}$ transaction	from the same source
#10	transaction \xrightarrow{byUser} user $\xrightarrow{userCountry}$ country $\xrightarrow{userCountry^{-1}}$ user $\xrightarrow{byUser^{-1}}$ transaction	from the same user country

A HIN is constructed by linking entities of interest from several selected databases. Transactions are the target instances on which fraud decisions are made, so each transaction ID is represented as a node in the network, and the set of transaction IDs compose a node type in the network schema. In addition, other entities that are directly or indirectly related to a transaction are considered here, and they compose other node types in the schema, including billing accounts, user accounts, game titles, IP addresses, *etc.* Links are added based on common semantics. For example, a transaction is linked with a user if the user placed the transaction, and a transaction is linked with an item if the transaction contains the item.

As a result, the constructed HIN is composed of over 400K nodes and 1.5M links. It integrates data involving $m = 12$ types of nodes, such as *transaction*, *user*, *item*, *title*, *currency*, *source*, *country*, *etc.* which are connected through $r = 15$ types of links, such as “transaction $\xrightarrow{containsItem}$ item”, “billing $\xrightarrow{billingIP}$ IP”, *etc.* Its network schema is shown in Figure 2. Note that *billing* and *country* are directly linked in two different semantics: country corresponding to the billing address, and country associated with the Bank Identification Number (BIN). Each billing account is designed to be associated with one country following each semantics.

IV. METHOD

In this paper, we present a broad learning approach, named HITFRAUD, for collective fraud detection in HINs by capturing inter-transaction dependency. The proposed framework is outlined as:

- 1) The aforementioned HIN is designed so that it contains entities of interest as nodes, *e.g.*, transactions, billing accounts, user accounts, game titles, and IP addresses. Two entities are linked if they appear together in a record.
- 2) Efficient ways of finding downsized meta-paths from transaction nodes are explored using breadth-first-search on the network schema. Meta-path based features are computed by pairing the pre-computed downsized meta-paths.

- 3) The predicted labels of test transactions are obtained from a trained classifier, and meta-path based features are updated using both the labels of training transactions and the predicted labels of test transactions, in alternation until convergence.

A. Capturing Inter-Transaction Dependency

First, we briefly review the concept of *meta-path* following previous work [1, 7, 9]. It has been demonstrated to be useful for mining social networks, bioinformatic networks, and bibliographic networks [1, 7, 8, 9].

In general, a meta-path corresponds to a type of path within the network schema, containing a certain sequence of link types. For example, in Figure 2, a meta-path “transaction $\xrightarrow{containsItem}$ item $\xrightarrow{isTitle}$ title $\xrightarrow{isTitle^{-1}}$ item $\xrightarrow{containsItem^{-1}}$ transaction” denotes a composite relation between transactions where $containsItem^{-1}$ represents the inverted relation of *containsItem*. The semantic meaning of this meta-path is that transactions contain items that belong to the same game title. Different meta-paths usually represent different semantic meanings between linked nodes. In this manner, various relationships among transactions can be described by a set of meta-paths. By capturing such inter-transaction dependency and aggregating the label information of the linked transactions, we could better detect correlated fraudulent behaviors. In other words, we could identify transactions with highly risky values in a categorical variable, *e.g.*, game title. Intuitively, when a recent launch of FIFA attracts many frauds to it, transactions that contain FIFA would be more suspicious than others. Similarly, transactions that are related to user accounts, billing accounts, IP addresses and currencies with high risk could also be identified. In Table I, we show several examples of meta-paths that are used in this work and their semantics.

The implementation of meta-paths is essentially a chain of matrix multiplications. Let’s denote a meta-path as $\mathcal{P} = \langle \mathcal{E}^{k_1}, \dots, \mathcal{E}^{k_l} \rangle$ where the source node of \mathcal{E}^{k_1} is of type s and the target node of \mathcal{E}^{k_l} is of type t . The semantic

meaning of this meta-path is mathematically described as $\mathbf{P} = \mathbf{A}^{k_1} \times \dots \times \mathbf{A}^{k_l} \in \mathbb{R}^{n_s \times n_t}$. Since \mathbf{P} is usually asymmetric, the meta-paths #5 and #6 in Table I are not identical. It is usually assumed that the strength of connection between v_p^s and v_q^t on such semantics is positively correlated with $\mathbf{P}[p, q]$, because $\mathbf{P}[p, q]$ is the (weighted) count of paths connecting v_p^s and v_q^t that follow the sequence of links in \mathcal{P} . Hereinafter, $\overline{\mathcal{P}}$ and \mathbf{P} will be used interchangeably when the meaning is clear from context.

B. Redundancy in Meta-Paths

Cardinality, which refers to the maximum number of times a node of the source node type can be linked with nodes of the target node type, is represented by the styling of a line and its endpoint in Figure 2. The notation style is similar to entity-relationship diagrams (ERDs) where a Crow’s foot shows many-to-one relationship.

Let’s consider a meta-path “transaction $\xrightarrow{\text{fromSource}}$ source $\xrightarrow{\text{fromSource}^{-1}}$ transaction” in Table I and denote it as $\mathcal{P} = \langle \mathcal{E}^{k_1}, \mathcal{E}^{k_2} \rangle$ where \mathcal{E}^{k_1} is the set of links indicating which source a transaction is from, and \mathcal{E}^{k_2} is its inverted relation. The adjacency matrix of \mathcal{E}^{k_1} is denoted as $\mathbf{A} \in \mathbb{R}^{n \times n_s}$ where $n = 130K$ is the number of transactions and n_s is the number of sources, then that of \mathcal{E}^{k_2} is \mathbf{A}^T , and $\mathbf{P} = \mathbf{A} \times \mathbf{A}^T$. Note that an inverted relation is described by the transpose of its original adjacency matrix rather than the inverse. Because each transaction is conducted on one of EA game stores, \mathcal{E}^{k_1} here is a many-to-one relation between transactions and sources. That is to say, \mathbf{A} is extremely sparse with one value per row, and its sparsity ratio is $1 - 1/n_s > 98\%$.

The computation of a meta-path through the matrix chain multiplication turns a sparse adjacency matrix \mathbf{A} into a dense matrix \mathbf{P} . Moreover, there are a lot of redundancy in \mathbf{P} . For this particular meta-path, if i -th transaction and the j -th transaction are from the same source k , *i.e.*, $\mathbf{A}[i, k] = \mathbf{A}[j, k] = 1$, they have exactly the same row and column in \mathbf{P} , *i.e.*, $\mathbf{P}[i, :] = \mathbf{P}[j, :] = \mathbf{P}[:, i] = \mathbf{P}[:, j]$.

C. Efficient Computation

Multiplying adjacency matrices of a meta-path in the natural sequence can be inefficient from the time perspective, considering the classic matrix chain multiplication problem. It can be optimized using dynamic programming in $O(l^3)$ where l is the length of a meta-path which is usually very small. In practice, however, the real difficulty in computing meta-paths lies in the space constraint. As we discussed above, turning sparse matrices into a full matrix is definitely not desirable because it makes the space cost from linear to quadratic with respect to the number of transactions.

Fortunately, we are not interested in the concrete form of a meta-path itself. For the purpose of obtaining features for fraud detection on transactions, the meta-paths that we need to compute should have the same source node type and target node type which is a transaction. Because each transaction may be linked with different number of transactions through a meta-path, aggregation functions are employed to combine

the label information of linked transactions in order to derive a fixed number of meta-path based features. For example, we can use the weighted label fraction of linked transactions as the feature $\mathbf{z} \in \mathbb{R}^n$ for each meta-path [7, 8]. It is formulated as follows:

$$\mathbf{z} = \mathbf{D} \times \mathbf{P} \times \mathbf{y} \quad (1)$$

where $\mathbf{D} \in \mathbb{R}^{n \times n}$ is a diagonal matrix and $\mathbf{D}[i, i] = 1/\sum_j \mathbf{P}[i, j]$. In this manner, z_i indicates the ratio of being frauds among transactions that are connected with the i -transaction through the meta-path.

1) *Many-to-One Cases*: Next, we will introduce how to obtain the feature \mathbf{z} without explicitly computing the meta-path \mathbf{P} . Since the concept of meta-path is defined as a sequence of link types, it can also be considered as a sequence of node types that are endpoints of these links (different link sequences can correspond to the same node sequence though), say, $\mathcal{P} = \langle \mathcal{V}^{k_0}, \dots, \mathcal{V}^{k_l} \rangle$. We have $k_0 = k_l = 1$ since we only consider meta-paths whose source node type and target node type are transaction. It is assumed that transaction is always the node set of the largest size, *i.e.*, $\text{argmax}_i \{|\mathcal{V}^i| \mid \mathcal{V}^i \in \mathcal{P}\} = 1$, and t denotes the node set of the smallest size in \mathcal{P} , *i.e.*, $\text{argmin}_i \{|\mathcal{V}^i| \mid \mathcal{V}^i \in \mathcal{P}\} = t$. Therefore, we can decompose a meta-path \mathcal{P} into two parts at the node type t , and its matrix form can be written as $\mathbf{P} = \mathbf{P}_1 \times \mathbf{P}_2^T$ where $\mathbf{P}_1, \mathbf{P}_2 \in \mathbb{R}^{n \times n_t}$.

DEFINITION 1 (Simple Meta-path): A meta-path is a simple meta-path if it is a sequence of many-to-one relations.

Note that a simple meta-path itself is a (composite) many-to-one relation. For example, “transaction $\xrightarrow{\text{byBilling}}$ billing $\xrightarrow{\text{isAccount}}$ account $\xrightarrow{\text{isType}}$ type” is a simple meta-path. All meta-paths in Table I are a concatenation of a simple meta-path and another meta-path except for the seventh and the eighth ones which will be discussed later. Their features can efficiently be computed.

LEMMA 1: Given a meta-path $\mathbf{P} = \mathbf{P}_1 \times \mathbf{P}_2^T$ where \mathbf{P}_1 is a simple meta-path, the computation of Eq. (1) on \mathbf{P} can be reduced to:

$$\mathbf{z} = \mathbf{P}_1 \times (\mathbf{D}_2 \times \mathbf{P}_2^T \times \mathbf{y}) \quad (2)$$

where $\mathbf{D}_2 \in \mathbb{R}^{n_t \times n_t}$ is a diagonal matrix and $\mathbf{D}_2[i, i] = 1/\sum_j \mathbf{P}_2[j, i]$.

PROOF 1: The detailed proof is presented in [2].

Eq. (2) is important because it enables us to obtain the weighted label fraction of linked nodes without extensive matrix operations. In this manner, we can effectively avoid computing the redundant full matrix of a meta-path as an intermediate result.

2) *Many-to-Many Cases*: Let’s consider the meta-path “transaction $\xrightarrow{\text{containsItem}}$ item $\xrightarrow{\text{isTitle}}$ title $\xrightarrow{\text{isTitle}^{-1}}$ item $\xrightarrow{\text{containsItem}^{-1}}$ transaction” in Table I. It involves a many-to-many link type, *i.e.*, “transaction $\xrightarrow{\text{containsItem}}$ item”, in Figure 2.

DEFINITION 2 (Complex Meta-path): A meta-path is a complex meta-path if it contains at least one many-to-many relation.

Note that a complex meta-path is “usually” a (composite) many-to-many relation. Now, we dichotomize all meta-paths into simple meta-paths and complex meta-paths. Given any nontrivial meta-path ($l > 1$), we still decompose it into two parts at the node type of the smallest size, $\mathbf{P} = \mathbf{P}_1 \times \mathbf{P}_2^T$, and we propose to compute its features as follows:

$$\mathbf{z} = \mathbf{D}_1 \times \mathbf{P}_1 \times (\mathbf{D}_2 \times \mathbf{P}_2^T \times \mathbf{y}) \quad (3)$$

where $\mathbf{D}_1 \in \mathbb{R}^{n \times n}$ is a diagonal matrix and $\mathbf{D}_1[i, i] = 1/\sum_j \mathbf{P}_1[i, j]$. Obviously, $\mathbf{D}_1 = \mathbf{I}_n$ where \mathbf{I}_n is an identity matrix when \mathbf{P}_1 is a simple meta-path. Therefore, Eq. (2) is a special case of Eq. (3) in the many-to-one scenarios.

However, Eq. (3) is not equivalent to Eq. (1) or Eq. (2) when \mathbf{P}_1 is a complex meta-path. For example, $\mathcal{P} = \text{“transaction} \xrightarrow{\text{containsItem}} \text{item} \xrightarrow{\text{isTitle}} \text{title} \xrightarrow{\text{isTitle}^{-1}} \text{item} \xrightarrow{\text{containsItem}^{-1}} \text{transaction} \text{”}$ can be decomposed as $\mathcal{P}_1 = \mathcal{P}_2 = \text{“transaction} \xrightarrow{\text{containsItem}} \text{item} \xrightarrow{\text{isTitle}} \text{title} \text{”}$ where \mathcal{P}_1 is a complex meta-path. In Eq. (1), we compute the average fraction of frauds in the linked transactions and each transaction is weighted by the number of common game titles in the current transaction. It does not, however, distinguish which game titles are shared. In other words, each shared game title is counted equally. In contrast, Eq. (3) counts a shared rare title more than a shared popular title because \mathbf{D}_2 accounts for a normalization step at the title level. The similarity between two transactions increases proportionally to the number of titles they share, but is offset by the popularity of the title.

Here, we present an analogy to term frequency-inverse document frequency (tf-idf) in information retrieval. Essentially, Eq. (1) computes the similarity between linked transactions in a term frequency manner, while Eq. (3) captures the tf-idf information. Most importantly, Eq. (1) is intractable when $\mathbf{P} \in \mathbb{R}^{n \times n}$ is too large and nearly full, while Eq. (3) only needs to deal with sparse matrices $\mathbf{P}_1, \mathbf{P}_2 \in \mathbb{R}^{n \times n_t}$ where n_t is typically much smaller than n .

D. Collective Fraud Detection

So far we have explored the inter-transaction dependency to capture correlated fraudulent behaviors where the correlations mainly exist between training transactions and test transactions. We further notice that fraudulent behaviors are fast evolving. For example, a batch of new transactions may be from the same new billing account but with different IP addresses, some of which are rather risky and others may be proxy. It is desirable to mark all these transactions as suspicious.

As stated in Section II, the inference problem for collective fraud detection in a HIN is to learn a predictive function $f : (\mathcal{V}, \mathcal{E}, \mathbf{X}) \rightarrow \mathbf{y}$. Conventional classification approaches usually make an independent and identically distributed (i.i.d.) assumption, and thus the probability of each transaction being frauds is inferred independently as $f(\mathbf{x}_i) \propto \Pr(y_i = 1 | \mathbf{x}_i)$. In addition to the given features \mathbf{X} , we include features $\{\mathbf{z}^1, \dots, \mathbf{z}^c\}$ that are derived from meta-paths $\{\mathcal{P}^1, \dots, \mathcal{P}^c\}$

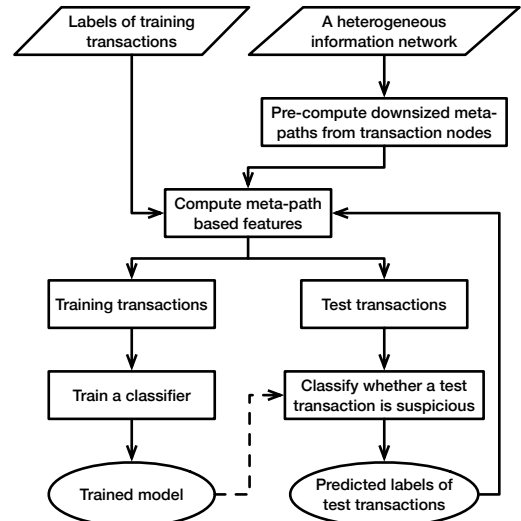


Fig. 3. The flowchart of HITFRAUD.

where c is the number of extracted meta-paths. Therefore, the target is to learn

$$f(\mathbf{x}_i) \propto \Pr(y_i = 1 | \mathbf{x}_i, \mathbf{z}_i^1, \dots, \mathbf{z}_i^c) \quad (4)$$

In this manner, however, the inference of different transactions is essentially not independent, because meta-path based features contain the label information of linked transactions in both training and test sets. It can be done in an iterative framework where the label of a transaction is inferred based on the labels of its linked transactions through its meta-path based features, and its predicted label will further be used to infer the labels of its linked transactions by updating their meta-path based features. It is similar to the framework of Heterogeneous Collective Classification (HCC) [8], and we improve it with a much more efficient way of computing meta-path based features.

DEFINITION 3 (Downsized Meta-path): Given the node sequence of a meta-path $\mathcal{P} = \langle \mathcal{V}^{k_0}, \dots, \mathcal{V}^{k_l} \rangle$, it is a downsized meta-path if $n_{k_0} > \dots > n_{k_l}$.

We design the process of meta-path exploration based on discussions in the last section which can be summarized into two facts: (1) Meta-paths that are used for feature computation in our task always start from and end at transaction nodes. (2) Each meta-path can be decomposed into two parts at the node type of the smallest size. Therefore, we can perform a breadth-first-search from transaction nodes to find all downsized meta-paths from \mathcal{V}^1 to each other node type. In the search procedure, say, the current meta-path \mathbf{P} is from \mathcal{V}^1 to \mathcal{V}^i , we enumerate link types “ $\mathcal{V}^i \xrightarrow{\mathcal{E}^k} \mathcal{V}^j$ ” in the network schema. If $|\mathcal{V}^i| > |\mathcal{V}^j|$, a new meta-path $\mathbf{P}' = \mathbf{P} \times \mathbf{A}^k$ is added into \mathcal{S}_j . Search will be expanded from the newly added meta-paths until all downsized meta-paths from \mathcal{V}^1 to \mathcal{V}^i have been included in \mathcal{S}_i .

In contrast to the original HCC framework where all meta-paths connecting the target entities are pre-computed in the initialization step which results in a lot of redundancy as

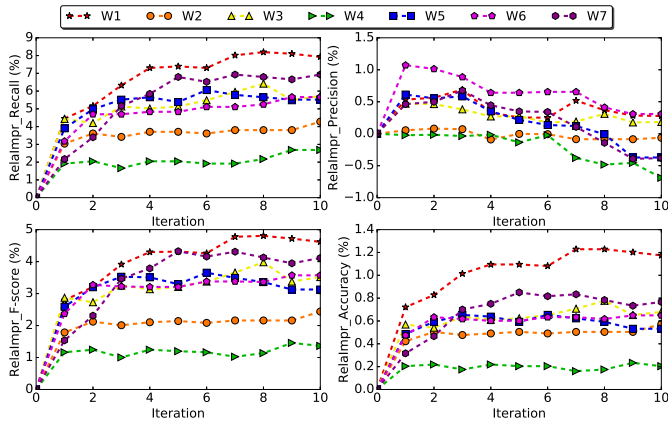


Fig. 4. Convergence analysis of HITFRAUD with random forest as the base classifier.

discussed in Section IV-B, we only pre-compute downsized meta-paths and then organically combine them using Eq. (3) in the feature computation step. In this manner, we have two bonus effects: (1) Redundant meta-paths are automatically avoided because a concatenation of two downsized meta-paths is for sure not a sub-sequence of another one. (2) There is no more need to manually tune the maximal length of meta-paths.

The flowchart of the proposed method HITFRAUD is outlined in Figure 3 where the predicted labels of test transactions are initialized from a classifier that is trained using the given features \mathbf{X} only. Then a new classifier is trained using additional meta-path based features which are iteratively updated based on both the labels of training transactions and the predicted labels of test transactions.

V. EXPERIMENTS

Experiments are conducted on EA payment transaction data that contains $n = 130K$ transactions with manual review labels. For training and evaluation purposes, we segment the data into two consecutive parts so that one-week data is used for testing models and the preceding weeks are used for training. Based on sliding windows, 7 data segmentations are created, each of which is denoted as $W1$ to $W7$. Four metrics are reported: recall, precision, F-score and accuracy.

Figure 4 shows the detailed changes of evaluation metrics on the predictions of test transactions during the iterative process where random forest is used as the base classifier. Note that the predicted labels of test transactions are initialized from the baseline, therefore we denote the baseline performance as the value at iteration 0. Since the absolute prediction performances are different across datasets, here we plot the relative improvement of HITFRAUD over the baseline.

It can be observed that recall, F-score and accuracy are significantly improved in the first few iterations and become stable in the late stages. In terms of precision, it is slightly degraded in some datasets. However, we consider this as a valuable trade-off, because we use no more than 1% decrease in precision to exchange for nearly 8% increase in recall, and

thereby always improve F-score and accuracy. More importantly, recall is more critical in practice when the predicted labels are used as assistance for manual review of suspicious transactions. More experimental results are available in [2].

VI. CONCLUSION

In this paper, we propose HITFRAUD, a collective fraud detection algorithm that captures the inter-transaction dependency. Meta-path based features are efficiently computed through the use of pre-computing downsized meta-paths. Suspicious transactions in the test set are collectively identified when they share common fraudulent behaviors. Experiments on EA payment transaction data demonstrate that the prediction performance is effectively boosted by HITFRAUD with fast convergence. It is validated that the correlated and fast evolving fraudulent behaviors can indeed be explored to more effectively capture fraud payment transactions.

VII. ACKNOWLEDGEMENTS

This work is supported in part by NSF through grants IIS-1526499, and CNS-1626432, and NSFC 61672313. We also thank Victor Chen, Shuang Chen, Ying Wang, Frank Yang, Alexandru Doroftei, Marco Sederquist, Thilmin Gee, and Helen Sulic for insightful discussions and valuable feedbacks.

REFERENCES

- [1] B. Cao, X. Kong, and P. S. Yu. Collective prediction of multiple types of links in heterogeneous information networks. In *ICDM*, pages 50–59. IEEE, 2014.
- [2] B. Cao, M. Mao, S. Viidu, and P. S. Yu. HitFraud: A broad learning approach for collective fraud detection in heterogeneous information networks. *arXiv*, 2017.
- [3] Z. Gyöngyi, H. Garcia-Molina, and J. Pedersen. Combating web spam with trustrank. In *VLDB*, pages 576–587, 2004.
- [4] B. Hooi, H. A. Song, A. Beutel, N. Shah, K. Shin, and C. Faloutsos. Fraudar: Bounding graph fraud in the face of camouflage. In *KDD*, pages 895–904. ACM, 2016.
- [5] M. Ji, J. Han, and M. Danilevsky. Ranking-based classification of heterogeneous information networks. In *KDD*, pages 1298–1306. ACM, 2011.
- [6] M. Jiang, P. Cui, A. Beutel, C. Faloutsos, and S. Yang. CatchSync: catching synchronized behavior in large directed graphs. In *KDD*, pages 941–950. ACM, 2014.
- [7] X. Kong, B. Cao, and P. S. Yu. Multi-label classification by mining label and instance correlations from heterogeneous information networks. In *KDD*, pages 614–622. ACM, 2013.
- [8] X. Kong, P. S. Yu, Y. Ding, and D. J. Wild. Meta path-based collective classification in heterogeneous information networks. In *CIKM*, pages 1567–1571. ACM, 2012.
- [9] Y. Sun, J. Han, X. Yan, P. S. Yu, and T. Wu. Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. In *VLDB*, pages 992–1003, 2011.
- [10] Y. Sun, Y. Yu, and J. Han. Ranking-based clustering of heterogeneous information networks with star network schema. In *KDD*, pages 797–806. ACM, 2009.