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





Jie Wang[1] Zheng Yan[1] Jiahe Lan[1] Xuyan Li[1] Elisa Bertino[2]

1 XIDIAN UNIVERSITY 2 PURDUE UNIVERSITY

Network and Distributed System Security (NDSS) Symposium, 25 Feb 2026

Background

- Trust: A complex and multifaceted concept
  - A subjective belief held by one entity (trustor) towards another (trustee)
  - Subjectivity, dynamicity, context-awareness, asymmetry, conditional transitivity, etc
- Trust Evaluation
  - Quantify trust by considering the factors that affect trust

Key Applications

Fraud detection

Intrusion detection

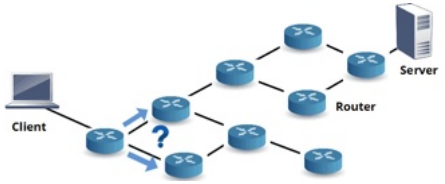

Trustworthy routing

Access control

Industry & Standards

Trustworthiness 6G White Paper

The 6G network will integrate various capabilities such as communication, sensing, computing, and intelligence, making it necessary to redefine the network architecture. The novel network architecture should support native trustworthiness and can be flexibly adapted for tasks such as collaborative sensing and distributed learning to proliferate AI applications on a large scale. Data, as well as the knowledge and intelligence derived from it, is the driving force behind 6G network architecture redesign, wherein new features will be developed to enable E2E native trustworthiness. These include new data governance architectures supporting data compliance and monetization, as well as advanced privacy protection and quantum attack defense technologies.

6.1 A conceptual model of trustworthy networking ITU-T Y.3053

In order to make networks trustable in heterogeneous communication environments, fundamental features beyond the secure communications are needed. First, an identity of a network element should be well defined (identification). Then it is necessary to check whether the identified element is trustworthy (trust evaluation). Finally, trustworthy communication between the peer network elements should be provided (trustworthy communication).

Background

- ❑ Taxonomy of trust evaluation methods
  - Statistics-based, inference-based, and learning-based (trust prediction)
- ❑ Graph Neural Network (GNN): A ML paradigm for graph-structured data
  - ✓ Trust relationships can be naturally modeled as graphs
  - ✓ GNN's message passing is compatible with basic trust properties
  - ✓ End-to-end evaluation manner

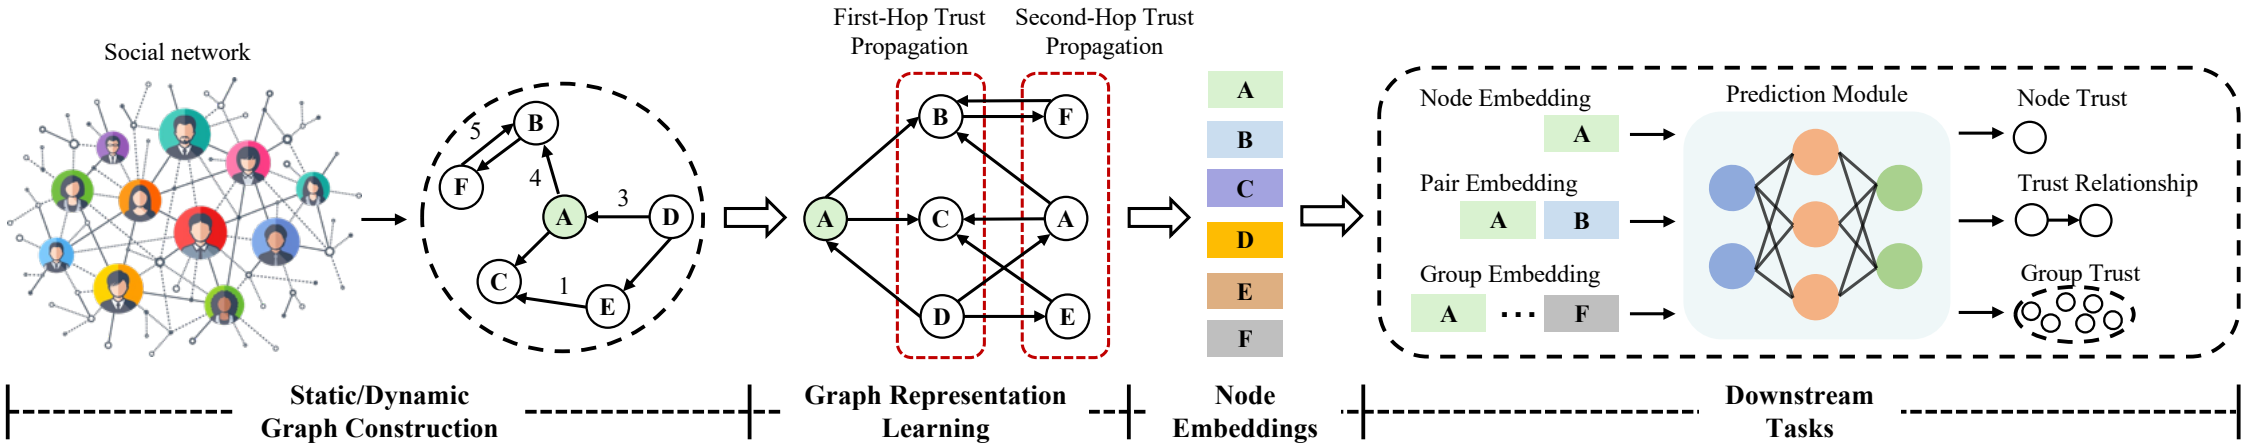

Motivations

TABLE I: Comparison of existing GNN-based trust prediction models.
●: it satisfies a criterion; ○: it does not satisfy a criterion; ◐: it partially satisfies a criterion.

Models	Dynamicity	Heterogeneity	Context-Awareness	Robustness
Guardian [INFOCOM'20]	○	○	○	○
Medley [INFOCOM'21]	●	○	○	○
GATrust [TKDE'23]	○	○	○	○
TrustGNN [TNNLS'24]	○	○	○	○
KGTrust [WWW'23]	○	●	○	○
DTrust [INFOCOM'23]	●	○	○	○
TrustGuard [TDSC'24]	●	○	○	◐

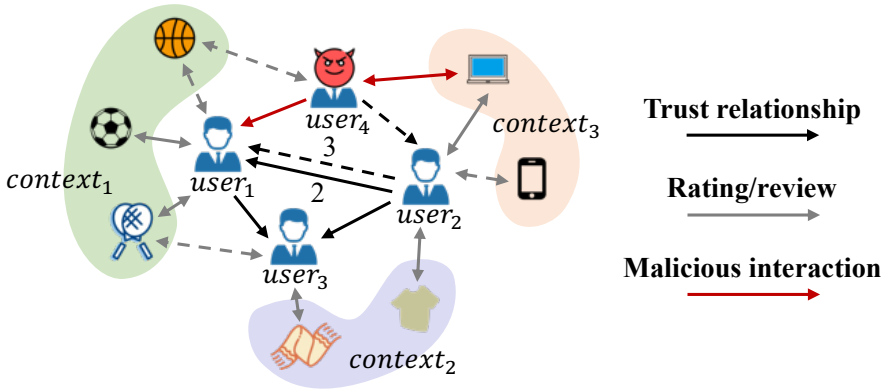

Dashed and solid arrows represent interactions at different time points

Motivations

TABLE I: Comparison of existing GNN-based trust prediction models.
●: it satisfies a criterion; ○: it does not satisfy a criterion; ◐: it partially satisfies a criterion.

Models	Dynamicity	Heterogeneity	Context-Awareness	Robustness
Guardian [INFOCOM'20]	○	○	○	○
Medley [INFOCOM'21]	●	○	○	○
GATrust [TKDE'23]	○	○	○	○
TrustGNN [TNNLS'24]	○	○	○	○
KGTrust [WWW'23]	○	●	○	○
DTrust [INFOCOM'23]	●	○	○	○
TrustGuard [TDSC'24]	●	○	○	◐

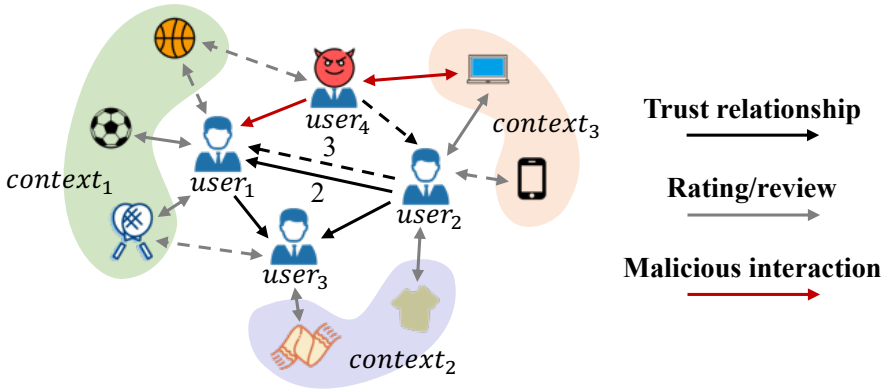

Dashed and solid arrows represent interactions at different time points

Dynamicity:

- User interactions evolve over time, and trust relationships often change accordingly
- Static models fail to capture such dynamics and may violate temporal causality

Static Model *(Random Split)*

Training Set	Testing Set
t_2, t_5, t_4	t_1, t_3

Using the future to predict the past

Dynamic Model *(Temporal Split)*

Training Set	Testing Set
t_1, t_2, t_3	t_4, t_5

Using the past to predict the future

Motivations

TABLE I: Comparison of existing GNN-based trust prediction models.
●: it satisfies a criterion; ○: it does not satisfy a criterion; ◐: it partially satisfies a criterion.

Models	Dynamicity	Heterogeneity	Context-Awareness	Robustness
Guardian [INFOCOM'20]	○	○	○	○
Medley [INFOCOM'21]	●	○	○	○
GATrust [TKDE'23]	○	○	○	○
TrustGNN [TNNLS'24]	○	○	○	○
KGTrust [WWW'23]	○	●	○	○
DTrust [INFOCOM'23]	●	○	○	○
TrustGuard [TDSC'24]	●	○	○	◐

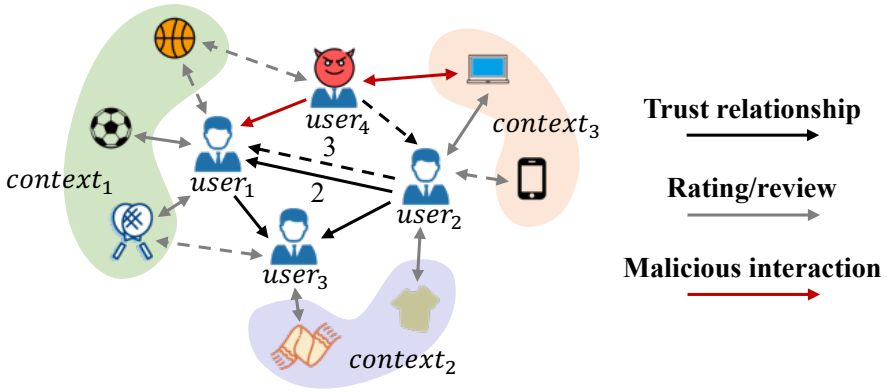

Dashed and solid arrows represent interactions at different time points

Heterogeneity:

- Real-world networks are characterized by different types of nodes and edges
- User–item interactions are semantically rich and denser than trust relationships

TABLE II: Statistics of real-world trust networks.

Feature	Epinions	FilmTrust	Flixster	Ciao	Douban
# users	40,163	1,508	5,213	7,375	129,490
# items	139,738	2,071	18,197	99,746	58,541
# ratings	664,824	35,497	409,803	280,391	16,830,839
rating density	**0.051%**	**1.14%**	**0.04%**	0.03%	**0.222%**
# trustors	33,960	609	47,029	6,792	129,490
# trustees	49,289	732	47,029	7,297	129,490
# trust relations	487,183	1,853	655,054	111,781	1,692,952
trust density	0.029%	0.42%	0.03%	**0.23%**	0.010%

Almost all models focus on homogeneous user-to-user networks

Motivations

TABLE I: Comparison of existing GNN-based trust prediction models.
●: it satisfies a criterion; ○: it does not satisfy a criterion; ◐: it partially satisfies a criterion.

Models	Dynamicity	Heterogeneity	Context-Awareness	Robustness
Guardian [INFOCOM'20]	○	○	○	○
Medley [INFOCOM'21]	●	○	○	○
GATrust [TKDE'23]	○	○	○	○
TrustGNN [TNNLS'24]	○	○	○	○
KGTrust [WWW'23]	○	●	○	○
DTrust [INFOCOM'23]	●	○	○	○
TrustGuard [TDSC'24]	●	○	○	◐

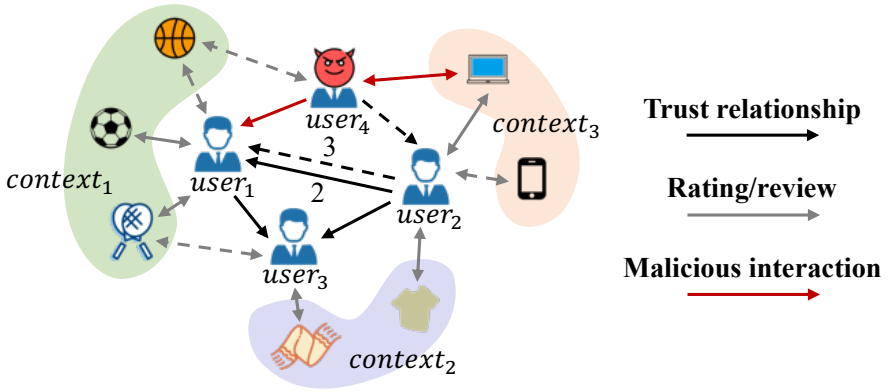

Dashed and solid arrows represent interactions at different time points

Context-Awareness:

- Trust varies across different contexts
- Context refers to any information that describes the specific situations in which a trust relationship is established

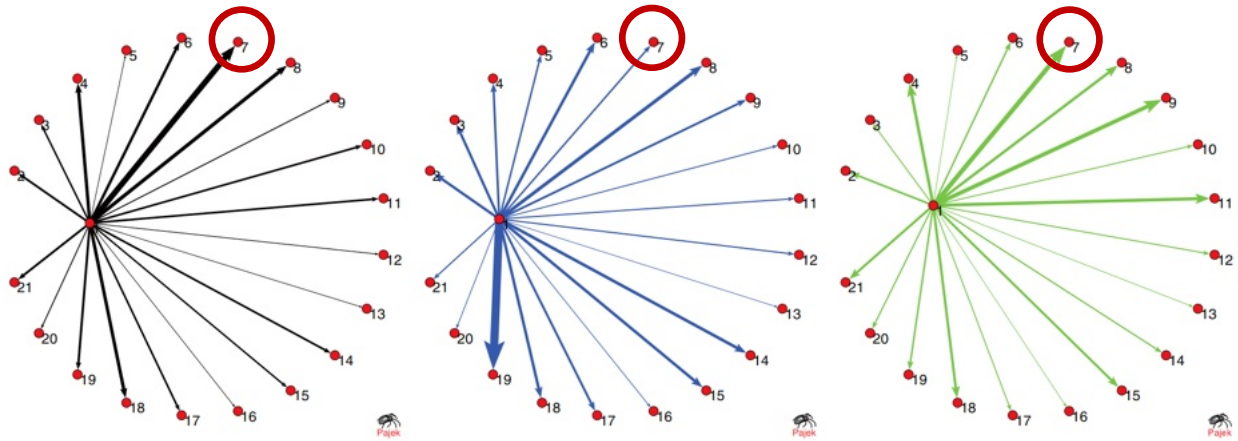

(a) Overall Trust (b) Trust in Home & Garden (c) Trust in Restaurants

Context-aware trust

Motivations

TABLE I: Comparison of existing GNN-based trust prediction models.
●: it satisfies a criterion; ○: it does not satisfy a criterion; ◐: it partially satisfies a criterion.

Models	Dynamicity	Heterogeneity	Context-Awareness	Robustness
Guardian [INFOCOM'20]	○	○	○	○
Medley [INFOCOM'21]	●	○	○	○
GATrust [TKDE'23]	○	○	○	○
TrustGNN [TNNLS'24]	○	○	○	○
KGTrust [WWW'23]	○	●	○	○
DTrust [INFOCOM'23]	●	○	○	○
TrustGuard [TDSC'24]	●	○	○	◐

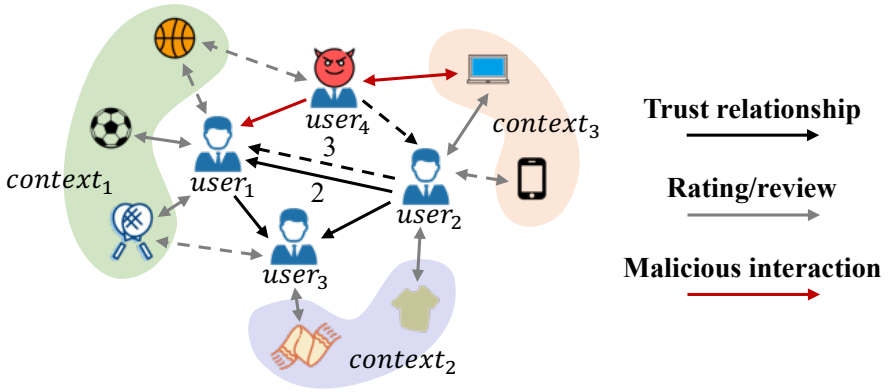

Dashed and solid arrows represent interactions at different time points

Robustness:

- Attackers can inject malicious interactions into training data to mislead model learning
- The attacks may target trust prediction (trust-oriented) and GNN itself (GNN-oriented)

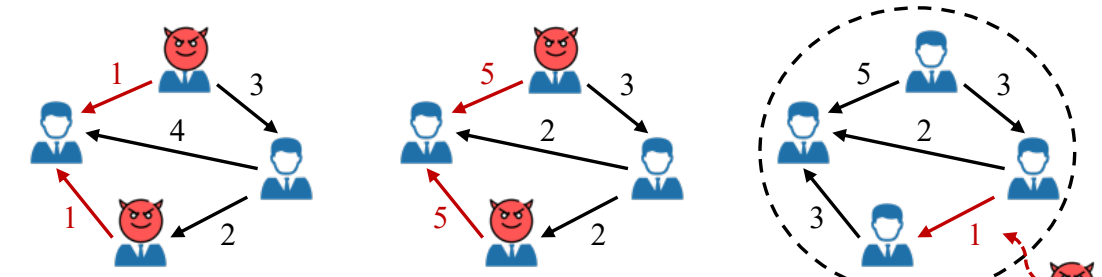

(a) Bad-mouthing Attack (b) Good-mouthing Attack (c) GNN-oriented Attack

Trust-oriented (Local Attack) ⇩ Global Attack

Motivations

Q1: Insufficient modeling of trust dynamicity

Q2: Limited consideration of real-world network heterogeneity

Q3: Lack of support for context-awareness

Q4: Limited attention to model robustness

Motivations

Q1: Insufficient modeling of trust dynamicity

Q2: Limited consideration of real-world network heterogeneity

Q3: Lack of support for context-awareness

Q4: Limited attention to model robustness

Is it possible to design a trust prediction model that can support all the required aspects?

CAT Overview

CAT: A Context-Aware GNN-based Trust prediction model

- Supports trust dynamicity
- Handles real-world network heterogeneity
- Is resilient to data poisoning attacks

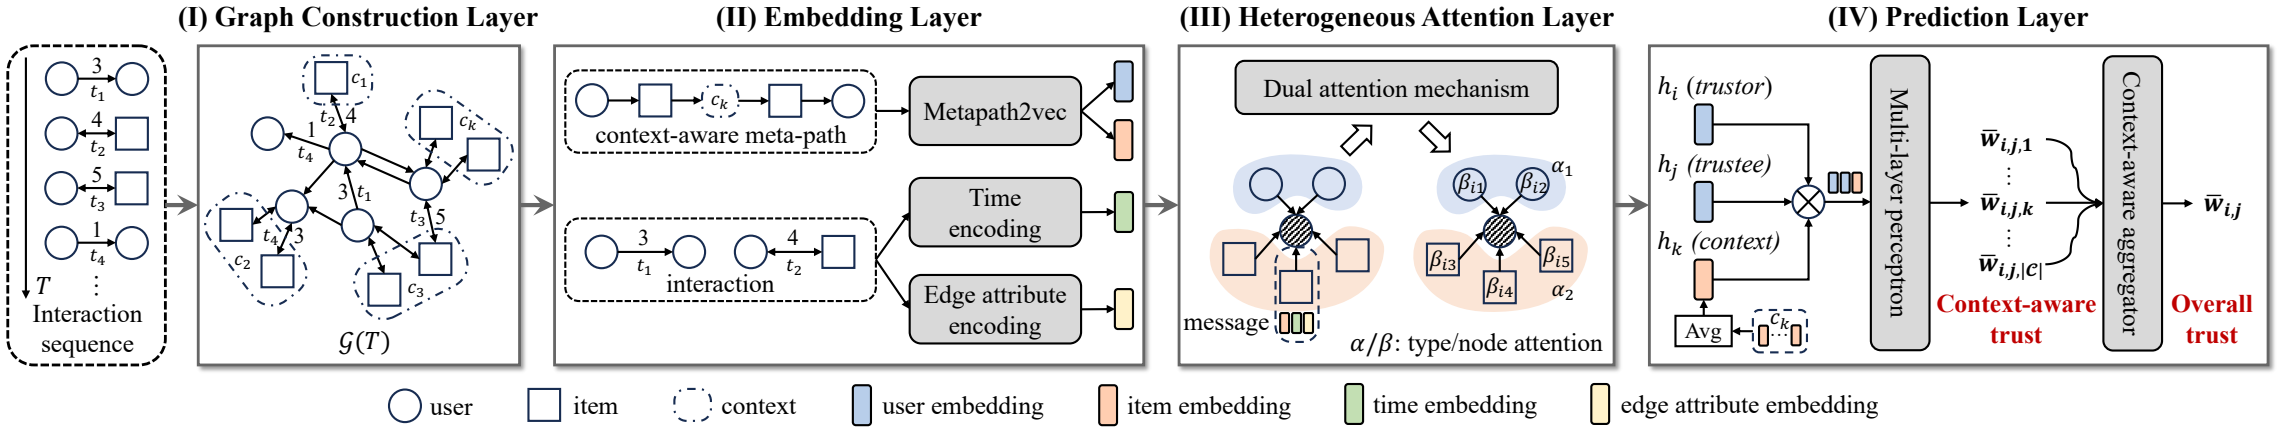

Challenges & Solutions

Challenge1: Modeling dynamicity with fine granularity conflicts with scalability

Challenges & Solutions

Challenge1: Modeling dynamicity with fine granularity conflicts with scalability

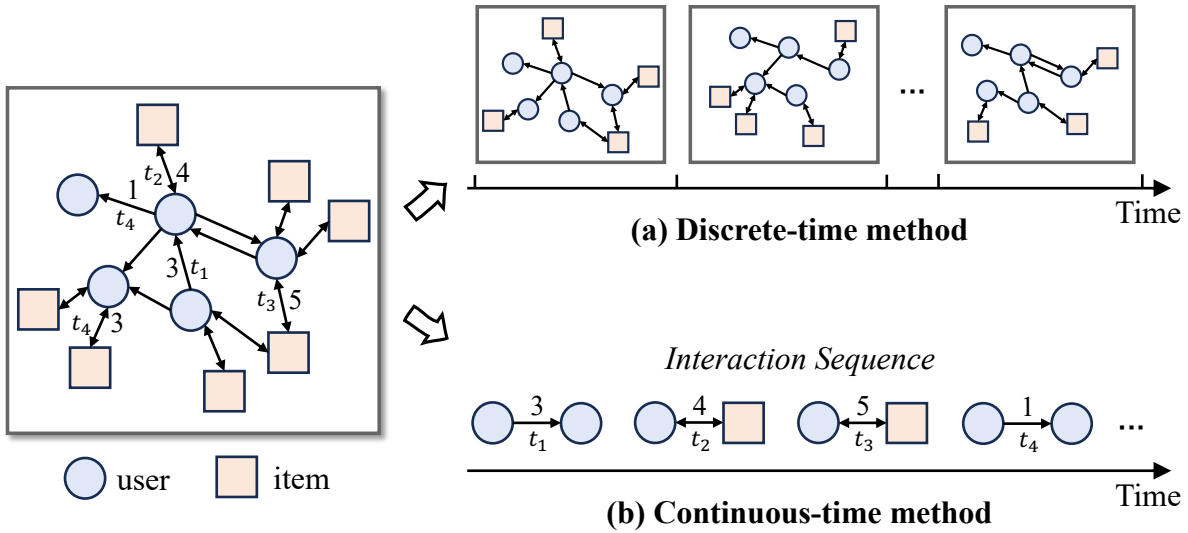

(a) Lose temporal information within snapshots

(b) Suffer from low scalability

Challenges & Solutions

Challenge1: Modeling dynamicity with fine granularity conflicts with scalability

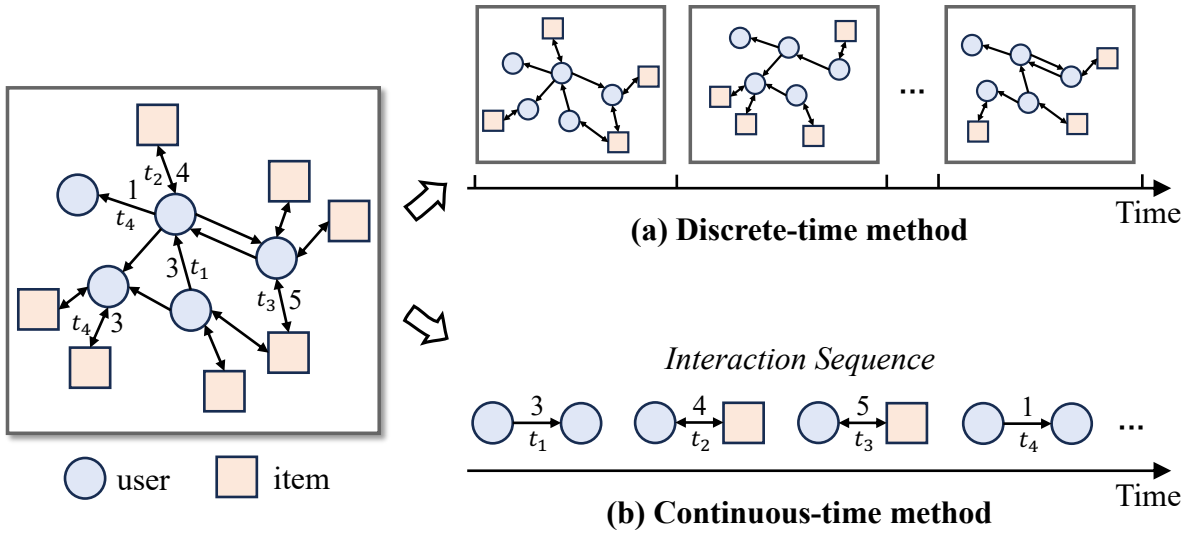

(a) Lose temporal information within snapshots

(b) Suffer from low scalability

1. Use time encoding to make full use of timestamped interactions

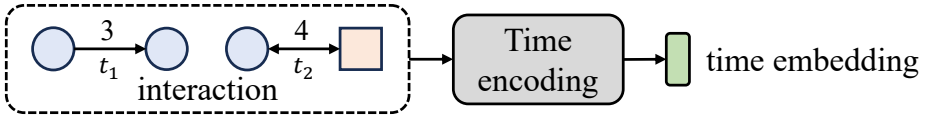

Challenges & Solutions

Challenge1: Modeling dynamicity with fine granularity conflicts with scalability

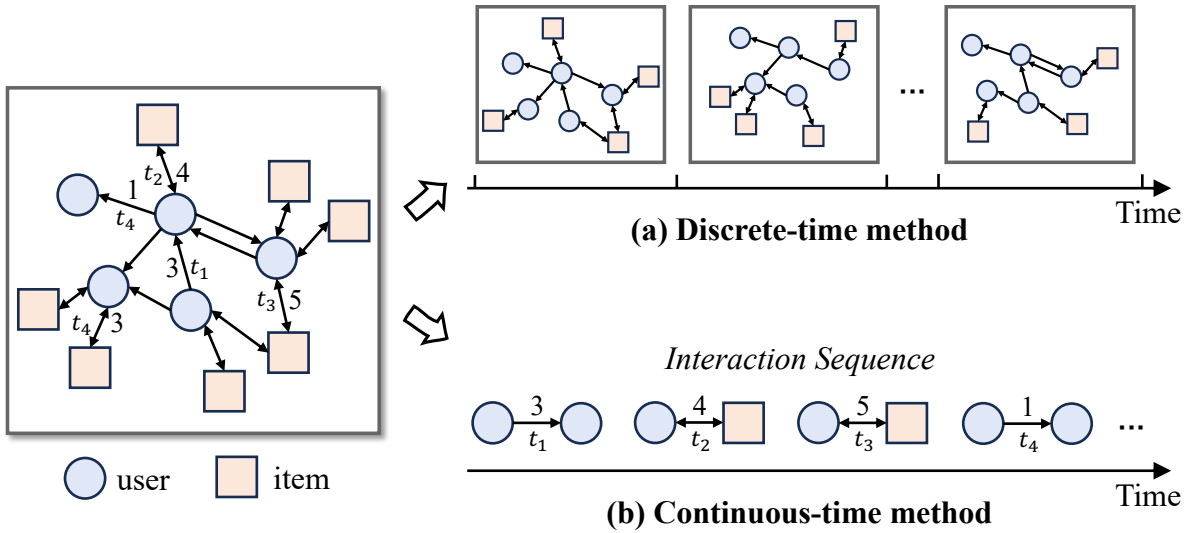

(a) Lose temporal information within snapshots

(b) Suffer from low scalability

1. Use time encoding to make full use of timestamped interactions

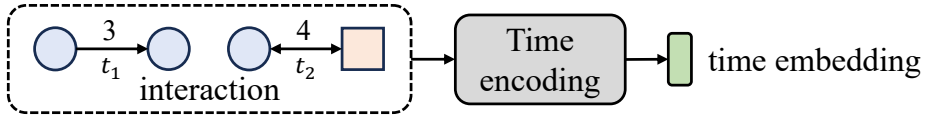

2. Focus on limited yet crucial interactions:

- Recent-time neighbor sampling
- One-hop trust propagation

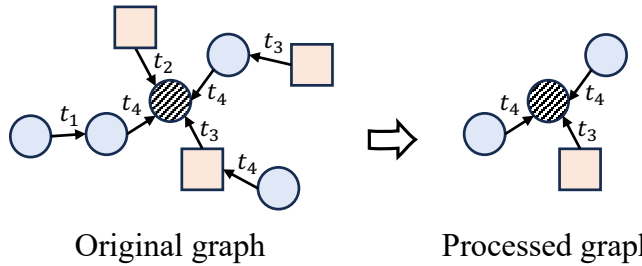

Challenges & Solutions

Challenge2: It is difficult to extract key information from heterogeneous graphs

Challenges & Solutions

Challenge2: It is difficult to extract key information from heterogeneous graphs

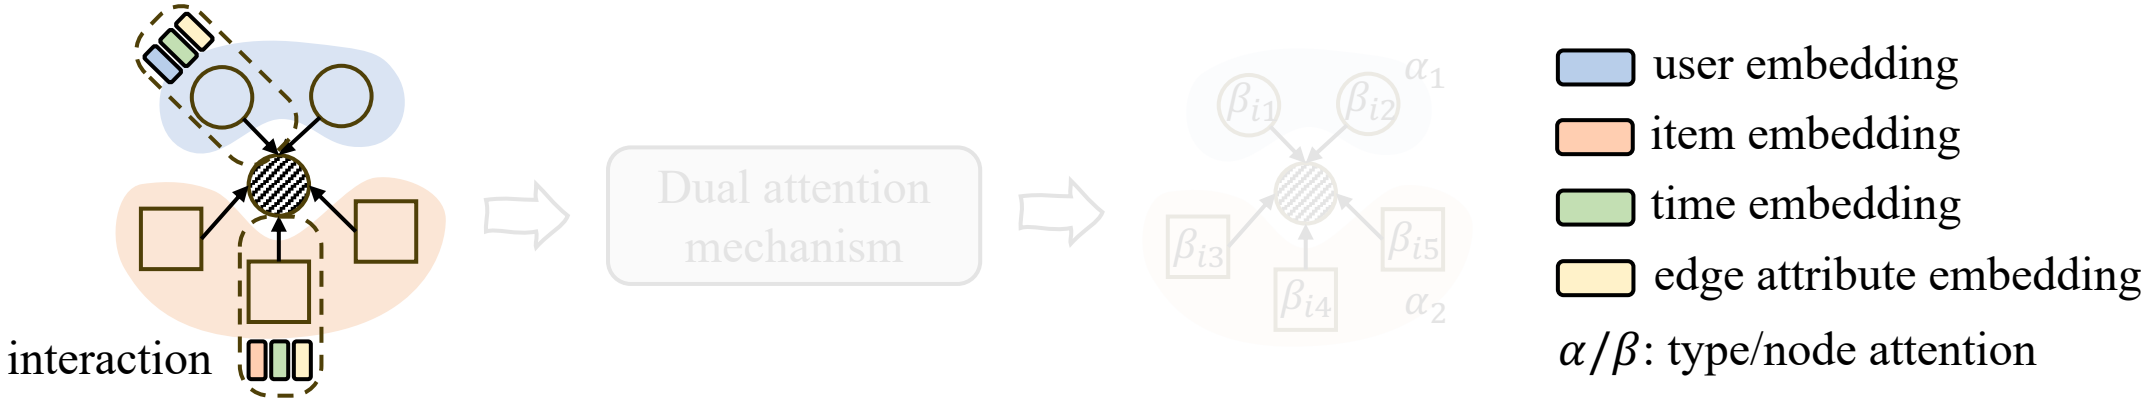

Heterogeneous interactions may introduce irrelevant information that undermine trust propagation

Challenges & Solutions

Challenge2: It is difficult to extract key information from heterogeneous graphs

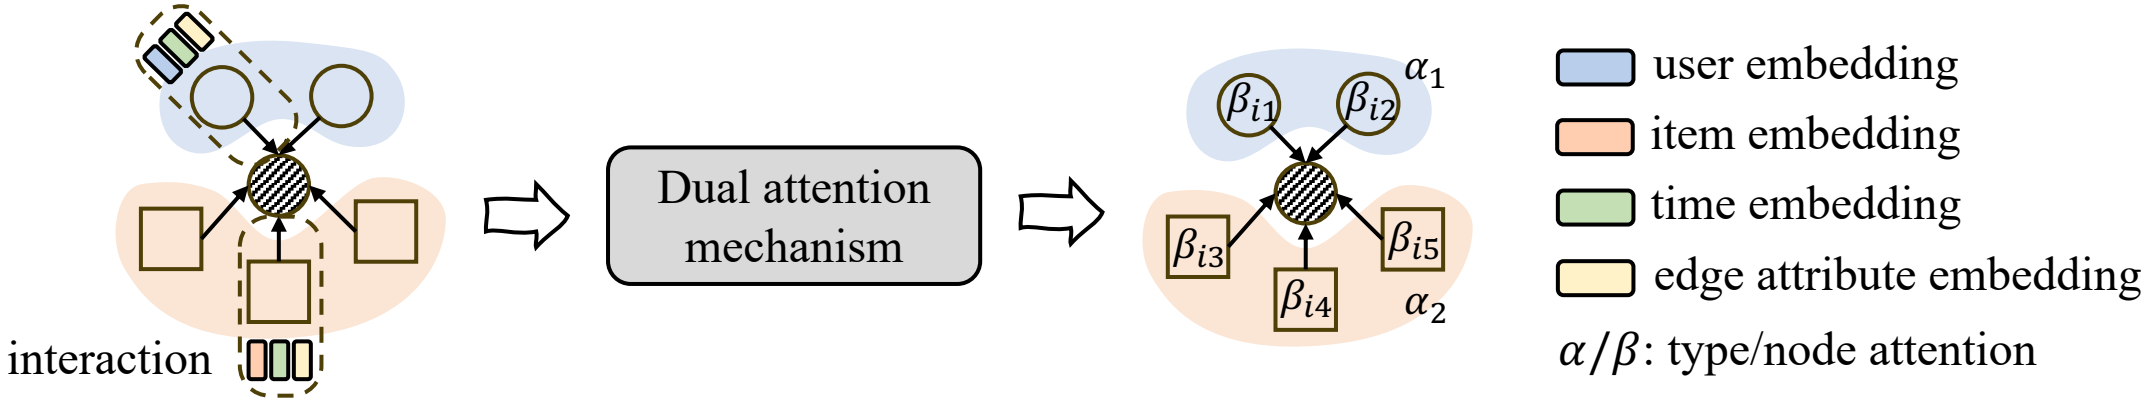

Heterogeneous interactions may introduce irrelevant information that undermine trust propagation

1. Use type attention to learn the importance of each node type (interaction type)

2. Use node attention to learn the importance of each node within the same type

Challenges & Solutions

Challenge3: Contexts are less explicit than nodes or edges, and existing datasets lack context-specific trust labels

Challenges & Solutions

Challenge3: Contexts are less explicit than nodes or edges, and existing datasets lack context-specific trust labels

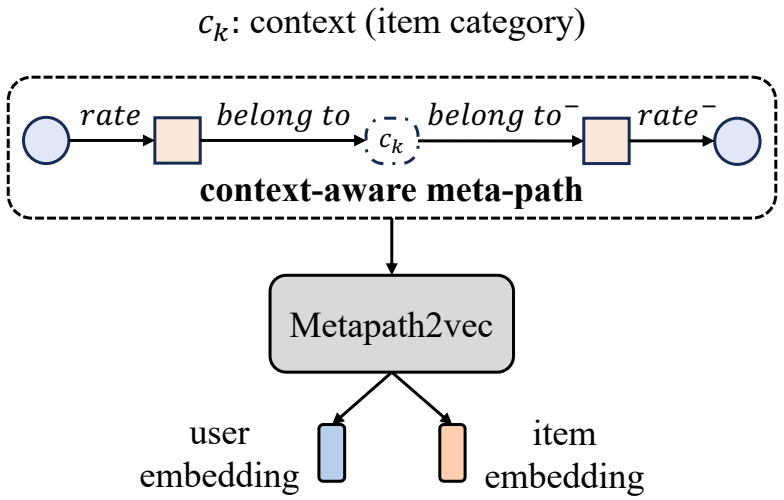

1. Introduce a context-aware meta-path to incorporate contextual information

Challenges & Solutions

Challenge3: Contexts are less explicit than nodes or edges, and existing datasets lack context-specific trust labels

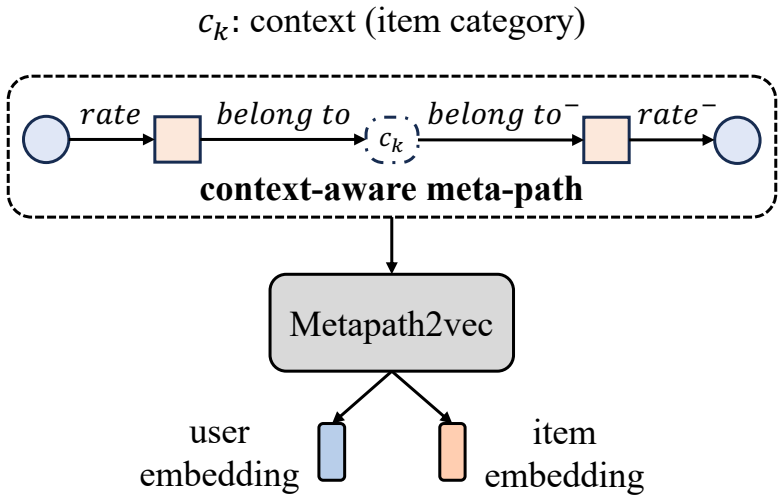

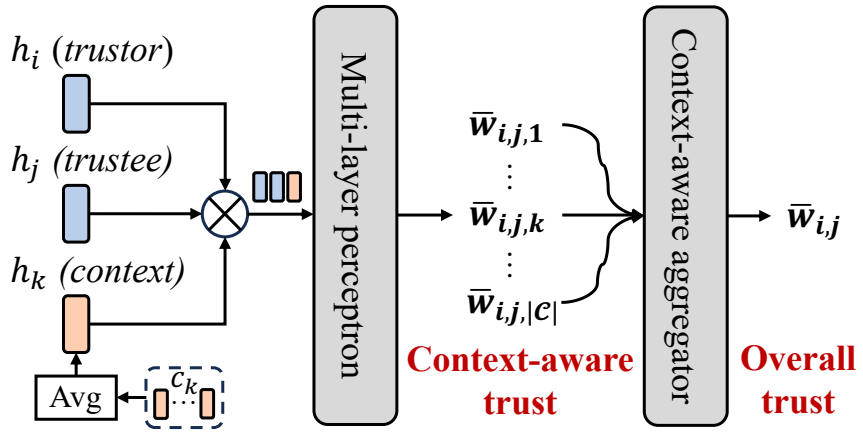

1. Introduce a context-aware meta-path to incorporate contextual information

2. Propose a context-aware aggregator to avoid the reliance on context-specific trust labels

Challenges & Solutions

Challenge3: Contexts are less explicit than nodes or edges, and existing datasets lack context-specific trust labels

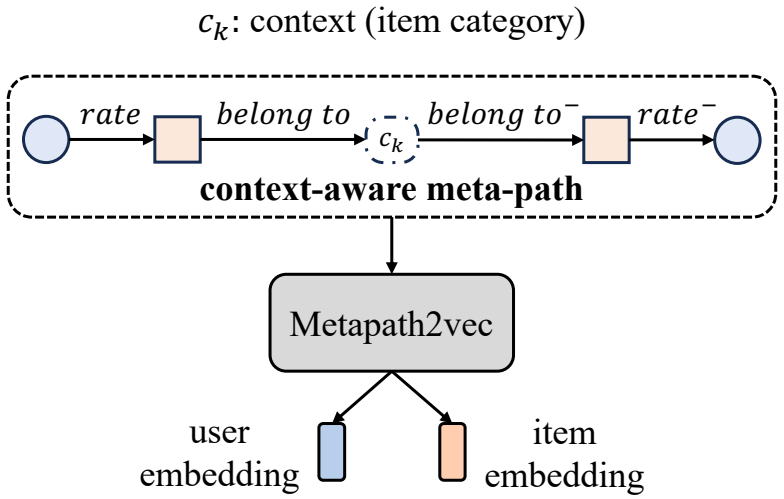

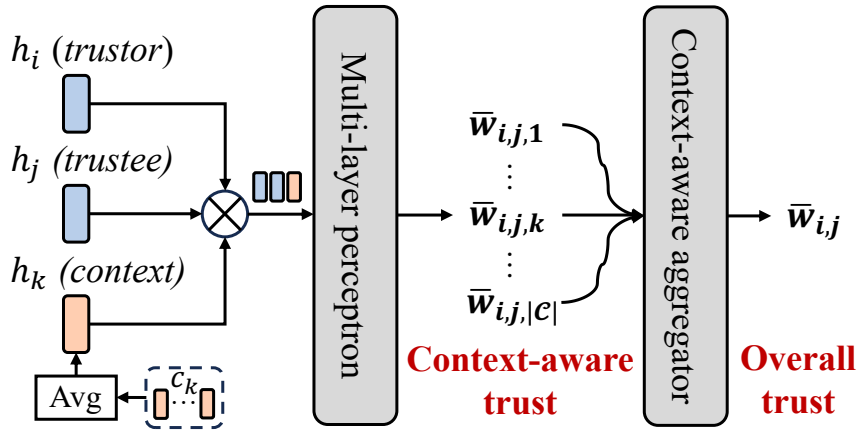

1. Introduce a context-aware meta-path to incorporate contextual information

2. Propose a context-aware aggregator to avoid the reliance on context-specific trust labels

By jointly addressing Challenge 2&3, CAT gains a robust semantic understanding

CAT Overview

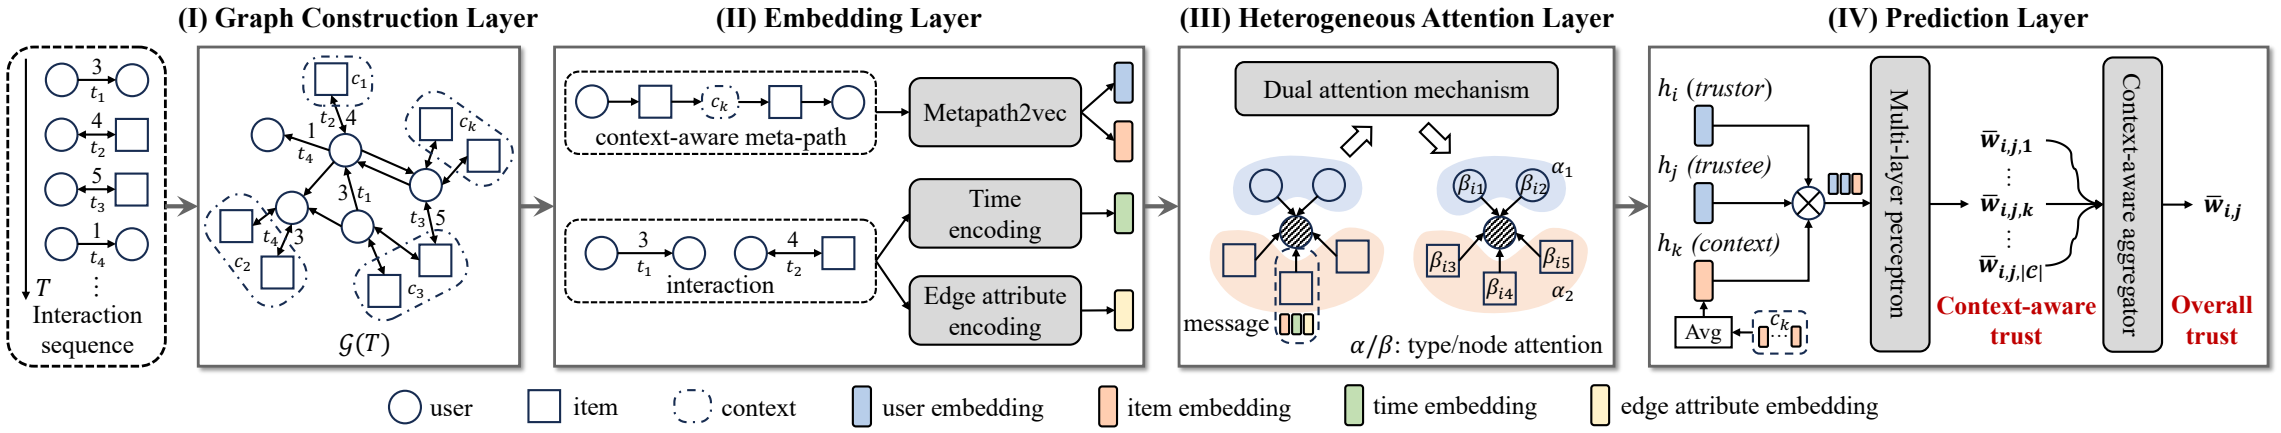

- Graph Construction Layer
  - Build a dynamic heterogeneous graph based on interaction sequence (continuous-time)
- Embedding Layer
  - Initialize embeddings for nodes, time, and edge attributes (e.g., trust levels)

CAT Overview

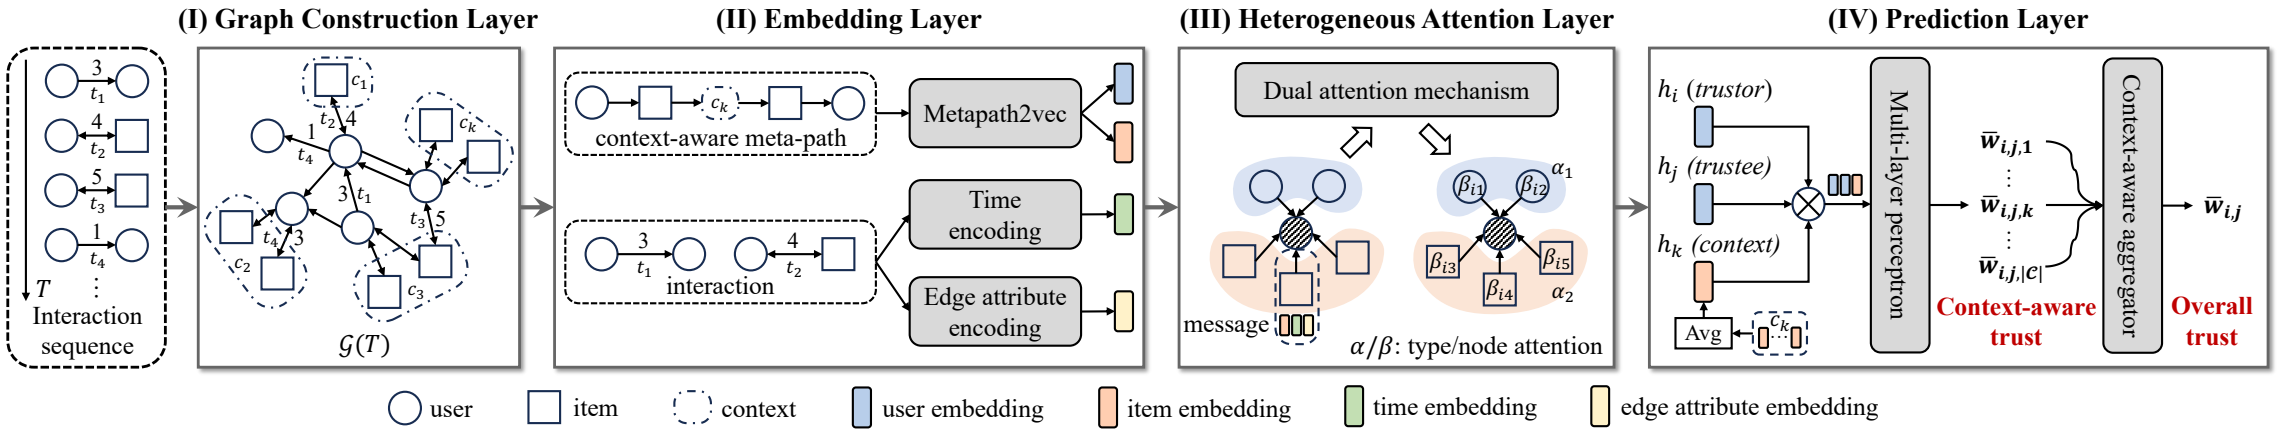

- ❑ Heterogeneous Attention Layer
  - Learn node embeddings by selectively propagating and aggregating trust information
- ❑ Prediction Layer
  - Predict the trust relationship between any two users under a specific context

CAT Overview

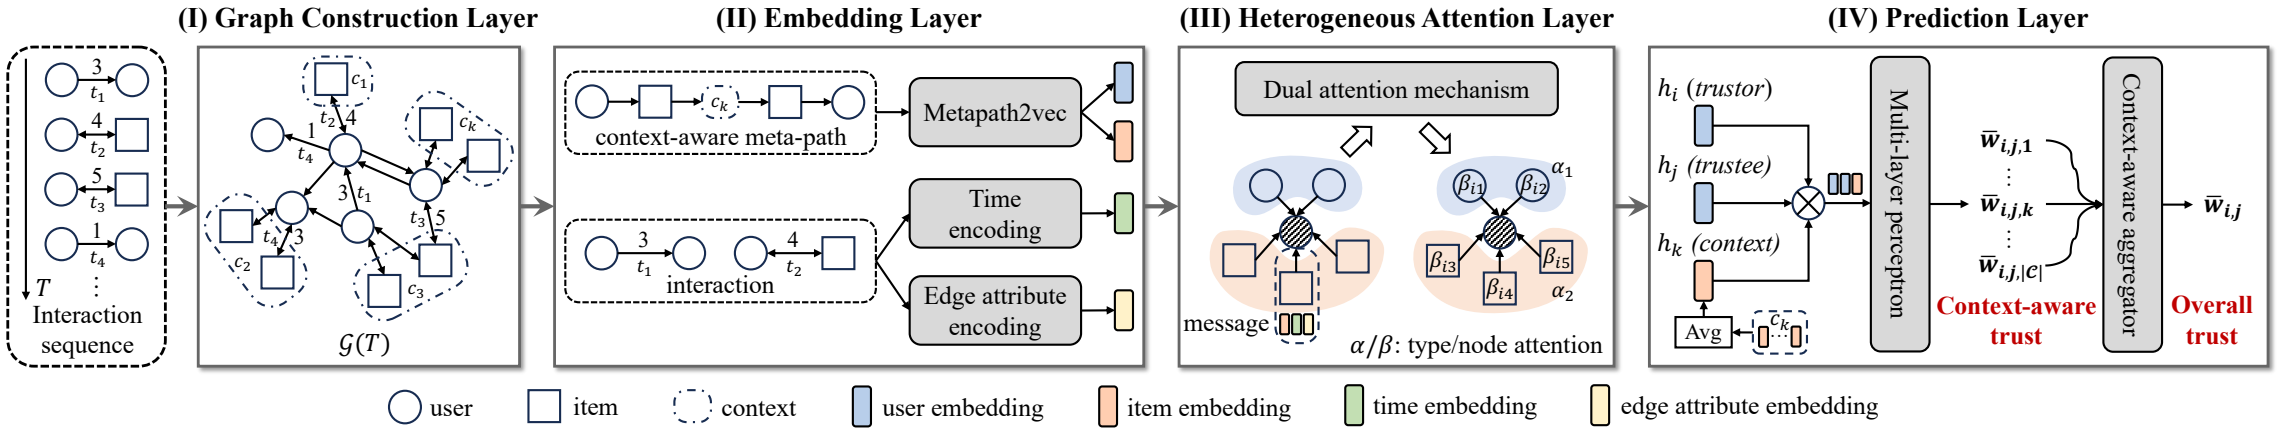

- ❑ Heterogeneous Attention Layer
  - Learn node embeddings by selectively propagating and aggregating trust information
- ❑ Prediction Layer
  - Predict the trust relationship between any two users under a specific context

CAT can predict both context-aware trust (even without context-specific labels) and overall trust!

Experimental Design

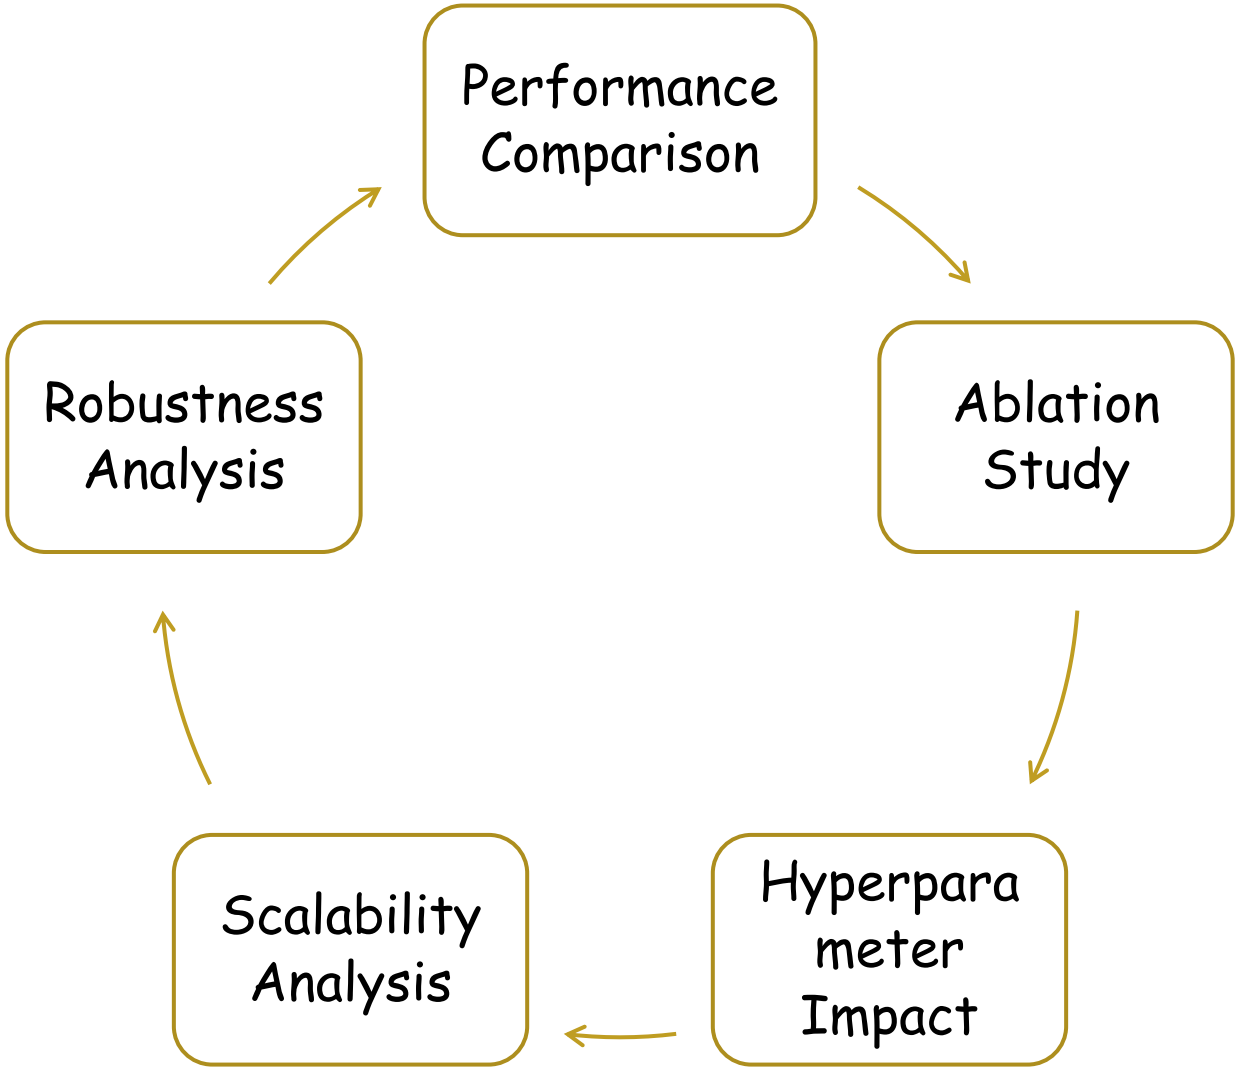

Experimental Setup

❑ Datasets

- Sourced from real-world social networking-based consumer review sites

TABLE III: 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	✓	✗

❑ Evaluation Metrics

- Mean Reciprocal Rank (MRR), Average Precision (AP), Area Under the ROC Curve (AUC)

❑ Baseline Models

- Four types of baselines that vary in their support for dynamicity and heterogeneity
- CAT → Heterogeneous dynamic model

Performance Comparison

❑ Overall Trust Prediction

- Epinions: Evaluate the model's ability to predict dynamic trust relationships
- Ciao & CiaoDVD: Evaluate model stability across different data split ratios

TABLE IV: 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
GATrust	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	0.5081	0.9151	0.9588	0.6168	0.9446	0.9675
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
GATrust	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	0.2707	0.9365	0.8773	0.2040	0.9387	0.8596
HGT	0.2693	0.9326	0.8603	0.2931	0.9504	0.8633
CAT	**0.4082**	**0.9527**	**0.8933**	**0.3987**	**0.9594**	**0.8799**

TABLE V: 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	MRR	AP	AUC	MRR	AP	AUC	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
	GATrust	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	0.2772	0.8470	0.8550	0.2775	0.8508	0.8600	0.2835	0.8559	0.8650	0.3008	0.8623	0.8693
	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
	GATrust	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	0.7606	0.9826	0.9826	0.7563	0.9827	0.9828	0.7481	0.9823	0.9829	0.7340	0.9814	0.9826
	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**

CAT outperforms all baselines across three datasets, particularly on Task 2, highlighting its potential to address the cold start issue

Performance Comparison

❑ Context-Aware Trust Prediction (only supported by CAT)

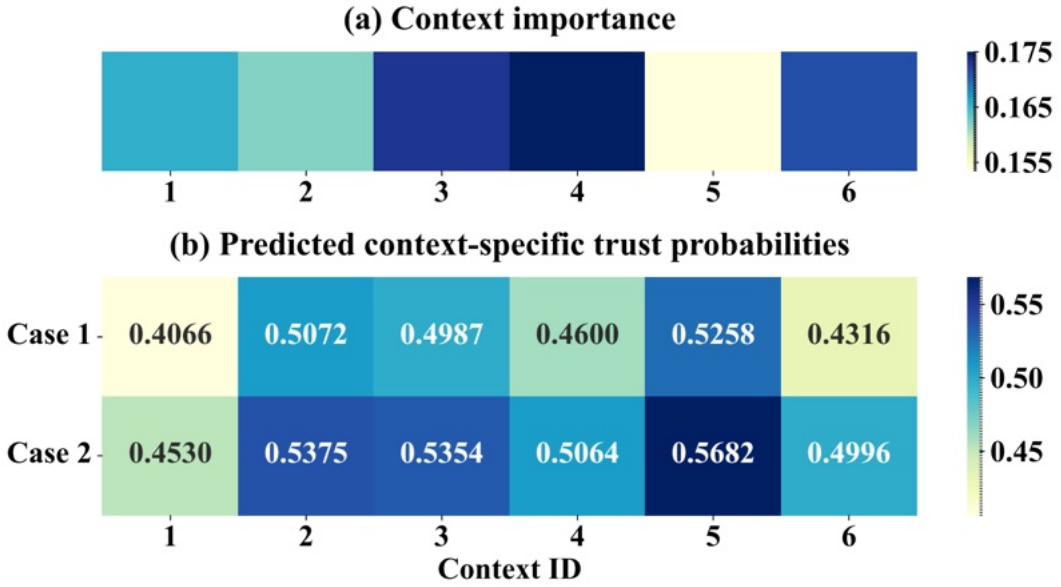

Fig. 3: Results of context-aware trust prediction on the Ciao dataset. Case 1 illustrates a user pair $\langle v_i, v_j \rangle$ with an overall distrusted relationship, whereas case 2 depicts a user pair with an overall trusted relationship.

- Different contexts contribute differently to trust establishment
- Trust probabilities can vary notably for the same user pair across contexts

Performance Comparison

❑ Context-Aware Trust Prediction (only supported by CAT)

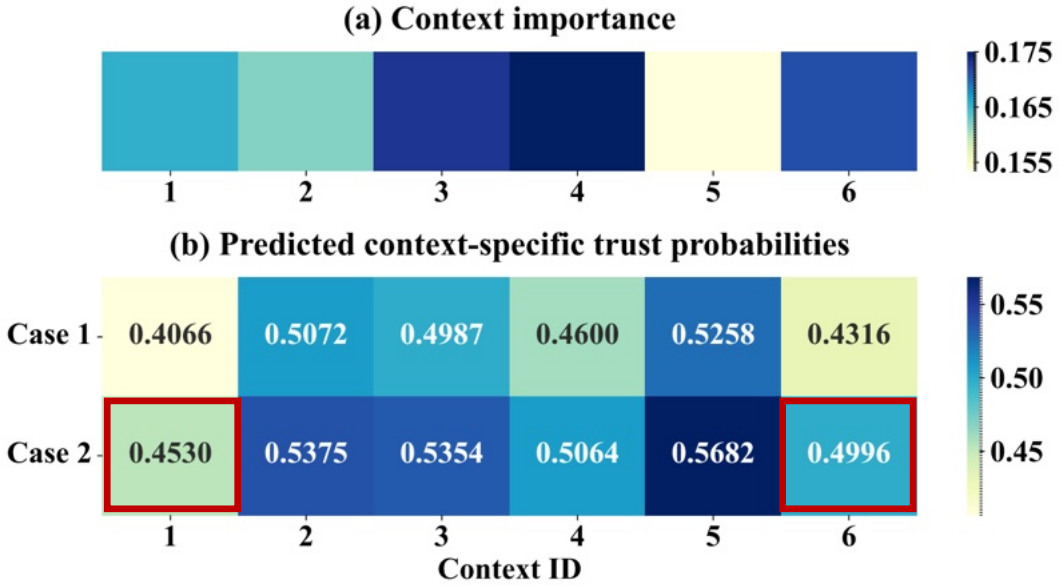

Fig. 3: Results of context-aware trust prediction on the Ciao dataset. Case 1 illustrates a user pair $\langle v_i, v_j \rangle$ with an overall distrusted relationship, whereas case 2 depicts a user pair with an overall trusted relationship.

- Different contexts contribute differently to trust establishment
- Trust probabilities can vary notably for the same user pair across contexts
- CAT's context-awareness capability enables flexible trust-based applications

Ablation Study

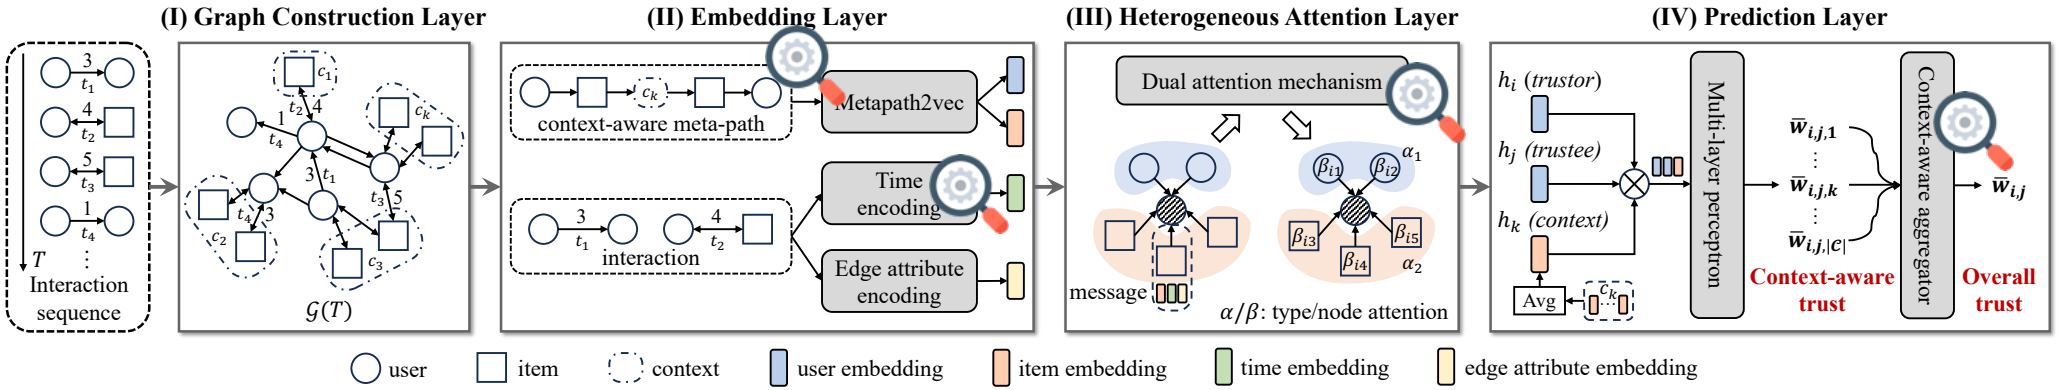

TABLE VI: Ablation studies on CAT.

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

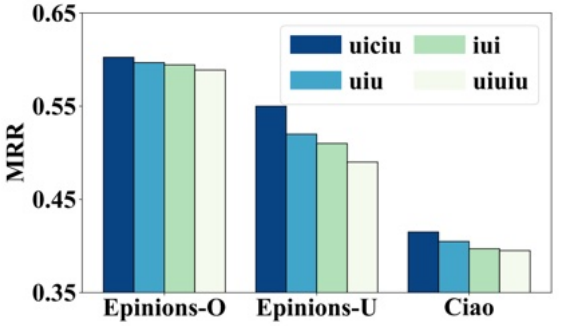

Fig. 4: Comparison of different meta-paths.

CAT consistently outperforms its five variants, highlighting the importance of modeling dynamicity, heterogeneity, and context-awareness

Hyperparameter Impact

- Embedding dimension, batch size, # sampled neighbors, trust propagation length

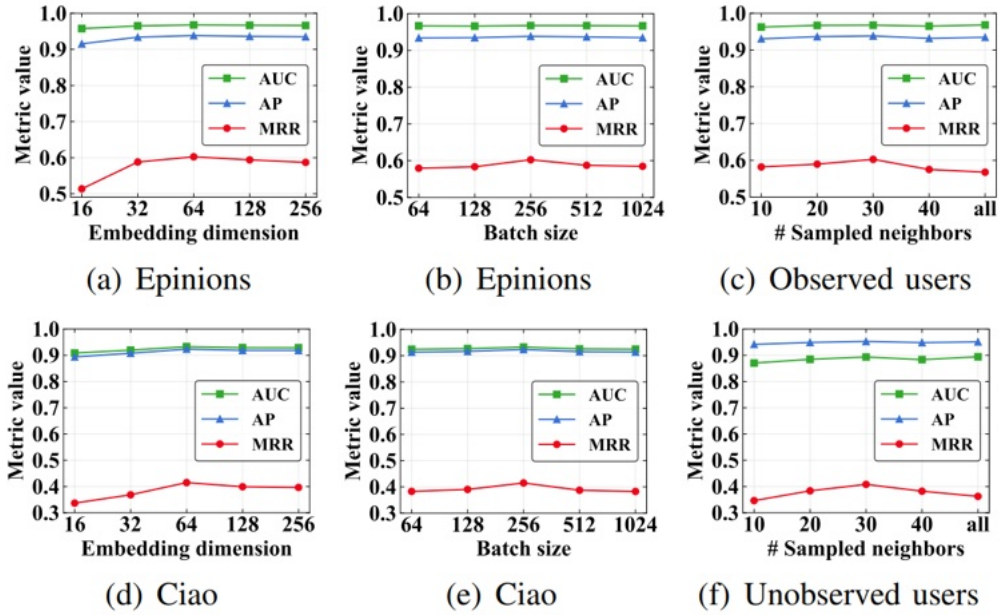

Fig. 5: Effects of embedding dimension, batch size, and number of sampled neighbors on model performance.

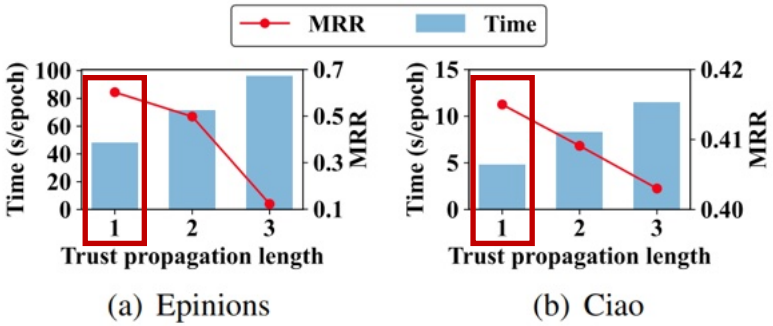

Fig. 6: Effect of trust propagation length on model performance.

As propagation length increases, CAT's performance degrades and runtime grows:

- Heterogeneous graphs introduce excessive irrelevant information
- Embeddings generated by Metapath2vec already encode high-hop information

Scalability Analysis

- We compare CAT's runtime with that of HGT (the best-performing baseline)

↓ 44.39% ↓ 40.79% ↓ 73.97%

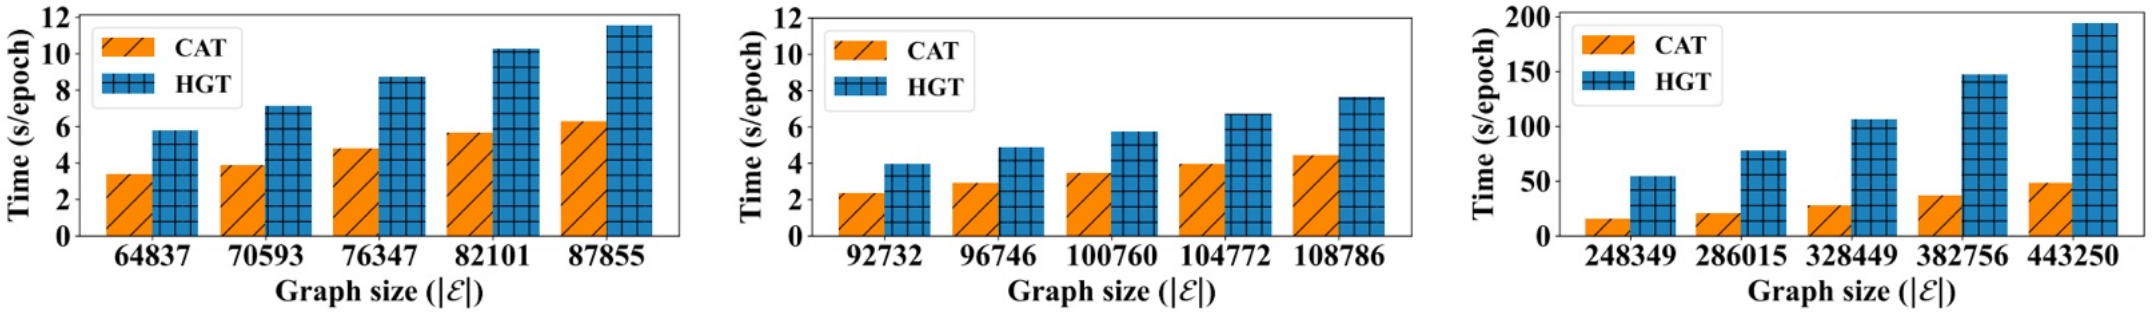

(a) Ciao (b) CiaoDVD (c) Epinions

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

CAT scales well by restricting trust propagation to one-hop neighbors and sampling only recent, critical interactions

Robustness Analysis

- We focus on both trust-oriented (local) and GNN-oriented attacks (global)

TABLE VII: 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%**

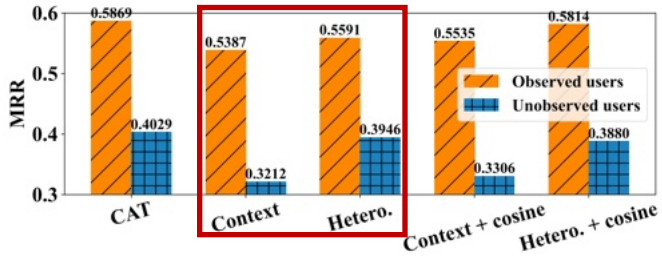

Fig. 8: Detailed evaluation and comparison of robustness against GNN-oriented attacks with 10% perturbation.

CAT demonstrates superior robustness over existing GNN-based trust prediction models, owing to its strong semantic understanding.

Robustness Analysis

- We focus on both trust-oriented (local) and GNN-oriented attacks (global)

TABLE VII: 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%**

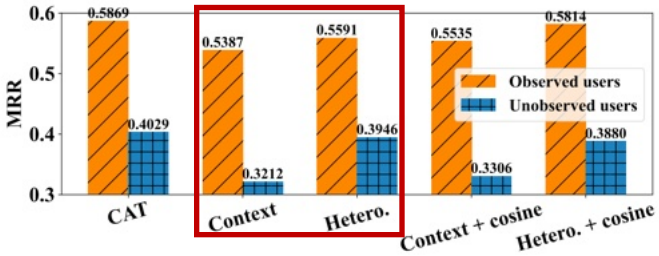

Fig. 8: Detailed evaluation and comparison of robustness against GNN-oriented attacks with 10% perturbation.

CAT demonstrates superior robustness over existing GNN-based trust prediction models, owing to its strong semantic understanding.

TABLE IX: 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					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%

Adaptive attacker: Blur the boundaries between two distinct contexts

Conclusion & Future Work

- CAT: A GNN-based Trust Prediction Model that
  - Supports trust dynamicity
  - Handles real-world network heterogeneity
  - Predicts both context-aware trust and overall trust
  - Is resilient to trust/GNN-oriented data poisoning attacks

- Future Work
  - How to incorporate textual information (e.g., user profiles and reviews) into CAT?
  - How to defend against evasion attacks that occur during the testing phase?

Paper

Code

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

Q & A


Jie Wang


Zheng Yan


Jiahe Lan


Xuyan Li


Elisa Bertino

Network and Distributed System Security (NDSS) Symposium, 25 Feb 2026