Fine-Grained Behavioral Modeling With Graph Neural Networks for Financial Identity Theft Detection

Min Gao, Qiongzan Ye, Yangbo Gao, Zhenhua Zhang, Yu Chen, Yupeng Li, *Member, IEEE*, Shutong Chen, Qingyuan Gong, Xin Wang, and Yang Chen, *Senior Member, IEEE*

***Abstract*—Online-to-Offline (O2O) e-commerce services and their users confront a spectrum of fraud risks, where financial identity theft is prevalent and severe. However, current approaches are inadequate to cover such fraud. To address this problem, we consider both environmental entity interactions and activity sequences to model more granular user behaviors. According to our preliminary study, we discovered that fraudulent users exhibit high aggregations of various environmental entities and fraudulent individuals using the same personal ID that features diverse interactions with different environmental entities. We further investigate the abnormal behaviors of individual fraudsters. Motivated by these discoveries, we propose a deep learning-based behavior modeling framework named EnvIT to capture the above behavior patterns. Therefore, EnvIT is sufficiently general to learn user representations for various e-commerce fraud situations. Extensive experiments are conducted on two real-world datasets provided by Meituan and Vesta, respectively. The results demonstrate the superiority of our method, with a 0.17%-13.50% improvement in AUC and 1.13%-22.57% in R@90%P on the Meituan dataset, and a 0.71%-11.94% improvement in AUC and 2.99%-21.19% in R@90%P on the Vesta dataset, respectively.**

***Index Terms*—Identity theft, financial fraud detection, graph neural networks, behavioral modeling.**

I. Introduction

THE advent of Online-to-Offline (O2O) e-commerce services has revolutionized the way we live. These services facilitate online purchases of goods and services [1], as well as the application for loans [2]. However, these conveniences are accompanied by various fraud risks. According to reports from the Federal Trade Commission (FTC) [3], identity theft cases topped the list of reported issues among all categories. The FTC defines *identity theft* as a pervasive fraud risk that refers to the unauthorized use of personal or financial information of an individual for fraud activities [4], [5]. In addition to the individuals' names, the private information contains their Social Security Numbers,[1] credit card numbers, and bank account details. Due to financial identity theft, both individuals and online platforms may suffer financial losses and reputation damage [6]. The risk of identity theft has evolved with the adaptation of various methods, including phone number recycling [7], phishing [8], and account hijacking [9]. Among these, phone number recycling [7] is a prevalent practice for fraudsters, as recycled mobile numbers are linked to authenticated accounts on e-commerce platforms.[2] The rapid growth of e-commerce has inadvertently contributed to the prevalence of financial identity theft [11]. For example, fraudsters may steal others' credit card information to make purchases or misuse others' identity information to take out loans. Therefore, it is imperative to detect identity theft fraud to guarantee the safety of online financial transactions within the realm of e-commerce.

Recent studies [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28] leverage deep learning techniques to detect fraud users by learning user representations from users' historical behaviors, social networks,

Received 12 January 2025; revised 15 September 2025; accepted 21 October 2025. Date of publication 31 October 2025; date of current version 16 January 2026. This work was supported in part by the National Natural Science Foundation of China under Grant 62072115, Grant 62102094, and Grant 62202402, in part by Shanghai Science and Technology Innovation Action Plan Project under Grant 22510713600, in part by Guangdong Basic and Applied Basic Research Foundation under Grant 2022A1515011583, in part by One-off Tier 2 Startup Grant 2020/2021 of Hong Kong Baptist University under Grant RCOFSGT2/20-21/COMM/002, in part by the Startup Grant (Tier 1) for New Academics AY2020/21 of Hong Kong Baptist University and the AI-Info Communication Study (AIS) Scheme 2021/22 under Reference AIS 21-22/06, in part by Guangdong and Hong Kong Universities "1+1+1" Joint Research Collaboration Scheme under Grant 2025A0505000001, in part by the Initiation Grant for Faculty Niche Research Areas 2023/24 under Grant RC-FNRA-IG/23-24/COMM/01, in part by the Research Grants Council of HKSAR under Grant HKBU 22202423 and Grant HKBU 12203425, and in part by Meituan. Recommended for acceptance by Dr. Muhammad Khurram Khan. *(Corresponding author: Yang Chen.)*

Min Gao, Qiongzan Ye, Shutong Chen, Xin Wang, and Yang Chen are with the College of Computer Science and Artificial Intelligence, Fudan University, Shanghai 200438, China, and also with the Shanghai Key Laboratory of Intelligent Information Processing, Fudan University, Shanghai 200438, China (e-mail: mgao21@m.fudan.edu.cn; qzye19@fudan.edu.cn; shutongchen17@fudan.edu.cn; xinw@fudan.edu.cn; chenyang@fudan.edu.cn).

Yangbo Gao, Zhenhua Zhang, and Yu Chen are with Meituan, Beijing 100102, China (e-mail: gaoyangbo02@meituan.com; zhangzhenhua02@meituan.com; chenyu17@meituan.com).

Yupeng Li is with the Department of Interactive Media, Hong Kong Baptist University, China (e-mail: ypengl@hkbu.edu.hk).

Qingyuan Gong is with the Research Institute of Intelligent Complex Systems, Fudan University, Shanghai 200433, China (e-mail: gongqingyuan@fudan.edu.cn).

Digital Object Identifier 10.1109/TNSE.2025.3627451

[1] https://www.ssa.gov/history/ssn/geocard.html

[2] Verified accounts refer to those users who have completed the identity verification process with identity cards or bank cards on e-commerce platforms. Recently, user authentication on several prominent platforms in China, such as Taobao and Meituan, relies on Short Message Service (SMS) One-Time Password messages [10]. Consequently, fraudsters can exploit reassigned mobile numbers to gain access and assume control of these verified accounts.

and user profile information. These techniques fall into two main categories: sequence-based techniques and graph neural network (GNN)-based techniques. Sequence-based techniques differentiate fraudulent accounts from normal ones by leveraging the history of sequential behavior [5], [12], [13], [14], [15], [16], [17], [18], [19], [20], [29], [30], [31]. However, they overlook the intricate interconnections between financial accounts, which are critical to identifying fraud gangs. GNN-based studies facilitate graph structures to model the relationships between users and other entities and learn their representations with GNNs [21], [22], [23], [24], [25], [26], [32]. Most approaches (e.g., [21], [22], [23]) utilize the aggregations of environmental entities, such as IP addresses or devices (e.g., Media Access Control Address and International Mobile Equipment Identity of the device) to detect fraud groups. However, these methods are based on the assumption that fraudsters typically operate within interconnected groups, linked by a limited number of devices or IP addresses. Consequently, these approaches fall short in effectively identifying individual fraudsters. Moreover, several leading e-commerce platforms such as Taobao and Meituan,[3] permit multiple verified accounts to be associated with a single personal ID[4] [33]. As a result, there might be differences in devices and IP addresses among verified accounts sharing the same personal ID. These distinctions could potentially serve as indicators to discern whether these accounts are linked to the same individual. However, current research has not yet fully exploited these significant distinctions in addressing the issue of detecting financial identity theft on e-commerce platforms.

We have investigated actual cases of financial identity theft on Meituan, a famous O2O e-commerce platform in China that provides services related to local life and finance [34], [35]. To understand fraud patterns at different levels, we investigate the suspicious activities exhibited by fraudulent accounts and the links between the accounts and the entities in the access environment. In this work, we consider the IP addresses, devices, and transmission addresses used by the accounts as access environment entities. To deeply understand the characteristics of fraudulent behaviors in financial identity theft, we summarize three types of typical patterns from real data and analyze the modeling challenges they cause and the modeling limitations of existing methods: (1) environment aggregation, i.e., multiple fraudulent accounts share the same devices, IPs, or addresses, which makes it difficult for sequence models to model inter-account relationships, and the existing GNN methods are difficult to differentiate between normal and abnormal sharing, which makes it easy to introduce noise; (2) environmental difference, which means that although multiple accounts share the same personal ID information, their associations with environmental entities, such as devices or other accounts, exhibit significantly different structural patterns. Some fraudulent accounts are primarily connected via a large number of device nodes, while others establish connections through shared identity nodes. These distinct patterns reflect different fraud strategies and are challenging for existing models to distinguish or interpret effectively; and (3) behavioral characteristics, which means that individual fraudulent accounts show concentrated or repetitive operations in a short period of time, and GNNs are unable to depict the temporal dynamics, while sequential models do not integrate the structural and identity contextual information. These challenges indicate that existing methods have significant shortcomings in modeling multidimensional fraudulent behaviors, which motivates us to propose a new approach that fuses graph structure and behavioral modeling to effectively identify such complex fraud patterns.

Guided by these observations, we model the complex interactions between accounts and various entities with different graph structures (a heterogeneous environment graph and a homogeneous account graph). We propose a framework named **EnvIT** for identity theft detection. Our framework takes into account both the associations across the environmental entities and the historical activities of the accounts. Specifically, EnvIT is composed of five key components: *environmental aggregation extractor*, *environmental difference extractor*, *behavior feature extractor*, *attentive feature fusion module*, and *detection layer*. The design of the environmental aggregation extractor is utilized to generate deep representations of environmental aggregation from the heterogeneous environment graph. This module has the capability to include fraud groups exhibiting the first fraud pattern: environmental aggregation. The design of the environmental difference extractor is employed to learn the embeddings of environmental differences in the accounts from the homogeneous account graph. This module could cover fraudulent accounts with the second fraud pattern: environmental difference. The behavior feature extractor module is designed to extract behavior representations from historical behavior sequences from the account. This module could detect individual fraudsters with the third fraud pattern of abnormal temporal behavior activities. With these three components, EnvIT could learn informative account representations under different e-commerce fraud situations. A previous version of this work [36] only considered heterogeneous and homogeneous graph structures. To achieve a robust fusion of various representations from these three modules covering complex situations, we introduce the self-attention mechanism and design a module named the attentive feature fusion module. As described in [16], [37], [38], the attention weights of this module could enhance the interpretability of the detection results. Finally, a detection layer determines the likelihood of an account being fraudulent.

To summarize, this work makes the following contributions:

- To our knowledge, this study is the first to tackle the issue of financial identity theft in O2O e-commerce platforms. It formulates the problem from various graph perspectives to capture intricate relationships between accounts and environmental entities.
- We analyze users' behavioral patterns according to the dataset from Meituan, one of China's top O2O e-commerce platforms. We study and summarize fraud patterns from different perspectives, including environmental associations and behavioral activities.
- We propose EnvIT, a novel framework designed to detect financial identity theft by leveraging both the associations

[3]https://about.meituan.com/en

[4]In this context, "ID" stands for the resident identity card, i.e., an official document used for personal identification in China's mainland.

across various environmental entities and the historical activities of accounts. EnvIT leverages three powerful components to capture comprehensive account representations in various e-commerce fraud scenarios, including environmental aggregation, environmental difference, and historical behavior. These representations are then fused using an attention mechanism to enable effective fraud detection.

- We evaluate the performance of EnvIT on two real-world datasets from two leading e-commerce platforms. Our experimental findings indicate that EnvIT surpasses several state-of-the-art methods and is capable of detecting more fraud than Meituan's current solution. Furthermore, we conduct an interpretability analysis of the attention weights across meta-paths on two datasets.

II. Related Work

Recently, there has been a surge in the development of deep learning approaches to detect various forms of financial fraud. Notably, approaches with sequence and graph learning have gained significant popularity in this domain.

Sequence-based methods employ neural networks like Convolutional Neural Network (CNN) and Recurrent Neural Network (RNN) to analyze users' historical behavior. For example, Liu et al. [18] presented Local Intention Calibrated Tree-LSTM (LIC Tree-LSTM) to capture the specific fraudulent intention and build a behavior tree according to the transaction sequences for each user. Branco et al. [17] employed Gated Recurrent Units (GRU) to model users' payment behaviors as interconnected sequences for the task of credit card fraud detection. Babaev et al. [14] also utilized GRU to predict credit scores for applicants based on their loan records. Cheng et al. [15] developed a framework for dynamic default prediction with Gated LSTM. This framework leveraged loan behaviors to forecast repayment delinquency in the context of networked-guarantee loans. Guo et al. [12] designed HAInt-LSTM with a modified forget gate, which can be used to recognize fraud patterns from intervals between consecutive time steps. Xiao et al. [39] presented a Multiview row-INteractive Time-aware framework to detect fraudulent behaviors from time-series data. Xie et al. [40] designed a time-aware gate and combined an LSTM to learn users' transaction representations. Gao et al. [41] introduced a temporal pattern encoder and a statistical feature encoder aimed at identifying accounts compromised due to phone number reassignment. Xie et al. [42] modeled user behaviors from both time-aware and location-aware perspectives and presented a spatial-temporal gated model to learn representations of new transactions based on historical transactions. These approaches capture temporal patterns effectively, making themselves particularly useful for detecting fraudulent behaviors. However, they typically model each account independently, failing to take advantage of the complex interactions between users and environmental entities (e.g., IPs, devices, and delivery addresses). As a result, they are limited in identifying coordinated fraud, such as fraud gangs operating multiple accounts with shared personal IDs or devices.

Numerous studies [23], [43], [44], [45], [46], [47], [48] have demonstrated the effectiveness of graph neural networks (GNNs) [49] in the detection of financial fraud. These studies encompass methods for identifying fraudulent transactions [50], [51], predicting loan defaults [52], [53], detecting insurance fraud [22], and identifying cash-out users [24]. Among them, Fang et al. [25] presented mHGNN to identify illicit transactions and constructed an attributed heterogeneous information network (AHIN) to capture intricate interactions among entities, including topics, products, and comments. Similarly, Zhong et al. [26] designed MAHINDER, a fraud detection model that relies on a multiview AHIN for the identification of credit default fraud. Hu et al. [24] developed HACUD, a detection model equipped with a hierarchical attention mechanism designed to identify cash-out users within credit card payment services. Xu et al. [54] presented a Spectrum-Enhanced and Environment-Constrained Graph Fraud Detector (SEC-GFD) to detect fraud by leveraging spectrum and label information. Liu et al. [21] presented GEM, a study focused on analyzing aggregation patterns derived from devices and temporal activities associated with malicious accounts in user-device graphs. Wang et al. [55] developed a semi-supervised graph attentive network specifically designed for fraud detection using a multi-view graph (i.e., users' relationships, attributes, and devices). Xiang et al. [43] presented a Gated Temporal Attention Network to detect credit card fraud. In the context of financial identity theft detection on e-commerce platforms, existing graph-based methods, including those that use heterogeneous graphs, have shown strong capabilities in modeling multiple relations. However, they often struggle to effectively capture fraud patterns related to environmental contexts. Specifically, in environmental aggregation scenarios, fraudulent accounts might share the same devices, IPs, or addresses, forming high-aggregation patterns, while in environmental difference cases, accounts linked to the same personal ID may appear in diverse environmental relations to evade detection. Current models do not account for these different fraud patterns.

To address these challenges, our method jointly models both the differences between verified accounts sharing the same personal ID and their interactions between environmental entities and temporal behaviors, enabling more effective detection of diverse and complex identity theft patterns. Table I offers a comparative summary of representative approaches in terms of their ability to model graph structures, heterogeneity, and behavioral sequences. As shown in the comparison, most existing methods capture only one or two of these dimensions, whereas our approach integrates all three, allowing for a more comprehensive and robust detection of identity theft behaviors.

III. Case Analysis

A. Dataset and Preliminary Exploration

Meituan offers a range of services related to daily life, including delivering food, booking taxis, and purchasing tickets online. Additionally, Meituan has developed a consumer finance business that offers consumer loans to enhance users' purchasing capabilities. For this study, a dataset was obtained from the

TABLE I
COMPARISON OF REPRESENTATIVE MODELS IN MODELING KEY ASPECTS OF FINANCIAL IDENTITY THEFT BEHAVIORS

Model	Graph Structure	Heterogeneity Information	Behavioral Sequence
LSTM	✗	✗	✓
Bi-LSTM	✗	✗	✓
GRU	✗	✗	✓
HAInt-LSTM	✗	✗	✓
TSF	✓	✗	✓
GraphSAGE	✓	✗	✗
GAT	✓	✗	✗
RGCN	✓	✓	✗
HAN	✓	✓	✗
HACUD	✓	✓	✗
SOBT	✗	✓	✗
SEC-GFD	✓	✓	✗
Ours	✓	✓	✓

consumer finance business department of Meituan. There are 15,350 sampled accounts that had applied for loans before June 2020 in the dataset. There are also 6,993 additional accounts in the dataset that underwent verification using the same IDs as the sampled accounts. We have anonymized personal ID-related information through a hash function and excluded sensitive data to maintain user privacy. There are three types of accounts in the dataset: normal, fraudulent, and unlabeled. Normal accounts exhibit regular loan behavior, characterized by punctual repayment and full settlement of all loans, and are manually reviewed to ensure stability and compliance. In contrast, fraudulent accounts are initially identified from user feedback. Users provide feedback in cases where they receive collection calls for unauthorized overdue loans when reviewing their credit reports. It is important to highlight that every fraud label is manually checked by a team of risk control specialists to guarantee their accuracy and reliability before being ultimately added to the dataset. In addition to the two categories of clearly labeled accounts mentioned above, the remaining accounts are classified as unlabeled accounts. Since these unlabeled accounts might contain potential fraudulent activities, we adopt a semi-supervised learning strategy. This approach not only uses the labeled accounts for supervised training but also fully involves the unlabeled accounts in the message passing mechanism to enhance our model's ability to capture potential fraudulent behaviors. The imbalanced distribution between normal and fraudulent accounts may result in suboptimal performance of fraud detection models, particularly for the minority class, i.e., the fraudulent accounts [56]. Therefore, we employ an under-sampling strategy to exclude a portion of normal accounts. Additionally, the dataset encompasses fundamental account details, loan application dates, and historical activities, including logins, payments, and delivery orders spanning from January 2018 to June 2020. For accounts that have applied for a loan, an approval process is conducted before the loan is disbursed. The objective is to determine whether identity theft is involved in the account approval process. Consequently, we exclude post-loan application activities for each account, while retaining the complete historical behaviors of accounts without a loan application. The statistical overview of the processed dataset is presented in Table II. To provide an intuitive understanding of the environmental aggregation and difference patterns of accounts, we visualize the relationships between accounts and environmental entities based on the Meituan dataset, as shown in Fig. 1.

TABLE II
STATISTICS OF THE MEITUAN DATASET

Account Type	Normal	Fraudulent	Unlabeled
# of account	14,156	1,194	6,993
Avg. # of login	11.04	25.48	9.05
Avg. # of payment	61.52	51.76	19.97
Avg. login interval	413.51 min	324.39 min	374.88 min
Avg. payment interval	624.58 min	626.19 min	609.77 min

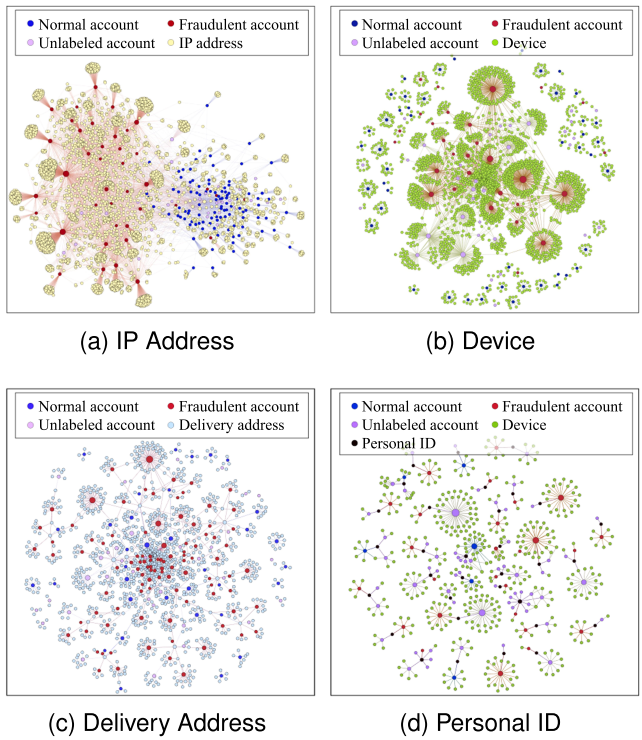

Fig. 1. The relationships between accounts and environmental entities. Node colors represent different types of nodes: blue for normal accounts, red for fraudulent accounts, purple for unlabeled accounts, yellow for IP addresses, green for devices, cyan for delivery addresses, and black for personal IDs. Node size is proportional to node degree. Edge color reflects the label of the source node, e.g., blue for edges emitted from normal accounts.

Environmental Aggregation Analysis: We obtained the IP address and device details associated with the user's account from the user's login record. Fig. 1(a), (b), and (c) illustrate subgraphs where account nodes are linked to IP addresses (nodes in yellow), devices (nodes in green), and delivery addresses (nodes in cyan), respectively. We find that these three subgraphs present similar results, i.e., fraudulent accounts (nodes in red) exhibit more IP addresses, devices, and delivery addresses associated with them than normal accounts (nodes in blue). These findings indicate that there are aggregations of IP addresses, devices, and delivery addresses in fraudulent accounts, which

suggests that these accounts may be constituted as fraudulent gangs. Therefore, all three types of environmentally associated entities can be used as potential information for capturing more sophisticated behaviors by fraudulent teams.

Environmental Difference Analysis: As shown in Fig. 1(d), personal IDs are shared with different accounts that are linked to environmental entities with large differences. Here, the differences are manifested in two ways: one part of the fraud accounts (nodes in red) establishes extensive association networks through a large number of devices (nodes in green), reflecting the characteristic of using devices as the main directly associated objects; the other part of the fraud nodes forms complex connections with other accounts more through personal IDs (nodes in black), exhibiting the characteristic of using personal information as the directly associated environmental nodes. It is important to distinguish between these two different environmental associations, as they reflect different fraud strategies.

Behavior Characteristics Analysis: We thoroughly examined login and payment activities, revealing that the accounts engaged in fraudulent activities exhibit notably unusual historical behaviors. In particular, certain accounts demonstrate a significantly elevated frequency of login or payment transactions within condensed timeframes, particularly during late-night hours. In addition, these accounts often exhibit login attempts originating from multiple disparate locations or devices within a short period of time, further distinguishing their behavior from that of legitimate users.

All these findings provide valuable insights and motivate us to design corresponding components for capturing specific fraud behaviors in our subsequent model design by modeling account behaviors and complex connections.

IV. EnvIT Framework

A. Problem Formulation

Based on information from environmental associations and historical activity sequences, we aim to ascertain whether a verified account is being operated by a fraudster. This objective can be viewed as a binary classification problem.

Given the associations between environmental entities and the activity details of user accounts, along with the hashed values of their verified personal IDs, we can build a heterogeneous environment graph $\mathcal{G}_e = (\mathcal{U}, \mathcal{D}, \mathcal{E}, \mathcal{R})$ and a homogeneous account graph $\mathcal{G}_a = (\mathcal{U}, \mathcal{E}_a)$. The set of account nodes is denoted as $\mathcal{U}$, and different environmental entity nodes (such as IP nodes and device nodes) are represented by $\mathcal{D}$. The $\mathcal{E}$ denotes the set of edges between two adjacent nodes. Each edge type refers to a particular relationship. And $\mathcal{R}$ denotes the set of relationships that occur between account nodes and environmental entity nodes. A relationship can be the interaction or usage between an account node and an environmental entity node. Each $e_{ij}^{r_m} \in \mathcal{E}$ represents the connection between an account node $i \in \mathcal{U}$ and an environmental entity node $j \in \mathcal{D}$, defined by the relationship $r_m \in \mathcal{R}$. The access environment (e.g., IP address and device) is characterized by multiple types of entities with a total of M types, i.e., $M = |\mathcal{R}|$. Additionally, characteristics such as gender and age are considered features for each account. The ground-truth labels for the account nodes are denoted as Y. The $\mathcal{E}_a$ represents the set of edges connecting neighboring nodes in $\mathcal{G}_a$. The accounts $u \in \mathcal{U}$ and $v \in \mathcal{U}$ are verified using the same personal ID for each edge $e_{uv}^a \in \mathcal{E}_a$.

From the historical behaviors of the accounts, we derive a behavior sequence $S = [s_1, s_2, \ldots, s_T]$. Within this sequence, $s_i (i \in \{1, 2, \ldots, T\})$ represents the embeddings of the behavior, and T refers to the duration of the sequence. Each behavior record encompasses various aspects, including the time interval, time slot, GPS location, and login IP address. The summary of main notations and definitions is listed in Table III.

B. The EnvIT Model

1) Model Overview: In this section, we provide a detailed introduction to the EnvIT model. Fig. 2 shows the overall framework of our proposed model. Specifically, EnvIT is made up of five modules: the environmental aggregation extractor, the environmental difference extractor, the behavior feature extractor, the attentive feature fusion module, and the detection layer. The initial three components leverage diverse data sources, including logs, behavioral sequences, and other relevant information, to comprehensively analyze and model the intricate relationships between the account and its surroundings. The first module focuses on capturing the representations derived from the association of the account with environmental factors. The second module concentrates on the account divergence within the environment. The third module draws upon the historical behavior patterns of the accounts. Subsequently, the module of attentive feature fusion is designed to fuse embeddings from the preceding three modules. Finally, the detection layer is responsible for forecasting the likelihood of an account being fraudulent.

2) Environmental Aggregation Extractor: To achieve a more accurate representation of accounts, we utilize environmental aggregation techniques that capture the nuanced behaviors and interactions within the social network. To generate the initial embeddings X for all account nodes, we embed their attribute features with a fully connected layer. The initial embedding of each environmental entity node j in the heterogeneous graph $\mathcal{G}_e$ is randomly initialized to x_j. Here, the account node i's embedding x_i and environmental entity node j's embedding x_j have the same dimension d. In our heterogeneous graph, we incorporate multiple meta-paths to capture diverse environmental relationships between accounts, devices, IPs, and locations. Specifically, we consider the following meta-paths to model heterogeneous relationships among accounts and their associated environmental entities:

- *Account-Device-Account (A-D-A):* This meta-path captures the interaction between accounts that share the same device, potentially indicating device-based aggregation activities.
- *Account-IP-Account (A-I-A):* It models the relationships between accounts that originate from the same IP prefix, revealing potential clusters of accounts associated with a common network environment.
- *Account-Location-Account (A-L-A):* This meta-path represents accounts that share the same geographic location, suggesting possible geographical correlations.

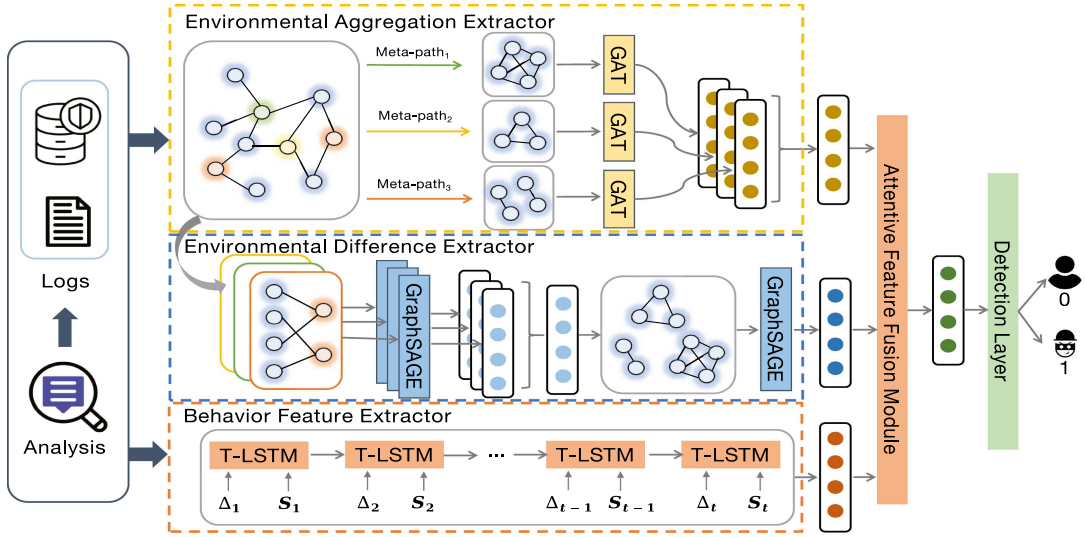

Fig. 2. An illustration of the EnvIT framework for detecting identity theft. EnvIT comprises five components: an environmental aggregation extractor, an environmental difference extractor, a behavior feature extractor, an attentive feature fusion module, and a detection layer. Firstly, the input graph data and user features are encoded via the first two components, and the input behavior sequences are encoded via the third component. Then, the encoded outputs from the top three components are fused via the attentive feature fusion module. Finally, in the rightmost part, the detection layer is used to detect fraudulent accounts.

- *Device-Account-Device (D-A-D):* This meta-path reflects the relationships among devices that are used by the same account and provides valuable insights into the device-based usage patterns of accounts.
- *IP-Account-IP (I-A-I):* This meta-path refers to the relationships between IP addresses that are shared with the same account. It captures the semantic pattern of an account being linked to multiple IP addresses. The I-A-I path is beneficial for detecting network-level anomalies, such as IP spoofing, which are commonly used in fraudulent activities.
- *Location-Account-Location (L-A-L):* This meta-path represents the relationships between locations that are associated with the same account. It models the semantic pattern of an account being linked to multiple locations. This meta-path is particularly useful in identifying accounts that exhibit suspicious location-based patterns.

These meta-paths are carefully selected based on domain knowledge in fraud detection, where shared environmental entities, such as devices, IP addresses, locations, and network information, often serve as indicators of coordinated behavior. Then, we denote the meta-paths [57] as $P_m = \mathcal{U} \rightarrow \mathcal{D}^m \rightarrow \mathcal{U}$ in the heterogeneous graph $\mathcal{G}_e$. By leveraging this representation, we can identify the meta-path-based neighbors of each account node. We recognize that each account node has a unique impact on the targeted account node i. To address this, we used the Graph Attention Network (GAT) [58] to update the representation of the targeted account node by aggregating messages from its neighboring nodes. The process can be summarized as $h_i^{mk} = \sigma(\sum_{v \in \mathcal{N}(i)^m} \alpha_{iv}^{mk} W^{mk} h_v^{mk-1})$, where the neighbors of node i under meta-path P_m, are represented as $\mathcal{N}(i)^m$ with $m \in \{1, \ldots, M\}$. The representation of account node i is denoted as h_i^{mk}, with W^{mk} representing a learnable weight matrix at the k-th hidden layer. Additionally, the representation of each account node i in the 0-th hidden layer is denoted as x_i. The ReLU activation function is represented by $\sigma(\cdot)$. Moreover, α_{iv}^{mk} denotes the attention score between node i and v at the k-th hidden layer under meta-path P_m. The attention score can be computed as follows:

$$\alpha_{iv}^{mk} = \text{softmax}\left(\text{LeakyReLU}\left(\vec{a}^T \left[z_i^{mk-1} \,||\, z_v^{mk-1}\right]\right)\right), \quad (1)$$

where $z_i^{mk-1} = W^{mk} h_i^{mk-1}$, and $\vec{a}$ refers to a vector of learnable parameters. The transposition and concatenation operations are represented by $\cdot^T$ and $||$, respectively. Here, the attention score reflects the relative importance of neighboring accounts connected via different meta-paths, which can reveal potential risk relationships. For example, if an account assigns higher weights to another account that shares the same device, this may indicate a coordinated fraudulent pattern through a shared device.

With the total number of hidden layers being K_a, the output of the final hidden layer of the account node i on the meta-path P_m is represented as $h_i^{mK_a}$, which can be simplified as h_i^m. Hence, the environmental aggregation of each account i can be calculated by $h_i^a = \sum_{m=1}^{M} h_i^m$.

3) Environmental Difference Extractor: According to the heterogeneous environment graph $\mathcal{G}_e$, we have retrieved the environmental representation for each account node. Subsequently, we compute the environmental differences among verified accounts that possess identical personal identification. To better capture the distinct relationships between account nodes and different types of environment entities, we partition the heterogeneous environment graph $\mathcal{G}_e$ into multiple bipartite sub-graphs $\mathcal{G}_e^{r_1}, \mathcal{G}_e^{r_2}, \ldots, \mathcal{G}_e^{r_M}$, where $r_1, r_2, \ldots, r_M \in \mathcal{R}$. Each sub-graph $\mathcal{G}_e^{r_m}$ $(m \in \{1, \ldots, M\})$ contains only the

TABLE III
SUMMARY OF KEY NOTATIONS AND DEFINITIONS

Notation	Definition
$\mathcal{G}_e$	A heterogeneous environment graph
$\mathcal{G}_a$	A homogeneous account graph
$\mathcal{U}$	Set of account nodes
$\mathcal{D}$	Set of multiple environmental entity nodes
$\mathcal{E}$	Set of edges
$\mathcal{R}$	Set of relationships
$e_{ij}^{r_m}$	A connection between an account node $i \in \mathcal{U}$ and an environment node $j \in \mathcal{D}$ with relation $r_m \in \mathcal{R}$
M	Number of multiple entity types within the access environment
e_{uv}^{a}	An edge between two account nodes u and v with the same personal ID
$S = [s_1, s_2, \cdots, s_T]$	A behavior sequence
T	The duration of the sequence
P_m	One meta-path in heterogeneous graph
$\mathcal{N}(i)^m$	The meta-path based neighbors of node i under meta-path P_m
h_i^{mk}	Representation of account i under meta-path P_m in the k -th hidden layer
W^{mk}	The learnable weight matrix at the k-th hidden layer
α_{iv}^{mk}	The attention score between node i and v in the k -th hidden layer under meta-path P_m
$\vec{a}$	A vector of learnable parameters
$.^T$	The transposition operation
$\|\|$	The concatenation operation
$h_i^{mK_a}$	The output of the final hidden layer of the account node i on the meta-path P_m
h_i^m	A simplified version of $h_i^{mK_a}$
h_i^a	The environmental aggregation of each account i
$\mathcal{G}_e^{r_m}$	Each bipartite sub-graph
MEAN	The MEAN function operates on each vector $h_j^{r_m k-1}$ in an element-wise mean manner
$W^{r_m k-1}$	The trainable parameters
K_e	The total number of hidden layers
$h_i^{r_m}$	The final representation from the hidden layer for account node i
h_i^e	The initial account representations with the learned environmental embedding
K_d	The total number of hidden layers
$h_i^{K_d}$	The output of the last hidden layer of the environmental difference extractor module
h_i^d	The environmental difference of each account i
s_t	The input of behavior event at time step t
h_t and h_{t-1}	The hidden states at time step t and the previous time step $t-1$, respectively
f_t, i_t, and o_t	The forget, input, and output gates, respectively
c_t and c_{t-1}	The cell memories at time step t and previous time step $t-1$, respectively
c_{t-1}^S and c_{t-1}^L	The short-term and long-term memory of the cell at previous time step $t-1$, respectively
$g(\Delta t)$	The discount function for the short-term memory of the cell in the previous time step
Δt	The time interval between two adjacent events
$\{W_d, b_d\}$	The learnable parameters of subspace decomposition
$\{W_c, U_c, b_c\}$	The learnable parameters of candidate values
h_i^b	The behavior representation of the account nodes
H_i	The combination the representations of account i
H_i^{Attn}	The attention function
Q_i	The query matrix
K_i	The key matrix
V_i	The value matrix
d_k	The dimension of key matrix K_i
d_v	The dimension of value matrix V_i
h_i^{Attn}	The final fused feature vector
W_i^f	The learnable weight matrix in the detection layer
$\hat{y_i}$ and $\hat{Y}$	The final output of account i and all account nodes, respectively
y_i and Y	The true label of account i and all account nodes, respectively
N	The total number of account nodes

edges linking account nodes to a specific category of environment entities. In our setting, we define three environment categories based on commonly observed contextual features in fraud detection scenarios: device identifiers (Device), IP addresses (IP), and delivery addresses (Address). Thus, we set $R = \{Device, IP, Address\}$, resulting in $M = 3$ bipartite sub-graphs.

Each bipartite sub-graphs on one particular type of environmental interaction, allowing our model to better capture potential behavioral patterns in that specific context. For example, the sharing of the same device across multiple accounts might indicate coordinated control, while the use of different devices under the same personal ID might imply attempts at identity theft. This decomposition enables our model to disentangle the semantics of different environmental relations, helping to avoid information noise caused by mixing heterogeneous node types in the message-passing process.

Inspired by GraphSAGE [59], a GNN variant known for its powerful inductive learning on various graphs, we utilize this module to update the embeddings of account nodes by aggregating information from the interconnected environmental entity nodes. This can be described as follows:

$$h_i^{r_m k} = \sigma\left(W^{r_m k-1} \cdot \text{MEAN}\left(\left\{h_j^{r_m k-1}\right\}\right)\right), \quad (2)$$

where $h_i^{r_m k}$ refers to the account node i's representation at the k-th hidden layer of sub-graph $\mathcal{G}_e^{r_m}$. The set of environmental entity nodes associated with account node i in the sub-graph $\mathcal{G}_e^{r_m}$ is denoted by $\mathcal{N}(i)^{r_m}$, where $j \in \mathcal{N}(i)^{r_m}$. We begin by setting $h_i^{r_m 0} = x_i$ and $h_j^{r_m 0} = x_j^{r_m}$. The MEAN function operates on each vector $h_j^{r_m k-1}$ in an element-wise mean manner, and $\sigma(\cdot)$ represents the ReLU activation function. The trainable parameters are denoted as $W^{r_m k-1}$. Here, K_e refers to the total number of hidden layers. For account node i, its final representation from the hidden layer is $h_i^{r_m K_e}$, denoted as $h_i^{r_m}$. To derive the environment embeddings for each account node, we compute the element-wise sum of its representations at the final hidden layers across all sub-graphs by $h_i^e = \sum_{m=1}^{M} h_i^{r_m}$.

For each node i in the homogeneous account graph $\mathcal{G}_a$, we initialize its representations with the learned environmental embedding h_i^e. GraphSAGE model is utilized on graph $\mathcal{G}_a$ to capture the environmental difference among the verified accounts sharing the same personal ID. As a result, for each account i, the representation of the environmental difference at the k-th hidden layer can be calculated by $h_i^k = \sigma(W^{k-1} \cdot \text{MEAN}(\{h_v^{k-1}\}))$, where $v \in \mathcal{N}(i)$ and $\mathcal{N}(i)$ refers to the account nodes sharing the same personal ID with the account node i. h_i^0 is initialized as h_i^e. The MEAN and $\sigma(\cdot)$ functions mentioned above remain the same. The total number of hidden layers is K_d. The output of the last hidden layer, denoted as h_i^d, can be simplified from $h_i^{K_d}$.

4) Behavior Feature Extractor: We design the behavior feature extractor to obtain a deep representation of the historical behavior sequences for each account. It is common for malicious behavior to occur suddenly and be concentrated within a short period of time in practical scenarios. As a result, the time interval between consecutive behavior events plays a crucial role in identifying abnormal behavior. To encode the behavior of accounts, we employ the Time-Aware LSTM (T-LSTM) [60]. This modified version of Long-Short Term Memory (LSTM) [61] takes into account the time interval between successive events within a sequence. The formalization of T-LSTM is as follows:

$$c_{t-1}^S = \tanh\left(W_d c_{t-1} + b_d\right) , \tag{3}$$

$$c_{t-1}^{S\prime} = g\left(\Delta t\right) * c_{t-1}^S , \tag{4}$$

$$c_{t-1}^L = c_{t-1} - c_{t-1}^S , \tag{5}$$

$$c_t = f_t * \left(c_{t-1}^L + c_{t-1}^{S\prime}\right) + i_t * \tanh\left(W_c s_t + U_c h_{t-1} + b_c\right) , \tag{6}$$

$$h_t = o_t * \tanh\left(c_t\right) , \tag{7}$$

where the cell memories at time step t and previous time step $t-1$ are represented as c_t and c_{t-1}, respectively. The c_{t-1}^S and c_{t-1}^L are the short-term and long-term memory of the cell at the previous time step $t-1$, respectively. s_t denotes the input of behavior event at time step t. The hidden states at time step t and previous time step $t-1$ are represented as h_t and h_{t-1}, respectively. Additionally, we have forget, input, and output gates, which are f_t, i_t, and o_t, respectively. In T-LSTM, a function $g(\Delta t)$ is utilized to discount the short-term memory of the cell at the previous time step. This function depends on the time interval Δt between two consecutive behavior events and can be expressed by $g(\Delta t) = \frac{1}{\log(\Delta t + e)}$. As a result of the function $g(\Delta t)$, the previous cell's short-term memory c_{t-1}^S is adjusted to $c_{t-1}^{S\prime}$. Importantly, the larger the time interval between two consecutive behavior events, the less short-term memory will be retained. In T-LSTM, $\{W_d, b_d\}$ are learnable parameters of subspace decomposition. $\{W_c, U_c, b_c\}$ are learnable parameters of candidate values. Here, the $\sigma(\cdot)$ denotes a sigmoid layer. To obtain the behavior representation of the account, we take the sum of all the hidden states of the cells at different time steps. This can be expressed as $h_i^b = \sum_{t=1}^{T} h_t$.

5) Attentive Feature Fusion Module: There may be differences in the significance of environment and behavior-related features when handling various fraudulent accounts. As a result, combining these features through simple concatenation or addition may not yield optimal results. To address this, attention mechanisms can be employed to focus more on important information. This enables better integration of the features and enhances the overall performance of the model. To learn the attention scores of different representations, we utilize the self-attention mechanism [62]. We combine the representations of account i to create a matrix, which serves as the input of the attentive feature fusion module. Let $H = \text{Concat}(h_i^d, h_i^a, h_i^b)^T$, where h_i^d, h_i^a, h_i^b are the learned embeddings from three extractors. The Concat denotes the concatenation function, and $\cdot^T$ is the transposition operation. We apply a scaled dot-product self-attention mechanism to compute the attention output matrix. This process can be denoted as

$$H_i^{\text{Attn}} = \text{softmax}\left(\frac{Q_i K_i^T}{\sqrt{d_k}}\right) V_i , \tag{8}$$

where $Q_i = H_i W_Q$, $K_i = H_i W_K$, and $V_i = H_i W_V$ are the query, key, and value matrices, respectively. W_Q, W_K, and W_V are learnable projection matrices, and d_k is the dimension of the key matrix. Then, we compute the final fused representation using $h_i^{\text{Attn}} = \sum_{j=1}^{3} H_i^{\text{Attn}}[j, :]$, achieved by summing over the rows of $H_i^{\mathbf{Attn}}$. Here, j refers to the j-th row of the learned attention output matrix. This operation aggregates the weighted semantic embeddings from different environmental relations into the final fused representation of accounts.

6) Detection Layer: The detection layer of EnvIT comprises a fully connected layer and a softmax layer. The final predicted fraud probability of the model for user i can be obtained by $\hat{y}_i = \text{softmax}(W_i^f h_i^{\text{Attn}} + b_i^f)$, where the output of the attentive module is projected into a two-dimensional space with the learnable weight matrix W_i^f and the learnable bias matrix b_i^f. We use $\hat{Y}$ denoting the predicted labels for all account nodes.

C. EnvIT Algorithm and Model Optimization

Algorithm 1 outlines the detailed process of EnvIT. Line 1 initializes the model by setting up its three main components and the necessary parameters. Lines 2-10 illustrate the training process of EnvIT. Specifically, Lines 3-5 describe the process of extracting node representations based on EnvIT's three essential components. In Line 6, these representations are fused using the

Algorithm 1: EnvIT Algorithm.

Input: A heterogeneous environment graph $\mathcal{G}_e$, a homogeneous account graph $\mathcal{G}_a$, a behavior sequence S, initial node features X, ground-truth labels Y, and number of epochs n_{epochs}

Output: Predicted account labels $\hat{Y}_{\text{test}}$

1 **Initialize** model parameters randomly
// --- Training Phase ---
2 **for** *epoch* $\leftarrow 1$ **to** n_{epochs} **do**
3 $h_i^a \leftarrow \texttt{AggregationExtractor}(\mathcal{G}_a, X)$
4 $h_i^d \leftarrow \texttt{DifferenceExtractor}(\mathcal{G}_e, X)$
5 $h_i^b \leftarrow \texttt{BehaviorExtractor}(S, X)$
6 $h_i^{\text{Attn}} \leftarrow \texttt{AttentionFusion}(h_i^a, h_i^d, h_i^b)$
7 $\hat{Y}_{\text{train}} \leftarrow \texttt{DetectionLayer}(h_i^{\text{Attn}})$
8 $\mathcal{L} \leftarrow \texttt{Loss}(\hat{Y}_{\text{train}}, Y_{\text{train}})$
9 Backpropagate and update model parameters
10 **end**
// --- Inference Phase ---
11 Load the best model parameters
12 $\hat{Y}_{\text{test}} \leftarrow \texttt{DetectionLayer}(h_i^{\text{Attn}})$
13 **return** $\hat{Y}_{test}$

AttentionFusion module. Line 7 generates the predicted output of EnvIT. Lines 8-9 handle the computation of the loss and the update of model parameters. After training converges, we save the optimized model as the best model, and the learned representations of the account nodes can be derived. Line 11 involves the loading of the best model, and Lines 12-13 perform inference on the test dataset, producing the final predicted labels.

To train our model, we employ the cross-entropy loss function, defined as $\mathcal{L} = -\sum_{i=1}^{N}(y_i \log(\hat{y}_i) + (1-y_i)\log(1-\hat{y}_i))$, where y_i represents the true label of account i, which indicates whether the account i is categorized as normal or fraudulent. The total number of accounts is denoted as N. It should be noted that practical scenarios may involve unlabeled accounts. Consequently, we focus on the loss function exclusively on labeled accounts and adopt a semi-supervised training strategy for the proposed EnvIT model.

V. Experiments

A. Experimental Setup

1) Dataset: We evaluate the detection performance of EnvIT on two datasets from different leading platforms. The first dataset, described in Section III-A, represents one of the scenarios of identity theft fraud on e-commerce platforms. In this particular scenario, fraudsters pretend to be others to take out loans in consumer financial services. We assign labels of 1, 0, and -1 to the fraudulent accounts, normal accounts, and unlabeled accounts, respectively. To create the heterogeneous graph, we extract detailed information such as the initial three bytes of the IP addresses, the device UUIDs derived from the login records, and the latitude and longitude values obtained from delivery addresses. This graph establishes connections between accounts and various entities within the access environment. Additionally, we build a homogeneous graph consisting of the verified accounts sharing the same personal ID.

The other dataset[5] is from Vesta Corporation,[6] one of the world's leading payment service company. This dataset (denoted as Vesta dataset) comprises a wide range of features from device type to product features of real-world e-commerce transactions. It is widely used in the research pertaining to the credit fraud detection [63], [64]. Therefore, we select this dataset to represent another scenario of identity theft fraud on e-commerce platforms where fraudsters steal others' credit cards to pay for orders. To create user accounts, we extract the unique user ID from each transaction record based on the bankcard, address, days related features. The transaction behavior sequences of users are sorted based on the account. If an account has one or more transaction records labeled as fraudulent, the account will be labeled as a fraudulent account (marked as 1). The account with all normal transactions will be labeled as a normal account (marked as 0). We combine all bankcard features in each transaction record to represent the verified personal ID and we consider the network information and device information as the entities of access environments. Similarly, we construct a heterogeneous graph linking accounts and other entities within the access environment. Additionally, we create a homogeneous graph for the verified accounts that share the same personal ID. The statistics of the graphs on the above two datasets are illustrated in Table IV.

TABLE IV
Summary of Two Datasets

	Meituan dataset		Vesta dataset	
# Nodes	Normal accounts	14,156		
	Fraud accounts	1,194	Normal accounts	72,646
	Unlabeled accounts	6,993	Fraud accounts	8,585
	IP	46,543	Network	65,225
	Device	84,033	Device	11,150
	Address	40,842		
# Edges	Account-IP	97,106		
	Account-Device	87,972	Account-Network	82,489
	Account-Address	49,620	Account-Device	20,101
	Account-Account	48,655	Account-Account	93,875

In our experiments, we divided each dataset into three parts: training set, validation set, and test set. The split was done randomly with a ratio of 3:1:1. During training, we utilized all the accounts in the training set and computed the cross-entropy loss based solely on the predicted results of the labeled accounts.

2) Baselines: Given the absence of existing methods for financial identity fraud detection, we conducted a comparative analysis of EnvIT against various representative methods. These methods are selected based on their suitability for fraud detection and can be categorized into three groups:

- *Sequence-based methods:* **LSTM** [61], **Bi-LSTM** [65], **GRU** [66], **HATnt-LSTM** [12], and **TSF** [41]. LSTM, Bi-LSTM, and GRU are well-established recurrent neural networks that are commonly employed for modeling fraudulent behavior sequences [67]. HATnt-LSTM is an approach for fraud detection that focuses on sequential behavior representations. It excels at capturing the temporal dynamics of fraudulent activities, incorporating the interval time between consecutive time steps. To generate behavior embeddings from the historical behavior of accounts, we employ HATnt-LSTM. Then, these behavior embeddings are fed into the final detection layer for identification of fraudulent accounts. TSF examines the issue of detecting compromised accounts caused by phone number reassignments. It incorporates a temporal pattern encoder and a

[5]https://www.kaggle.com/c/ieee-fraud-detection/overview
[6]https://www.vesta.io/

statistical feature encoder to capture behavioral evolution and significant operational features.

- *Homogeneous GNNs:* **GraphSAGE** [59] and **GAT** [58]. Both of these graph neural networks are commonly adopted for graph fraud detection. We streamline the classification task by transforming the heterogeneous environment graph into homogeneous environment graphs. Subsequently, we learn the embeddings of account nodes in each environment graph, as well as the account graph. Finally, we concatenate the embeddings of each graph together for the final detection.
- *Heterogeneous GNNs:* **R-GCN** [68], **HAN** [69], **HACUD** [24], **SOBT** [5], and **SEC-GFD** [54]. R-GCN is a specialized graph neural network designed for modeling complex relations among nodes in heterogeneous graphs. HAN considers meta-paths from neighbors and employs node-level and semantic-level attentions to learn informative node embeddings. Both R-GCN and HAN are suitable for identifying the fraud nodes. In the context of financial services, HACUD emerges as a cash-out fraud detection model that considers diverse attributes and relations among various entities based on users' preferences. We leverage them to model the account embeddings for the final detection. SOBT is a credit card fraud detection model that uses a fraud feature-boosting mechanism. It incorporates a spiral oversampling balancing technique to balance the ratio of legitimate to fraudulent transactions and improve the model's discrimination capability. SEC-GFD designs a Spectrum-Enhanced and Environment-Constrainted Graph Fraud Detector (SEC-GFD) to comprehensively incorporate spectrum and label information into a fraud detector.

Furthermore, we compare EnvIT with two of its variants, namely **EnvIT-Graph** and **EnvIT-Beh**. The former exclusively integrates environmental difference and environmental aggregation as inputs to the final detection layer. The latter relies solely on behavioral features to generate behavioral embeddings, which are subsequently utilized in the final detection layer.

3) Metrics: We use the Area Under Curve (AUC) [70] metric to evaluate the detection performance. This metric represents the area under the receiver operating characteristic (ROC) curve, assessing the pairwise ranking performance of the classification results between positive and negative instances. The AUC is a crucial metric in fraud detection, particularly in imbalanced classification scenarios. To prevent normal accounts from being misclassified as fraudulent, we strive to achieve high precision in practice. However, it is equally important to identify as many fraudulent accounts as possible, thus achieving a high recall. Following previous studies [71], [72], we employ *Recall@T%Precision* as another metric. This metric is defined as the recall value obtained when the model achieves a precision score of $T\%$. To compare the detection performance of all models, we set T to be 80, 85, and 90, respectively.

4) Implementation Details: All models are based on PyTorch [73] and Deep Graph Library (DGL) [74]. For a fair comparison, the learning rate, batch size, and embedding size are set to 1×10^{-3}, 256, and 32, respectively, to maintain consistency across all models. Adam is chosen as the optimizer. For sequence-based methods and the behavior feature extractor in EnvIT, we set the number of time steps to 40. Furthermore, we employ 2 propagation layers for all GNNs. Particularly, GraphSAGE, EnvIT-Graph, and EnvIT adopt a "mean-pooling" aggregation strategy. We sample a maximum of 100 neighbors. Additionally, the GAT, HAN, EnvIT-Graph, and EnvIT incorporate 8 attention heads. For the Meituan dataset, we leverage three meta-paths based on IP address, device, and delivery address for HAN, HACUD, EnvIT-Graph, and EnvIT. On the Vesta dataset, we select two meta-paths based on the network information and device for HAN, HACUD, EnvIT-Graph, and EnvIT.

When training EnvIT and other baseline models to evaluate the overall performance, there is some randomness in the experiments (e.g., the partition of the dataset and the training process of each neural network). Therefore, we use a set of random seeds to control the randomness. They are generated by numpy.random, a Python-based random number routine. We run 5 rounds to train all models with the same set of random seeds, and the results are averaged. Additionally, the early stopping mechanism is employed to avoid over-fitting. The training process will stop if the AUC value in the validation set does not increase for 15 epochs.

B. Performance Evaluation

1) Overall Performance: Based on the experimental results (see Table V), a number of conclusions can be drawn: (1) The inclusion of time intervals allows EnvIT-Beh to outperform all sequence-based methods except for TSF, showcasing its superior performance by considering the varying time intervals. (2) Incorporating the connections among verified accounts sharing the same personal ID (i.e., the account graph) enhances the performance of most GNNs. (3) EnvIT-Graph outperforms other GNNs, particularly in terms of *Recall@T%Precision*. Among verified accounts sharing the same personal ID, this finding supports the presence of environmental differences. (4) On Meituan dataset, GAT surpasses GraphSAGE and R-GCN. This can be attributed to GAT's attention mechanism, which effectively aggregates information from different neighbors while considering their relative importance. However, the performance of GAT is worse than that of GraphSAGE and R-GCN on Vesta dataset. This is mainly attributed to the fact that the importance of different neighbor nodes has little significance, because the average degree of the account nodes in Vesta dataset is so small. (5) HAN goes beyond accounting for the importance of different neighbors and incorporates the importance of different meta-paths. This enables HAN to outperform GAT on both datasets. (6) HACUD demonstrates an inferior performance on both datasets, demonstrating the distinction between the cash-out and the financial identity theft as different types of financial fraud. (7) SOBT outperforms many sequence-based models and several baselines on both datasets, likely due to its advanced feature selection and oversampling techniques. Similarly, SEC-GFD performs better than most graph-based models, which could be

TABLE V
PERFORMANCE COMPARISON OF ENVIT AND BASELINES. THE W/O ID INDICATES THE MODEL EXCLUSIVELY LEVERAGING THE HETEROGENEOUS ENVIRONMENT GRAPH WITHOUT THE ACCOUNT GRAPH, R@T%P IS SHORT FOR RECALL@T%PRECISION. THE BEST IS BOLD.

Model	Meituan dataset					
	AUC	R@80%P	R@85%P	R@90%P	Recall	F1-Score
LSTM	0.7893	0.2149	0.1904	0.1470	0.2650	0.3882
Bi-LSTM	0.7681	0.2144	0.1887	0.1608	0.2691	0.3927
GRU	0.7774	0.2344	0.2009	0.1716	0.2837	0.4119
HAInt-LSTM	0.8044	0.2695	0.2243	0.1940	0.3054	0.4333
TSF	0.9014	0.3195	0.2765	0.2361	0.3047	0.4276
GraphSAGE (w/o ID)	0.8680	0.2409	0.2031	0.1415	0.3011	0.4293
GraphSAGE	0.8801	0.3417	0.2787	0.2283	0.3445	0.4814
GAT (w/o ID)	0.8835	0.3978	0.3740	0.3501	0.3810	0.5222
GAT	0.8932	0.4622	0.4048	0.3473	0.4146	0.5534
R-GCN (w/o ID)	0.8462	0.2235	0.1921	0.1593	0.2579	0.3804
R-GCN	0.8637	0.2478	0.2282	0.2239	0.3017	0.4268
HAN (w/o ID)	0.8792	0.3638	0.3145	0.2898	0.3464	0.4868
HAN	0.8916	0.4247	0.3507	0.3159	0.3696	0.5151
HACUD (w/o ID)	0.8716	0.2777	0.2506	0.2158	0.3125	0.4463
HACUD	0.8836	0.3776	0.3367	0.2909	0.3652	0.5059
SOBT	0.8910	0.3221	0.2815	0.2115	0.4146	0.5206
SEC-GFD	0.8815	0.3474	0.3053	0.2812	0.3870	0.5045
EnvIT-Graph	0.8921	0.4462	0.4076	0.3559	0.4475	0.5801
EnvIT-Beh	0.8006	0.2786	0.2191	0.1966	0.3243	0.4503
EnvIT	**0.9031**	**0.4619**	**0.4171**	**0.3672**	**0.4503**	**0.5805**
Model	Vesta dataset					
	AUC	R@80%P	R@85%P	R@90%P	Recall	F1-Score
LSTM	0.7993	0.3516	0.3120	0.3011	0.1773	0.2709
Bi-LSTM	0.7993	0.3609	0.3366	0.3063	0.1443	0.2301
GRU	0.8111	0.4066	0.3526	0.3109	0.1230	0.1969
HAInt-LSTM	0.8121	0.4097	0.3777	0.3246	0.1705	0.2667
TSF	0.8757	0.3858	0.3322	0.2808	0.3404	0.5204
GraphSAGE (w/o ID)	0.8745	0.4695	0.4108	0.3727	0.1226	0.4469
GraphSAGE	0.9003	0.5498	0.4980	0.4124	0.3347	0.4635
GAT (w/o ID)	0.8523	0.4057	0.3769	0.3687	0.2051	0.2036
GAT	0.8913	0.4828	0.4155	0.3767	0.2292	0.3476
R-GCN (w/o ID)	0.8704	0.4597	0.4087	0.3593	0.2051	0.2898
R-GCN	0.8899	0.4999	0.4544	0.3865	0.3121	0.4372
HAN (w/o ID)	0.8721	0.4758	0.4249	0.3458	0.2232	0.3392
HAN	0.8919	0.5335	0.4901	0.3947	0.3325	0.4581
HACUD (w/o ID)	0.8589	0.4373	0.4016	0.3632	0.3065	0.5047
HACUD	0.8674	0.4436	0.4047	0.3650	0.4002	0.4408
SOBT	0.8558	0.4957	0.4481	0.3962	0.4771	0.6012
SEC-GFD	0.8723	0.3400	0.3066	0.2663	**0.5804**	0.5488
EnvIT-Graph	0.9116	0.5894	0.5228	0.4483	0.4160	0.5384
EnvIT-Beh	0.8273	0.4181	0.3949	0.3603	0.1797	0.2744
EnvIT	**0.9187**	**0.6172**	**0.5648**	**0.4782**	0.5402	**0.6355**

due to its incorporation of spectrum and label information for fraud detection.

Furthermore, EnvIT-Graph surpasses EnvIT-Beh, emphasizing the greater importance of the associations across environmental entities over the historical activities. Nonetheless, neither approach outperforms EnvIT. This highlights the importance of simultaneously incorporating both the associations across environmental entities and historical activity sequences into the detection model. We examine the significance of performances of models with Welch's *t*-test [75] to strengthen the experimental validity. We calculate the corresponding p-values to investigate the difference between the performance of our model and the baselines. By performing Welch's *t*-test, we find that the value of p is less than 0.05. This verifies the significance of the detection results in distinguishing fraudulent accounts from normal accounts.

Taking the Meituan dataset as an example, the Precision-Recall curve and the ROC curve of EnvIT and baselines are displayed in Fig. 3. The results clearly demonstrate that EnvIT consistently exhibits a superior performance in terms of both precision and recall values. This underscores the effectiveness of EnvIT in fraud detection.

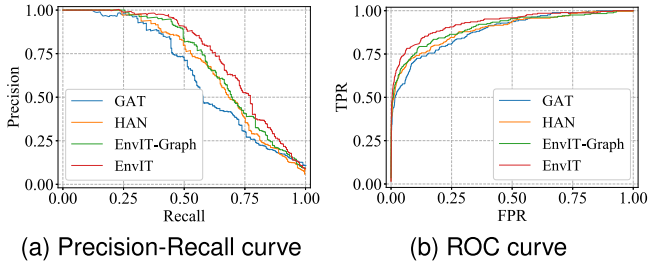

Fig. 3. Precision-Recall and ROC curves of models on Meituan dataset.

TABLE VI
ABLATION STUDY OF ENVIT (R@T%P IS SHORT FOR RECALL@T%PRECISION)

Dataset	Model	AUC	R@80%P	R@85%P	R@90%P
Meituan dataset	EnvIT (w/o Env_agg)	0.8502	0.3038	0.2785	0.2152
	EnvIT (w/o Env_diff)	0.8859	0.3797	0.3586	0.3333
	EnvIT (w/o Beh.)	0.8921	0.4462	0.4076	0.3559
	EnvIT (w/o Attn.)	0.8972	0.4515	0.3165	0.2869
	EnvIT	**0.9031**	**0.4619**	**0.4171**	**0.3672**
Vesta dataset	EnvIT (w/o Env_agg)	0.8990	0.5375	0.4852	0.4071
	EnvIT (w/o Env_diff)	0.8561	0.4456	0.3991	0.3529
	EnvIT (w/o Beh.)	0.9116	0.5894	0.5228	0.4483
	EnvIT (w/o Attn.)	0.9052	0.5642	0.4891	0.3914
	EnvIT	**0.9187**	**0.6172**	**0.5648**	**0.4782**

We attribute this superior performance to EnvIT's unique design, which jointly models heterogeneous environmental associations and historical activity sequences. In contrast, existing state-of-the-art models typically focus on one aspect (e.g., only temporal modeling or only graph structure) and fail to capture the intricate dynamics between users and their surrounding environment. These results confirm that the combination of complex social relationships and temporal information within a unified architecture leads to a more expressive and effective fraud detection model.

2) Ablation Study: To gain insights into the impact of different components on detection performance, we performed an ablation study. Table VI presents the results for four variants of EnvIT.

- EnvIT (w/o Env_agg): EnvIT removes the environmental aggregation extractor;
- EnvIT (w/o Env_diff): EnvIT eliminates the environmental difference extractor;
- EnvIT (w/o Beh.): EnvIT removes the behavior feature extractor, i.e., EnvIT-Graph;
- EnvIT (w/o Attn.): EnvIT replaces the attentive feature fusion module by employing an averaging manner to fuse different features.

According to Table VI, we observe that the importance of different modules varies across the two datasets. The environmental aggregation extractor contributes the most to the Meituan dataset, while the environmental difference extractor contributes the most to the Vesta dataset. This discrepancy may be due to the different fraud behaviors in the two datasets. In the Meituan dataset, fraudulent accounts tend to form clusters through shared environmental entities, making the aggregation patterns more informative. In contrast, the Vesta dataset contains more dispersed and heterogeneous fraud activities, where frequent changes in environments are used to evade detection. The complementary nature of the two modules enhances the robustness of the model against various fraud strategies. For both datasets, historical

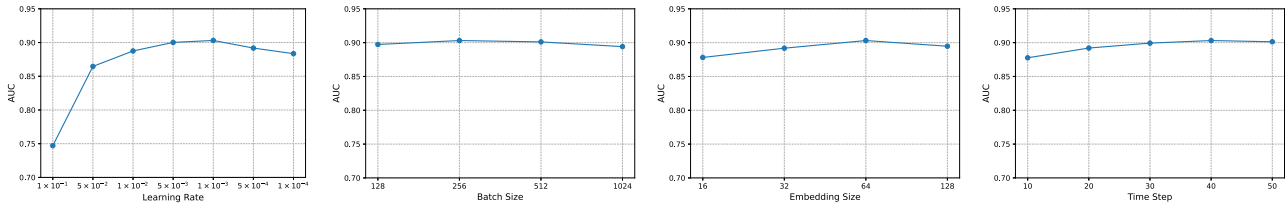

(a) Learning Rate (b) Batch Size (c) Embedding Size (d) Time Step

Fig. 4. Experimental results of different hyper-parameters. (a)-(d) are different hyper-parameter settings for learning rates, batch sizes, embedding sizes, and time steps, respectively.

TABLE VII
EFFECTS OF DIFFERENT ENVIRONMENTAL SUBGRAPHS

Dataset	Model	AUC	R@80%P	R@85%P	R@90%P
Meituan dataset	EnvIT (M1)	0.8930	0.4219	0.3797	0.3629
	EnvIT (M2)	0.9007	0.4515	0.4051	0.3840
	EnvIT (M3)	0.8839	0.4051	0.3502	0.3333
	EnvIT	**0.9031**	**0.4619**	**0.4171**	**0.3672**
Vesta dataset	EnvIT (M4)	0.9075	0.5862	0.5177	0.4359
	EnvIT (M2)	0.9091	0.6063	0.5247	0.4454
	EnvIT	**0.9187**	**0.6172**	**0.5648**	**0.4782**

TABLE VIII
EFFECTS OF DIFFERENT ENVIRONMENTAL METAPATHS

Dataset	Model	AUC	R@80%P	R@85%P	R@90%P
Meituan dataset	EnvIT (M5)	0.8899	0.4557	0.3839	0.3598
	EnvIT (M6)	0.8691	0.3966	0.2827	0.2110
	EnvIT (M7)	0.8996	0.4557	0.4093	0.3882
	EnvIT	**0.9031**	**0.4619**	**0.4171**	**0.3672**
Vesta dataset	EnvIT (M8)	0.9036	0.5557	0.4885	0.3908
	EnvIT (M6)	0.9096	0.5770	0.4983	0.4201
	EnvIT	**0.9187**	**0.6172**	**0.5648**	**0.4782**

behavior is the least important. Notably, the performance of EnvIT (w/o Attn.) is inferior to EnvIT, which suggests the need for a well-designed fusion approach for the three representations. This observation confirms the effectiveness of the attentive module we have incorporated.

3) Evaluation on Different Environmental Features: By varying the type of sub-graphs and meta-paths for evaluation, we can observe the impact of environment-related sub-graphs and meta-paths on the overall model's detection performance.

The performance results of EnvIT without different sub-graphs of environmental differences are shown in Table VII. "w/o ip_subgraph" (M1), "w/o device_subgraph" (M2), "w/o address_subgraph" (M3), and "w/o network_subgraph" (M4) represent deleting the account-IP address sub-graph, account-device sub-graph, account-delivery address sub-graph, and account-network sub-graph on the basis of the original model, respectively. According to the results, EnvIT's prediction performance drops the most for the Meituan dataset when the environmental difference feature of the delivery address is excluded. For the Vesta dataset, the environmental difference feature of network information used by verified accounts with the same personal ID is more significant than the differences between devices.

The performance results of EnvIT without different features for environmental aggregation are shown in Table VIII . Specifically, "w/o ip_meta-path" (M5), "w/o device_meta-path" (M6), "w/o address_meta-path" (M7), and "w/o network_meta-path" (M8) indicate that EnvIT does not consider meta-paths based on IP address, device, delivery address, and network information, respectively. In the Meituan dataset, removing the device meta-path (M6) decreases AUC from 0.9031 to 0.8691 ($\downarrow 3.4\%$), R@80%P from 0.4619 to 0.3966 ($\downarrow 6.53\%$), R@85%P from 0.4171 to 0.2827 ($\downarrow 13.44\%$), and R@95%P from 0.3672 to 0.2110 ($\downarrow 15.62\%$). One could observe that the device aggregation is the most important feature. For the Vesta dataset, removing the network meta-path (M8) lowers AUC from 0.9187 to 0.9036 ($\downarrow 1.51\%$), R@80%P from 0.6172 to 0.5557 ($\downarrow 6.15\%$), R@85%P from 0.5648 to 0.4885 ($\downarrow 7.63\%$), and R@95%P from 0.4782 to 0.3908 ($\downarrow 8.74\%$), indicating that network information has the greatest impact. We also observe that different metrics respond differently. Specifically, AUC is relatively stable across different meta-path settings, reflecting the overall classification capability of models, while R@T%P metrics are more sensitive because they emphasize high values of the precision, where missing key environmental information could reduce the value of the recall. This highlights the importance of specific environmental relationships under strict constraints of the precision.

4) Hyper-Parameter Study: We choose a set of values to train the model for each hyper-parameter. We select the value where the AUC achieves the highest on the validation set. Specifically, we tune the learning rate, batch size, embedding size, and number of time steps in $[1 \times 10^{-1}, 5 \times 10^{-2}, 1 \times 10^{-2}, 5 \times 10^{-3}, 1 \times 10^{-3}, 5 \times 10^{-4}, 1 \times 10^{-4}]$, [128, 256, 512, 1024], [16, 32, 64, 128], and [10, 20, 30, 40, 50], respectively. We present the detection results of EnvIT (AUC metric) with different learning rates and batch sizes in Fig. 4(a) and (b), respectively. Fig. 4(c) shows the results of EnvIT with embedding sizes of hidden states in the environmental difference extractor and the environmental aggregation extractor. Similarly, the results of EnvIT with different time steps of the behavior feature extractor are presented in Fig. 4(d).

According to Fig. 4(a), the change in learning rate has a great impact on EnvIT's detection performance. As the learning rate decreases, the AUC obtained by EnvIT first increases and then decreases. The best result is achieved when the learning rate is 1×10^{-3}, indicating that choosing 1×10^{-3} is better than choosing other values. From Fig. 4(b), we observe that EnvIT performs the best on AUC when the batch size is 256. One can see that different embedding sizes have little influence on the experimental results in Fig. 4(c). Under different embedding sizes, the AUC of EnvIT can reach more than 0.90. Choosing an embedding size of 64 is optimal in our experimental settings.

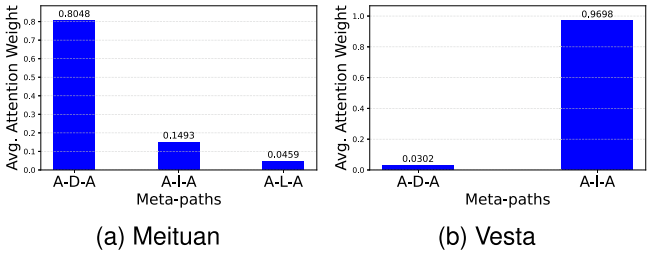

Fig. 5. Averaged attention weights over meta-paths for fraudulent accounts on two datasets: (a) Meituan and (b) Vesta. In the Meituan dataset, the A-D-A meta-path receives the highest attention weight, while the A–I–A meta-path plays a significant role in the Vesta dataset.

Generally, the longer the behavior sequence, the more information the behavior feature extractor can learn. However, each account has a different length of the behavior sequence. We use the zero-padding technique when the sequence length of an account is less than the number of time steps. Therefore, it is not appropriate to choose a very long sequence, which will result in too many zeros in the input of the behavior feature extractor. Fig. 4(d) shows that 40 is the optimal value of the number of time steps that achieves an AUC greater than 0.90.

C. Visualization of Attention Weights Based on Meta-Paths

To better understand how our model leverages the environmental structure for financial identity theft detection, we visualize the averaged attention weights across meta-paths on two datasets: account-device-account (A-D-A), account-IP (Network)-account (A-I-A), and account-location-account (A-L-A) in Fig. 5. The Meituan dataset shows that the meta-path A-D-A carries the highest significance with a weight of 0.8048, as displayed in Fig. 5(a). In contrast, the meta-paths A-I-A and A-L-A have little impact (0.1493 and 0.0459, respectively). However, for the Vesta dataset, the A–I–A meta-path receives the highest weight (0.9698), while the A–D–A meta-path contributes minimally (0.0302), as shown in Fig. 5(b). Such varied weighting across datasets indicates that our model emphasizes different interactions, which indicates suspicious behavior for different platforms. By assigning higher weights to semantically informative paths, the model effectively captures structural signals relevant to fraud detection. Therefore, the attention mechanism used in EnvIT not only reveals the global importance of meta-paths but also provides case-level interpretability. For each account, the attention weights highlight which environmental relations (e.g., device, IP, or address) and behavioral patterns (e.g., login sequences) contributed most to the prediction. These insights can guide fraud analysts in focusing on key features and interactions, providing a foundation for further investigation.

VI. Limitations and Discussion

In this section, we listed some limitations of our method and discussed some key issues.

A. Limitations

The above experimental results demonstrate EnvIT's superior performance in the detection of identity theft fraud accounts. However, there are some limitations of this work. We will discuss them and provide a vision for future work.

- First, EnvIT assumes that the historical behaviors (i.e., login and payment records) of the account exists when detecting whether an account is fraudulent. This premise allows us to extract the devices, IP addresses and other entities within the access environment used by the account from the historical behaviors. However, some fraudulent accounts do exhibit few activities. They are isolated nodes in the heterogeneous environment graph. Also, most of the time steps of their behavior sequence would be padded by zero vectors. Therefore, it is difficult to learn the fraud patterns of these accounts. To overcome this limitation, we will employ the "cross-site linking" function [76], [77] to import social relationships from other websites. Since the accounts used by the same user on different online websites can be linked by the "cross-site linking" function, it is possible to use the social relationships on other websites to address the problem of the sparsity of accounts' relationships on O2O e-commerce platforms.
- Additionally, EnvIT incorporates unlabeled accounts through message passing in the heterogeneous graph, enabling structural and behavioral information from labeled accounts to propagate to unlabeled nodes. Although we have adopted a semi-supervised learning approach for model training, we recognize that some advanced fraudulent accounts may not have been identified. Since unlabeled accounts only participate in the message passing process, but do not generate supervised signals, there is a risk of model learning bias. Therefore, future work might consider investigating the background nodes [78], [79] to advance our method and introducing methods such as pseudo-labeling [80], [81] or contrastive learning [82], [83] to better adapt to such scenarios. Pseudo-labeling could serve to expand the effective labeled accounts, and contrastive learning might aid in distinguishing subtle differences between fraudulent and normal activities. Employing these strategies could alleviate label bias and enhance the detection of silent or undetected fraud. Moreover, we only adopt under-sampling to deal with the class imbalance problem of identity theft fraud. Under-sampling is applied effectively in feature-based methods. While in graph-based methods, under-sampling may remove important structural relationships of the graph, thus limiting the model's ability to generalize in real online deployment. The study of the class imbalance problem on graphs is promising [56], [84]. In the future, we also plan to design other samplers to solve the class imbalance problem of identity theft fraud.
- Finally, EnvIT's performance on identity theft detection has demonstrated its superiority of capturing fraud patterns from access environment and behavior sequences. As we all know, many factors lead to identity theft fraud, including phishing attacks [85], malware software, social

media fraud [86], and loss of items. Further differentiation among the different types is a valuable direction. In the future, we aim to assign finer-grained labels to each account based on user feedback. We plan to extend our approach to distinguish finer-grained types that cause identity theft fraud.

B. Practical Deployment Challenges

We discussed some practical challenges in deploying EnvIT in real-world fraud detection systems. First, low-latency inference is crucial in high-traffic environments, which requires further optimization of our model to meet real-time demands. Second, fraud patterns evolve rapidly in practice, so the model must adapt over time to remain effective. Lastly, interpretability is vital for compliance and trust, especially in financial domains where decisions must be explainable. In future work, we plan to develop visualization tools to make these case-level explanations more accessible and actionable for analysts.

C. Model Robustness

Although EnvIT demonstrates strong performance on both benchmark datasets and in real-world deployment, it currently does not explicitly address adversarial evasion tactics such as synthetic identity creation or device/VPN spoofing. However, fine-grained behavior modeling and graph aggregation may implicitly capture inconsistencies associated with such behaviors. To further enhance robustness against evolving fraud strategies, future work will explore techniques such as adversarial behavior simulation, and self-supervised pre-training on evolving user activity to better adapt to real-time data drift and adversarial attacks.

VII. Conclusion

This paper investigates the problem of detecting financial identity theft in O2O e-commerce services. We analyze real-world cases from Meituan. Our findings reveal some key fraud patterns at different levels, such as associations across environmental entities and behavioral activities exhibited by accounts. Based on these important discoveries, we introduce EnvIT, a deep learning-based framework to detect financial identity theft fraud by considering both the environmental association factors and the historical activities of accounts. The effectiveness of EnvIT and the interpretability of the detection results are demonstrated by experimental results on two datasets from Meituan and Vesta. We empirically show that factors from the associated access environment and temporal behavior activities are critical for the detection of financial identity theft fraud.

References

[1] F. Xu, Z. Han, J. Piao, and Y. Li, "'I think you'll like it': Modelling the online purchase behavior in social E-commerce," in *Proc. ACM Hum. Comput. Interact.*, 2019, vol. 3, no. 65, pp. 1–23.

[2] T. Liang et al., "Credit risk and limits forecasting in E-commerce consumer lending service via multi-view-aware mixture-of-experts nets," in *Proc. 14th ACM Int. Conf. Web Search Data Mining*, 2021, pp. 229–237.

[3] Federal Trade Commission (FTC), "Consumer sentinel network data book 2021," FTC, 2022. [Online]. Available: https://www.ftc.gov/reports/consumer-sentinel-network-data-book-2021

[4] Federal Trade Commission (FTC), "What to know about identity theft," FTC, 2021. [Online]. Available: https://www.consumer.ftc.gov/articles/what-know-about-identity-theft

[5] L. Ni, J. Li, H. Xu, X. Wang, and J. Zhang, "Fraud feature boosting mechanism and spiral oversampling balancing technique for credit card fraud detection," *IEEE Trans. Comput. Social Syst.*, vol. 11, no. 2, pp. 1615–1630, Apr. 2024.

[6] A. Mutemi and F. Bacao, "E-commerce fraud detection based on machine learning techniques: Systematic literature review," *Big Data Mining Analytics*, vol. 7, no. 2, pp. 419–444, 2024.

[7] A. McDonald, C. Sugatan, T. Guberek, and F. Schaub, "The annoying, the disturbing, and the weird: Challenges with phone numbers as identifiers and phone number recycling," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2021, pp. 1–14.

[8] K. Thomas et al., "Data breaches, phishing, or malware? Understanding the risks of stolen credentials," in *Proc. 2017 ACM SIGSAC Conf. Comput. Commun. Secur.*, 2017, pp. 1421–1434.

[9] A. Mirian, J. DeBlasio, S. Savage, G. M. Voelker, and K. Thomas, "Hack for hire: Exploring the emerging market for account hijacking," in *Proc. World Wide Web Conf.*, 2019, pp. 1279–1289.

[10] Z. Lei, Y. Nan, Y. Fratantonio, and A. Bianchi, "On the insecurity of SMS one-time password messages against local attackers in modern mobile devices," in *Proc. Netw. Distrib. System Secur. Symp.*, 2021, pp. 580–597.

[11] D. Hummer and D. J. Rebovich, "Identity theft and financial loss," in *Handbook on Crime and Technology*. Cheltenham, U.K.: Edward Elgar Publishing, 2023, pp. 38–53.

[12] J. Guo, G. Liu, Y. Zuo, and J. Wu, "Learning sequential behavior representations for fraud detection," in *Proc. IEEE Int. Conf. Data Mining*, 2018, pp. 127–136.

[13] Q. Gong et al., "DeepScan: Exploiting deep learning for malicious account detection in location-based social networks," *IEEE Commun. Mag.*, vol. 56, no. 11, pp. 21–27, Nov. 2018.

[14] D. Babaev, M. Savchenko, A. Tuzhilin, and D. Umerenkov, "E.T.-RNN: Applying deep learning to credit loan applications," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 2183–2190.

[15] D. Cheng, Y. Zhang, F. Yang, Y. Tu, Z. Niu, and L. Zhang, "A dynamic default prediction framework for networked-guarantee loans," in *Proc. 28th ACM Int. Conf. Inf. Knowl. Manage.*, 2019, pp. 2547–2555.

[16] Y. Zhu et al., "Modeling users' behavior sequences with hierarchical explainable network for cross-domain fraud detection," in *Proc. Web Conf.*, 2020, pp. 928–938.

[17] B. Branco, P. Abreu, A. S. Gomes, M. S. Almeida, J. T. Ascensão, and P. Bizarro, "Interleaved sequence RNNs for fraud detection," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 3101–3109.

[18] C. Liu et al., "Fraud transactions detection via behavior tree with local intention calibration," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 3035–3043.

[19] W. Lin et al., "Online credit payment fraud detection via structure-aware hierarchical recurrent neural network," in *Proc. 30th Int. Joint Conf. Artif. Intell.*, 2021, pp. 3670–3676.

[20] X. He, Q. Gong, Y. Chen, Y. Zhang, X. Wang, and X. Fu, "DatingSec: Detecting malicious accounts in dating apps using a content-based attention network," *IEEE Trans. Dependable Secure Comput.*, vol. 18, no. 5, pp. 2193–2208, Sep./Oct. 2021.

[21] Z. Liu, C. Chen, X. Yang, J. Zhou, X. Li, and L. Song, "Heterogeneous graph neural networks for malicious account detection," in *Proc. 27th ACM Int. Conf. Inf. Knowl. Manage.*, 2018, pp. 2077–2085.

[22] C. Liang et al., "Uncovering insurance fraud conspiracy with network learning," in *Proc. 42nd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2019, pp. 1181–1184.

[23] Z. Liu et al., "GeniePath: Graph neural networks with adaptive receptive paths," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 4424–4431.

[24] B. Hu, Z. Zhang, C. Shi, J. Zhou, X. Li, and Y. Qi, "Cash-out user detection based on attributed heterogeneous information network with a hierarchical attention mechanism," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 946–953.

[25] Y. Fan et al., "Metagraph aggregated heterogeneous graph neural network for illicit traded product identification in underground market," in *Proc. IEEE Int. Conf. Data Mining*, 2020, pp. 132–141.

[26] Q. Zhong et al., "Financial defaulter detection on online credit payment via multi-view attributed heterogeneous information network," in *Proc. Web Conf.*, 2020, pp. 785–795.

[27] M. Mendoza, M. Tesconi, and S. Cresci, "Bots in social and interaction networks: Detection and impact estimation," *ACM Trans. Inf. Syst.*, vol. 39, no. 1, pp. 1–32, 2020, doi: 10.1145/3419369.

[28] H. Peng, J. Zhao, L. Li, Y. Ren, and S. Zhao, "One-class adversarial fraud detection nets with class specific representations," *IEEE Trans. Netw. Sci. Eng.*, vol. 10, no. 6, pp. 3793–3803, Nov./Dec. 2023.

[29] J. Lian, X. Wang, X. Lin, Z. Wu, S. Wang, and W. Guo, "Graph anomaly detection via multi-view discriminative awareness learning," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 6, pp. 6623–6635, Nov./Dec. 2024.

[30] W. Yu, Y. Wang, L. Liu, Y. An, B. Yuan, and J. Panneerselvam, "A multiperspective fraud detection method for multiparticipant e-commerce transactions," *IEEE Trans. Computat. Social Syst.*, vol. 11, no. 2, pp. 1564–1576, Apr. 2024.

[31] H. Zhu, M. Zhou, G. Liu, Y. Xie, S. Liu, and C. Guo, "NUS: Noisy-sample-removed undersampling scheme for imbalanced classification and application to credit card fraud detection," *IEEE Trans. Computat. Social Syst.*, vol. 11, no. 2, pp. 1793–1804, Apr. 2024.

[32] H. Peng, R. Zhang, Y. Dou, R. Yang, J. Zhang, and P. S. Yu, "Reinforced neighborhood selection guided multi-relational graph neural networks," *ACM Trans. Inf. Syst.*, vol. 40, no. 4, pp. 1–46, 2021, doi: 10.1145/3490181.

[33] Alipay, "Multiple alipay accounts are allowed to be bound to the same personal ID," Alipay, 2019. [Online]. Available: https://docs.alipayplus.com/alipayplus/alipayplus/faq_acq/account_management

[34] X. Ding et al., "Delivery scope: A new way of restaurant retrieval for on-demand food delivery service," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 3026–3034.

[35] Y. Ping et al., "User consumption intention prediction in Meituan," in *Proc. 27th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2021, pp. 3472–3482.

[36] Q. Ye et al., "Modeling access environment and behavior sequence for financial identity theft detection in E-commerce services," in *Proc. Int. Joint Conf. Neural Netw.*, 2022, pp. 1–8.

[37] L. Cui, H. Seo, M. Tabar, F. Ma, S. Wang, and D. Lee, "DETERRENT: Knowledge guided graph attention network for detecting healthcare misinformation," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 492–502.

[38] J. Luo, M. Ye, C. Xiao, and F. Ma, "HiTANet: Hierarchical time-aware attention networks for risk prediction on electronic health records," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 647–656.

[39] F. Xiao, Y. Wu, M. Zhang, G. Chen, and B. C. Ooi, "MINT: Detecting fraudulent behaviors from time-series relational data," *Proc. VLDB Endowment*, vol. 16, no. 12, pp. 3610–3623, 2023.

[40] Y. Xie, G. Liu, C. Yan, C. Jiang, M. Zhou, and M. Li, "Learning transactional behavioral representations for credit card fraud detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 5735–5748, Apr. 2024.

[41] M. Gao et al., "Detecting compromised accounts caused by phone number recycling on E-commerce platforms: Taking Meituan as an example," *Front. Inf. Technol. Electron. Eng.*, vol. 25, no. 8, pp. 1077–1095, 2024.

[42] Y. Xie et al., "A spatial–temporal gated network for credit card fraud detection by learning transactional representations," *IEEE Trans. Automat. Sci. Eng.*, vol. 21, no. 4, pp. 6978–6991, Oct. 2024.

[43] S. Xiang et al., "Semi-supervised credit card fraud detection via attribute-driven graph representation," in *Proc. AAAI Conf. Artif. Intell.*, 2023, vol. 37, no. 12, pp. 14557–14 565.

[44] G. Zhang et al., "eFraudCom: An e-commerce fraud detection system via competitive graph neural networks," *ACM Trans. Inf. Syst.*, vol. 40, no. 3, pp. 1–29, 2022, doi: 10.1145/3474379.

[45] Y. Dou, Z. Liu, L. Sun, Y. Deng, H. Peng, and P. S. Yu, "Enhancing graph neural network-based fraud detectors against camouflaged fraudsters," in *Proc. 29th ACM Int. Conf. Inf. Knowl. Manage.*, 2020, pp. 315–324.

[46] S. Yang et al., "Financial risk analysis for SMEs with graph-based supply chain mining," in *Proc. 29th Int. Joint Conf. Artif. Intell.*, 2020, pp. 4661–4667.

[47] H. Wang, C. Zhou, J. Wu, W. Dang, X. Zhu, and J. Wang, "Deep structure learning for fraud detection," in *Proc. IEEE Int. Conf. Data Mining*, 2018, pp. 567–576.

[48] J. Zhang and Q. Xu, "Attention-aware heterogeneous graph neural network," *Big Data Mining Analytics*, vol. 4, no. 4, pp. 233–241, 2021.

[49] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 1, pp. 4–24, Jan. 2021.

[50] C. Liu, L. Sun, X. Ao, J. Feng, Q. He, and H. Yang, "Intention-aware heterogeneous graph attention networks for fraud transactions detection," in *Proc. 27th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2021, pp. 3280–3288.

[51] H. Wang et al., "Collaboration based multi-label propagation for fraud detection," in *Proc. 29th Int. Joint Conf. Artif. Intell.*, 2020, pp. 2477–2483.

[52] B. Hu et al., "Loan default analysis with multiplex graph learning," in *Proc. 29th ACM Int. Conf. Inf. Knowl. Manage.*, 2020, pp. 2525–2532.

[53] B. Xu, H. Shen, B. Sun, R. An, Q. Cao, and X. Cheng, "Towards consumer loan fraud detection: Graph neural networks with role-constrained conditional random field," in *Proc. AAAI Conf. Artif. Intell.*, 2021, pp. 4537–4545.

[54] F. Xu, N. Wang, H. Wu, X. Wen, X. Zhao, and H. Wan, "Revisiting graph-based fraud detection in sight of heterophily and spectrum," in *Proc. AAAI Conf. Artif. Intell.*, 2024, vol. 38, no. 8, pp. 9214–9222.

[55] D. Wang et al., "A semi-supervised graph attentive network for financial fraud detection," in *Proc. IEEE Int. Conf. Data Mining*, 2019, pp. 598–607.

[56] Y. Liu et al., "Pick and choose: A GNN-based imbalanced learning approach for fraud detection," in *Proc. Web Conf.*, 2021, pp. 3168–3177.

[57] Y. Sun, J. Han, X. Yan, P. S. Yu, and T. Wu, "PathSim: Meta path-based top-K similarity search in heterogeneous information networks," *Proc. VLDB Endowment*, vol. 4, no. 11, pp. 992–1003, 2011.

[58] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *Proc. 6th Int. Conf. Learn. Representations (Poster)*, 2018.

[59] W. L. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 1025–1035.

[60] I. M. Baytas, C. Xiao, X. Zhang, F. Wang, A. K. Jain, and J. Zhou, "Patient subtyping via time-aware LSTM networks," in *Proc. 23 rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 65–74.

[61] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.

[62] A. Vaswani et al., "Attention is all you need," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.

[63] H. Najadat, O. Altiti, A. A. Aqouleh, and M. Younes, "Credit card fraud detection based on machine and deep learning," in *Proc. 11th Int. Conf. Inf. Commun. Syst.*, 2020, pp. 204–208.

[64] K. I. Alkhatib, A. I. Al-Aiad, M. H. Almahmoud, and O. N. Elayan, "Credit card fraud detection based on deep neural network approach," in *Proc. 12th Int. Conf. Inf. Commun. Syst.*, 2021, pp. 153–156.

[65] A. Graves and J. Schmidhuber, "Framewise phoneme classification with bidirectional LSTM and other neural network architectures," *Neural Netw.*, vol. 18, no. 5/6, pp. 602–610, 2005.

[66] K. Cho et al., "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2014, pp. 1724–1734.

[67] J. Li, "Cyber security meets artificial intelligence: A survey," *Front. Inf. Technol. Electron. Eng.*, vol. 19, no. 12, pp. 1462–1474, 2018.

[68] M. S. Schlichtkrull, T. N. Kipf, P. Bloem, R. van den Berg, I. Titov, and M. Welling, "Modeling relational data with graph convolutional networks," in *Proc. Eur. Semantic Web Conf.*, 2018, pp. 593–607.

[69] X. Wang et al., "Heterogeneous graph attention network," in *Proc. World Wide Web Conf.*, 2019, pp. 2022–2032.

[70] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, 2006.

[71] A. Li, Z. Qin, R. Liu, Y. Yang, and D. Li, "Spam review detection with graph convolutional networks," in *Proc. 28th ACM Int. Conf. Inf. Knowl. Manage.*, 2019, pp. 2703–2711.

[72] T. Xu et al., "Deep entity classification: Abusive account detection for online social networks," in *Proc. 30th USENIX Secur. Symposium. USENIX*, 2021, pp. 4097–4114.

[73] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 8024–8035.

[74] M. Wang et al., "Deep graph library: A graph-centric, highly-performant package for graph neural networks," 2019, *arXiv:1909.01315*.

[75] B. L. Welch, "On the comparison of several mean values: An alternative approach," *Biometrika*, vol. 38, no. 3/4, pp. 330–336, 1951.

[76] Q. Gong, Y. Chen, J. Hu, Q. Cao, P. Hui, and X. Wang, "Understanding cross-site linking in online social networks," *ACM Trans. Web*, vol. 12, no. 4, 2018, Art. no. 25.

[77] Q. Gong et al., "Cross-site prediction on social influence for cold-start users in online social networks," *ACM Trans. Web*, vol. 15, no. 2, 2021, Art. no. 6.
[78] X. Huang et al., "DGraph: A large-scale financial dataset for graph anomaly detection," in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, vol. 35, pp. 22765–22777.
[79] L. Cao, H. Deng, Y. Yang, C. Wang, and L. Chen, "Graph-skeleton: 1% nodes are sufficient to represent billion-scale graph," in *Proc. ACM Web Conf.*, 2024, pp. 570–581.
[80] B. Zhang et al., "FlexMatch: Boosting semi-supervised learning with curriculum pseudo labeling," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, vol. 34, pp. 18408–18419.
[81] B. Wang et al., "Deep insights into noisy pseudo labeling on graph data," in *Proc. NeurIPS*, 2023, pp. 76214–76 228.
[82] S. Zhang, Z. Hu, A. Subramonian, and Y. Sun, "Motif-driven contrastive learning of graph representations," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 8, pp. 4063–4075, Aug. 2024.
[83] Y. Chen, J. Frias, and Y. R. Gel, "TopoGCL: Topological graph contrastive learning," in *Proc. AAAI Conf. Artif. Intell.*, 2024, vol. 38, no. 10, pp. 11453–11 461.
[84] M. Shi, Y. Tang, X. Zhu, D. A. Wilson, and J. Liu, "Multi-class imbalanced graph convolutional network learning," in *Proc. 29th Int. Joint Conf. Artif. Intell.*, 2020, pp. 2879–2885.
[85] M. Alsharnouby, F. Alaca, and S. Chiasson, "Why phishing still works: User strategies for combating phishing attacks," *Int. J. Hum.- Comput. Stud.*, vol. 82, pp. 69–82, 2015.
[86] D. Kuchhal and F. Li, "A view into YouTube view fraud," in *Proc. ACM Web Conf.*, 2022, pp. 555–563.

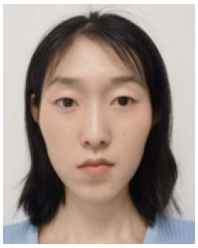

Min Gao received the B.S. degree from the School of Mathematics and Computer Science, Anshan Normal University, Anshan, China, in 2018, and the M.S. degree from the School of Mathematics and Information, Fujian Normal University, Fuzhou, China, in 2021. She is currently working toward the Ph.D. degree with the College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China. Her research interests include social network analysis, graph learning, and LLMs.

Qiongzan Ye received the B.E. degree from the School of Cyber Science and Engineering, Wuhan University, Wuhan, China, in 2019, and the M.S. degree from the School of Computer Science, Fudan University, Shanghai, China, in 2023. She was a Visiting Student with the School of Computer Science and Engineering, Beihang University, Beijing, China, in 2019. Her research interests include machine learning, data mining, network security, and privacy preservation.

Yangbo Gao received the B.E. and M.S. degrees from the Huazhong University of Science and Technology, Wuhan, China, in 2013 and 2016, respectively. He is currently working with Meituan Corp. His research interests include data mining and machine learning.

Zhenhua Zhang received the M.S. degree from the School of Information, Liaoning University, Shenyang, China. He is currently a Senior Technical Expert with Meituan Corp. He has more than 10 years of experience in risk control research and development, and holds 12 technical invention patents in the field of risk control. His research interests include machine learning and financial risk control.

Yu Chen received the B.E. and Ph.D. degrees from the Department of Computer Science and Technology, Tsinghua University, Beijing, China, in 2004 and 2009, respectively. He is currently a Technical Director of Meituan Corp. His research interests include data management, data mining, and machine learning.

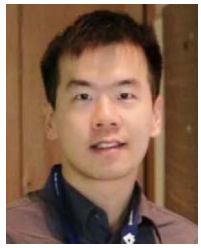

Yupeng Li (Member, IEEE) received the Ph.D. degree in computer science from The University of Hong Kong, Hong Kong. He was a Postdoctoral Researcher with the University of Toronto, Toronto, ON, Canada. He is currently an Assistant Professor with Hong Kong Baptist University, Hong Kong. He has authored or coauthored articles in prestigious venues, such as IEEE INFOCOM, ACM MobiHoc, IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, IEEE/ACM TRANSACTIONS ON NETWORKING. His research interests include network science and in particular, algorithmic decision making and machine learning problems, which arise in networked systems, such as information networks, and ride-sharing platforms. He is also excited about interdisciplinary research that applies algorithmic techniques to edging problems. He serves on the technical committees of some top conferences, for example, IEEE INFOCOM. He is a Member of ACM.

Shutong Chen received the B.E. degree in computer science from Fudan University, Shanghai, China, in 2021.She is currently working toward the master's degree in data science with Tsinghua-Berkeley Shenzhen Institute, Shenzhen, China. Her research interests include machine learning and social network mining.

Qingyuan Gong received the Ph.D. degree in computer science from the School of Computer Science, Fudan University, Shanghai, China, in 2020. She was a Visiting Student with the University of Göttingen, Göttingen, Germany, in 2015 and 2019, also the University of Chicago, Chicago, USA, in 2018. She was a Postdoctoral Researcher with the School of Computer Science, Fudan University until 2022. She is currently a pretenure Associate Professor with the Research Institute of Intelligent Complex Systems, Fudan University. She has authored or coauthored referred papers in *IEEE Communications Magazine*, *ACM Transactions on the Web*, IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, IEEE TRANSACTIONS ON MOBILE COMPUTING, Springer WWW Journal, ACM CIKM, and ICPP. Her research interests include network security, user behavior analysis, and computational social systems.

Xin Wang received the B.S. degree in information theory and the M.S. degree in communication and electronic systems from Xidian University, Xi'an, China, in 1994 and 1997, respectively, and the Ph.D. degree in computer science from Shizuoka University, Shizuoka, Japan, in 2002. He is currently a Professor with Fudan University, Shanghai, China. His research interests include quality of network service, next-generation network architecture, mobile Internet, and network coding.

Yang Chen (Senior Member, IEEE) received the B.S. and Ph.D. degrees from the Department of Electronic Engineering, Tsinghua University, Beijing, China, in 2004 and 2009, respectively. During 2006–2008, he was a Visiting Student with Microsoft Research Asia, and visited Stanford University, Stanford, CA, USA, in 2007. From 2009 to 2011, he was a Research Associate and the Deputy Head of Computer Networks Group, Institute of Computer Science, University of Göttingen, Göttingen, Germany. From 2011 to 2014, he was a Postdoctoral Associate with the Department of Computer Science, Duke University, Durham, NC, USA, where he served as Senior Personnel in the NSF Mobility First project. He was a Nokia Visiting Professor with Aalto University, Espoo, Finland, in 2019. He is currently a Professor with the College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China, and leads the Big Data and Networking (DataNET) group, Fudan University. His research interests include online social networks, Internet architecture, and mobile computing. He serves as a Senior Associate Editor of *ACM Transactions on Social Computing*, an Associate Editor-in-Chief of *Journal of Social Computing*, and an Editorial Board Member of Elsevier *Computer Communications*. He served as a OC/SPC/PC Member for many international conferences, including SOSP, SIGCOMM, WWW, IMC, IJCAI, AAAI, ECAI, DASFAA, IWQoS, ICCCN, GLOBECOM, and ICC. He is a Senior Member of ACM, and a Fellow of the Royal Geographical Society.