

HiLoMix: Robust High- and Low-Frequency Graph Learning Framework for Mixing Address Association

Xiaofan Tu^{1,2}, Tiantian Duan³, Shuyi Miao^{1,2}, Hanwen Zhang^{3,4}, Yi Sun^{3,4,5*}

¹School of Artificial Intelligence, Beihang University

²Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing

³Institute of Computing Technology, Chinese Academy of Sciences

⁴University of Chinese Academy of Sciences

⁵Shandong Key Laboratory of Blockchain Finance

{xiaofantu, shuyimiao}@buaa.edu.cn, {duantiantian, hwzhang, sunyi}@ict.ac.cn

Abstract

As mixing services are increasingly being exploited by malicious actors for illicit transactions, mixing address association has emerged as a critical research task. A range of approaches have been explored, with graph-based models standing out for their ability to capture structural patterns in transaction networks. However, these approaches face two main challenges: label noise and label scarcity, leading to suboptimal performance and limited generalization. To address these, we propose HiLoMix, a graph-based learning framework specifically designed for mixing address association. First, we construct the Heterogeneous Attributed Mixing Interaction Graph (HAMIG) to enrich the topological structure. Second, we introduce frequency-aware graph contrastive learning that captures complementary structural signals from high- and low-frequency graph views. Third, we employ weak supervised learning that assigns confidence-based weights to noisy labels. Then, we jointly train high-pass and low-pass GNNs using both unsupervised contrastive signals and confidence-based supervision to learn robust node representations. Finally, we adopt a stacking framework to fuse predictions from multiple heterogeneous models, further improving generalization and robustness. Experimental results demonstrate that HiLoMix outperforms existing methods in mixing address association.

Introduction

Web3.0 represents a paradigm shift toward a decentralized and user-centric internet, with Ethereum serving as its core infrastructure. Identities on Ethereum are represented by account addresses that are decoupled from real-world identities; however, this pseudonymity remains vulnerable to existing deanonymization attacks (Miao et al. 2025; Shen et al. 2021; Zhou et al. 2022). To mitigate these risks, users often turn to mixing services, among which **Tornado Cash** is the most widely used Ethereum mixer. As illustrated in Figure 1, it enables users to deposit assets into a funding pool and withdraw them without revealing which specific deposit it corresponds to, thus obscuring transaction trails and enhancing anonymity. However, Tornado Cash has also raised significant security concerns. According to Elliptic reports,

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

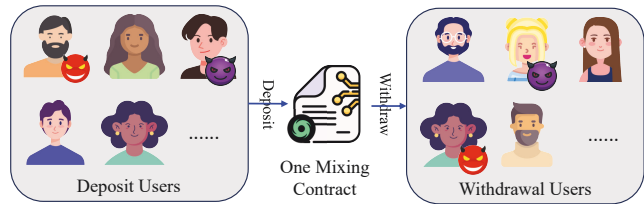


Figure 1: Mixing process of Tornado Cash. Characters marked with a devil icon represent illicit users. Two characters sharing the same icon indicate that they have transferred funds through mixer.

Tornado Cash facilitated over \$ 1.5 billion in illicit transfers prior to being sanctioned by OFAC in August 2022. With the recent lifting of this sanction, its potential misuse for money laundering may escalate. These developments highlight the urgent need to deanonymize Tornado Cash by uncovering associations among addresses involved in mixing.

Several approaches have been proposed to establish associations among mixing addresses, including heuristic and machine learning-based approaches. Heuristic approaches (Béres et al. 2021; Wang et al. 2023) rely on empirical patterns and domain-specific assumptions, but struggle to detect malicious actors who deliberately evade detection and suffer from poor scalability. In contrast, machine learning approaches automate the association of mixing addresses by training predictive models that learn implicit transaction patterns from blockchain data, offering greater generalization and scalability. However, such approaches face two fundamental challenges in practice:

- **Label noise:** As available labels are often derived from heuristic rules, they inevitably contain noise, introducing ambiguity during training and impairing the model’s generalization ability.
- **Label scarcity:** The limited availability of ground-truth associations prevents the model from establishing comprehensive decision boundaries, significantly degrading its performance and generalization ability.

(Hu et al. 2023, 2024) employed a pre-trained Transformer (Vaswani et al. 2017) to reduce reliance on labeled data. However, their approach overlooks the structural re-

relationships among accounts, which are crucial for mixing address association, as true associations are often embedded in complex multi-hop interaction structures, which sequence modeling alone cannot capture. In contrast, (Béres et al. 2021; Du et al. 2024) modeled account interactions as graphs, enabling the capture of topological structures and interaction dependencies between accounts, and attempted to alleviate the label scarcity problem through techniques such as synthetic labeling. However, these challenges remain largely unresolved and existing methods still suffer from the following **limitations**:

- 1) **Sparse graph structure.** The scarcity of associations results in highly sparse graph structures, leaving many nodes isolated and limiting the efficiency of message passing and feature aggregation.
- 2) **Limited supervision.** Limited associations hinder models from learning comprehensive decision boundaries, leading to poor generalization performance in large-scale real-world transaction environments.
- 3) **Overlooked label noise.** Heuristic labels are inherently noisy, which not only compromises training and evaluation, but also perturbs the graph structure, distorting connectivity patterns and neighborhood information, thereby impairing learning effectiveness.

To address the above limitations, we propose **HiLoMix**, a robust **H**igh-frequency and **L**ow-frequency graph learning framework tailored for mixing address association. First, to alleviate graph sparsity, we construct the *Heterogeneous Attributed Mixing Interaction Graph* (HAMIG), which encodes both account associations and transaction interactions as edges, enriching the graph topology and enhancing message propagation. Second, to mitigate limited supervision, we introduce a frequency-aware graph contrastive learning paradigm that captures complementary structural signals by contrasting node representations from high- and low-frequency views. Third, to handle label noise, we apply confidence-based label-weighted supervised learning, which assigns reliability-aware weights to noisy labels, enabling the model to leverage weak supervision without being dominated by noise. We jointly train a high-pass and low-pass GNN using both unsupervised contrastive signals and weighted supervision to learn robust node representations. Finally, we adopt the stacking framework that integrates the predictions of three heterogeneous base models, logistic regression, random forest, and multilayer perceptron, to further enhance generalization and robustness.

Our main **contributions** are summarized as follows:

- We propose HiLoMix, a robust high- and low-frequency graph learning framework tailored for the mixing address association task.
- We adopt heterogeneous graph modeling, frequency-aware graph contrastive learning, and confidence-based label-weighted supervised learning to address the above limitations of existing approaches.
- HiLoMix achieves state-of-the-art performance, outperforming the best baseline by 5.69%, 7.34% and 15.61%

in F_1 , AUC and MRR, respectively. Furthermore, we curate an up-to-date ground-truth dataset to facilitate future research on mixing address association.

Background and Related Work

Ethereum and Tornado Cash

Unlike Bitcoin, which supports only limited scripting capabilities for deterministic verification, Ethereum enables the deployment and execution of smart contracts through its Turing-complete Ethereum Virtual Machine (EVM). These smart contracts are immutable self-executing code deployed on-chain, capable of autonomously managing digital assets and enforcing application logic without centralized control. This decentralized computational infrastructure underpins a wide range of applications, including decentralized finance (DeFi), digital identity systems, decentralized autonomous organizations, and non-fungible tokens (NFTs), thereby fostering a trustless and transparent digital ecosystem.

Tornado Cash is a decentralized, non-custodial protocol designed to enhance transaction privacy on Ethereum by operating as a smart contract-based coin mixer. It employs Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge (zk-SNARKs) to conceal the on-chain linkage between source and destination addresses, thereby significantly complicating efforts to trace the flow of funds. By cryptographically severing the deterministic connection between sender and receiver, Tornado Cash preserves transaction anonymity while maintaining on-chain verifiability.

Summary of Existing Studies

Heuristic approaches are rule-based approaches that infer mixing address associations from empirical patterns and domain-specific assumptions, rather than formal models or labeled data. (Hong et al. 2018) first explored such associations by linking addresses in the Helix Bitcoin mixer, based on the observation that user deposits were approximately equal to their withdrawals. (Béres et al. 2021; Wang et al. 2023) extended these ideas by leveraging more transaction-level features, such as timing, distinctive gas prices, multi-denomination behavior, and TORN mining, to improve address linkage. Despite extensive exploration, heuristic approaches face inherent limitations. First, they rely on low-level transaction features (e.g. time, gas price) and require exhaustive pattern matching across all transactions, which limits scalability (Du et al. 2024). Second, rule-based approaches have limited detection capability, as only a small fraction of users match the predefined patterns. In many cases, they fail to identify malicious actors who actively evade detection and instead capture only inadvertent or care-less users.

Machine learning-based approaches automate the association of mixing addresses by training predictive models that learn implicit transaction patterns from blockchain data. (Hu et al. 2023) employed a pre-trained Transformer (Vaswani et al. 2017) to encode transaction sequences of Ethereum accounts, capturing temporal and contextual behavior patterns. Subsequently, (Hu et al. 2024) achieved both

parameter and computational efficiency using frequency-aware compression and asymmetric training compared (Hu et al. 2023). However, their approaches overlook the structural relationships among accounts, which are crucial for mixing address association. (Du et al. 2024) modeled the address interactions as a Mixing Interaction Graph (MIG) and formulated the task as a node-pair link prediction problem. By applying GraphSAGE (Hamilton, Ying, and Leskovec 2017), their approach captures the topological structures and interaction dependencies among accounts. Furthermore, (Béres et al. 2021) applied classic graph representation learning for address associations. (Zhong and Mueen 2024) developed a deep neural network to link disjoint Bitcoin address clusters in a self-supervised manner. Although these approaches are not specifically designed for mixing services, they can be readily adapted to our task due to the similarity in task formulation. However, machine learning models heavily rely on labeled data, which is unavailable for the mixing address association task. As a workaround, heuristic rules are employed to generate address associations but inevitably introduce noise, compromising both training and evaluation. To the best of our knowledge, the impact of such label noise has not been formally investigated in this context.

Methodology

Overview

As depicted in Figure 2, HiLoMix comprises four key components: (1) **HAMIG construction**, which encodes diverse account interactions into a heterogeneous graph to enrich the topological structure; (2) **Frequency-aware graph contrast learning**, which contrasts node representations from high- and low-frequency views to capture robust structural patterns; (3) **Confidence-based supervision**, which assigns adaptive weights to noisy labels based on estimated reliability to mitigate supervision noise; (4) **Mixing address association**, which employs a stacking framework to integrate the knowledge of multiple models for final prediction.

HAMIG Construction

(Du et al. 2024) introduced the Mixing Interaction Graph (MIG) to capture the topological structure of account associations, where each account is modeled as a node and the address associations are represented as edges. Although effective in describing the basic interaction topology, MIG suffers from structural sparsity due to limited address associations, which restricts effective message propagation in graph-based learning. To address this limitation, we extend MIG by incorporating mixing transaction interactions as additional edges, thereby constructing a heterogeneous account interaction graph that significantly enriches the structural connectivity and semantic information of MIG. In addition, we extract a node feature matrix $\mathcal{X} \in \mathbb{R}^{n \times d}$, where n denotes the number of account nodes and d the feature dimension. Each feature vector encodes an account’s interaction behaviors with Tornado Cash mixing contracts, including its interaction frequency, timestamps, and gas price statistics. We define the resulting enhanced structure as the Heterogeneous Attributed Mixing Interaction Graph (HAMIG),

which jointly captures both the topological dependencies among accounts and the attribute characteristics of individual nodes. Unlike MIG, HAMIG is formulated as an undirected heterogeneous graph $\mathcal{G} = (\mathcal{V}_a, \mathcal{V}_t, \mathcal{E}_{at}, \mathcal{E}_{aa}, \mathcal{X})$, where \mathcal{V}_a denotes the set of account addresses interacting with Tornado Cash, \mathcal{V}_t denotes the set of smart contracts within Tornado Cash, $\mathcal{E}_{at} = \{(v_i, v_j) \mid v_i \in \mathcal{V}_a, v_j \in \mathcal{V}_t\}$ represents mixing transaction edges, $\mathcal{E}_{aa} = \{(v_i, v_j) \mid v_i, v_j \in \mathcal{V}_a\}$ represents address association edges, and \mathcal{X} is the node feature matrix.

Frequency-aware Graph Contrast Learning

To mitigate insufficient supervision caused by label scarcity, we employ contrastive learning, which enables model training without ground-truth labels by leveraging instance-level similarity. Frequency-aware graph contrastive learning (GCL) offers a principled augmentation strategy by semantically meaningful and structurally coherent views based on the spectral characteristics of graph signals (Wan et al. 2024). However, directly applying biased-pass filters to the original graph signals may lead to entangled views with overlapping frequency components. To address this, inspired by (Wu et al. 2024), we first decompose the graph signal into high- and low-frequency components before filtering. This separation enables each GNN branch to focus on a distinct frequency view, facilitating representation learning while minimizing cross-view interference. Specifically, we employ a multilayer perceptron (MLP) to estimate the edge smoothness probability s_{ij} between node v_i and v_j , where both node features are taken into account. Given the node feature matrix $\mathcal{H}_0 = \text{MLP}(\mathcal{X})$, the edge-wise smoothness is computed as follows:

$$s_{ij} = \text{Sigmoid}(\text{Linear}(h_i \parallel h_j) + \text{Linear}(h_j \parallel h_i)) \quad (1)$$

where h_i indicates the feature vector of the i -th node, \parallel denotes vector concatenation, and the sigmoid function ensures that s_{ij} is bounded within $[0, 1]$. Based on estimated edge smoothness scores, HAMIG \mathcal{G} is decomposed into two frequency-specific graph views: a low-frequency view $\mathcal{G}^{(LF)}$ and a high-frequency view $\mathcal{G}^{(HF)}$. Both views share the same node set, while the adjacency matrix A is partitioned into $A^{(LF)}$ and $A^{(HF)}$ as follows:

$$\begin{aligned} A_{ij}^{(LF)} &= s_{ij}, & A_{ij}^{(HF)} &= 1 - s_{ij} & \text{where } A_{ij} &= 1 \\ A_{ij}^{(LF)} &= A_{ij}^{(HF)} &= 0 & & \text{where } A_{ij} &= 0 \end{aligned} \quad (2)$$

To capture the complementary structural semantics embedded in different frequency bands of the graph signal, we design a dual-branch frequency-aware architecture to process low- and high-frequency components separately. This design is based on the principle that graph frequency characterizes the rate of variation in the node features across the graph topology: low-frequency components capture smooth, global structures, while high-frequency components highlight sharp local changes and structural anomalies.

For the low-frequency view $\mathcal{G}^{(LF)}$, we employ a low-pass graph neural network to propagate and smooth features across structurally similar neighborhoods, thereby capturing

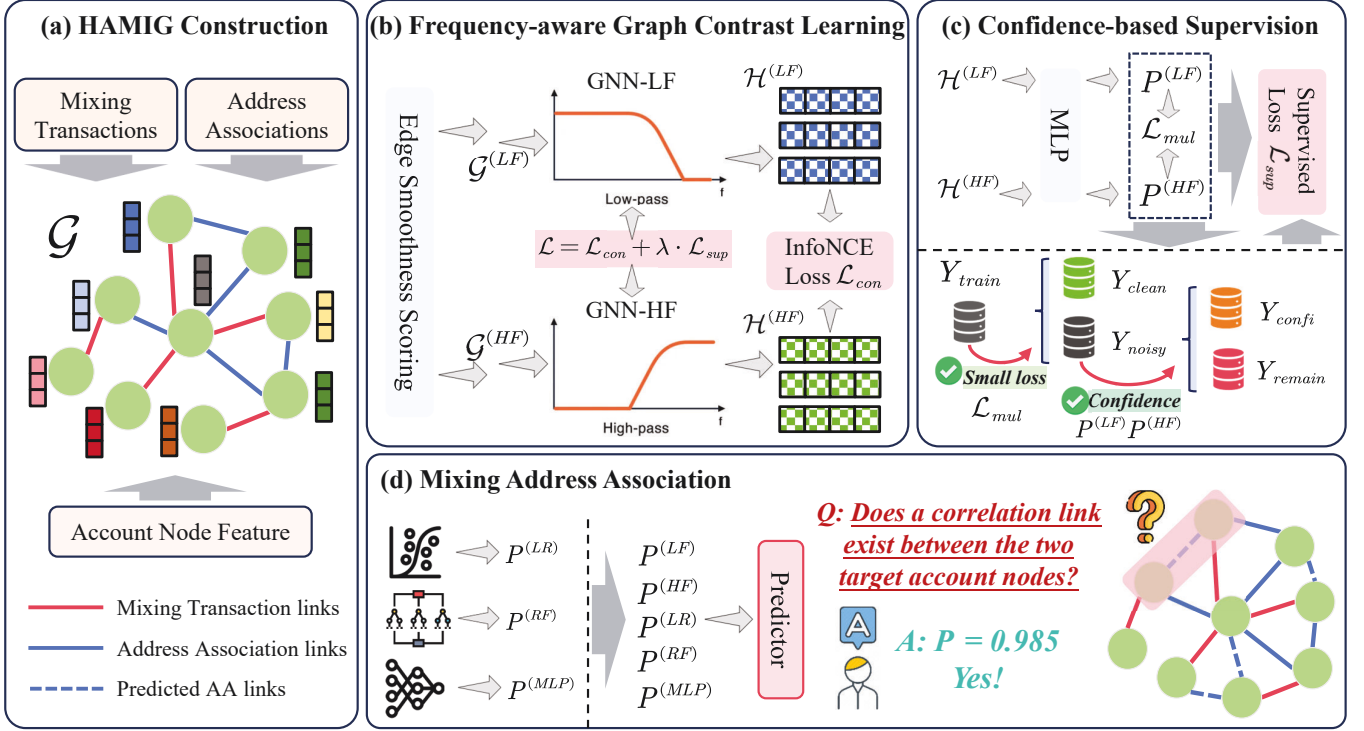


Figure 2: The overview framework of our HiLoMix.

consistent global semantics.

$$\mathcal{H}_l^{(LF)} = \left(I + \tilde{A}^{(LF)} \right) \mathcal{H}_{l-1}^{(LF)}, \quad l \in \{1, \dots, L\} \quad (3)$$

In contrast, for the high-frequency view $\mathcal{G}^{(HF)}$, we utilize a high-pass GNN to preserve discriminative signals and amplify local heterogeneity based on the normalized graph Laplacian matrix.

$$\mathcal{H}_l^{(HF)} = \left(I - \alpha \cdot \tilde{A}^{(HF)} \right) \mathcal{H}_{l-1}^{(HF)}, \quad l \in \{1, \dots, L\} \quad (4)$$

Here, $\tilde{A}^{(LF)}$ and $\tilde{A}^{(HF)}$ denote the normalized adjacency matrices for the low- and high-frequency graph views, respectively; α controls the high-pass filter intensity; and $\mathcal{H}_L^{(LF)}$ and $\mathcal{H}_L^{(HF)}$ (i.e., $\mathcal{H}^{(LF)}$ and $\mathcal{H}^{(HF)}$) are the final node embeddings from the two views.

By applying low-pass and high-pass GNNs to the low- and high-frequency graph views, respectively, we obtain node representations that emphasize distinct frequency semantics. These representations, $\mathcal{H}^{(LF)}$ and $\mathcal{H}^{(HF)}$, are then contrasted using InfoNCE loss (van den Oord, Li, and Vinyals 2019) to facilitate node representation learning.

$$\mathcal{L}^{(LF)} \left(h_i^{(LF)}, h_i^{(HF)} \right) = -\log \frac{s \left(h_i^{(LF)}, h_i^{(HF)} \right)}{\sum_{j=1}^N s \left(h_i^{(LF)}, h_j^{(HF)} \right)} \quad (5)$$

$$\mathcal{L}^{(HF)} \left(h_i^{(HF)}, h_i^{(LF)} \right) = -\log \frac{s \left(h_i^{(HF)}, h_i^{(LF)} \right)}{\sum_{j=1}^N s \left(h_i^{(HF)}, h_j^{(LF)} \right)} \quad (6)$$

where $s \left(h_i^{(LF)}, h_i^{(HF)} \right) = \exp \left(\omega \left(h_i^{(LF)}, h_i^{(HF)} \right) / \tau \right)$, with ω denoting cosine similarity and τ representing the temperature coefficient. This formulation encourages the model to produce semantically aligned embeddings across the two frequency-specific views. The overall contrastive objective is then defined as:

$$\mathcal{L}_{con} = \frac{1}{2|\mathcal{E}_{aa}|} \sum_{e \in \mathcal{E}_{aa}} \left[\mathcal{L}^{(LF)} + \mathcal{L}^{(HF)} \right] \quad (7)$$

Confidence-based Supervision

Although unsupervised paradigms such as contrastive learning help reduce dependence on labels and alleviate the impact of noise, they often fall short in performance compared to supervised methods. To address this limitation, we adopt a hybrid training strategy that integrates supervised learning signals to guide model optimization more effectively. Specifically, we compute the predicted probabilities of the association edge e_{ij} from the two frequency-specific networks using the formulation $p_{ij} = \text{Linear} \left(h_i \parallel h_j \right)$. Based on these predictions, the mutual loss (Qian et al. 2023) for edge e_{ij} is defined as:

$$\begin{aligned} \mathcal{L}_{mul}^{ij} &= -y_{ij} \left[\log \left(p_{ij}^{(LF)} \right) + \log \left(p_{ij}^{(HF)} \right) \right] \\ &= -y_{ij} \log \left(p_{ij}^{(LF)} \cdot p_{ij}^{(HF)} \right) \end{aligned} \quad (8)$$

Mutual loss gauges predictive confidence: \mathcal{L}_{mul}^{ij} has a low value when both networks confidently and correctly predict the presence of edge e_{ij} .

Clean Label Set Extraction Deep neural networks exhibit a memorization effect, where clean and simple patterns are learned in the early stages of training, while noisy labels are gradually overfitted in later epochs (Arpit et al. 2017). To adaptively manage label noise during training, we adopt the small-loss criterion (Han et al. 2018), which selects clean samples based on training losses. In epoch t , we define a threshold to distinguish clean from noisy labels as follows:

$$\mathcal{L}_{thre}^t = \text{Percentile} \left(\mathcal{L}_{mul}^{ij}, 1 - \frac{t}{2 \cdot T_{\max}} \right) \quad (9)$$

where T_{\max} denotes the total number of training epochs, $0.5 \leq \left(1 - \frac{t}{2 \cdot T_{\max}}\right) < 1$ ensures that at least half of the labels are retained as clean, $\text{Percentile}(\mathcal{L}, p)$ returns the value below which $100 * p\%$ of the loss values in \mathcal{L} fall. This dynamic threshold allows more samples to be included during the early stages for sufficient training and gradually reduces the number of selected instances as training progresses.

We define $\mathcal{L}_{thre}^{avg} = \text{Average}(\mathcal{L}_{mul}^{ij})$ to further ensure that small loss samples are clean regardless of their relative rank. Then clean and noisy label sets can be divided as follows:

$$\mathcal{E}_{cl} = \left\{ e_{ij} \in \mathcal{E}_{aa} \mid \mathcal{L}_{mul}^{ij} < \max(\mathcal{L}_{thre}^t, \mathcal{L}_{thre}^{avg}) \right\} \quad (10)$$

Confident Label Set Extraction Given that \mathcal{E}_{cl} denotes the clean label set, we define the remaining potentially noisy labels as $\mathcal{E}_{ns} = \mathcal{E}_{aa} \setminus \mathcal{E}_{cl}$. Inspired by (Qian et al. 2023), we further extract a subset $\mathcal{E}_{cf} \subseteq \mathcal{E}_{ns}$, where two predictions, $p_{ij}^{(LF)}$ and $p_{ij}^{(HF)}$, confidently agree on the same prediction that differs from the observed label y_{ij} . Formally, this subset is characterized as follows:

$$c = \arg \max p_{ij}^{(LF)} = \arg \max p_{ij}^{(HF)} \neq y_{ij} \quad (11)$$

$$\mathcal{E}_{cf} = \left\{ e_{ij} \in \mathcal{E}_{ns} \mid \mu(e_{ij}) > Thre = 1 - \frac{t}{2 \cdot T_{\max}} \right\} \quad (12)$$

where the confidence score $\mu(e_{ij})$ is defined as $\mu(e_{ij}) = \sqrt{p_{ij,c}^{(LF)} \cdot p_{ij,c}^{(HF)}}$. A higher value of $\mu(e_{ij})$ suggests a greater likelihood that e_{ij} is mislabeled, yet correctly predicted. The threshold $Thre$ is designed to be dynamic: stricter in the early stages of training and gradually relaxed as the training progresses.

Given the previously defined sets \mathcal{E}_{cl} and \mathcal{E}_{cf} , we denote the remaining labels as \mathcal{E}_{re} , from which no useful information can be reliably extracted.

$$\mathcal{E}_{re} = \mathcal{E}_{aa} \setminus \mathcal{E}_{cl} \setminus \mathcal{E}_{cf} \quad (13)$$

After partitioning the labels into \mathcal{E}_{cl} , \mathcal{E}_{cf} , and \mathcal{E}_{re} , the confidence-based supervision loss for mixing address association is defined as follows:

$$\mathcal{L}_{sup} = -\frac{1}{|\mathcal{E}_{aa}|} \sum_{e_{ij} \in \mathcal{E}_{aa}} \xi(e_{ij}) (\log p_{ij, \hat{y}_{ij}}^{(LF)} + \log p_{ij, \hat{y}_{ij}}^{(HF)}) \quad (14)$$

where

$$\begin{cases} \xi(e_{ij}) = 1, & \hat{y}_{ij} = y_{ij}, & \text{if } e_{ij} \in \mathcal{E}_{cl} \\ \xi(e_{ij}) = \mu(i), & \hat{y}_{ij} = c, & \text{if } e_{ij} \in \mathcal{E}_{cf} \\ \xi(e_{ij}) = 0.5, & \hat{y}_{ij} = y_{ij}, & \text{if } e_{ij} \in \mathcal{E}_{re} \end{cases} \quad (15)$$

where $\xi(\cdot)$ serves as a confidence score, assigning each label a different weight to guide the model towards more reliable supervision. Finally, the total loss used to update both the high-pass and low-pass GNNs is formulated as follows:

$$\mathcal{L} = \mathcal{L}_{con} + \lambda \cdot \mathcal{L}_{sup} \quad (16)$$

Mixing Address Association

After training the high-pass and low-pass graph neural networks, instead of simply selecting the better one or applying confidence-based voting, we adopt the stacking method (Wolpert 1992) to integrate their knowledge for the final prediction. Stacking aims to exploit the strengths of multiple base models by training a meta-learner on their predictions. In general, heterogeneous models, differing in learning algorithms, hypothesis spaces, or inductive biases, promote diversity and enhance ensemble performance. However, in our case, high-pass and low-pass GNNs are inherently homogeneous, limiting their ability to provide complementary information and even resulting in performance degeneration, as shown in Table 2. To overcome this limitation, we introduce three additional heterogeneous models, random forest, logistic regression, and multilayer perceptron, into the stacking ensemble, enriching it with diverse decision boundaries and representational capabilities.

Experiments

Experiment Setup

Dataset: We collect 554,589 raw Tornado Cash mixing transactions spanning from 2019-12-16 to 2025-01-17. After data cleaning and preprocessing, 371,181 valid transactions are retained. Following the heuristic rules proposed in (Wu et al. 2022; Wang et al. 2023), we identify 4,074 unique address associations as labels. By modeling these associations and transaction interactions as edges, we construct HAMIG comprising 106,982 account nodes and 375,255 edges of two types.

Although our dataset is constructed solely from Tornado Cash, we note that this does not compromise generality. Tornado Cash remains the largest and most representative Ethereum-based mixer, accounting for over 95% of the total transaction volume and total value locked (TVL) across all Ethereum mixing services according to (Akopova 2024). Therefore, focusing on Tornado Cash already captures the predominant behavioral and structural patterns of Ethereum-based mixing activities.

Evaluation metrics: To comprehensively evaluate the performance of HiLoMix, we adopt four metrics across two categories: classification and ranking. Specifically, F1-score and AUC are used to evaluate classification performance, while MRR and Hits@K assess ranking effectiveness. The formal definitions of MRR and Hits@K are provided:

$$\text{MRR} = \frac{1}{N} \sum_{i=1}^N \frac{1}{\text{rank}_i} \quad (17)$$

$$\text{Hits@K} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\text{rank}_i \leq K) \quad (18)$$

Methods	F ₁	AUC	MRR	Hits@3	Hits@5	Hits@10
DeepWalk	0.6844±0.0030	0.7078±0.0067	0.0866±0.0056	0.0568±0.0081	0.0932±0.0043	0.1940±0.0057
Node2Vec	0.6970±0.0040	0.7251±0.0033	0.0850±0.0027	0.0553±0.0046	0.0926±0.0079	0.1870±0.0019
GCN	0.7639±0.0035	0.7795±0.0063	0.1011±0.0090	0.0542±0.0114	0.1233±0.0239	0.3151±0.0369
GAT	0.7674±0.0064	0.7828±0.0090	0.1025±0.0047	0.0569±0.0054	0.1217±0.0106	0.3179±0.0093
GIN	0.7869±0.0036	0.8167±0.0068	0.1415±0.0096	0.1143±0.0135	0.1900±0.0210	0.3802±0.0327
GraphSAGE	0.7867±0.0028	0.8138±0.0022	0.1187±0.0127	0.0794±0.0178	0.1535±0.0235	0.3568±0.0211
Co-teaching	0.7859±0.0057	0.8011±0.0028	0.0961±0.0106	0.0460±0.0160	0.0927±0.0255	0.2578±0.0557
Co-teaching+	0.7690±0.0014	0.8230±0.0040	0.2712±0.0025	0.2994±0.0007	0.4281±0.0036	0.6169±0.0026
NRGL	0.7866±0.0058	<u>0.8513±0.0029</u>	<u>0.3989±0.0058</u>	<u>0.4784±0.0079</u>	<u>0.5896±0.0078</u>	<u>0.7263±0.0038</u>
BERT4ETH	0.6990±0.0083	0.7897±0.0215	0.3366±0.0656	0.3878±0.0872	0.4694±0.0991	0.5714±0.0949
MixBroker	<u>0.7930±0.0024</u>	0.8182±0.0049	0.1062±0.0089	0.0585±0.0117	0.1231±0.0170	0.3323±0.0292
HiLoMix	0.8382±0.0044	0.9137±0.0033	0.4612±0.0087	0.5173±0.0086	0.6556±0.0114	0.8037±0.0111
%Improv.	5.69%	7.34%	15.61%	8.13%	11.18%	10.65%

Table 1: Overall Performance Comparison. **Bold** and underline represent the best and second best performance, respectively. We repeat the experiment with 3 random seeds and report the average metrics with standard deviation.

where N is the number of positive samples, rank_i denotes the ranking position of the i -th positive sample in the candidate list based on predicted probabilities, and $\mathbb{I}(\cdot)$ is the indicator function.

Baselines: We compare HiLoMix with 11 baseline models, grouped into three categories: (1) Graph representation learning methods: DeepWalk (Perozzi, Al-Rfou, and Skiena 2014), Node2Vec (Grover and Leskovec 2016), GCN (Kipf and Welling 2016), GAT (Velickovic et al. 2017), GIN (Xu et al. 2018), and GraphSAGE (Hamilton, Ying, and Leskovec 2017). Since the mixing address association is formulated as a node-pair link prediction task in HAMIG, classical graph representation learning models serve as strong baselines. (2) Robust learning methods: Co-teaching (Han et al. 2018), Co-teaching+ (Yu et al. 2019) and NRGL (Wu et al. 2024). Given the central limitation of label noise in our task, we include both classical and state-of-the-art robust learning approaches for comparison. (3) Mixing address association methods: BERT4ETH (Hu et al. 2023) and MixBroker (Du et al. 2024). These methods are designed specifically for the mixing address association task.

Implementation Details: All implementations are conducted using PyTorch 2.3.1 and DGL 2.4.0. The model is trained for 50 epochs with a batch size of 128, using the Adam optimizer with a learning rate of 0.003. All experiments are performed on a single NVIDIA A100 GPU. For contrastive loss, we use a batch of 128 nodes and set the temperature coefficient at 0.5. For MRR and Hits@K metrics, 50 negative samples are generated for each positive instance. All baselines are implemented following their original papers and fine-tuned to achieve optimal performance.

Performance Comparison

We conduct a systematic comparison between HiLoMix and the baseline methods. The results of the mixing address association performance are presented in Table 1. The im-

provement (%Improv.) is calculated as the relative gain over the second-best model.

Mixing address association under the classification setting. HiLoMix achieves an AUC of 0.9137 and a F₁ score of 0.8382, outperforming all baselines with relative improvements ranging from 7.34% to 29.09% in AUC and 5.69% to 22.47% in F₁, demonstrating its effectiveness in the classification setting. Classic graph representation learning methods perform the worst, primarily due to their vulnerability to label noise and graph sparsity resulting from scarce associations. NRGL achieves the second highest AUC, validating the effectiveness of explicit label noise handling. MixBroker attains the second-best F₁ score, benefiting from task-specific graph modeling that captures user interaction structures. In contrast, BERT4ETH underperforms standard GNNs, highlighting the importance of incorporating structural information in address association.

Mixing address association under the ranking setting. HiLoMix outperforms the second-best model by 15.61%, 8.13%, 11.18% and 10.65% in MRR, Hits@3, Hits@5 and Hits@10, respectively. These results indicate that HiLoMix effectively ranks true address associations near the top among all candidates, demonstrating a strong generalization to unseen links. In line with the classification results, the classic graph representation learning methods perform the worst. NRGL ranks second across all ranking metrics, further validating the effectiveness of robust learning strategies for mixing address association task. BERT4ETH leverages sequence modeling to generate expressive account representations; however, its disregard for structural dependencies leads to suboptimal performance. MixBroker, while incorporating user interaction structures through graph modeling, suffers from graph sparsity and structural perturbations.

Ablation Study

To systematically validate the effectiveness of each component in HiLoMix, we conduct ablation studies by remov-

Methods	F ₁	AUC	MRR
w/o hetero models	0.7393	0.7948	0.2977
w/o HiLo GNNs	0.8075	0.8761	0.3507
w/o Hi GNN	0.8169	0.8977	0.4202
w/o Lo GNN	0.8286	0.9046	0.4129
w/o hetero graph	0.8158	0.8931	0.4419
w/o HiLo learning	0.8083	0.8911	0.4228
w/o label division	0.8238	0.8938	0.4260
w/o stacking	0.7882	0.8547	0.4061
HiLoMix	0.8382	0.9137	0.4612

Table 2: Ablation Study Results.

ing or modifying specific modules. Specifically, “w/o hetero models” excludes all additional heterogeneous base models from the stacking ensemble. “w/o HiLo GNNs” removes both high-pass and low-pass GNNs from the ensemble. “w/o Hi GNN” and “w/o Lo GNN” remove high-pass and low-pass GNN, respectively. “w/o hetero graph” removes all mixing transaction edges, retaining only address association edges in HAMIG. “w/o HiLo learning” replaces the high- and low-pass GNNs with two GCNs. “w/o label division” disables the confidence-based label division mechanism. “w/o stacking” discards the stacking framework and instead selects the better result between the high- and low-pass GNNs.

As shown in Table 2, removing any component results in a consistent drop in both classification and ranking metrics, confirming the complementary contributions of all modules. Among all variants, excluding heterogeneous base models leads to the greatest degradation, with F₁ decreasing by 11.80% and MRR by 35.45%. Although the drop is substantial, this result is expected: in the *w/o hetero models* setting, the high-pass and low-pass GNNs are inherently homogeneous, limiting their ability to provide complementary information. In contrast, noticeable drops occur when the core GNNs are removed from the stacking ensemble, confirming their pivotal role in HiLoMix’s effectiveness. The remaining variants yield moderate performance decreases, further validating the necessity of each component. These smaller drops also demonstrate the robustness of HiLoMix and indicate that the stacking framework effectively integrates noise-resistant and complementary knowledge from heterogeneous base models, maintaining stable performance even when certain components are weakened.

Feature Importance Evaluation

We conduct a comparative study across seven feature variants by progressively removing less important dimensions. Additionally, we incorporate node features from MixBroker for further comparison. As shown in Table 3, HiLoMix achieves the best performance across all metrics when using our 197-dimensional account node features, outperforming the second-best features, MixBroker features. Although the 221-dimensional feature set contains more attributes, its performance is inferior to that of the 197-dimensional set due to reduced interpretability. These results suggest that our 197-

Methods	F ₁	AUC	MRR
28-Dim Feature	0.8092	0.8780	0.3834
53-Dim Feature	0.7984	0.8686	0.3691
105-Dim Feature	0.8131	0.8761	0.3926
141-Dim Feature	0.8067	0.8861	0.4261
167-Dim Feature	0.8208	0.8957	0.4249
197-Dim Feature	0.8382	0.9137	0.4612
221-Dim Feature	0.8183	0.8948	0.4408
MixBroker Feature	0.8248	0.9025	0.4423

Table 3: Performance Comparison of Different Features.

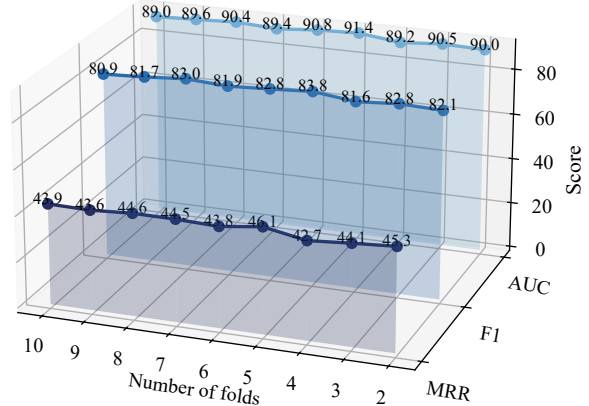


Figure 3: Performance of HiLoMix when training the base models in the stacking framework using different K of K-fold cross-validation.

dimensional features not only enhance model explainability but also reduce HAMIG’s memory footprint without compromising performance.

Hyper-parameter Sensitivity Analysis

We perform a sensitivity analysis on key hyper-parameters in HiLoMix. We begin by evaluating model stability with respect to the number of folds used in cross-validation within the stacking ensemble. As shown in Figure 3, HiLoMix performance remains generally stable across different fold settings. The best overall performance is observed with 5-fold cross-validation, which is therefore adopted as the default setting. Both fewer and greater numbers of folds lead to slight performance degradation, potentially due to overfitting or unstable ensemble integration.

We further investigate two additional hyper-parameters: loss weighting λ , which balances the supervised and contrastive learning objectives, and the high-pass filter strength α , which modulates the emphasis on high-frequency components. As illustrated in Figure 4, HiLoMix demonstrates robust performance across a wide range of λ and α values. Both metrics peak around $\lambda = 2$ and $\alpha = 0.5$, suggesting that a moderate balance between the supervised and contrastive learning objectives and the well-calibrated filter strength are critical for optimal performance.

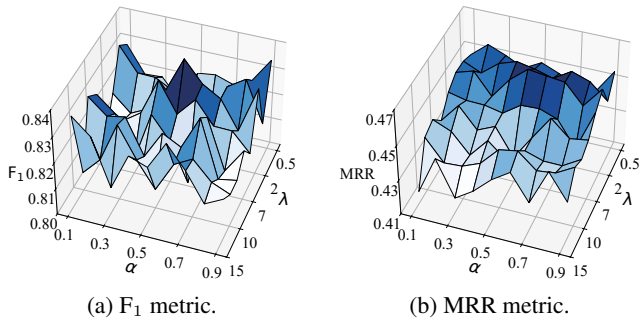


Figure 4: Performance of HiLoMix under different settings of hyperparameters α and λ .

Method	Training time (s)
HiLoMix	6.91
GIN	0.26
NRGL	15.27
MixBroker	0.26

Table 4: Comparison of training efficiency.

Efficiency Analysis

Although HiLoMix introduces additional components compared to simple baselines, most of them are computationally lightweight and do not incur substantial overhead. We compare HiLoMix with representative baselines, i.e., the best-performing model from each baseline category. The average training time per epoch is reported in Table 4.

As shown in Table 4, HiLoMix incurs moderate additional computation compared to simple GNNs but remains substantially faster than NRGL. In contrast, DeepWalk/Node2Vec and BERT4ETH require several hours for random walk generation or pretraining to obtain node embeddings, making them far less efficient than our model. All models in our experiments, including HiLoMix and the baselines, complete inference on the entire test set in under 0.01 seconds. In general, HiLoMix achieves a favorable balance between performance and computational cost.

Structural and Label Dynamics Analysis

We quantitatively compare MIG and HAMIG in terms of key structural properties. As shown in Table 5, HAMIG effectively alleviates the structural sparsity of MIG, increasing the average node degree by approximately $45\times$, and also improving the clustering coefficient and graph density by several orders of magnitude. This enriched topology facilitates more efficient message propagation among nodes.

We track the number of *clean*, *confident-flipped*, and *remaining* labels across epochs, as illustrated in Figure 5. The evolution process can be roughly divided into two stages: early training (0–30 epochs) and later training (30–50 epochs). During the early stage, the number of clean labels gradually decreases, while remaining labels increase, as the model transitions from fitting clean samples to handling potentially noisy ones. At this point, confident-flipped labels

Metric	MIG	HAMIG
Average node degree	0.0565	2.561
Average clustering coefficient	0.0004	0.0390
Graph density	$4.93e-7$	$2.23e-5$

Table 5: Comparison of graph structural properties between MIG and HAMIG.

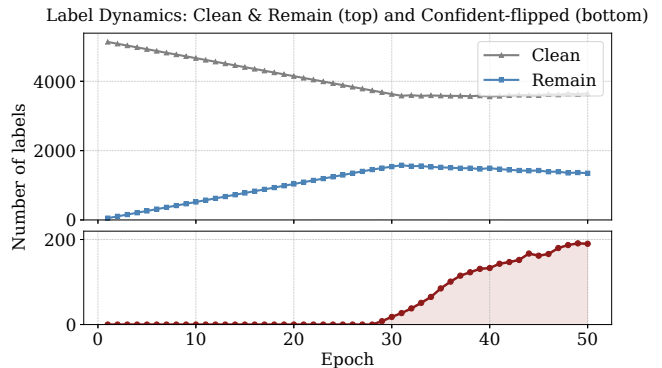


Figure 5: Dynamics of three label sets during training.

remain nearly zero, indicating that the model has not yet learned sufficiently reliable structural patterns. In the later stage, the number of confident-flipped labels steadily increases, showing that the model begins to capture meaningful structural dependencies and distinguish true associations from incorrect ones. Meanwhile, clean labels stabilize, and the model progressively filters mislabeled samples from the remaining set, demonstrating improved robustness to noisy supervision. This dynamic evolution process confirms that confidence-based supervision effectively governs label noise and guides the model toward reliable learning.

Conclusion

In this work, we propose HiLoMix, a novel graph-based learning framework tailored for mixing address association. By modeling diverse account interactions in Tornado Cash as HAMIG, we effectively address the issue of graph sparsity. To alleviate the limited supervision and label noise problem, we introduce frequency-aware graph contrastive learning alongside confidence-based label-weighted supervised learning, jointly optimizing high-pass and low-pass GNNs for robust node representations. Extensive experiments demonstrate that HiLoMix consistently outperforms existing approaches on the mixing address association task. Furthermore, we curate an up-to-date ground-truth dataset to support and facilitate future research in this field.

Acknowledgements

This work was supported by Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing and the National Natural Science Foundation of China (U22B2032).

References

- Akopova, A. 2024. Detecting Ethereum Mixers. <https://cse.aua.am/files/2024/05/DetectingEthereumMixers.pdf>. Accessed: 2025-10-12.
- Arpit, D.; Jastrzundefinedbski, S.; Ballas, N.; Krueger, D.; Bengio, E.; Kanwal, M. S.; Maharaj, T.; Fischer, A.; Courville, A.; Bengio, Y.; and Lacoste-Julien, S. 2017. A closer look at memorization in deep networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, 233–242.
- Béres, F.; Seres, I. A.; Benczúr, A. A.; and Quinyne-Collins, M. 2021. Blockchain is Watching You: Profiling and De-anonymizing Ethereum Users. In *2021 IEEE International Conference on Decentralized Applications and Infrastructures (DAPPS)*, 69–78.
- Du, H.; Che, Z.; Shen, M.; Zhu, L.; and Hu, J. 2024. Breaking the Anonymity of Ethereum Mixing Services Using Graph Feature Learning. *Trans. Info. For. Sec.*, 19: 616–631.
- Grover, A.; and Leskovec, J. 2016. node2vec: Scalable Feature Learning for Networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 855–864.
- Hamilton, W. L.; Ying, R.; and Leskovec, J. 2017. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 1025–1035.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I. W.; and Sugiyama, M. 2018. Co-teaching: robust training of deep neural networks with extremely noisy labels. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 8536–8546.
- Hong, Y.; Kwon, H.; Lee, J.; and Hur, J. 2018. A Practical De-mixing Algorithm for Bitcoin Mixing Services. In *Proceedings of the 2nd ACM Workshop on Blockchains, Cryptocurrencies, and Contracts*, 15–20.
- Hu, S.; Huang, T.; Chow, K.-H.; Wei, W.; Wu, Y.; and Liu, L. 2024. ZipZap: Efficient Training of Language Models for Large-Scale Fraud Detection on Blockchain. In *Proceedings of the ACM Web Conference 2024*, 2807–2816.
- Hu, S.; Zhang, Z.; Luo, B.; Lu, S.; He, B.; and Liu, L. 2023. BERT4ETH: A Pre-trained Transformer for Ethereum Fraud Detection. In *Proceedings of the ACM Web Conference 2023*, 2189–2197.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Miao, S.; Qiu, W.; Zheng, H.; Zhang, Q.; Tu, X.; Liu, X.; Liu, Y.; Dong, J.; and Zheng, Z. 2025. Know Your Account: Double Graph Inference-Based Account De-Anonymization on Ethereum. In *2025 IEEE 41st International Conference on Data Engineering (ICDE)*, 1305–1319.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. DeepWalk: online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 701–710.
- Qian, S.; Ying, H.; Hu, R.; Zhou, J.; Chen, J.; Chen, D. Z.; and Wu, J. 2023. Robust Training of Graph Neural Networks via Noise Governance. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, 607–615.
- Shen, M.; Zhang, J.; Zhu, L.; Xu, K.; and Du, X. 2021. Accurate Decentralized Application Identification via Encrypted Traffic Analysis Using Graph Neural Networks. *IEEE Transactions on Information Forensics and Security*, 16: 2367–2380.
- van den Oord, A.; Li, Y.; and Vinyals, O. 2019. Representation Learning with Contrastive Predictive Coding. *arXiv:1807.03748*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 6000–6010.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; Bengio, Y.; et al. 2017. Graph attention networks. *stat*, 1050.
- Wan, G.; Tian, Y.; Huang, W.; Chawla, N. V.; and Ye, M. 2024. S3GCL: spectral, swift, spatial graph contrastive learning. In *Proceedings of the 41st International Conference on Machine Learning*.
- Wang, Z.; Chaliasos, S.; Qin, K.; Zhou, L.; Gao, L.; Berrang, P.; Livshits, B.; and Gervais, A. 2023. On How Zero-Knowledge Proof Blockchain Mixers Improve, and Worsen User Privacy. In *Proceedings of the ACM Web Conference 2023*, 2022–2032.
- Wolpert, D. H. 1992. Stacked generalization. *Neural networks*, 5: 241–259.
- Wu, J.; Hu, R.; Li, D.; Huang, Z.; Ren, L.; and Zang, Y. 2024. Robust Heterophilic Graph Learning against Label Noise for Anomaly Detection. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 2451–2459.
- Wu, M.; McTighe, W.; Wang, K.; Seres, I. A.; Bax, N.; Puebla, M.; Mendez, M.; Carrone, F.; De Matthey, T.; Demaestri, H. O.; et al. 2022. Tutela: An open-source tool for assessing user-privacy on ethereum and tornado cash. *arXiv preprint arXiv:2201.06811*.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- Yu, X.; Han, B.; Yao, J.; Niu, G.; Tsang, I.; and Sugiyama, M. 2019. How does disagreement help generalization against label corruption? In *Proceedings of the 36th International Conference on Machine Learning*, 7164–7173.
- Zhong, S.; and Mueen, A. 2024. BitLINK: Temporal Linkage of Address Clusters in Bitcoin Blockchain. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4583–4594.
- Zhou, J.; Hu, C.; Chi, J.; Wu, J.; Shen, M.; and Xuan, Q. 2022. Behavior-Aware Account De-Anonymization on Ethereum Interaction Graph. *IEEE Transactions on Information Forensics and Security*, 17: 3433–3448.