

CAT: Can Trust be Predicted with Context-Awareness in Dynamic Heterogeneous Networks?

Jie Wang*, Zheng Yan*[†], Jiahe Lan*, Xuyan Li[†], and Elisa Bertino[‡]

*State Key Laboratory of Integrated Services Networks, School of Cyber Engineering, Xidian University

[†]Hangzhou Institute of Technology, Xidian University [‡]Department of Computer Science, Purdue University

jwang1997@stu.xidian.edu.cn, zyan@xidian.edu.cn, jhlan16@stu.xidian.edu.cn, xli77144@gmail.com, bertino@purdue.edu

Abstract—Trust prediction provides valuable support for decision-making, risk mitigation, and system security enhancement. Recently, Graph Neural Networks (GNNs) have emerged as a promising approach for trust prediction, owing to their ability to learn expressive node representations that capture intricate trust relationships within a network. However, current GNN-based trust prediction models face several limitations: (i) Most of them fail to capture trust dynamicity, leading to questionable inferences. (ii) They rarely consider the heterogeneous nature of real-world networks, resulting in a loss of rich semantics. (iii) None of them support context-awareness, a basic property of trust, making prediction results coarse-grained.

To this end, we propose CAT, the first Context-Aware GNN-based Trust prediction model that supports trust dynamicity and accurately represents real-world heterogeneity. CAT consists of a graph construction layer, an embedding layer, a heterogeneous attention layer, and a prediction layer. It handles dynamic graphs using continuous-time representations and captures temporal information through a time encoding function. To model graph heterogeneity and leverage semantic information, CAT employs a dual attention mechanism that identifies the importance of different node types and nodes within each type. For context-awareness, we introduce a new notion of meta-paths to extract contextual features. By constructing context embeddings and integrating a context-aware aggregator, CAT can predict both context-aware trust and overall trust. Extensive experiments on three real-world datasets demonstrate that CAT outperforms five groups of baselines in trust prediction, while exhibiting strong scalability to large-scale graphs and robustness against both trust-oriented and GNN-oriented attacks.

I. INTRODUCTION

Trust is a complex and multifaceted concept, referring to a subjective view held by one entity towards another within a specific context. It is characterized by subjectivity, dynamicity, context-awareness, asymmetry, and conditional transitivity [1]. Trust evaluation, which quantifies trust by considering factors affecting it, is one of the important approaches for achieving cyber trust. It has been applied to various fields for such purposes as fraud detection, intrusion detection, service management, and access control [2], [3]. For example, social networks are often affected by fraudsters and dishonest ratings,

as studies have reported that 10% of Twitter users are fake [4] and that 16% of restaurant ratings on Yelp are dishonest [5]. In such cases, trust evaluation helps identify trustworthy entities and filter out misleading information, thereby improving the trustworthiness of the network. Overall, trust evaluation plays a crucial role in facilitating decision-making, mitigating risks, and enhancing system security.

Graph Neural Networks (GNNs), an emerging Machine Learning (ML) paradigm designed for graph-structured data, have demonstrated strong performance in various tasks, such as anomaly detection [12], [13], [14], social recommendation [15], [16], and malware analysis [17], [18]. Their strength lies in the ability to learn expressive node representations (i.e., embeddings) through information propagation and aggregation, fully leveraging the structural information inherent in graphs. This capability has motivated researchers to apply GNNs to trust evaluation, also known as trust prediction¹. Compared with statistical [19], reasoning [20], and traditional ML approaches [21], using GNNs for trust prediction offers several distinct advantages. First, trust relationships and other interactions between nodes can be naturally modeled as graphs, while GNNs are well-suited for processing graph data. Second, the message-passing mechanism of GNNs is compatible with basic trust properties, including conditional transitivity and composability, thus enhancing prediction accuracy. Third, unlike traditional ML approaches that require heavy and complex feature engineering, GNNs provide an end-to-end prediction mode, greatly simplifying the process of trust prediction. Consequently, GNNs present a promising approach for intelligent and precise trust prediction.

Existing Approaches and Motivations. Table I summarizes existing GNN-based trust prediction models in terms of dynamicity, heterogeneity, context-awareness, and robustness. We use social networks as an example to illustrate the practical significance of these properties, as shown in Fig. 1. Based on our review, we make the following observations:

(i) Most models [6], [8], [9], [10] do not account for trust dynamicity, resulting in questionable inferences and poor prediction accuracy [22], [23]. As illustrated in Fig. 1, user interactions evolve over time, and trust relationships often change accordingly; for example, u_2 has different trust levels

TABLE I: Comparison between CAT and existing GNN-based trust prediction models.

●: it satisfies a criterion; ○: it does not satisfy a criterion; ◉: it partially satisfies a criterion.

Models	Guardian [6]	Medley [7]	GATrust [8]	TrustGNN [9]	KGTrust [10]	DTrust [11]	TrustGuard [2]	CAT
Dynamicity	○	●	○	○	○	●	●	●
Heterogeneity	○	○	○	○	●	○	○	●
Context-Awareness	○	○	○	○	○	○	○	●
Robustness	○	○	○	○	○	○	●	●

towards u_1 at different time points. Static models fail to capture such dynamics and may mistakenly utilize later interactions to predict early trust relationships, thereby violating temporal causality.

(ii) Almost all models [2], [6], [7], [8], [9], [11] overlook the heterogeneous nature of real-world networks. Fig. 1 presents a typical heterogeneous network with various types of nodes (i.e., users and items) and edges (i.e., trust relationships between users and user ratings on items). Most trust prediction models focus solely on homogeneous user-to-user networks and thus miss valuable semantic information, such as user preferences or item characteristics, which have been shown to enhance trust prediction [10], [24].

(iii) None of the existing models support context-awareness, which is essential for several reasons. First, trust is inherently context-aware, as a study [25] found that fewer than 1% of people trust their friends in all contexts in real-world networks. Here, *context* refers to any information that describes the specific situations in which a trust relationship is established, such as item categories (as illustrated in Fig. 1), interaction domains, or task types. For example, u_2 may trust u_1 in sport (c_1) recommendations due to u_1 's expertise, but not in clothing (c_2) recommendations. Second, since different contexts influence trust in distinct ways, modeling trust within each context can improve overall trust prediction. Third, evaluating trust in one context helps predict trust in similar yet unobserved contexts. Last, context-aware models enable flexible trust-based applications by offering fine-grained evaluation results. For instance, one can identify malicious behaviors of others in a specific context based on context-aware trust ².

(iv) Resilience to data poisoning attacks is rarely evaluated in existing models [6], [7], [8], [9], [10], [11], hindering their practical adoption. While TrustGuard [2] addresses trust-oriented attacks that operate at the node level, it overlooks more sophisticated GNN-oriented attacks that operate at the graph-structural level. In practice, attackers can inject malicious interactions (red lines in Fig. 1) into training data to mislead model learning. Such attacks are prevalent in real-world systems; for example, Taobao has recognized service boosting via malicious ratings as a threat to its platform [5].

Therefore, it is critical to develop a trust prediction model that can capture trust dynamicity, represent real-world situations, support context-awareness, and be resilient to data poisoning attacks.

²A context-aware trust level reflects a trustor's trust in a trustee within a specific context, while an overall trust level aggregates these context-specific levels to provide a holistic view of the trustor's general trust in the trustee.

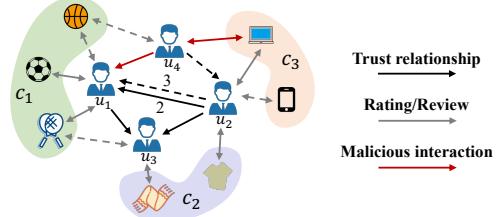


Fig. 1: A motivating example of dynamic heterogeneous networks represented by a real-world social network with interactions involving users $u_1 \sim u_4$ and items across various contexts $c_1 \sim c_3$, including sports, electronics, and clothing. Dashed and solid arrows indicate interactions occurring at different time points.

Technical Challenges. However, designing such a model is non-trivial as it requires addressing the following challenges:

TC1: Modeling dynamicity with fine granularity while maintaining scalability is a challenging task. Time is a continuous variable that requires precise representation and careful handling. Some trust prediction models [2], [11] use discrete-time methods to represent dynamic trust networks. While efficient, they may miss critical temporal patterns across snapshots. Continuous-time methods [7] address this issue by considering all timestamps but suffer from low scalability. Therefore, how to achieve both accurate dynamicity modeling and scalability calls for an effective solution.

TC2: It is difficult to extract key information from heterogeneous graphs. Real-world networks often contain multiple types of nodes and interactions [26], which provide rich semantics for accurate trust prediction. However, this diversity may also introduce excessive noise that undermines trust propagation and aggregation. Thus, how to determine the proper importance of interactions formed by different types of nodes requires a thorough investigation.

TC3: The absence of labeled, context-specific trust data poses a significant challenge for predicting trust across varied contexts. Existing datasets only provide labels for overall trust, making direct context-aware trust prediction infeasible. Moreover, in graph structures, contexts are less obvious than nodes and edges. Hence, how to incorporate contextual information (**TC3-1**) and predict context-aware trust from available overall trust data (**TC3-2**) needs ingenious designs.

Our Proposal. In this paper, we propose CAT, the first GNN-based trust prediction model that simultaneously captures trust dynamicity, handles real-world network heterogeneity, supports context-awareness, and is resilient to data poisoning attacks. CAT adopts a layered architecture that includes a graph construction layer, an embedding layer, a heterogeneous

attention layer, and a prediction layer. The graph construction layer builds a dynamic heterogeneous graph using timestamped interactions, enabling fine-grained modeling of dynamism. The embedding layer initializes embeddings for nodes, time, and edge attributes. To address **TC3-1**, we introduce a new notion: context-aware meta-path, which incorporates contextual information while enhancing semantic representation. Meanwhile, a time encoding function maps the continuous-time domain into a vector space. The heterogeneous attention layer propagates and aggregates trust information to refine embeddings. It employs a dual attention mechanism to learn the importance of node types and the nodes within each type, addressing **TC2**. To resolve the contradiction in **TC1**, we adopt recent-time neighbor sampling and one-hop trust propagation strategies, focusing on limited yet crucial interactions. Last, the prediction layer predicts the trust relationship between any two users under a specific context. To tackle **TC3-2**, we treat item categories as contexts and create context embeddings by averaging the embeddings of items within the same context. This enables CAT to predict context-aware trust. To avoid dependency on context-specific trust labels, we further propose a context-aware aggregator that links context-aware trust with overall trust. By jointly addressing **TC2** and **TC3**, CAT gains a robust semantic understanding, forcing attackers to consider multiple factors for a successful attack.

Evaluation. We conduct extensive experiments to evaluate CAT’s effectiveness, scalability, and robustness on three real-world social network datasets: Epinions [27], Ciao [27], and CiaoDVD [28]. Experimental results show that CAT outperforms five groups of baselines, especially in predicting trust for unobserved users that do not have interactions at the training stage, achieving a 50.79% Mean Reciprocal Rank (MRR) improvement on Epinions over the best baseline. Notably, CAT can predict context-aware trust, which is not supported by all baseline models. It also scales well to large graphs, reducing average running time by 73.97% on Epinions. Furthermore, CAT shows strong robustness, with maximum performance drops of only 0.95% and 3.39% under trust-oriented [2] and GNN-oriented attacks [29], respectively, when predicting trust for observed users.

Contributions. The main contributions are as follows:

- We propose CAT, the first context-aware GNN-based trust prediction model, which also uniquely addresses the dynamic nature of trust in heterogeneous networks.
- We introduce the notion of context-aware meta-paths that incorporate contextual features and generalize across domains. By constructing context embeddings, CAT can predict context-aware trust. A context-aware aggregator is further proposed to avoid reliance on context-specific trust labels and enable CAT to predict overall trust.
- We model trust dynamism using a time encoding function and extract key semantic information through a dual attention mechanism. To enhance scalability, we further adopt recent-time neighbor sampling and one-hop trust propagation strategies.
- We conduct a comprehensive evaluation on three real-

world datasets. The results show that CAT effectively predicts trust for both observed and unobserved users, scales well to large graphs, and remains robust against both trust-oriented and GNN-oriented attacks.

II. RELATED WORK

This section briefly reviews recent advances in dynamic heterogeneous GNNs and GNN-based trust prediction.

A. Dynamic Heterogeneous GNNs

Many research efforts have focused on either dynamic GNNs [30] or heterogeneous GNNs [31]. However, research on integrating both dynamicity and heterogeneity is relatively limited [32]. DyHATR [33] employs node-level and edge-level attentions to capture heterogeneous information, and applies Recurrent Neural Networks (RNNs) with a self-attention mechanism to discern temporal patterns. HTGNN [34] adopts a hierarchical aggregation mechanism to consider both heterogeneous spatial and temporal features. While DyHATR and HTGNN support dynamicity by modeling a sequence of static snapshots, they overlook the fine granularity of time. In contrast, HGT [35] incorporates fine-grained timestamps using a relative temporal encoding technique. Additionally, it leverages meta-relations with Transformer-based attention [36] to handle graph heterogeneity. However, HGT requires stacking multiple GNN layers to capture high-order information, leading to over-smoothing and poor scalability [37].

Discussion. Although these models achieve impressive performance across several tasks, their effectiveness in trust prediction remains uncertain. Specifically, basic trust properties (e.g., context-awareness and asymmetry) and trust-related attributes (e.g., trust levels and ratings) are not explicitly modeled, which may compromise prediction accuracy. Moreover, the joint modeling of dynamicity and heterogeneity often raises scalability concerns. Therefore, how to adapt existing dynamic heterogeneous GNNs to trust prediction while ensuring scalability requires further extensive investigation.

B. GNN-based Trust Prediction

Generally, GNN-based trust prediction models can be classified into static models and dynamic models, depending on whether they incorporate temporal information.

Static Models. Lin *et al.* [6] proposed Guardian, the first GNN-based trust prediction model for online social networks. Guardian leverages the concepts of in-degree and out-degree to model trust asymmetry, and utilizes localized graph convolutions for trust propagation. GATrust [8] improves Guardian by fusing multi-faceted user properties, including user features, known trust relationships, and network structure, using learnable weights. Unlike the above two models, TrustGNN [9] models trust by specifying different trust chains and then applies an attention mechanism to discriminate their importance. However, the above models focus on homogeneous user-to-user trust networks, ignoring the heterogeneous interactions between different node types found in real-world networks. KGTrust [10] addresses this issue by initializing embeddings

with type-specific semantics and employing a discriminative convolutional mechanism for embedding learning. Despite these advances, all these models operate on static snapshots and thus fail to capture the dynamic nature of trust. This oversight leads to information loss and questionable inferences [22], [23], as they miss the sequence in which trust relationships are established. Moreover, the lack of support for context-awareness and robustness further limits their effectiveness in practice.

Dynamic Models. Current dynamic models adopt two approaches for modeling time: discrete-time and continuous-time. DTrust [11] and TrustGuard [2] use the discrete-time method that treats a dynamic graph as a series of time-ordered snapshots. They first learn static spatial features within a snapshot, and then capture temporal patterns across snapshots through Gated Recurrent Units (GRUs) and a self-attention mechanism, respectively. While efficient, this method may lose substantial structural dependencies across snapshots [35]. In addition, DTrust lacks robustness considerations, whereas TrustGuard employs a robust aggregator to counter trust-oriented attacks based on the network theory of homophily. However, TrustGuard’s efficacy under GNN-oriented attacks remains unclear. Medley [7] is a continuous-time method that fully utilizes each timestamp through a time encoding function. It achieves high prediction accuracy but suffers from expensive costs and lacks robustness. Moreover, all these models do not support heterogeneity and context-awareness.

Discussion. Despite remarkable progress in GNN-based trust prediction, several critical issues remain unresolved, as shown in Table I. First, insufficient attention is given to trust dynamicity modeling, particularly the efficient encoding of fine-grained timestamps. Second, heterogeneity is not well addressed in existing work. Third, none of the models support context-awareness, a basic nature of trust. Last, the resilience of these models to both trust-oriented and GNN-oriented attacks has been largely overlooked.

III. PROBLEM STATEMENT

In this section, we first introduce key definitions and formulate the problem we aim to address. Then, we describe the threat model of CAT.

A. Problem Definition

Definition 1. Heterogeneous Network. A heterogeneous network, also known as Heterogeneous Information Network (HIN), is represented as a six-tuple $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{T}, \mathcal{R}, \Phi, \Psi)$, where each node $v \in \mathcal{V}$ and each edge $e \in \mathcal{E}$ are associated with their type mapping functions $\Phi : \mathcal{V} \rightarrow \mathcal{T}$ and $\Psi : \mathcal{E} \rightarrow \mathcal{R}$, respectively. A network is considered heterogeneous when $|\mathcal{T}| + |\mathcal{R}| > 2$, while it is homogeneous when $|\mathcal{T}| = |\mathcal{R}| = 1$.

Definition 2. Dynamic Heterogeneous Network. A dynamic heterogeneous network observed within a time interval $[0, T]$ can be expressed as $\mathcal{G}(T) = (\mathcal{V}(T), \mathcal{E}(T), \mathcal{T}, \mathcal{R}, \Phi, \Psi)$, where $\mathcal{V}(T)$ and $\mathcal{E}(T)$ are time-varying sets of nodes and edges, respectively. $\mathcal{E}(T) = \{(e_{i,j}(t), r) \mid v_i(t), v_j(t) \in \mathcal{V}(T), r \in$

$\mathcal{R}, t \leq T\}$, where $e_{i,j}(t)$ is an edge between nodes v_i and v_j built at timestamp t , and r is the edge type.

Definition 3. Meta-path. A meta-path p of length s is defined as a path in the form of $\tau_1 \xrightarrow{r_1} \tau_2 \xrightarrow{r_2} \dots \xrightarrow{r_s} \tau_{s+1}$ with node types $\tau_1, \tau_2, \dots, \tau_{s+1} \in \mathcal{T}$ and edge types $r_1, r_2, \dots, r_s \in \mathcal{R}$. Note that any two node types and any two edge types in this path can be the same. The meta-path p describes a composite relation $r_1 \circ r_2 \circ \dots \circ r_s$ between node types τ_1 and τ_{s+1} , capturing rich semantics.

Take a heterogeneous social network (Fig. 1) as an example, the meta-path $user \xrightarrow{\text{rate}} item \xrightarrow{\text{rate}} user$, where rate- means “is rated by”, describes the relationship “two users have rated the same item”, indicating similar preferences or potential trust between them.

Definition 4. Contextual Trust Network. A contextual trust network is a dynamic heterogeneous network with contexts involved, denoted as $\mathcal{G}(T) = (\mathcal{V}(T), \mathcal{E}(T), \mathcal{T}, \mathcal{R}, \Phi, \Psi, \mathcal{C})$. Here, $\mathcal{T} = \{user, item\}$, $\mathcal{R} = \{\langle user, user \rangle, \langle user, item \rangle, \langle item, user \rangle\}$, and \mathcal{C} denotes the set of contexts.

In Definition 4, the term *user* refers to either a trustor or a trustee — two entities involved in a trust relationship, while the term *item* has application-specific meanings, such as products in social networks or assets in financial networks. Based on the datasets employed in this paper, each context $c \in \mathcal{C}$ corresponds to an item category. Among the edge types, $\langle user, user \rangle$ represents a directed trust relationship and is inherently asymmetric, while $\langle user, item \rangle$ and $\langle item, user \rangle$ represent inverse relationships, such as “rate” and “is rated by”. The absence of $\langle item, item \rangle$ relationships is based on the assumption that the primary interactions of interest occur between *users* and *items* or between *users* themselves. Additionally, the set of context-aware trust relationships is represented as $\{(\langle v_i, v_j \rangle, c_k, t, w_{i,j,k}, r) \mid e_{i,j}(t) \in \mathcal{E}(T), c_k \in \mathcal{C}, r = \langle user, user \rangle\}$, where $w_{i,j,k}$ denotes the trust level of the trustor-trustee pair $\langle v_i, v_j \rangle$ under context c_k . The value of $w_{i,j,k}$ can be binary or multi-level depending on specific scenarios. Notably, due to the asymmetric nature of trust, $w_{i,j,k}$ is not necessarily equal to $w_{j,i,k}$.

Definition 5. Trust Prediction. Given a contextual trust network $\mathcal{G}(T)$ observed in the time interval $[0, T]$, the goal of trust prediction is to learn a model $f(\cdot)$ that can predict, at time $T + \Delta(T)$, both the context-aware trust level $\bar{w}_{i,j,k}$ for any trustor-trustee pair $\langle v_i, v_j \rangle$ within a specific context c_k and the overall trust level $\bar{w}_{i,j}$ (aggregated across all contexts) of this pair. Here, $v_i, v_j \in \mathcal{V}(T)$ represent user nodes, $v_i \neq v_j$, and $c_k \in \mathcal{C}$.

The trust prediction task may vary depending on the type of available supervision. If only labels for trusted relationships are provided, with distrusted ones implicitly represented by the absence of links, the task reduces to link prediction: predicting whether a trusted relationship (link) exists between two users. Conversely, if labels for trust strength (e.g., numerical scores) are available, the task becomes edge classification: predicting the trust level (edge weight) between users.

B. Threat Model

An important requirement for designing a GNN-based trust prediction model is that the model can withstand both trust-oriented [2] and GNN-oriented attacks [29]. Trust-oriented attacks operate at the local, node level, targeting individual trustworthiness by manipulating specific trust relationships. For example, attackers may destroy or boost the trustworthiness of certain nodes by launching bad-mouthing or good-mouthing attacks. In contrast, GNN-oriented attacks typically operate at the global, graph-structural level, exploiting inherent vulnerabilities of GNNs and highlighting the adversarial relationship between the attacker and the model. In this case, attackers may have access to the full input graph and strategically manipulate its structural patterns. Both types of attacks are prevalent in real-world scenarios [5], [29] and usually manifest as data poisoning attacks, where attackers degrade model performance by manipulating links (e.g., trust relationships) during training. Therefore, this paper analyzes model robustness under both attacks.

We make the following assumptions. For trust-oriented attacks, we follow the approach proposed in [2] and focus on collaborative bad-mouthing attacks (more disruptive than good-mouthing attacks), where a group of malicious nodes collude to destroy the trustworthiness of target nodes. For GNN-oriented attacks, we consider a powerful attacker who can access and analyze the contextual trust network $\mathcal{G}(T)$, including its statistics and temporal evolution patterns. The attacker can manipulate the network structure by strategically adding or removing links, e.g., by creating a surrogate model to guide the attack. However, the attacker lacks knowledge of the architecture or internal parameters of the trust prediction model and cannot query it. These settings are realistic and consistent with existing work [29]. In practice, attackers typically do not know the details of a victim model, yet they can observe the input data (e.g., public social networks) and formulate attack strategies using surrogate models. We further assume that during the test phase, the input to the trust prediction model is a pair of nodes indicating the roles of trustor and trustee. Thus, evasion attacks are not applicable. Instead, data poisoning attacks may occur during the training phase to degrade the model's overall performance.

Beyond these two attacks, we also consider an adaptive data poisoning attack that specifically targets a context-aware trust prediction model by manipulating interactions across contexts to blur contextual boundaries (see Section VI).

IV. CAT DESIGN

In this section, we first provide an overview of CAT and then elaborate on its technical details.

A. Overview

As illustrated in Fig. 2, CAT consists of four layers: a graph construction layer, an embedding layer, a heterogeneous attention layer, and a prediction layer. Specifically, the graph construction layer builds a contextual trust graph using streaming user-item and user-user interactions. In the embedding

layer, a meta-path covering rich semantics and contextual features is first defined, and then node embeddings are initialized via Metapath2vec [38]. Meanwhile, time information and edge attributes (e.g., ratings from users on items) are encoded to facilitate information propagation and aggregation. The heterogeneous attention layer employs a dual attention mechanism to discriminate the importance of node types and the importance of nodes within the same type. Leveraging the message-passing mechanism of GNNs, CAT achieves trust propagation with varying attentions granted on different interactions. To enhance scalability, we adopt recent-time neighbor sampling and one-hop trust propagation strategies, focusing on limited yet crucial interactions. Finally, in the prediction layer, we first generate a context embedding by averaging item embeddings within the same context. Then, two user embeddings together with the context embedding are fed into a Multi-Layer Perceptron (MLP) to predict the latent trust level between the two users within that context. To overcome the lack of context-specific trust labels, we propose a context-aware aggregator to link context-aware trust with overall trust, weighted by context importance. The integration of these layers gives CAT a comprehensive semantic understanding, which hinders potential attacks, as attackers have to consider multiple contextual and semantic factors.

B. Graph Construction Layer

This layer aims to construct a contextual trust graph using user-user and user-item interactions, each associated with attributes such as timestamps, trust levels, and ratings. To effectively model the temporal dynamics of these interactions, we employ continuous-time representations, which describe the evolution of the graph as a function of continuous time. Unlike discrete-time representations [2], [11], [39] that divide time into fixed intervals and may lose fine-grained temporal patterns, continuous-time representations allow interactions and updates to occur at arbitrary time points, enabling a more accurate and realistic modeling of trust dynamics. In addition, items belonging to the same category are grouped under the same context.

C. Embedding Layer

This layer aims to initialize embeddings for nodes, time, and edge attributes.

Node Embedding. Typically, we can apply two transformation matrices for users and items to characterize their heterogeneity. While this method is straightforward, it fails to capture rich semantics embedded in diverse relations. Meta-paths have emerged as a promising solution to this limitation. However, conventional meta-paths, which focus on explicit links, lack the capacity to incorporate contextual information, a crucial factor for context-aware trust prediction. To address this, we extend the concept of meta-paths by introducing a new context-aware meta-path, defined as: $user \xrightarrow{\text{rate}} item \xrightarrow{\text{belong to}} context \xrightarrow{\text{belong to}^-} item \xrightarrow{\text{rate}^-} user$, simplified as $uiciu$. It is important to clarify that the “context” herein is not an actual node type but rather a

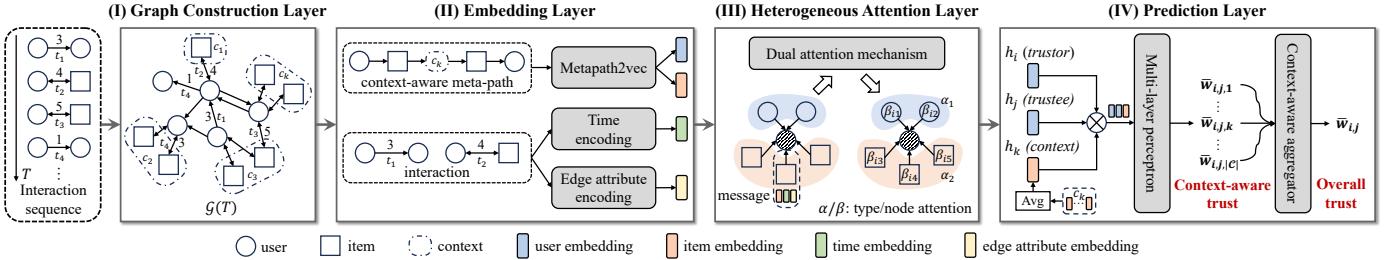


Fig. 2: CAT overview.

latent factor that reflects item characteristics. This meta-path describes the relationship “two users have interacted with two items belonging to the same context”, implying that these two users share similar preferences in that context. Then, we apply the Metapath2vec model [38] to generate semantically-rich embeddings $h_i, h_j \in \mathbb{R}^{d_v}$, where d_v is the embedding dimension. Unlike previous approaches [8], [40], [41] that manually define contextual factors (e.g., decomposing neighbor features into normal and abnormal components), our method incorporates contextual information through relational structures via context-aware meta-paths. This design offers three advantages: (i) It automatically captures latent contextual semantics from user-item interactions. (ii) It leverages meta-paths to connect different types of nodes and characterize diverse relations. (iii) It generalizes across networks by eliminating manual context engineering (as discussed in Section VI).

Time Embedding. Metapath2vec has a key limitation: it does not retain the temporal order of interactions within a meta-path. To address this problem and model trust dynamics, we introduce a time encoding function $\varphi : \Delta t \rightarrow \mathbb{R}^{d_t}$. Herein, Δt represents the time gap between an interaction and its associated target node, and d_t is the embedding dimension. We use relative time gaps instead of absolute timestamps to reduce inconsistencies that arise from the varying ranges of absolute timestamps (e.g., Unix timestamps) across different datasets and between training and testing sets. Additionally, unlike previous approaches [42], [43] that focus on the target timestamp of each node, we assign each node’s timestamp as the time at which its earliest interaction occurs. The rationale behind this choice is twofold: (i) The first interaction marks a node’s entry into a network, which holds historical significance. (ii) It simplifies time encoding by using a single representative timestamp per node, thereby avoiding redundant computation. The purpose of φ is to encode continuous time gaps into d_t -dimensional embeddings h_t , ensuring that temporally close gaps are also close in the embedding space. Formally, the time encoding function is expressed as:

$$\varphi(\Delta t) = \cos(\Delta t \cdot \omega + \delta), \quad (1)$$

where ω and δ are learnable parameters. The cosine function is chosen due to (i) its inherent periodicity, making it ideal to capture periodic temporal patterns, and (ii) its continuity and smoothness, which ensure that small changes in time result in small changes in the output of the function.

Edge Attribute Embedding. Except for timestamps, each edge may have such attributes as trust levels and user-item ratings, which can improve prediction accuracy but are often ignored by existing dynamic heterogeneous GNNs. To address this, we encode these attributes into embedding vectors. For instance, if the attributes refer to trust levels w and $w \in \{Trust, Distrust\}$, we first model them as $[0, 1]^\top$ and $[1, 0]^\top$ using one-hot encoding. Given that different datasets have different levels of trust, we then transform these one-hot vectors into embeddings $h_e \in \mathbb{R}^{d_e}$ of a specified dimension via a learnable matrix.

D. Heterogeneous Attention Layer

This layer aims to learn node embeddings by selectively propagating and aggregating trust information that encodes trust strength, relevance, and directionality. Specifically, we design a dual attention mechanism that assigns varying weights to node types (type attention) and nodes within each type (node attention), enabling CAT to discriminate the significance of heterogeneous interactions. Compared to existing attention mechanisms [31], [35] for handling graph heterogeneity, the proposed mechanism offers two advantages: (i) It accounts for both temporal information and edge attributes, which are essential for accurate trust prediction but often overlooked in prior work. (ii) It does not rely on any meta-path and focuses on node type importance, reducing manual effort and improving generality. Furthermore, to enhance scalability, we adopt recent-time neighbor sampling and one-hop trust propagation strategies for controlling the amount of trust information and the range of its propagation. In this subsection, we denote v_j as the source node and v_i the target node.

To begin with, we define a message $m_{i,j}$ as the trust information that needs to be propagated from v_j to v_i :

$$m_{i,j} = h_j + h_t + h_e, \quad (2)$$

where h_j , h_t , and h_e represent the embeddings of v_j , relative time gaps, and edge attributes, respectively. The messages associated with v_i serve as the basis for mining its latent features, such as the objective trust given by others. Note that we use an additive method rather than concatenation for message construction, as this avoids the high dimensionality introduced by concatenation and is thus more efficient.

Type Attention. Given a specific node v_i , it can engage in various interactions with different types of nodes, each having distinct impacts. For example, for a user node, a neighboring

user node from a $\langle user, user \rangle$ interaction propagates information about trust relationships, while a neighboring item node from a $\langle item, user \rangle$ interaction conveys information such as user preferences. Therefore, we employ a type attention mechanism to discriminate the significance of different node types or relations. First, we calculate the embedding for type τ by aggregating messages from nodes of type τ as:

$$h_\tau = \sum_{j \in \mathcal{N}_{v_i}^\tau} \frac{1}{\sqrt{d_i d_j}} m_{i,j}, \quad (3)$$

where $\mathcal{N}_{v_i}^\tau$ denotes the set of neighboring nodes of type τ for v_i . d_i and d_j are the degrees of v_i and v_j , used to normalize message influence based on node connectivity.

Then, we calculate the type attention score based on the node type embedding h_τ and the target node's embedding h_i . To obtain the relative importance of each type, we further normalize the attention scores across all the types using the softmax function. The above process is shown in Eq. 4.

$$\begin{aligned} a_\tau &= \sigma(\gamma_\tau^\top \cdot [h_i || h_\tau]), \\ \alpha_\tau &= \frac{\exp(a_\tau)}{\sum_{\tau' \in \mathcal{T}} \exp(a_{\tau'})}, \end{aligned} \quad (4)$$

where γ_τ is the attention vector specific to type τ , $||$ is a concatenation operation, and $\sigma(\cdot)$ is a non-linear activation function such as LeakyReLU.

Node Attention. Falling into the same type, different nodes can influence a target node differently. Therefore, we further calculate node attention scores with the consideration of type importance. Specifically, given a node v_i with type τ and its neighboring node v_j with type τ' (where τ and τ' may be the same or different), the normalized node attention score is calculated using Eq. 5.

$$\begin{aligned} b_{i,j} &= \sigma(\eta^\top \cdot \alpha_{\tau'}[h_i || m_{i,j}]), \\ \beta_{i,j} &= \frac{\exp(b_{i,j})}{\sum_{j' \in \mathcal{N}_{v_i}} \exp(b_{i,j'})}, \end{aligned} \quad (5)$$

where η is the attention vector for nodes, and $\alpha_{\tau'}$ refers to the significance of type τ' . The node attention score $\beta_{i,j}$ indicates the importance of v_j to v_i . Note that the effects of time and edge attributes are also learned through the dual attention mechanism since they are encoded in the messages.

Finally, to derive the embedding of node v_i , we aggregate its related messages $m_{i,j}$, each weighted by $\beta_{i,j}$ (computed via sequential type and node attention mechanisms). After that, a fully-connected layer is applied to extract high-level features. The formal process is described below:

$$h_i = \sigma\left(\left(\sum_{j \in \mathcal{N}_{v_i}} \beta_{i,j} m_{i,j}\right) \cdot W + B\right), \quad (6)$$

where W and B are learnable parameters.

Neighbor Sampling. A real-world heterogeneous network is typically very large and contains rich information. While this information benefits embedding learning, it may also introduce noise that undermines trust propagation and aggregation. Considering this issue and CAT's scalability, we

adopt a neighbor sampling strategy [30], [44] that selects a subset of neighbors (or interactions) for trust propagation and aggregation. There are two sampling strategies: uniform sampling, which uniformly samples neighbors of a target node, and recent-time sampling, which selects neighbors having the newest interactions with the target node. We adopt the recent-time sampling strategy since it is more efficient, and new interactions are generally more important than historical ones [7]. To the best of our knowledge, this strategy is applied for the first time in GNN-based dynamic trust prediction. A comparison of the two strategies is provided in Appendix E.

One-Hop Trust Propagation. Through the above process, latent trust information embedded in messages can be propagated and aggregated, guided by the graph structure. Specifically, if CAT is equipped with L heterogeneous attention layers, h_i fuses the information about trust from the L -hop neighbors of v_i . In other words, we can control the range of trust propagation by setting L . However, increasing L introduces several issues. First, a large L often leads to over-smoothing [45], where node embeddings become indistinguishable. Second, deeper architectures significantly increase computational overhead, making them less scalable in large-scale trust networks. Moreover, previous work has shown that trust is conditionally transitive and tends to decay with distance [9], [46], highlighting the importance of one-hop neighbors over multi-hop ones. Empirical results (Section V-D) further support the above statements: CAT performs best when $L = 1$, indicating that trust information from one-hop neighbors is sufficient for accurate and efficient prediction. Therefore, we adopt a one-hop trust propagation strategy to enhance scalability without compromising accuracy.

E. Prediction Layer

This layer aims to predict the latent trust between any two user nodes within a specific context. Since the previous layer already generates expressive embeddings through the dual attention mechanism, this layer adopts a simple two-layer MLP as the prediction model. The input to the MLP is a combination of three types of embeddings: those of the trustor, the trustee, and the context. While embeddings of the trustor and trustee are directly obtained from the previous layer, the context embedding is not readily accessible. To enable context-aware trust prediction, we calculate a context embedding by averaging the embeddings of items within the same context. The reason is that items within the same context collectively characterize that context from multiple aspects, and averaging effectively preserves the general characteristics of the context. Alternative methods for generating context embeddings are discussed in Appendix F. Formally, the prediction process is as follows:

$$\tilde{w}_{i,j,k} = \text{softmax}(\text{MLP}(h_i || h_j || h_k)), \quad (7)$$

where h_i , h_j , and h_k denote the embeddings of v_i , v_j , and the context c_k , respectively. $\tilde{w}_{i,j,k}$ is the probabilistic prediction vector of the trust from v_i to v_j within c_k . The predicted trust level is thus computed by $\bar{w}_{i,j,k} = \arg \max_q(\tilde{w}_{i,j,k})$, where

q refers to the index of the maximum value in $\tilde{w}_{i,j,k}$. It is important to note that $\tilde{w}_{i,j,k}$ may not equal $\tilde{w}_{j,i,k}$ due to the asymmetric nature of trust.

Context-Aware Aggregation. By the above steps, CAT has the potential to predict trust between users within any context, which is more fine-grained than the overall trust used in prior approaches. However, we are facing the lack of ground truth for trust levels within specific contexts. To address this challenge, we establish a connection between context-aware trust and overall trust. A solution is to average or sum trust levels across all contexts to derive an overall trust level [24]. This solution is efficient but does not consider the importance of different contexts. We thus introduce a context-aware aggregator to automatically determine the significance of each context in forming the overall trust level:

$$\tilde{w}_{i,j} = \sum_{k=1}^{|C|} g_k \cdot \tilde{w}_{i,j,k}, \quad (8)$$

where $|C|$ refers to the number of contexts. $G_k = [g_1, \dots, g_k, \dots, g_{|C|}]$ is a learnable vector, in which each element g_k indicates the importance of a corresponding context, and the sum of the elements equals one via the softmax function. The output $\tilde{w}_{i,j}$ is the probabilistic prediction vector of the overall trust from v_i to v_j . Similarly, we can compute the overall trust level using $\bar{w}_{i,j} = \arg \max_q(\tilde{w}_{i,j})$. The procedure of CAT is presented in Algorithm 1.

To summarize, *CAT is the first GNN-based trust prediction model that not only predicts the overall trust but also provides insights into the contexts in which trust relationships are more likely to form, despite the absence of context-specific trust labels*. This is a key contribution that distinguishes CAT from previous approaches [8], [40], [41], which consider multiple contextual factors to guide embedding learning but cannot explicitly predict context-aware trust. This capability is significant for enabling flexible trust-based applications, such as fraud detection and access control.

F. Model Training

To train CAT, we adopt the cross-entropy loss function to measure the discrepancy between the predicted overall trust and the ground truth. The loss function is formulated as:

$$\mathcal{L} = - \sum_{e_{i,j} \in \mathcal{E}(T)} \log \tilde{w}_{i,j; w_{i,j}} + \lambda \cdot \|\Theta\|_2^2, \quad (9)$$

where $w_{i,j}$ denotes the ground truth, Θ denotes all learnable parameters of the model, and λ controls the L_2 regularization to mitigate overfitting. Additionally, we use the Adam optimizer [47] for updating the model parameters.

V. EXPERIMENTAL EVALUATION

In this section, we first introduce experimental settings and then answer the following research questions: **RQ1:** How does CAT perform compared with the representative models regarding overall trust and context-aware trust? **RQ2:** How does each component of CAT contribute to its performance? **RQ3:**

Algorithm 1: CAT Procedure

```

Input: A dynamic graph  $\mathcal{G}(T)$ ;
Output: Context-aware trust level  $\tilde{w}_{i,j,k}$ , overall trust level  $\bar{w}_{i,j}$ ;
1   Generate node embeddings  $h_i, h_j$  for all nodes, time embeddings  $h_t$  and edge attribute embeddings  $h_e$  for all interactions; ▷ Embedding Initialization
2   for  $v_i \in \mathcal{V}(T)$  do
3     for  $j \in \text{Sampling}(\mathcal{N}_{v_i})$  do
4       Construct a message  $m_{i,j}$  using Eq. 2;
5       Calculate the type embedding  $h_\tau$  using Eq. 3;
6       Calculate the type attention score  $\alpha_\tau$  using Eq. 4;
7       Calculate the node attention score  $\beta_{i,j}$  using Eq. 5;
8     end
9     Form  $v_i$ 's embedding  $h_i$  using Eq. 6;
10   end
11   ▷ Trust Relationship Prediction
12   for  $e_{i,j} \in \mathcal{E}(T)$  do
13     Calculate the context-aware trust vector  $\tilde{w}_{i,j,k}$  using Eq. 7;
14     Calculate the overall trust vector  $\tilde{w}_{i,j}$  by aggregating trust across all contexts through Eq. 8;
15     Obtain the context-aware trust level  $\bar{w}_{i,j,k}$  via  $\arg \max_q(\tilde{w}_{i,j,k})$ ;
16     Obtain the overall trust level  $\bar{w}_{i,j}$  via  $\arg \max_q(\tilde{w}_{i,j})$ ;
17   end
18   ▷ Overall Trust Level Prediction
19 
```

How do different hyperparameters affect the performance of CAT? **RQ4:** Is CAT efficient and scalable compared with the representative models? **RQ5:** Is CAT robust against data poisoning attacks compared with the representative models? We also visualize the learned attention scores and embeddings to interpret CAT's behavior (Appendix G and H).

A. Experimental Settings

Datasets. Our experiments are conducted on three real-world datasets: Epinions [27], Ciao [27], and CiaoDVD [28], all sourced from social networking-based consumer review sites. These datasets are widely used benchmarks for trust prediction and, to the best of our knowledge, are the only publicly available ones that provide temporal information, heterogeneous interactions, and item categories, making them well-suited for our study. Specifically, they include: (i) user-item ratings ranging from 1 to 5; (ii) overall trust relationships between users, e.g., user A trusts user B regardless of specific contexts; (iii) item categories, a type of contextual information; and (iv) interaction timestamps. The statistics of the datasets are given in Table II.

Data preparation. For Epinions, we chronologically split the trust relationships into 70%-15%-15% and 80%-10%-10% for training, validation, and testing, two common split ratios used in dynamic graph analysis. For Ciao and CiaoDVD, where the trust relationships lack timestamps, we randomly split them into training, validation, and testing sets with ratios of x , $\frac{1-x}{2}$, and $\frac{1-x}{2}$, respectively ($x = 50\%, 60\%, 70\%, 80\%$). In addition, we initialize the time embeddings of trust relationships as zero vectors to ensure CAT works correctly on these two datasets. Since only positive links (i.e., trusted relationships) exist in all datasets, an equal number of unlinked user pairs is randomly selected to form a distrusted set for training, validation, and testing. This method, known

TABLE II: Statistics of Epinions, Ciao, and CiaoDVD datasets.

Datasets	# Users	# Items	# Ratings	# Trust relationships	# Contexts	Timestamps of ratings	Timestamps of trust relationships
Epinions	9163	12573	265189	311158	25	✓	✓
Ciao	2378	16861	36065	57544	6	✓	✗
CiaoDVD	19533	16121	72665	40133	17	✓	✗

as negative sampling, is commonly used in current research to create a balanced dataset for effective model training [48].

Baseline Models. To conduct a comprehensive comparison, we select four types of GNN models as baselines that vary in their support for dynamicity and heterogeneity. (i) *Homogeneous static model*. Guardian [6] is the first GNN model for trust prediction, which has demonstrated superiority over three types of non-GNN-based approaches. GATruth [8] utilizes attention mechanisms to incorporate multiple trust-related factors. (ii) *Homogeneous dynamic model*. Medley [7] and TrustGuard [2] are two state-of-the-art GNN-based trust prediction models with dynamicity support. They employ continuous-time and discrete-time methods, respectively, to predict time-aware trust relationships. (iii) *Heterogeneous static model*. HAN [31] introduces a hierarchical attention mechanism to learn the importance of meta-path-based neighbors and different meta-paths. KGTrust [10] is excluded from our evaluation due to its closed-source nature and limited improvements over Guardian. (iv) *Heterogeneous dynamic model*. Research on GNN-based trust prediction has not yet evolved to this model type, while CAT belongs to it. To validate the effectiveness of CAT, we adapt HGT [35] to the trust prediction task. HGT is a continuous-time model that captures graph heterogeneity through meta-relations and learns mutual attention across them. In addition to these GNN-based models, we include Linear, a simple yet strong baseline that uses only linear layers but has been shown to outperform many complex GNNs [49]. For this baseline, node features are initialized using Metapath2vec with node-type information, given the absence of textual attributes in the datasets.

Evaluation Metrics. We evaluate the effectiveness of CAT and baselines using *Mean Reciprocal Rank (MRR)*, *Average Precision (AP)*, and *Area Under the ROC Curve (AUC)*. These metrics are widely used to evaluate GNN-based link prediction [50], [51], [52], making them well-suited for our task. We do not use accuracy and F1 score, as they are more suitable for edge classification with explicit trust strength labels, whereas link prediction focuses on ranking positive links (trusted relationships) higher than negative ones (distrusted relationships). All metrics range from 0 to 1, with higher values indicating better performance. Detailed definitions are given in Appendix B. We also include *running time* to assess model efficiency and scalability.

Implementation Details. We implemented CAT ³ and baseline models using PyTorch on a server equipped with an Intel Xeon Platinum 8352V CPU and an RTX 4090 GPU. For training, we set the maximum number of epochs to 50 and adopt an early stopping strategy: training stops if AP

³<https://github.com/Jieerbobo/CAT>

TABLE III: Performance comparison between CAT and different baselines on the Epinions dataset.

In each column, the best result is highlighted in **bold**, while the second-best result is underlined.

Models	70%-15%-15%			80%-10%-10%		
	MRR	AP	AUC	MRR	AP	AUC
Task 1: Trust prediction for observed users						
Linear	0.3866	0.8247	0.8873	0.3622	0.8305	0.8907
Guardian	0.3097	0.8237	0.9233	0.4979	0.9080	0.9444
GATruth	0.4380	0.8851	0.9431	0.5168	0.9135	0.9460
Medley	0.4762	0.8944	0.9440	0.5577	0.9336	0.9628
TrustGuard	0.4955	0.8919	0.9382	0.5390	0.9214	0.9542
HAN	0.4054	0.8731	0.9396	0.4100	0.8869	0.9415
HGT	<u>0.5081</u>	<u>0.9151</u>	<u>0.9588</u>	<u>0.6168</u>	<u>0.9446</u>	<u>0.9675</u>
CAT	0.6025	0.9383	0.9677	0.6778	0.9603	0.9773
Task 2: Trust prediction for unobserved users						
Linear	0.2520	0.9067	0.8047	0.2675	0.9364	0.8297
Guardian	0.1497	0.8034	0.5902	0.1658	0.8573	0.6395
GATruth	0.1917	0.8137	0.6017	0.1770	0.8514	0.6140
Medley	0.1979	0.8884	0.7806	0.2203	0.9339	0.8381
TrustGuard	0.2571	0.8950	0.7312	0.2634	0.9320	0.7891
HAN	<u>0.2707</u>	<u>0.9365</u>	<u>0.8773</u>	0.2040	0.9387	0.8596
HGT	0.2693	0.9326	0.8603	<u>0.2931</u>	<u>0.9504</u>	<u>0.8633</u>
CAT	0.4082	0.9527	0.8933	0.3987	0.9594	0.8799

on the validation set does not improve for 10 consecutive epochs. Regarding hyperparameters of CAT, we conducted a grid search and set the learning rate to 0.001, dropout to 0, L_2 regularization coefficient to 0.0005, embedding dimension to 64, batch size to 256, and trust propagation length to 1 by default. Baseline models were similarly tuned to achieve their best performance in trust prediction, except that their trust propagation length was fixed to 2 (a common choice in the literature [7]). Considering the scales of the datasets, we set the number of sampled neighbors to 30 for Epinions and use all neighbors for Ciao and CiaoDVD. We report the average results obtained from 5 runs for each experiment.

B. Performance Comparison (RQ1)

In this subsection, we first evaluate the performance of CAT and baselines on predicting overall trust across three datasets, as these datasets only have labels for the overall trust. Subsequently, we illustrate how CAT outputs context-aware trust and discuss its practical significance.

Overall Trust. For the Epinions dataset, where timestamps of trust relationships are available, we design two tasks to assess model performance on dynamic graphs: trust prediction for observed users (task 1) and unobserved users (task 2). A key characteristic of dynamic graphs is that users may exist (have interactions) or disappear (without interactions) as time evolves [53]. To capture this, we define unobserved users as those who only exist in the validation and testing

TABLE IV: Performance comparison between CAT and different baselines on the Ciao and CiaoDVD datasets.
In each column, the best result is highlighted in **bold**, while the second-best result is underlined.

Datasets	Models	50%-25%-25%			60%-20%-20%			70%-15%-15%			80%-10%-10%		
		MRR	AP	AUC									
Ciao	Linear	0.1978	0.7676	0.7688	0.2118	0.7836	0.7847	0.2026	0.7730	0.7736	0.2098	0.7817	0.7824
	Guardian	0.2177	0.8126	0.8318	0.2397	0.8198	0.8312	0.2222	0.8123	0.8290	0.2223	0.8135	0.8311
	GATrue	0.2475	0.8292	0.8432	0.2440	0.8285	0.8440	0.2438	0.8313	0.8473	0.2293	0.8236	0.8433
	Medley	0.2092	0.8049	0.8257	0.2356	0.8270	0.8438	0.2573	0.8406	0.8526	0.2484	0.8358	0.8513
	TrustGuard	0.2706	0.8327	0.8354	0.2593	0.8281	0.8332	0.2806	0.8416	0.8446	0.2760	0.8403	0.8436
	HAN	0.2411	0.8195	0.8299	0.2529	0.8298	0.8393	0.2575	0.8313	0.8408	0.2643	0.8341	0.8413
	HGT	<u>0.2772</u>	<u>0.8470</u>	<u>0.8550</u>	<u>0.2775</u>	<u>0.8508</u>	<u>0.8600</u>	<u>0.2835</u>	<u>0.8559</u>	<u>0.8650</u>	<u>0.3008</u>	<u>0.8623</u>	<u>0.8693</u>
	CAT	0.3716	0.9097	0.9215	0.3881	0.9149	0.9257	0.4150	0.9234	0.9327	0.4042	0.9200	0.9298
CiaoDVD	Linear	0.4119	0.9168	0.9226	0.3995	0.9114	0.9176	0.4097	0.9156	0.9213	0.4157	0.9174	0.9231
	Guardian	0.3571	0.9146	0.9358	0.3509	0.9131	0.9351	0.3748	0.9177	0.9366	0.3876	0.9180	0.9362
	GATrue	0.3848	0.9199	0.9390	0.3796	0.9187	0.9379	0.3891	0.9193	0.9394	0.3660	0.9126	0.9331
	Medley	0.4318	0.9320	0.9440	0.4605	0.9394	0.9492	0.4578	0.9392	0.9496	0.4742	0.9410	0.9497
	TrustGuard	0.5820	0.9571	0.9601	0.5786	0.9577	0.9613	0.5872	0.9566	0.9586	0.5618	0.9539	0.9577
	HAN	0.7258	0.9796	0.9816	0.7350	0.9804	0.9818	0.7138	0.9794	0.9819	0.7173	0.9801	0.9824
	HGT	<u>0.7606</u>	<u>0.9826</u>	<u>0.9826</u>	<u>0.7563</u>	<u>0.9827</u>	<u>0.9828</u>	<u>0.7481</u>	<u>0.9823</u>	<u>0.9829</u>	<u>0.7340</u>	<u>0.9814</u>	<u>0.9826</u>
	CAT	0.8225	0.9871	0.9861	0.8259	0.9872	0.9862	0.8413	0.9890	0.9883	0.8406	0.9890	0.9880

stages, with no interactions during training. As shown in Table III, CAT consistently performs better than baseline models across both tasks, regardless of data split ratios. The Linear baseline, though surpassing some homogeneous static models, lags far behind CAT and other baselines. These results highlight the necessity of incorporating heterogeneity and dynamicity into the model design. While HGT also considers these two aspects, it lacks support for trust properties and trust-related attributes, making it not fully compatible with trust prediction. Its complex architecture also risks overfitting, resulting in suboptimal performance. We also observe a clear performance decay for all models in task 2 compared with task 1, indicating that task 2 is more difficult. This is reasonable as task 2 requires that the patterns learned by the trust prediction models can be generalized to unobserved users. Despite the challenging nature of task 2, CAT improves MRR by 50.79% and 36.03% under the settings of 70%-15%-15% and 80%-10%-10%, respectively, compared with the best baseline. These findings showcase that CAT effectively handles dynamic graphs and has the potential to address the cold start issue [54], a long-standing issue in trust prediction. Additionally, we find that HAN performs worse than Medley and TrustGuard in task 1 but outperforms it in task 2. This suggests that modeling heterogeneity can deal with sparse trust relationships between unobserved users.

For the Ciao and CiaoDVD datasets, where timestamps of trust relationships are missing, we evaluate the stability of CAT and baselines by varying data split ratios. Table IV shows that CAT consistently outperforms baselines across different split ratios, with average MRR improvements of 38.67% on Ciao and 11.08% on CiaoDVD over the best baseline. Among the baselines, HGT performs best, highlighting the significance of modeling both dynamicity and heterogeneity. Notably, heterogeneous GNNs show superiority over homogeneous GNNs, suggesting that heterogeneous graphs offer rich information for accurate trust prediction. Furthermore, the Linear baseline shows substantial performance gaps com-

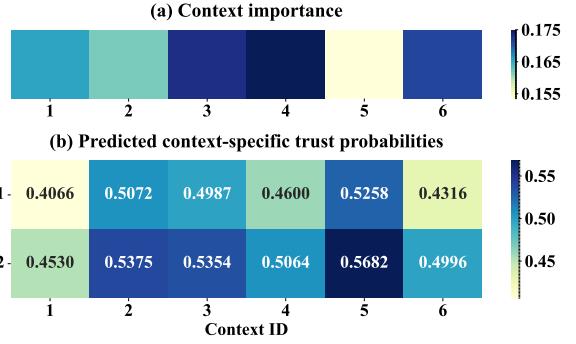


Fig. 3: Results of context-aware trust prediction on the Ciao dataset. Case 1 illustrates a user pair $\langle v_i, v_j \rangle$ with an overall distrusted relationship, whereas case 2 depicts a user pair with an overall trusted relationship.

pared with most GNN models. This may be attributed to its overly simple architecture, which tends to underfit in dynamic heterogeneous networks. In contrast, CAT effectively captures diverse semantics through its advanced design, making it well-suited for trust prediction in such complex networks.

Context-Aware Trust and a Real-World Use Case.

Among existing GNN-based trust prediction models, CAT is unique in its ability to predict context-aware trust. Since other models lack this capability and current datasets do not provide context-specific trust labels, a direct quantitative comparison is not feasible. Instead, we present case studies to show how CAT outputs context-aware trust levels and their practical significance. First, CAT can explain how different contexts influence an overall trust level and provide a detailed trust probability for each context. Fig. 3(a) shows the importance of each context learned by the context-aware aggregator, revealing that different contexts contribute differently to trust establishment. Fig. 3(b) further illustrates that trust probabilities can vary notably for the same user pair across contexts, which aligns with the definition of context-awareness [1]. For instance, in case 1, user v_i generally distrusts v_j , yet exhibits trust

TABLE V: Ablation studies on CAT.

CAT and variants	Epinions: Observed users			Epinions: Unobserved users			Ciao		
	MRR	AP	AUC	MRR	AP	AUC	MRR	AP	AUC
CAT	0.6025	0.9383	0.9677	0.4082	0.9527	0.8933	0.4150	0.9234	0.9327
w/o Time Embedding	0.5575	0.9240	0.9597	0.2441	0.9040	0.8045	0.3702	0.9050	0.9142
w/o Type Attention	0.5941	0.9368	0.9677	0.3033	0.9393	0.8776	0.4035	0.9216	0.9322
w/o Node Attention	0.5781	0.9328	0.9653	0.3975	0.9501	0.8851	0.4042	0.9194	0.9293
w/o Ca Meta-path	0.5763	0.9274	0.9600	0.2924	0.9328	0.8592	0.3959	0.9157	0.9256
w/o Ca Aggregator	0.5642	0.9285	0.9636	0.3859	0.9480	0.8817	0.3993	0.9196	0.9302

in contexts 2 and 5 (trust probability > 0.5). Similarly, in case 2, v_i generally trusts v_j and has high trust probabilities in several contexts, but still distrusts v_j in contexts 1 and 6. These examples highlight the importance of modeling context-awareness. However, baseline models that predict only overall trust fail to capture such nuanced contextual distinctions.

Second, CAT’s context-awareness enables flexible trust-based applications, such as fraud detection on social platforms, where users may behave normally in some contexts but maliciously in others (e.g., posting benign reviews for electronics but fake reviews for clothing). CAT models the social network as a contextual trust graph and predicts trust for each user pair within each context. As shown in Fig. 3(b), when CAT’s predicted trust probability for v_j falls below a threshold in certain contexts, this indicates potential fraudulent behavior in those contexts. Such signals allow the platform to warn v_i or restrict v_j ’s activities there. The threshold can be customized based on individual trust preferences, thus providing flexibility.

C. Ablation Study (RQ2)

In this subsection, we investigate the necessity and rationality of each component in CAT using a fixed 70%-15%-15% split ratio across Epinions and Ciao datasets, which have the largest and smallest sizes, respectively.

To begin with, we construct five CAT variants by removing key components: (i) CAT w/o Time Embedding, which ignores temporal information and involves changes in both embedding and heterogeneous attention layers. (ii) CAT w/o Type Attention, which assumes equal importance across node types. (iii) CAT w/o Node Attention, which assumes equal importance among nodes within the same type. (iv) CAT w/o Ca Meta-path, where user and item embeddings are initialized with random vectors instead of using the context-aware meta-path. (v) CAT w/o Ca Aggregator, which replaces the context-aware aggregator with a mean aggregator that averages trust across all contexts to produce overall trust.

It can be concluded from Table V that: (i) CAT is consistently better than its five variants, validating the importance of integrating trust dynamicity, type and node importance under graph heterogeneity, and context-awareness. This joint consideration is a unique contribution of our approach, making CAT effective for trust prediction and applicable to real-world networks. (ii) Time embeddings and context-aware meta-paths contribute the most to model performance, indicating that basic trust properties are crucial for accurate trust prediction. (iii) Neither “w/o Node Attention” nor “w/o Type Attention”

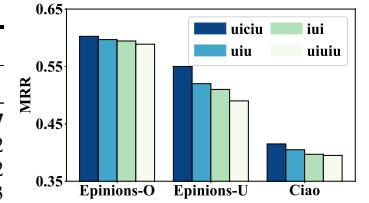


Fig. 4: Comparison of different meta-paths.

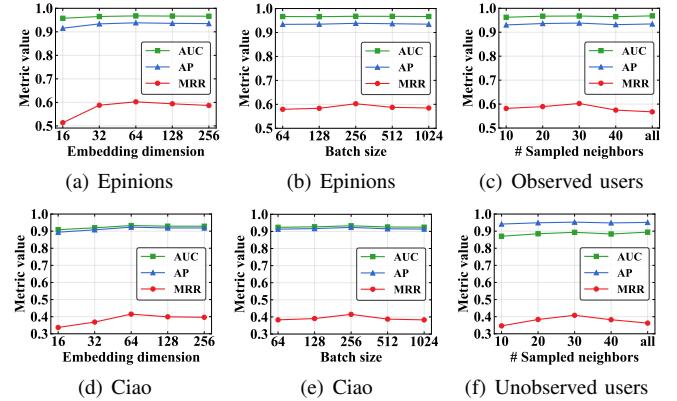


Fig. 5: Effects of embedding dimension, batch size, and number of sampled neighbors on model performance.

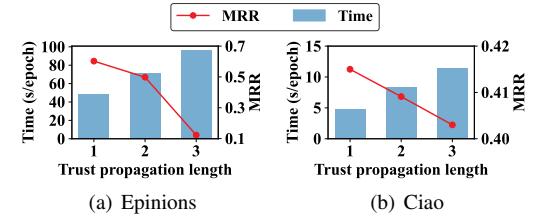


Fig. 6: Effect of trust propagation length on model performance.

is competitive to CAT, highlighting the significance of node distinctions and semantic information in forming effective embeddings. (iv) Removing the context-aware aggregator reduces performance, confirming that different contexts influence trust differently. Effectively modeling such contextual differences can improve overall trust prediction.

We further justify the introduction of the new context-aware meta-path *uiciu* in Fig. 4. The results show that *uiciu* has superiority over other meta-paths thanks to its incorporation of contextual information. Additionally, *uiuiu*, which captures high-hop trust information, is less effective than *uiu*, suggesting that irrelevant information from distant connections may undermine trust propagation and aggregation.

D. Hyperparameter Analysis (RQ3)

In this subsection, we analyze the impact of embedding dimension, batch size, number of sampled neighbors, and trust propagation length on CAT’s performance using a fixed 70%-15%-15% split ratio across Epinions and Ciao.

Embedding Dimension. As shown in Fig. 5(a) and Fig. 5(d), CAT’s performance improves as the embedding

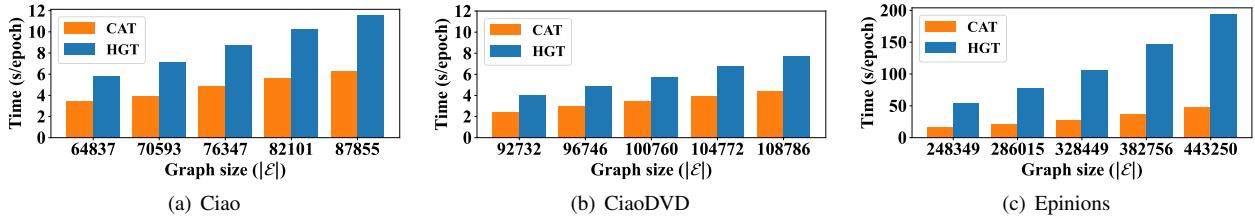


Fig. 7: Running time of CAT and HGT on various graph sizes.

dimension increases. However, excessively increasing the embedding dimension results in a performance decline. This outcome is as expected since a small embedding dimension limits the model’s learning ability, while a large dimension increases the risk of overfitting.

Batch Size. A real-world network is often extremely large, making it impractical to train a GNN-based trust prediction model by directly inputting the entire network. A solution is to divide the large network into subgraphs and process several subgraphs at a time. The batch size, referring to the number of subgraphs processed together, may impact the model’s performance. A small batch size results in a partial view of node behaviors, leading to reduced prediction accuracy. Instead, a large batch size enhances the model performance by covering more information. Fig. 5(b) and Fig. 5(e) show that CAT is not sensitive to batch size.

Number of Sampled Neighbors. The number of sampled neighbors determines subgraph size and affects scalability. We study its impact on model performance using the Epinions dataset, since the smaller Ciao and CiaoDVD datasets do not incur significant computational costs when all neighbors are used. As shown in Fig. 5(c) and Fig. 5(f), model performance initially improves with more sampled neighbors but eventually decreases. This is reasonable because sampling too few neighbors leads to information loss, whereas sampling too many neighbors may introduce noise that degrades performance.

Trust Propagation Length. Trust propagation length refers to the number of neighbor hops used to form a node’s embedding. As we can see in Fig. 6(a) and Fig. 6(b), CAT’s performance significantly decreases while running time increases with longer propagation lengths, differing from prior findings [2], [9] where 2 or 3 hops are preferred. The reasons are two-fold: (i) From a data perspective, we focus on a heterogeneous graph while those studies use homogeneous ones. Although the heterogeneous graph provides rich semantics, it also introduces noise. Compared with one-hop neighbors, high-hop neighbors are less crucial and may mislead model training [43]. (ii) From a model perspective, the embeddings initialized by Metapath2vec already cover high-hop information. Thus, there is no need to stack multiple heterogeneous attention layers for trust propagation. Despite using one-hop neighbors, CAT still has high accuracy (see Section V-B).

E. Scalability Analysis (RQ4)

In this subsection, we evaluate CAT’s scalability by comparing its running time (i.e., average training time per epoch) with

that of HGT, the best baseline that also considers dynamicity and heterogeneity, across varying graph sizes. A detailed time complexity analysis of CAT is provided in Appendix A.

Fig. 7 illustrates that as the graph size increases, the running time of both models increases for all datasets. Notably, CAT achieves average reductions in running time of 44.39% and 40.79% on the Ciao and CiaoDVD datasets, respectively, compared to HGT. Both models run faster on CiaoDVD than on Ciao because Ciao contains more trust relationships, which are the target links for prediction. On the Epinions dataset, CAT exhibits a more pronounced efficiency advantage, requiring about 48s per epoch for training on a graph with 443K edges, compared to 194s for HGT, yielding an average reduction of 73.97% across different graph sizes. This high efficiency stems from CAT’s one-hop trust propagation and recent-time neighbor sampling strategies. The former restricts trust propagation to immediate neighbors, reducing computation while maintaining high accuracy (see Fig. 6). The latter focuses on limited yet crucial interactions. For comparison, considering interactions with all neighbors makes CAT nearly three times slower. Together, these strategies greatly reduce the number of messages to be propagated and improve scalability, making CAT suitable for practical deployment.

F. Robustness Analysis (RQ5)

In this subsection, we evaluate the robustness of CAT and two state-of-the-art GNN-based trust prediction models (i.e., Medley [7] and TrustGuard [2]) against data poisoning attacks. Specifically, we follow TrustGuard [2] to construct trust-oriented attacks and T-Spear [29], the only recent data poisoning attack targeting continuous-time GNNs, to perform GNN-oriented attacks. Experiments are conducted on the Epinions dataset, which contains timestamped trust relationships and is thus compatible with T-Spear. Due to space constraints, we report only MRR results. Please refer to Appendix C for comprehensive results and detailed attack settings.

As shown in Table VI, all models exhibit performance declines as the perturbation rate increases, but CAT outperforms the baselines under both attacks. For example, under trust-oriented attacks, CAT experiences a maximum performance drop of 0.95% in task ①, compared to 14.34% for Medley and 8.62% for TrustGuard. Similarly, in task ②, CAT’s maximum performance drop is 4.63%, smaller than Medley’s 13.34% and TrustGuard’s 12.49%. We observe that CAT is more robust against trust-oriented attacks than GNN-oriented ones, while Medley and TrustGuard show the opposite trend. This may be

TABLE VI: Robustness comparison under trust-oriented and GNN-oriented attacks (MRR results).

①/②: trust predictions for observed/unobserved users. p : perturbation rate (ratio of added adversarial links to original links). MDR: Maximum Drop Rate (lower is better).

Tasks	Models	Clean	Trust-oriented Attacks					GNN-oriented Attacks				
			$p=5\%$	$p=10\%$	$p=15\%$	$p=20\%$	MDR ↓	$p=5\%$	$p=10\%$	$p=15\%$	$p=20\%$	MDR ↓
①	Medley	0.4762	0.4552	0.4367	0.4155	0.4079	14.34%	0.4650	0.4576	0.4477	0.4393	7.75%
	TrustGuard	0.4955	0.4831	0.4820	0.4565	0.4528	8.62%	0.4813	0.4711	0.4664	0.4721	5.87%
	CAT	0.6025	0.5968	0.6046	0.6144	0.6070	0.95%	0.5999	0.5869	0.5842	0.5821	3.39%
②	Medley	0.1979	0.1894	0.1791	0.1756	0.1715	13.34%	0.1909	0.1890	0.1846	0.1737	12.23%
	TrustGuard	0.2571	0.2465	0.2431	0.2253	0.2250	12.49%	0.2477	0.2296	0.2348	0.2289	10.97%
	CAT	0.4082	0.3988	0.3924	0.3911	0.3893	4.63%	0.4022	0.4029	0.3918	0.3858	5.49%

because the surrogate model used in T-Spear is closer to CAT’s architecture than those of the other two models. Additionally, we note an interesting phenomenon: under task ① of trust-oriented attacks, CAT’s performance initially declines but later improves, even surpassing the clean setting. We analyze this result at a 15% perturbation rate and find that this attack may act as data augmentation for two reasons: (i) It substantially reduces the number of connected components (from 100 to 2) with a small change in the clustering coefficient (0.133 vs. 0.131). This enhances graph connectivity and thus facilitates trust propagation. (ii) Similarity distribution analysis shows a 25.68% overlap between adversarial and original links, implying that some adversarial links are structurally benign.

To summarize, CAT demonstrates superior robustness over existing GNN-based trust prediction models, owing to its strong semantic understanding. By learning from a heterogeneous graph and incorporating advanced designs, CAT gains a comprehensive view of intrinsic interactions and node semantics, thereby increasing the difficulty for attackers to succeed. For instance, T-Spear’s surrogate model overlooks such semantic diversity. To validate this viewpoint, we construct two variants: (i) “Context”, which disables heterogeneity modeling, and (ii) “Hetero.”, which weakens contextual understanding. We also develop a defense mechanism for comparison because of the limited work in this area. Inspired by previous work [2], [55], we employ cosine similarity to filter out suspicious links. Fig. 8 shows the MRR results under GNN-oriented attacks at the perturbation rate of 10%. The results indicate that both heterogeneity and context-awareness are critical to CAT’s robustness, with the ability to handle heterogeneity being particularly significant. While cosine similarity offers slight gains, it remains inferior to CAT.

VI. DISCUSSION

Generality to Different Networks. While we only demonstrated CAT’s effectiveness in social networks, it is also applicable to other real-world networks that can be abstracted into *user*, *item*, and *item category* (i.e., *context*). Examples include financial and employment networks. The reasons are as follows: (i) Trust issues are prevalent in these networks, making trust prediction critical. For instance, in financial networks, fraudulent activities are a serious concern. CAT can analyze complex node interactions to uncover latent trust relationships and mitigate security threats. (ii) These networks

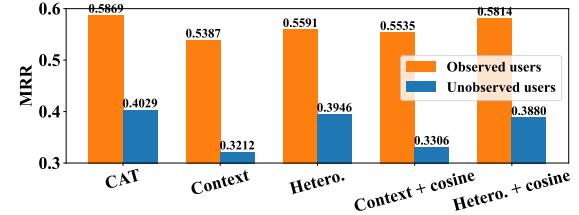


Fig. 8: Detailed evaluation and comparison of robustness against GNN-oriented attacks with 10% perturbation.

are highly dynamic, with interactions evolving over time. CAT is effective at capturing such dynamics to predict time-aware trust relationships. (iii) These networks are heterogeneous in nature, with diverse node and edge types. In financial networks, *user* can be investors or financial advisors, *item* can be stocks or cryptocurrencies, and *context* denotes the type of financial product. CAT can effectively handle such heterogeneity and predict the trust an investor has in an advisor for recommending a specific product type. Furthermore, the proposed context-aware meta-path remains applicable to these networks, reducing the need for additional human effort.

Other Types of Adversarial Attacks. We evaluated CAT’s robustness against two data poisoning attacks in Section V-F. Here, we further discuss its vulnerabilities to adaptive attacks and evasion attacks. First, consider an adaptive attacker who is aware that CAT derives its robustness from contextual differentiation and aims to blur contextual boundaries, making the model misinterpret distinct contexts as similar. Specifically, this adaptive data poisoning attack allows colluding attackers to manipulate interactions across contexts, e.g., by injecting malicious user-item interactions with timestamps crafted to mimic normal behavior. We evaluate CAT against this strong attack in Appendix D. The results show that while this attack is more disruptive than trust-oriented and GNN-oriented attacks, CAT remains robust, particularly when predicting trust for observed users. This robustness stems from its context-aware aggregation mechanism, which adaptively weights contexts to mitigate noise and preserve contextual distinctions.

Second, evasion attacks can cause a model to misclassify a pair’s trust relationship. Edge filtering [56] based on Jaccard or cosine similarities, as well as low-rank based defenses [57], [58], can be employed to filter out suspicious links before they are fed into CAT for prediction. Additionally, adversarial training [59] improves model robustness by augmenting the

training data with adversarial examples. In summary, these defenses can be easily integrated into CAT to further enhance its resilience against various adversarial attacks.

Diverse Trust Prediction Tasks. Beyond predicting trust relationships between trustor-trustee pairs, which corresponds to edge classification or link prediction in the field of GNNs, CAT also supports node-level trust prediction (i.e., node classification) and community-level trust prediction (i.e., graph classification). For node-level trust prediction, CAT’s prediction layer takes a node’s embedding along with a context embedding as input and outputs the node’s trust within the given context. For community-level trust prediction, it takes a group of nodes’ embeddings along with a context embedding as input and outputs the group’s trust within that context. This support for diverse tasks makes CAT flexible in practice.

Limitations and Future Work. While CAT extracts rich semantics through Metapath2vec and a dual attention mechanism, it currently overlooks textual attributes, such as user reviews and profiles, which may also exist in real-world networks. These texts can reveal user preferences and characteristics that are closely related to trust establishment. We did not include this information primarily due to scalability concerns and the limited availability of high-quality textual data in existing datasets. Nevertheless, CAT’s modular design allows easy integration of plug-in text encoders, including lightweight pre-trained models and LLM-based feature extractors, enabling future incorporation of textual information. For example, recent years have witnessed a surge in applying LLMs across various domains owing to their extensive knowledge bases and strong reasoning capabilities [60], [61]. By prompting LLMs to extract crucial information from the texts, we can obtain informative embeddings. This approach offers two advantages: (i) It enables CAT to comprehensively model real-world networks, improving its practicality. (ii) LLM-enhanced node embeddings offer superior robustness against both poisoning and evasion attacks [62], [63], further enhancing CAT’s robustness.

VII. CONCLUSION

In this paper, we proposed CAT, a GNN-based trust prediction model that captures trust dynamicity, represents real-world heterogeneity, supports context-awareness, and is robust against data poisoning attacks. To incorporate contextual and temporal information, we introduced a context-aware metapath and utilized a time encoding function, respectively. A dual attention mechanism identifies the importance of diverse interactions, while recent-time neighbor sampling and one-hop trust propagation strategies enhance scalability. By generating context embeddings and linking context-aware trust with overall trust, CAT can predict both types of trust. Extensive experiments on three real-world datasets demonstrate CAT’s effectiveness, scalability, and robustness.

ACKNOWLEDGMENT

This work is supported in part by the National Natural Science Foundation of China under Grant U23A20300; in

part by the Key Research Project of Shaanxi Natural Science Foundation under Grant 2023-JC-ZD-35; in part by the Concept Verification Funding of Hangzhou Institute of Technology of Xidian University under Grant GNYZ2024XX007; in part by the 111 Center under Grant B16037; and in part by the Fundamental Research Funds for the Central Universities under Grant YJSJ25011.

REFERENCES

- [1] W. Sherchan, S. Nepal, and C. Paris, “A survey of trust in social networks,” *ACM Computing Surveys*, vol. 45, no. 4, pp. 1–33, 2013.
- [2] J. Wang, Z. Yan, J. Lan, E. Bertino, and W. Pedrycz, “Trustguard: Gnn-based robust and explainable trust evaluation with dynamicity support,” *IEEE Transactions on Dependable and Secure Computing*, pp. 1–18, 2024.
- [3] B. Wang, J. Jia, and N. Z. Gong, “Graph-based security and privacy analytics via collective classification with joint weight learning and propagation,” in *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS)*, 2019.
- [4] B. Wang, L. Zhang, and N. Z. Gong, “Syllscar: Sybil detection in online social networks via local rule based propagation,” in *Proceedings of the IEEE Conference on Computer Communications*, 2017, pp. 1–9.
- [5] Y. Liu, W. Zhou, and H. Chen, “Efficiently promoting product online outcome: An iterative rating attack utilizing product and market property,” *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 6, pp. 1444–1457, 2017.
- [6] W. Lin, Z. Gao, and B. Li, “Guardian: Evaluating trust in online social networks with graph convolutional networks,” in *Proceedings of the IEEE Conference on Computer Communications*, 2020, pp. 914–923.
- [7] W. Lin and B. Li, “Medley: Predicting social trust in time-varying online social networks,” in *Proceedings of the IEEE Conference on Computer Communications*, 2021, pp. 1–10.
- [8] N. Jiang, W. Jie, J. Li, X. Liu, and D. Jin, “Gatrust: A multi-aspect graph attention network model for trust assessment in osns,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 6, pp. 5865–5878, 2023.
- [9] C. Huo, D. He, C. Liang, D. Jin, T. Qiu, and L. Wu, “Trustggn: Graph neural network-based trust evaluation via learnable propagative and composable nature,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–13, 2023.
- [10] Z. Yu, D. Jin, C. Huo, Z. Wang, X. Liu, H. Qi, J. Wu, and L. Wu, “Kgtrust: Evaluating trustworthiness of siot via knowledge enhanced graph neural networks,” in *Proceedings of the ACM Web Conference*, 2023, pp. 727–736.
- [11] J. Wen, N. Jiang, J. Li, X. Liu, H. Chen, Y. Ren, Z. Yuan, and Z. Tu, “Dtrust: Toward dynamic trust levels assessment in time-varying online social networks,” in *Proceedings of the IEEE Conference on Computer Communications*, 2023, pp. 1–10.
- [12] F. Yang, J. Xu, C. Xiong, Z. Li, and K. Zhang, “Prographer: An anomaly detection system based on provenance graph embedding,” in *Proceedings of the 32nd USENIX Security Symposium (USENIX Security)*, 2023, pp. 4355–4372.
- [13] M. U. Rehman, H. Ahmadi, and W. U. Hassan, “Flash: A comprehensive approach to intrusion detection via provenance graph representation learning,” in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2024, pp. 1–19.
- [14] K. Hsieh, M. Wong, S. Segarra, S. K. Mani, T. Eberl, A. Panasyuk, R. Netravali, R. Chandra, and S. Kandula, “{NetVigil}: Robust and {Low-Cost} anomaly detection for {East-West} data center security,” in *Proceedings of the 21st USENIX Symposium on Networked Systems Design and Implementation*, 2024, pp. 1771–1789.
- [15] W. Fan, Y. Ma, Q. Li, Y. He, E. Zhao, J. Tang, and D. Yin, “Graph neural networks for social recommendation,” in *Proceedings of the World Wide Web Conference*, 2019, pp. 417–426.
- [16] L. Yang, Z. Liu, Y. Dou, J. Ma, and P. S. Yu, “Consisrec: Enhancing gnn for social recommendation via consistent neighbor aggregation,” in *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021, pp. 2141–2145.

[17] Y.-H. Chen, S.-C. Lin, S.-C. Huang, C.-L. Lei, and C.-Y. Huang, “Guided malware sample analysis based on graph neural networks,” *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 4128–4143, 2023.

[18] J. Yan, G. Yan, and D. Jin, “Classifying malware represented as control flow graphs using deep graph convolutional neural network,” in *Proceedings of the 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2019, pp. 52–63.

[19] I.-R. Chen, F. Bao, and J. Guo, “Trust-based service management for social internet of things systems,” *IEEE Transactions on Dependable and Secure Computing*, vol. 13, no. 6, pp. 684–696, 2016.

[20] G. Liu, Q. Yang, H. Wang, and A. X. Liu, “Trust assessment in online social networks,” *IEEE Transactions on Dependable and Secure Computing*, vol. 18, no. 2, pp. 994–1007, 2019.

[21] G. Liu, C. Li, and Q. Yang, “Neuralwalk: Trust assessment in online social networks with neural networks,” in *Proceedings of the IEEE Conference on Computer Communications*, 2019, pp. 1999–2007.

[22] J. Khoury, D. Klisura, H. Zanddizari, G. D. L. T. Parra, P. Najafirad, and E. Bou-Harb, “Jbeil: Temporal graph-based inductive learning to infer lateral movement in evolving enterprise networks,” in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2024, pp. 1–17.

[23] X. Han, T. Pasquier, A. Bates, J. Mickens, and M. Seltzer, “Unicorn: Runtime provenance-based detector for advanced persistent threats,” in *Proceedings of the 27th Annual Network and Distributed System Security Symposium (NDSS)*, 2020.

[24] Q. Wang, W. Zhao, J. Yang, J. Wu, S. Xue, Q. Xing, and P. S. Yu, “C-deeptrust: A context-aware deep trust prediction model in online social networks,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 6, pp. 2767–2780, 2023.

[25] J. Tang, H. Gao, and H. Liu, “mtrust: Discerning multi-faceted trust in a connected world,” in *Proceedings of the fifth ACM International Conference on Web Search and Data Mining*, 2012, pp. 93–102.

[26] T. Cui, G. Gou, G. Xiong, Z. Li, M. Cui, and C. Liu, “Siamhan: Ipv6 address correlation attacks on tls encrypted traffic via siamese heterogeneous graph attention network,” in *Proceedings of the 30th USENIX Security Symposium (USENIX Security)*, 2021, pp. 4329–4346.

[27] J. Tang and H. Liu, *Trust in social media*. Morgan & Claypool Publishers, 2015.

[28] G. Guo, J. Zhang, D. Thalmann, and N. Yorke-Smith, “Etaf: An extended trust antecedents framework for trust prediction,” in *Proceedings of the International Conference on Advances in Social Networks Analysis and Mining*, 2014, pp. 540–547.

[29] D. Lee, J. Lee, and K. Shin, “Spear and shield: Adversarial attacks and defense methods for model-based link prediction on continuous-time dynamic graphs,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024, pp. 13374–13382.

[30] E. Rossi, B. Chamberlain, F. Frasca, D. Eynard, F. Monti, and M. Bronstein, “Temporal graph networks for deep learning on dynamic graphs,” in *ICML 2020 Workshop on Graph Representation Learning*, 2020.

[31] X. Wang, H. Ji, C. Shi, B. Wang, Y. Ye, P. Cui, and P. S. Yu, “Heterogeneous graph attention network,” in *Proceedings of the World Wide Web Conference*, 2019, pp. 2022–2032.

[32] Z. Zhang, Z. Zhang, X. Wang, Y. Qin, Z. Qin, and W. Zhu, “Dynamic heterogeneous graph attention neural architecture search,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023, pp. 11307–11315.

[33] H. Xue, L. Yang, W. Jiang, Y. Wei, Y. Hu, and Y. Lin, “Modeling dynamic heterogeneous network for link prediction using hierarchical attention with temporal rnn,” in *Proceedings of Machine Learning and Knowledge Discovery in Databases: European Conference*, 2021, pp. 282–298.

[34] Y. Fan, M. Ju, C. Zhang, and Y. Ye, “Heterogeneous temporal graph neural network,” in *Proceedings of the 2022 SIAM International Conference on Data Mining*, 2022, pp. 657–665.

[35] Z. Hu, Y. Dong, K. Wang, and Y. Sun, “Heterogeneous graph transformer,” in *Proceedings of the Web Conference*, 2020, pp. 2704–2710.

[36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[37] J. Liu, L. Song, G. Wang, and X. Shang, “Meta-hgt: Metapath-aware hypergraph transformer for heterogeneous information network embedding,” *Neural Networks*, vol. 157, pp. 65–76, 2023.

[38] Y. Dong, N. V. Chawla, and A. Swami, “Metapath2vec: Scalable representation learning for heterogeneous networks,” in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, pp. 135–144.

[39] I. J. King and H. H. Huang, “Euler: Detecting network lateral movement via scalable temporal graph link prediction,” in *Proceedings of the 29th Annual Network and Distributed System Security Symposium (NDSS)*, 2022.

[40] S. M. Ghafari, A. Joshi, A. Beheshti, C. Paris, S. Yakhchi, and M. Orgun, “Dcat: A deep context-aware trust prediction approach for online social networks,” in *Proceedings of the 17th International Conference on Advances in Mobile Computing & Multimedia*, 2019, pp. 20–27.

[41] P. Li, H. Yu, and X. Luo, “Context-aware graph neural network for graph-based fraud detection with extremely limited labels,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 11, 2025, pp. 12112–12120.

[42] D. Xu, C. Ruan, E. Korpeoglu, S. Kumar, and K. Achan, “Inductive representation learning on temporal graphs,” in *Proceedings of the 8th International Conference on Learning Representations*, 2020.

[43] W. Cong, S. Zhang, J. Kang, B. Yuan, H. Wu, X. Zhou, H. Tong, and M. Mahdavi, “Do we really need complicated model architectures for temporal networks?” in *Proceedings of the 11th International Conference on Learning Representations*, 2023.

[44] W. Hamilton, Z. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[45] D. Chen, Y. Lin, W. Li, P. Li, J. Zhou, and X. Sun, “Measuring and relieving the over-smoothing problem for graph neural networks from the topological view,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 3438–3445.

[46] A. Jøsang, R. Ismail, and C. Boyd, “A survey of trust and reputation systems for online service provision,” *Decision Support Systems*, vol. 43, no. 2, pp. 618–644, 2007.

[47] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proceedings of the 3rd International Conference on Learning Representations*, 2015.

[48] Z. Jia, Y. Xiong, Y. Nan, Y. Zhang, J. Zhao, and M. Wen, “Magic: Detecting advanced persistent threats via masked graph representation learning,” in *Proceedings of the 33rd USENIX Security Symposium (USENIX Security)*, 2024, pp. 1–18.

[49] T. Bilot, B. Jiang, Z. Li, N. El Madhoun, K. Al Agha, A. Zouaoui, and T. Pasquier, “Sometimes simpler is better: A comprehensive analysis of {State-of-the-Art}{Provenance-Based} intrusion detection systems,” in *34th USENIX Security Symposium (USENIX Security 25)*, 2025, pp. 7193–7212.

[50] T. N. Kipf and M. Welling, “Variational graph auto-encoders,” *arXiv preprint arXiv:1611.07308*, 2016.

[51] M. Zhang and Y. Chen, “Link prediction based on graph neural networks,” *Advances in Neural Information Processing Systems*, vol. 31, 2018.

[52] J. Li, H. Shomer, H. Mao, S. Zeng, Y. Ma, N. Shah, J. Tang, and D. Yin, “Evaluating graph neural networks for link prediction: Current pitfalls and new benchmarking,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[53] J. Lou, X. Yuan, R. Zhang, X. Yuan, N. Gong, and N.-F. Tzeng, “Grid: Protecting training graph from link stealing attacks on gnn models,” in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2025, pp. 1–19.

[54] J. Wang, X. Jing, Z. Yan, Y. Fu, W. Pedrycz, and L. T. Yang, “A survey on trust evaluation based on machine learning,” *ACM Computing Surveys*, vol. 53, no. 5, pp. 1–36, 2020.

[55] X. Zhang and M. Zitnik, “Gnnguard: Defending graph neural networks against adversarial attacks,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 9263–9275, 2020.

[56] H. Wu, C. Wang, Y. Tyshetskiy, A. Docherty, K. Lu, and L. Zhu, “Adversarial examples for graph data: Deep insights into attack and defense,” in *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 2019, pp. 4816–4823.

[57] N. Entezari, S. A. Al-Sayouri, A. Darvishzadeh, and E. E. Papalexakis, “All you need is low (rank) defending against adversarial attacks on graphs,” in *Proceedings of the 13th International Conference on Web Search and Data Mining*, 2020, pp. 169–177.

[58] J. Mu, B. Wang, Q. Li, K. Sun, M. Xu, and Z. Liu, “A hard label black-box adversarial attack against graph neural networks,” in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2021, pp. 108–125.

- [59] F. Tramer and D. Boneh, “Adversarial training and robustness for multiple perturbations,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [60] P. Liu, J. Liu, L. Fu, K. Lu, Y. Xia, X. Zhang, W. Chen, H. Weng, S. Ji, and W. Wang, “Exploring chatgpt’s capabilities on vulnerability management,” in *Proceedings of the 33rd USENIX Security Symposium (USENIX Security)*, 2024, pp. 811–828.
- [61] X. He, X. Bresson, T. Laurent, A. Perold, Y. LeCun, and B. Hooi, “Harnessing explanations: Llm-to-lm interpreter for enhanced text-attributed graph representation learning,” in *Proceedings of the 12th International Conference on Learning Representations*, 2024.
- [62] K. Guo, Z. Liu, Z. Chen, H. Wen, W. Jin, J. Tang, and Y. Chang, “Learning on graphs with large language models (llms): A deep dive into model robustness,” *arXiv preprint arXiv:2407.12068*, 2024.
- [63] Z. Zhang, X. Wang, H. Zhou, Y. Yu, M. Zhang, C. Yang, and C. Shi, “Can large language models improve the adversarial robustness of graph neural networks?” *arXiv preprint arXiv:2408.08685*, 2024.
- [64] W. Xia, M. Lai, C. Shan, Y. Zhang, X. Dai, X. Li, and D. Li, “Explaining temporal graph models through an explorer-navigator framework,” in *Proceedings of the 11th International Conference on Learning Representations*, 2023.
- [65] D. Han, Z. Wang, W. Chen, K. Wang, R. Yu, S. Wang, H. Zhang, Z. Wang, M. Jin, J. Yang, X. Shi, and X. Yin, “Anomaly detection in the open world: Normality shift detection, explanation, and adaptation,” in *Proceedings of the 30th Annual Network and Distributed System Security Symposium (NDSS)*, 2023.
- [66] L. Van der Maaten and G. Hinton, “Visualizing data using t-sne.” *Journal of Machine Learning Research*, vol. 9, no. 11, 2008.

APPENDIX A TIME COMPLEXITY OF CAT

We analyze the time complexity of the heterogeneous attention layer, as it is the dominant component in CAT. Let $|\mathcal{V}|$ be the number of nodes, $|\mathcal{E}|$ the number of edges, $|\mathcal{T}|$ the number of node types, and d the embedding dimension. At first, the time complexities for calculating type attention and node attention are $O(|\mathcal{T}|d + |\mathcal{E}|d)$ and $O(|\mathcal{E}|d)$, respectively. Then, the message aggregation for forming node embeddings, as shown in Eq. 6, costs about $O(|\mathcal{V}|d^2 + |\mathcal{E}|d)$. Therefore, the overall time complexity of CAT is about $O(|\mathcal{V}|d^2 + |\mathcal{E}|d + |\mathcal{T}|d)$. Since we adopt a neighbor sampling strategy, the actual used edges (i.e., interactions) would be smaller than $|\mathcal{E}|$. Additionally, the one-hop trust propagation strategy makes the complexity independent of the number of GNN layers L . Given that $|\mathcal{T}| \ll |\mathcal{E}|$, the total time complexity of CAT is correlated to $|\mathcal{V}|$, $|\mathcal{E}|$, and d .

APPENDIX B DETAILS OF EVALUATION METRICS

For clarity, we define a positive link as a trusted relationship between a trustor-trustee pair and a negative link as a distrusted relationship in the following metrics:

- **MRR:** MRR is the mean of the reciprocal rank over all positive links.

$$\text{MRR} = \frac{1}{N} \sum_{i=1}^N \frac{1}{\text{rank}_i},$$

where N is the total number of positive links, and rank_i is the rank of the i -th positive link according to its prediction score. To calculate MRR, we randomly sample 100 negative links for each positive link and rank the positive link among these negatives.

- **AP:** AP summarizes a precision-recall curve that shows the trade-off between precision and recall for different classification thresholds.

$$\text{AP} = \sum_n (R_n - R_{n-1}) P_n,$$

where P_n and R_n are the precision and recall at the n -th classification threshold.

- **AUC:** AUC measures the likelihood that a positive link is ranked higher than a random negative link.

$$\text{AUC} = \frac{\sum_{i \in \mathcal{S}^1} \sum_{j \in \mathcal{S}^2} \mathbf{1}(\text{rank}_i < \text{rank}_j)}{|\mathcal{S}^1| \cdot |\mathcal{S}^2|},$$

where \mathcal{S}^1 and \mathcal{S}^2 are sets of positive and negative links, respectively. The indicator function $\mathbf{1}$ outputs 1 if $\text{rank}_i < \text{rank}_j$, and 0 otherwise.

APPENDIX C ROBUSTNESS EVALUATION ON TRUST-ORIENTED AND GNN-ORIENTED ATTACKS

Attack Settings. For GNN-oriented data poisoning attacks, to the best of our knowledge, T-Spear [29] is the only recent method targeting continuous-time GNNs. It formulates the attack by considering four constraints to achieve unnoticeability, which has proven effective in several popular continuous-time GNNs. Following the setting of T-Spear, we inject different numbers of adversarial links (i.e., malicious trust relationships) into the training and validation stages by adjusting perturbation rates. After training with poisoned data, we calculate metrics based on the ground truth trust relationships in the test set.

For trust-oriented data poisoning attacks, we follow the setting of TrustGuard [2] to inject different proportions of malicious trust relationships. Since TrustGuard operates as a discrete-time model that does not provide timestamps for injected links, we generate timestamps by sampling from the temporal distribution of the original links. This approach, consistent with T-Spear [29], helps maintain attack stealthiness.

Results. Fig. VII presents the AP and AUC results of CAT and baseline models under two types of attacks. CAT exhibits the highest robustness, with a maximum performance drop of only 2.07%. In contrast, Medley shows limited robustness due to its lack of defense mechanisms, especially in task ②. Notably, all models experience minor performance drops in most cases. The reason is that both attacks simulate realistic threat scenarios, where attackers introduce some constraints to make their attacks stealthy and unnoticeable, thus limiting their attack strength.

APPENDIX D ROBUSTNESS EVALUATION ON ADAPTIVE ATTACKS

Attack Settings. The adaptive attack, referred to as the cross-context bridging attack, aims to blur the boundaries between two distinct contexts by injecting malicious user-item interactions, causing CAT to confuse them. The attack setting is as follows: (i) *Context selection:* We randomly select two distinct item categories from the Epinions dataset as contexts

TABLE VII: Robustness comparison under trust-oriented and GNN-oriented attacks (AP and AUC results).

①/②: trust predictions for observed/unobserved users. p : perturbation rate (ratio of added adversarial links to original links). MDR: Maximum Drop Rate (lower is better).

Metrics	Tasks	Models	Clean	Trust-oriented Attacks						GNN-oriented Attacks					
				$p=5\%$	$p=10\%$	$p=15\%$	$p=20\%$	MDR ↓	$p=5\%$	$p=10\%$	$p=15\%$	$p=20\%$	MDR ↓		
AP	①	Medley	0.8944	0.8862	0.8855	0.8728	0.8613	3.70%	0.8923	0.8838	0.8827	0.8810	1.50%		
		TrustGuard	0.8919	0.8860	0.8855	0.8724	0.8685	2.62%	0.8853	0.8816	0.8826	0.8814	1.18%		
		CAT	0.9383	0.9394	0.9391	0.9415	0.9401	-0.09%	0.9399	0.9350	0.9326	0.9335	0.61%		
	②	Medley	0.8884	0.8714	0.8473	0.8430	0.8328	6.26%	0.8799	0.8758	0.8780	0.8652	2.61%		
		TrustGuard	0.8950	0.8918	0.8918	0.8824	0.8802	1.65%	0.8901	0.8881	0.8898	0.8881	0.77%		
		CAT	0.9527	0.9515	0.9467	0.9461	0.9469	0.69%	0.9521	0.9507	0.9503	0.9499	0.29%		
AUC	①	Medley	0.9440	0.9399	0.9431	0.9350	0.9239	2.13%	0.9428	0.9347	0.9384	0.9364	0.99%		
		TrustGuard	0.9382	0.9340	0.9342	0.9249	0.9210	1.83%	0.9338	0.9329	0.9341	0.9322	0.64%		
		CAT	0.9677	0.9689	0.9673	0.9683	0.9681	0.04%	0.9695	0.9663	0.9648	0.9660	0.30%		
	②	Medley	0.7806	0.7474	0.6941	0.6820	0.6596	15.50%	0.7636	0.7570	0.7619	0.7383	5.42%		
		TrustGuard	0.7312	0.7252	0.7261	0.7079	0.7030	3.86%	0.7218	0.7216	0.7234	0.7216	1.31%		
		CAT	0.8933	0.8898	0.8765	0.8748	0.8776	2.07%	0.8913	0.8867	0.8875	0.8870	0.74%		

TABLE VIII: Robustness comparison under the adaptive attack (MRR results).

①/②: trust predictions for observed/unobserved users. p : perturbation rate (ratio of added adversarial links to original links). MDR: Maximum Drop Rate (lower is better).

Tasks	Clean	CAT					MDR ↓	w/o Context-aware Aggregation				
		$p=5\%$	$p=10\%$	$p=15\%$	$p=20\%$	MDR ↓		$p=5\%$	$p=10\%$	$p=15\%$	$p=20\%$	MDR ↓
①	0.6025	0.5855	0.5801	0.6025	0.5878	3.72%		0.5748	0.5335	0.5654	0.5569	11.45%
②	0.4082	0.3885	0.3754	0.3774	0.3742	8.33%		0.3841	0.3683	0.3763	0.3713	9.77%

c_k and $c_{k'}$. (ii) *Colluders selection*: We choose highly-active users with substantial interactions in c_k as colluders. These users are chosen because they capture the most representative behavioral patterns of c_k . Thus, when they engage extensively in $c_{k'}$, their behavior introduces misleading cross-context signals that make CAT confuse distinct contexts as similar, thereby distorting its learned contextual boundaries. (3) *Attack injection*: Each colluder injects new interactions with randomly chosen items in $c_{k'}$. The timestamps of these injected interactions are uniformly distributed across the dataset's historical time span to mimic normal temporal patterns and avoid detection. The total number of injected links is controlled by a perturbation rate, and a balanced allocation strategy specifies both the number of colluders and the number of links injected by each colluder.

We evaluate CAT's robustness against this strong attack under perturbation rates of 5%, 10%, 15%, and 20%. Since baseline models do not incorporate contextual information, the attack is not applicable to them. Instead, we construct a variant by removing CAT's context-aware aggregation mechanism to assess its contribution to model robustness. We present only MRR results, as AUC and AP exhibit similar trends.

Results. As shown in Table VIII, CAT experiences a maximum performance drop of 3.72% in task ①, compared with 0.95% under trust-oriented attacks and 3.39% under GNN-oriented attacks. For task ②, the drop reaches 8.33%, higher than 4.63% and 5.49% under the two respective attacks. These results indicate that the cross-context bridging attack is more effective than trust-oriented and GNN-oriented attacks, as it directly targets CAT's context-aware design. Nevertheless,

CAT remains robust in task ①, while its performance degrades more noticeably in task ②. This is because the attack makes the semantic boundaries between contexts less sharp, making it more difficult for CAT to learn distinguishable patterns that generalize to unobserved users.

When the context-aware aggregation mechanism is removed, CAT's robustness declines substantially, with the maximum performance drops increasing to 11.45% in task ① and 9.77% in task ②. This confirms the critical role of context-aware aggregation in enhancing model robustness, as it mitigates noise and preserves contextual distinctions by adaptively weighting contexts.

APPENDIX E COMPARISON OF TWO NEIGHBOR SAMPLING STRATEGIES

CAT adopts a recent-time neighbor sampling strategy to enhance scalability. In this section, we compare this strategy with the uniform neighbor sampling strategy, which uniformly samples nodes and is commonly used in static graph analysis [44]. The results in Table IX are based on the Epinions dataset using a 70%-15%-15% split ratio. From the table, we can see that recent-time neighbor sampling is more effective and efficient than uniform neighbor sampling, indicating that recent interactions are more informative and important than historical ones in dynamic graphs.

APPENDIX F COMPARISON OF CONTEXT EMBEDDING GENERATORS

In addition to the average generator used in CAT, we develop two alternative generators for context embeddings: (i) Long Short-Term Memory (LSTM) generator: An LSTM

TABLE IX: Comparison of uniform and recent-time neighbor sampling strategies.

①/②: trust predictions for observed/unobserved users.

Strategies	Tasks	MRR	AP	AUC	Time (s/epoch)
Uniform	①	0.5907	0.9368	0.9678	52.73
	②	0.3735	0.9487	0.8861	
Recent-time (CAT)	①	0.6025	0.9383	0.9677	48.04
	②	0.4082	0.9527	0.8933	

TABLE X: Comparison of different generators for context embeddings.

①/②: trust predictions for observed/unobserved users.

Generators	Tasks	MRR	AP	AUC	Time (s/epoch)
LSTM	①	0.5191	0.9168	0.9608	405.08
	②	0.3510	0.9458	0.8824	
MLP	①	0.4794	0.9091	0.9586	77.47
	②	0.3255	0.9448	0.8841	
Avg (CAT)	①	0.6025	0.9383	0.9677	48.04
	②	0.4082	0.9527	0.8933	

is used to generate a context embedding by processing an unordered set of item embeddings within the same context. (ii) MLP generator: An MLP aggregates item embeddings with the same context, followed by an elementwise max-pooling operation to obtain a context embedding. Results in Table X show that our simple, efficient average generator outperforms both alternatives. One possible reason is that the LSTM and MLP architectures introduce unnecessary complexity, increasing the risk of overfitting.

APPENDIX G

VISUALIZATION OF DUAL ATTENTION SCORES

CAT comprises several components, making it difficult to understand how it works and whether its functionality aligns with human expectations. To address this, we visualize two key components of CAT: the dual attention mechanism and the context-awareness aggregation mechanism, to explain its prediction process. Since the latter was investigated in Section V-B, this section focuses on the dual attention mechanism.

As shown in Fig. 9, v_0 represents the target user who interacts with both items and other users. We observe that different weights are assigned to the user type (i.e., $\langle user, user \rangle$ interactions) and the item type (i.e., $\langle item, user \rangle$ interactions), with the user type receiving higher importance. This is reasonable since our task is trust prediction, where trust relationships between user pairs directly affect the outcomes. Additionally, the node attention scores vary across nodes, and this variation is particularly evident within the item type. For example, the importance of v_6 to v_0 is 0.37, whereas the importance of v_4 to v_0 is only 0.31. Overall, this visualization suggests that the dual attention mechanism can effectively handle heterogeneous graphs by distinguishing both interaction types and individual nodes. However, the granularity of explanations could be further improved by investigating specific explanation

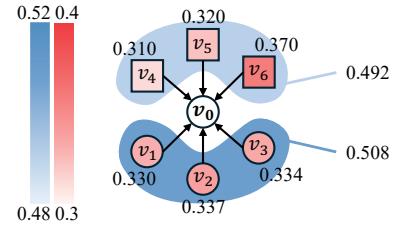


Fig. 9: Visualization of the dual attention scores, including type attention scores (shown in blue) and node attention scores (shown in red). $v_0 \sim v_3$ denote users and $v_4 \sim v_6$ denote items, with v_0 being the target node.

