Financial Crime Detection:

Utility, Robustness, Decentralization & Privacy

November 15, 2026
Financial AI Security, Privacy & Safety
An ICAIF 2026 Workshop

Kubilay Atasu
Associate Professor, Software Technology Dept.

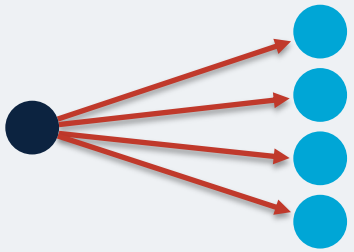
TU Delft Delft University of Technology



IBM AMLWorld: Realistic Synthetic Transaction Graphs

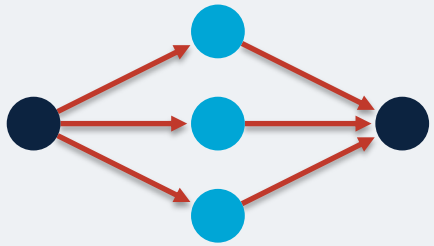
- Multi-agent simulation (normal and criminal agents)
- Known laundering patterns mixed into normal activity

Fan-Out



One account splits money across many accounts

Scatter-Gather



Money spreads through intermediaries and regroups

HI
Higher Illicit Ratio

HI-Small	HI-Medium	HI-Large
5M Txns	32M Txns	180M Txns

LI
Lower Illicit Ratio

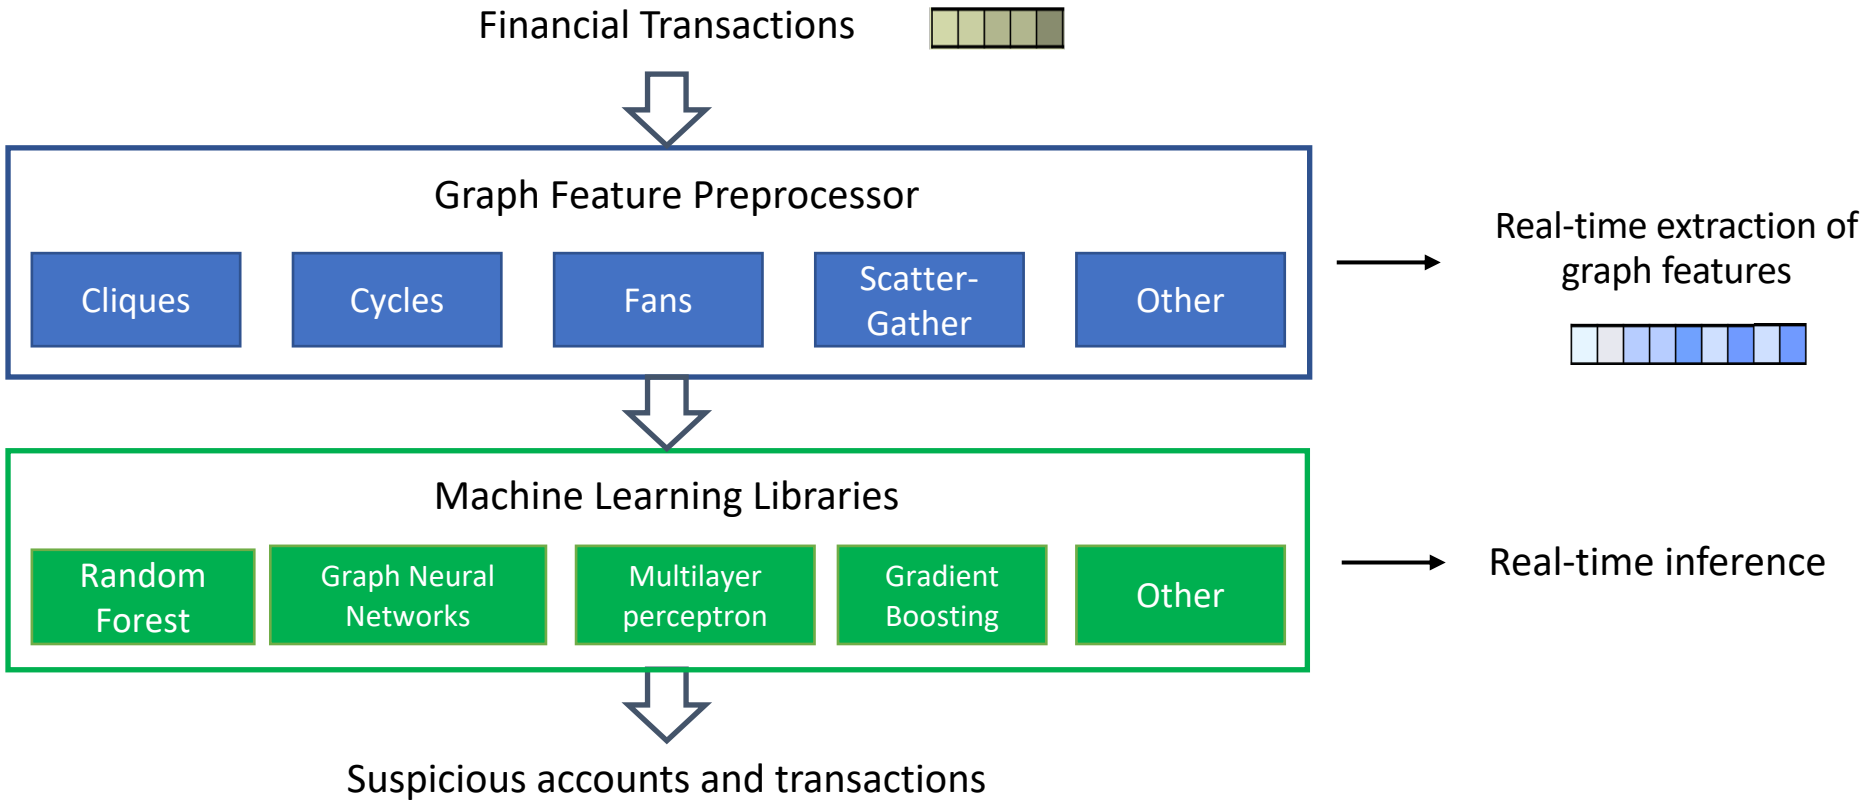
LI-Small	LI-Medium	LI-Large
7M Txns	31M Txns	176M Txns

Six Datasets
0.05–0.12% illicit, 10 to 97 days of activity

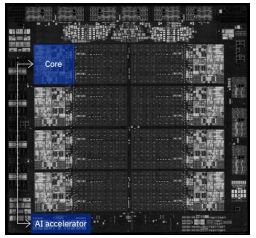
Datasets: *Erik R. Altman, Jovan Blanusa, Luc von Niederhäusern, Beni Egressy, Andreea Anghel, Kubilay Atasu:* ***Realistic Synthetic Financial Transactions for Anti-Money Laundering Models.*** *NeurIPS 2023*

Graph Machine Learning in IBM AI Toolkit for Z

Monitoring suspicious account activities in real time!

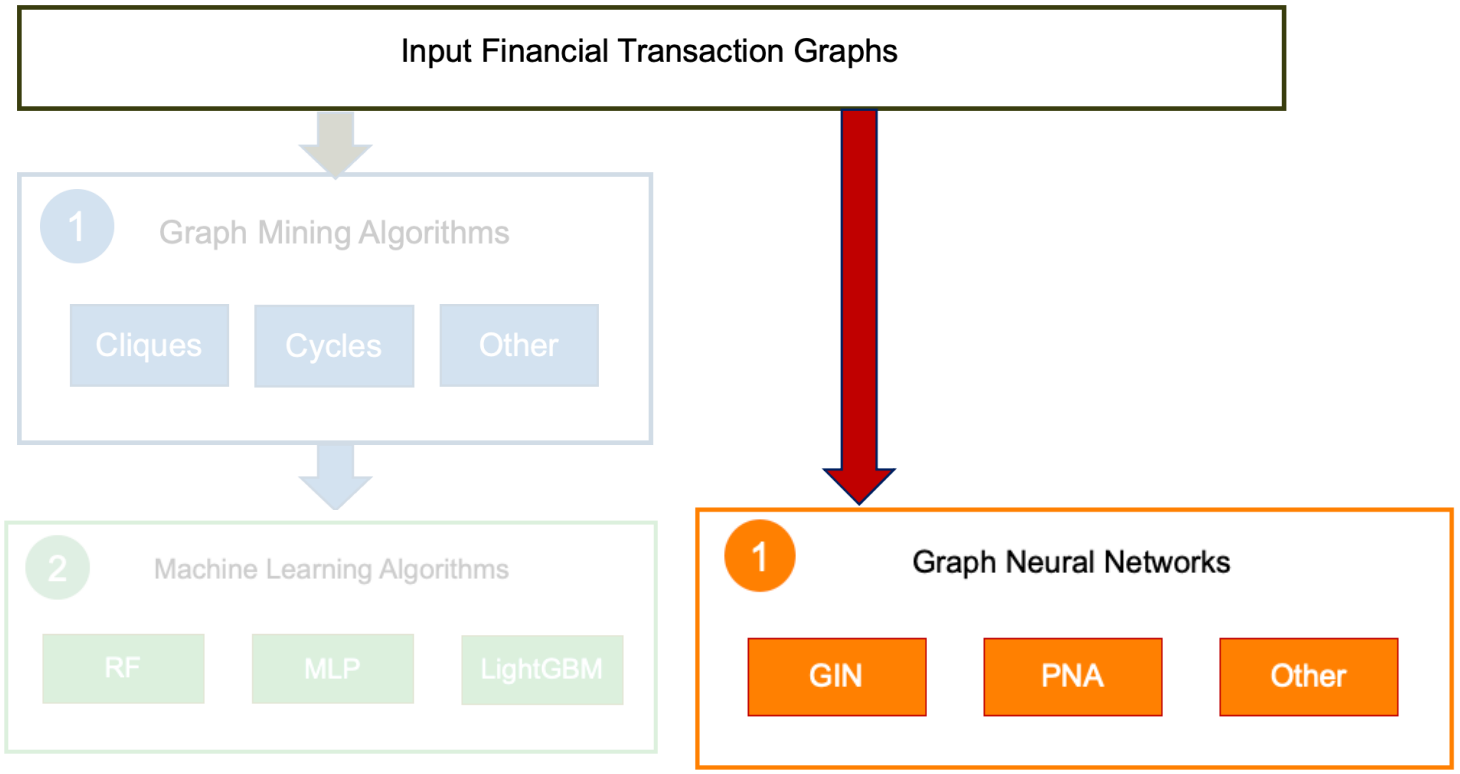






J. Blanusa et al.: Graph Feature Preprocessor: Real-time Subgraph-based Feature Extraction for Financial Crime Detection, ICAIF 2024.

Can Graph Neural Networks Help?



Why Graph Neural Networks?

Automation No feature engineering

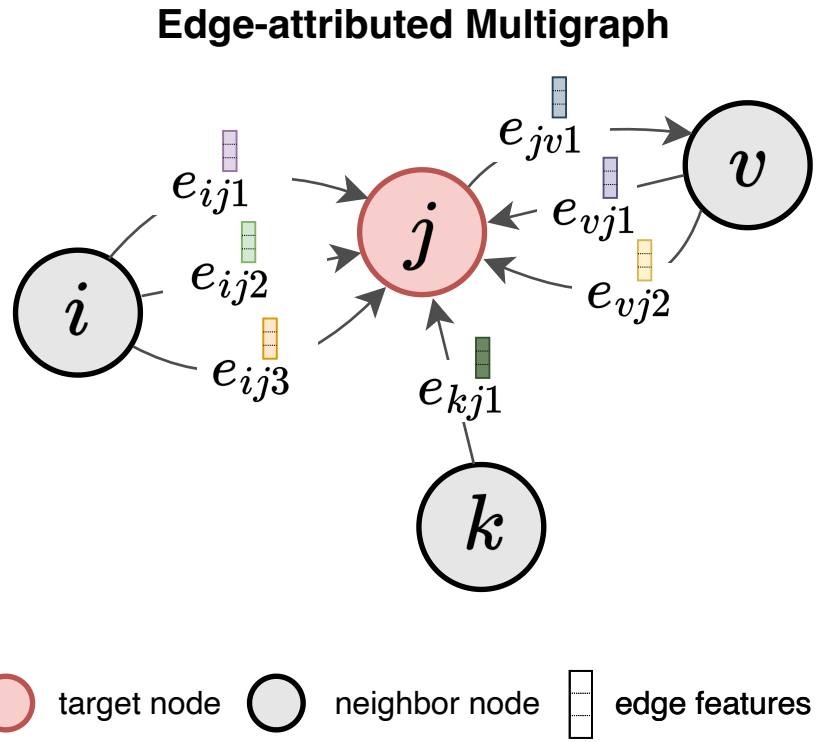
No (less) **domain knowledge** required

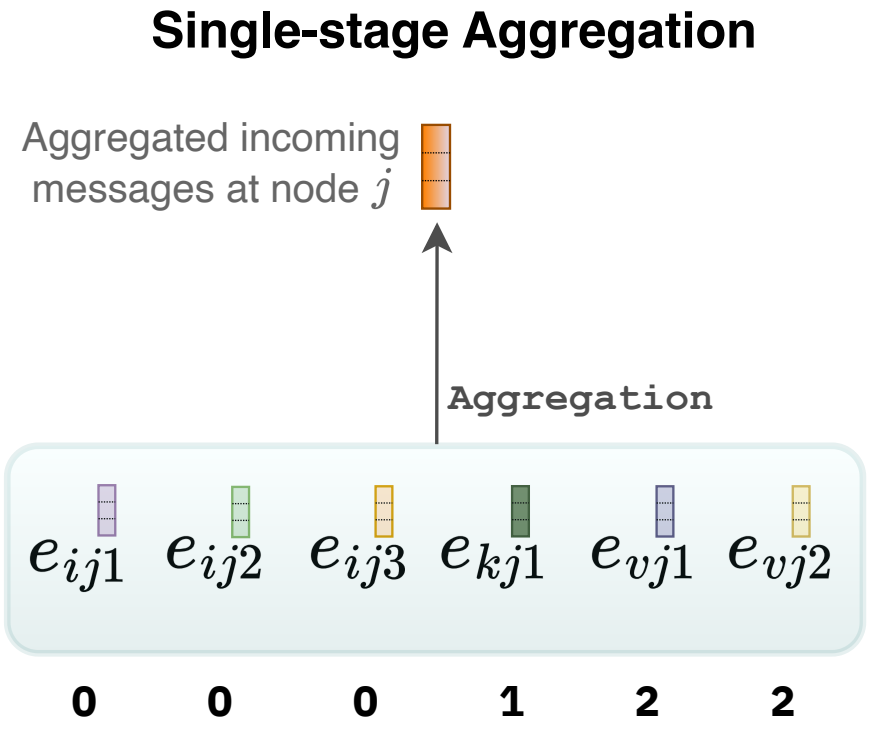
Can detect "**unestablished**" patterns

Differentiable – connect with LLMs

Challenge: Financial Transaction Graphs are Multigraphs with Rich Edge Attributes!

Standard GNNs on Edge-Attributed Multigraphs





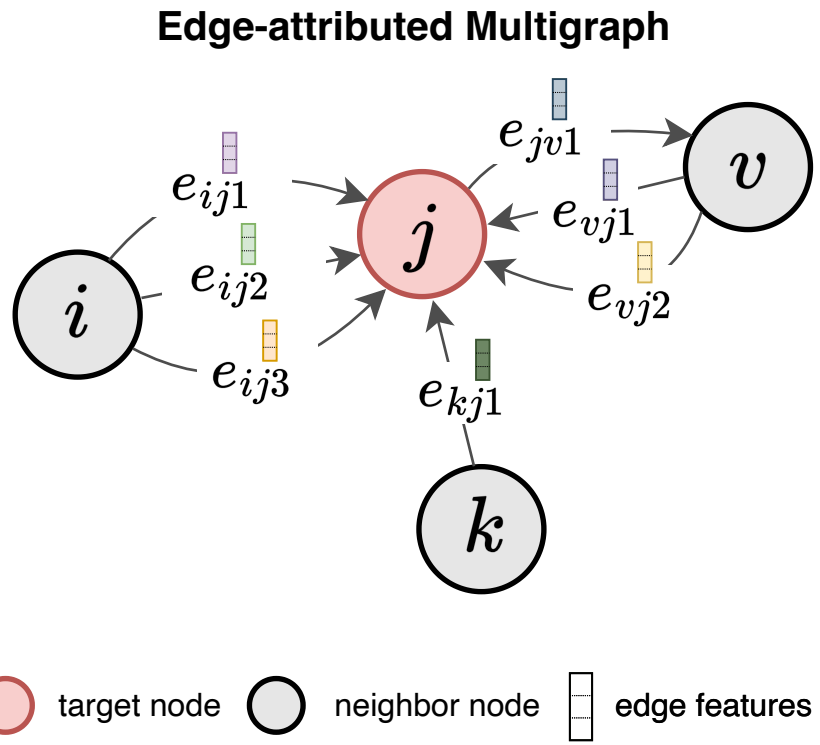
- That's not good enough!
- Computes in-degree, but not fan-in
- In-degre=6, but fan-in=3 for node j

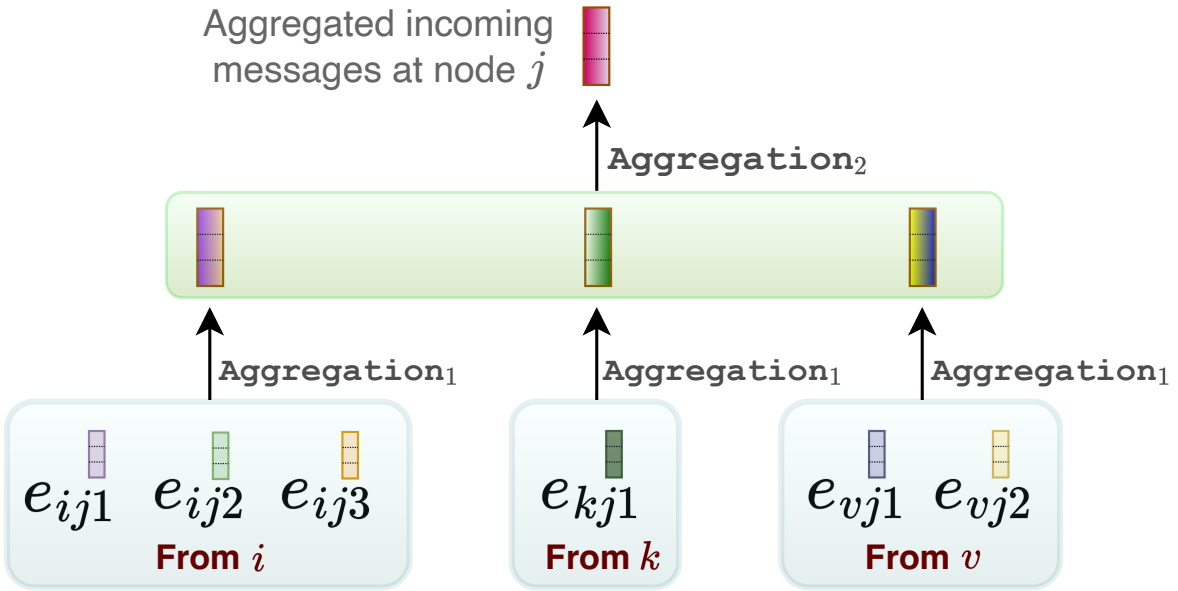
Multi-GNN: precomputed multi-graph port IDs

- State-of-the-art performance!
- Permutation equivariance lost!

Multi-GNN: B. Egressy et al.: Provably Powerful Graph Neural Networks for Directed Multigraphs. AAAI 2024

MEGA-GNN: Neighbor-aware Aggregation



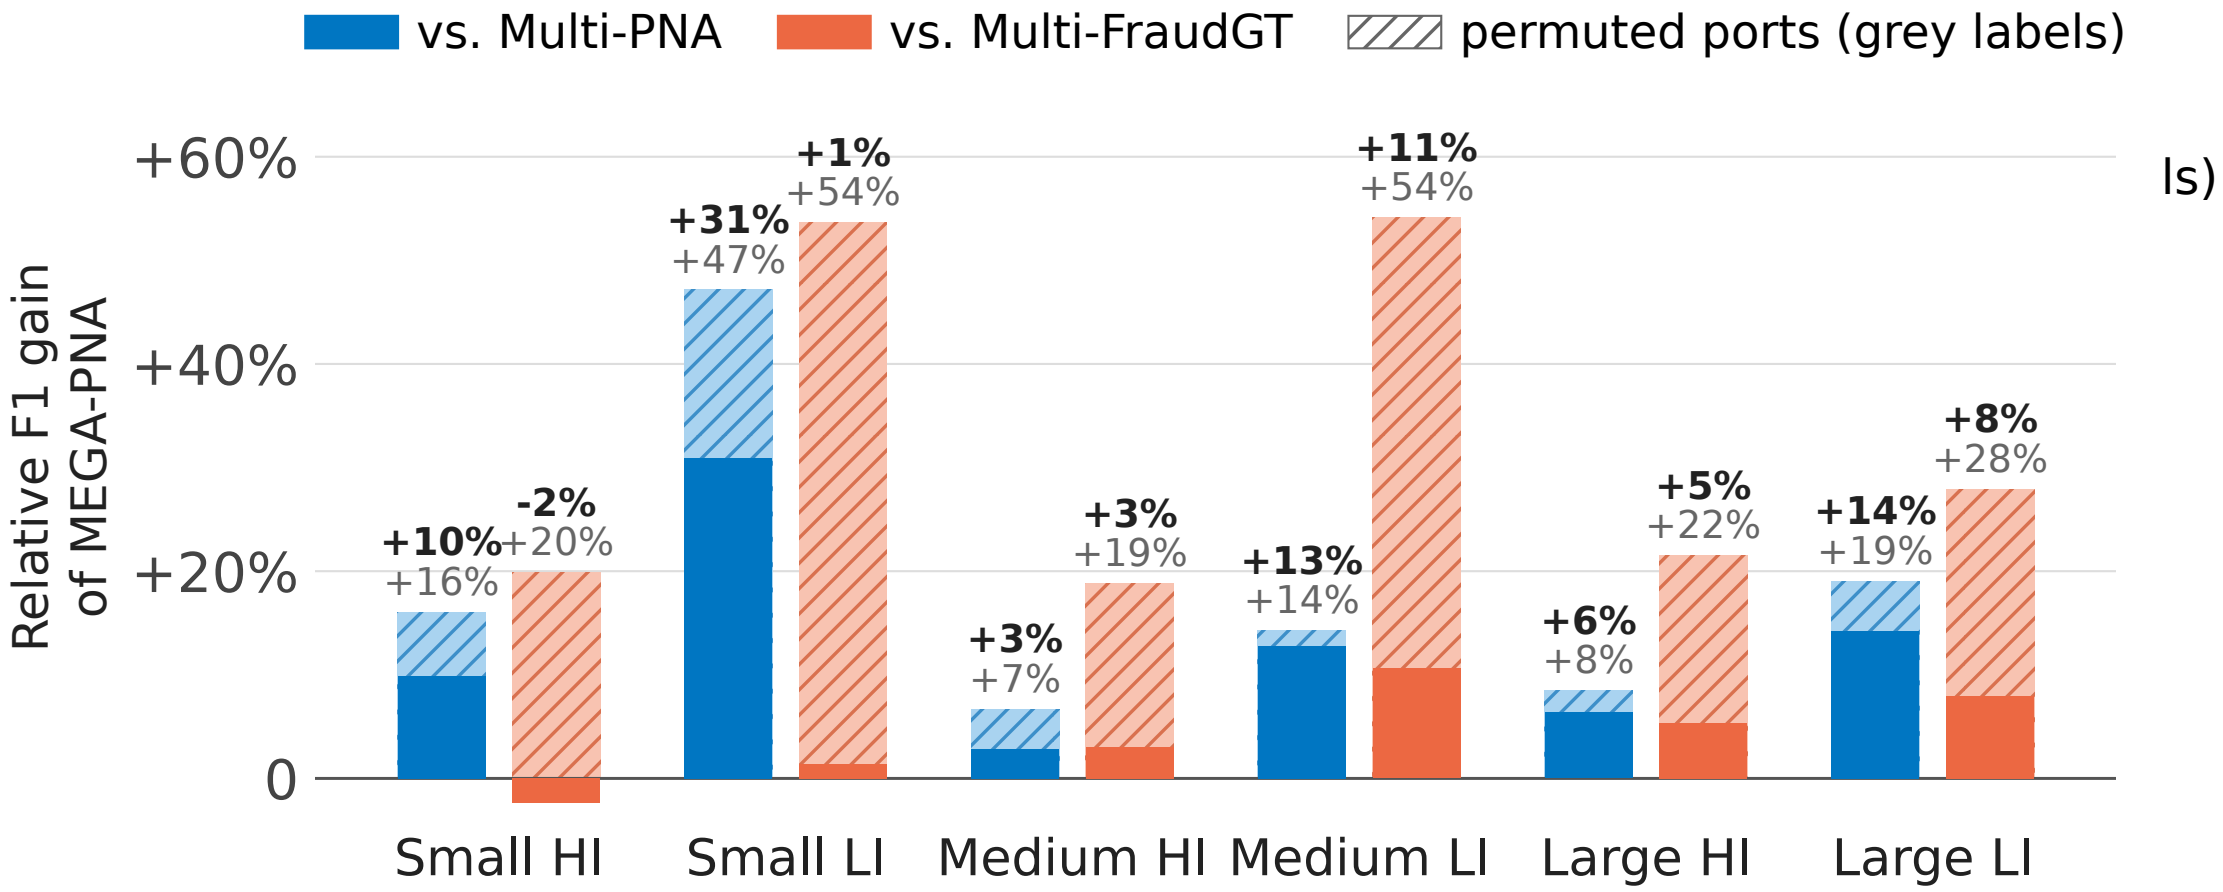


First aggregates multi-edges per neighbor, then aggregates across neighbors

- ✔ *Strictly more expressive than single-stage aggregation on edge-attributed multigraphs*
- ✔ *MEGA-GNN: permutation equivariant, same asymptotic complexity as standard GNNs*

MEGA-GNN: H. C. Bilgi and K. Atasu: Message Passing Neural Networks for Multigraphs with EdGe Attributes. TMLR 2026.

Effect Of Permuting Port IDs at Test Time



Multi-PNA and Multi-FraudGT use precomputed port IDs.

MEGA-PNA does not rely on port IDs.

MEGA-GNN: H. C. Bilgi and K. Atasu: Message Passing Neural Networks for Multigraphs with EdGe Attributes. TMLR 2026.

What have we ignored so far? (I)

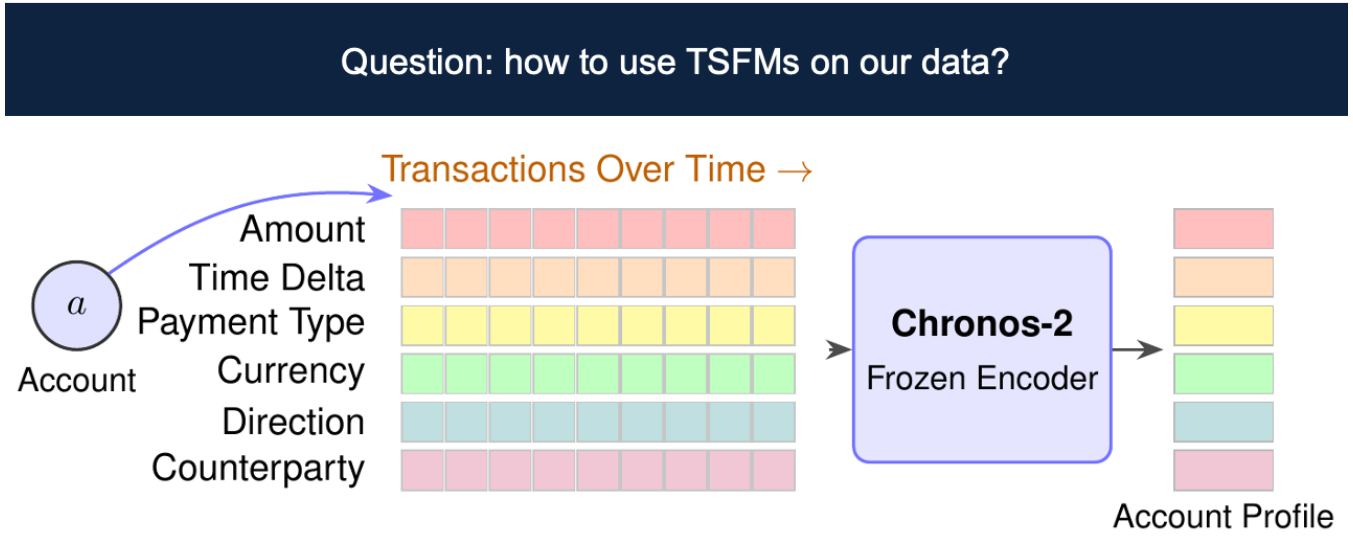
We have focused on the graph topology and ignored the time dimension.

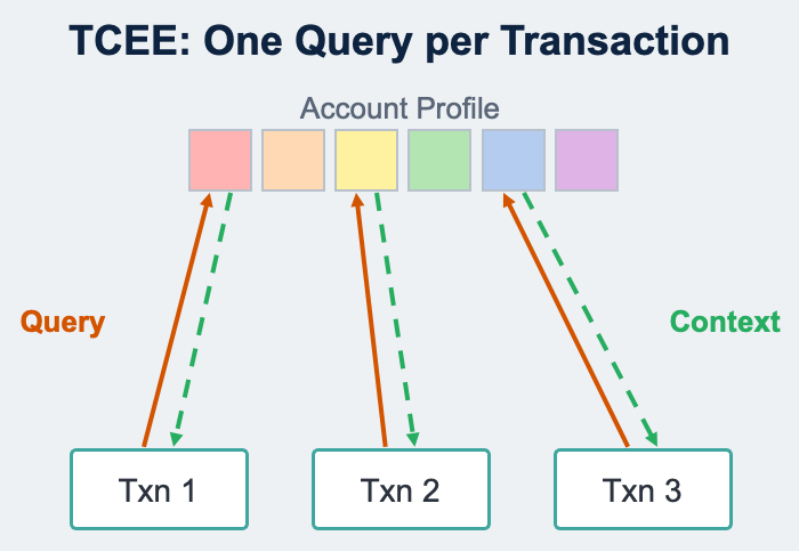
Financial transaction data can be modeled as a multivariate time series.

Can we improve the model utility by leveraging time series foundation models?

How can we combine TSFMs and GNNs to find both temporal and topological patterns?

Leveraging Time-Series Foundation Models

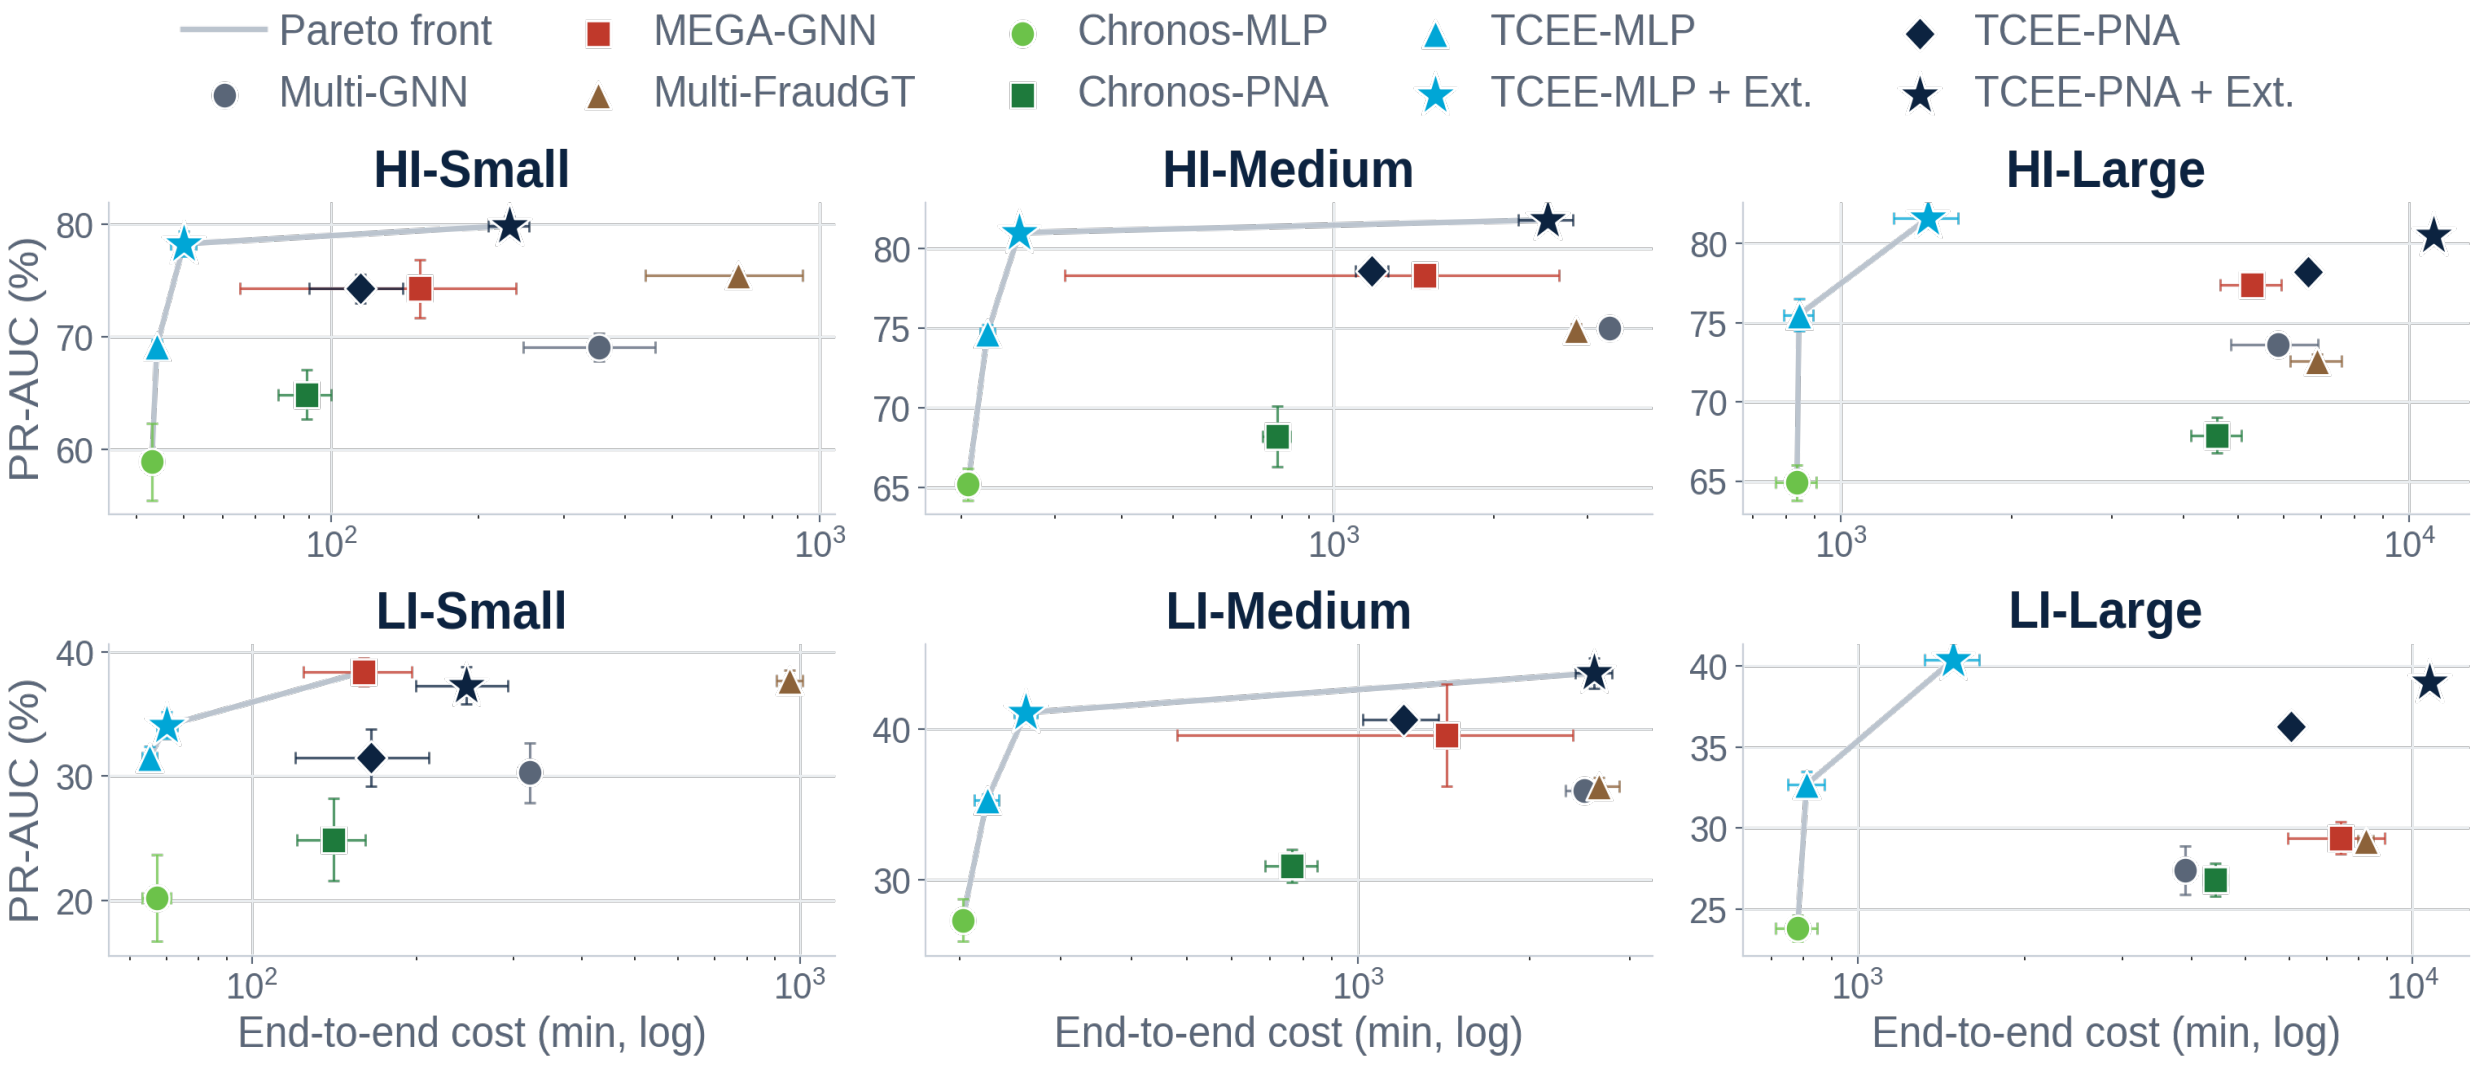




Idea: let each transaction query the account profiles of its accounts
→ Temporal-Context Edge Enrichment (TCEE)

Danae Savvidi et al. “Leveraging Time Series Foundation Models for Financial Crime Detection”, ICAIF 2026.

Leveraging Time-Series Foundation Models



Danae Savvidi et al. “Leveraging Time Series Foundation Models for Financial Crime Detection”, ICAIF 2026.

What have we ignored so far? (II)

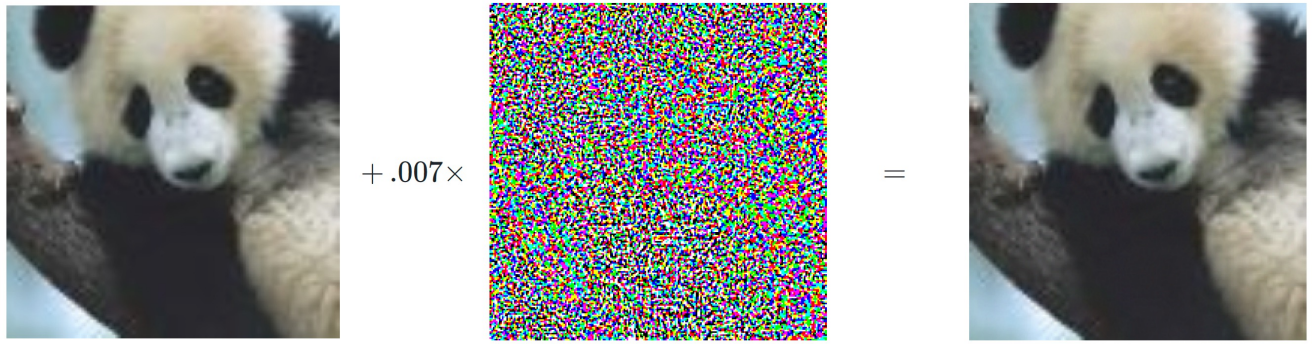
Criminal networks can identify and exploit the blind spots of financial crime models.

Standard GNN training is not robust against adversarial attacks.

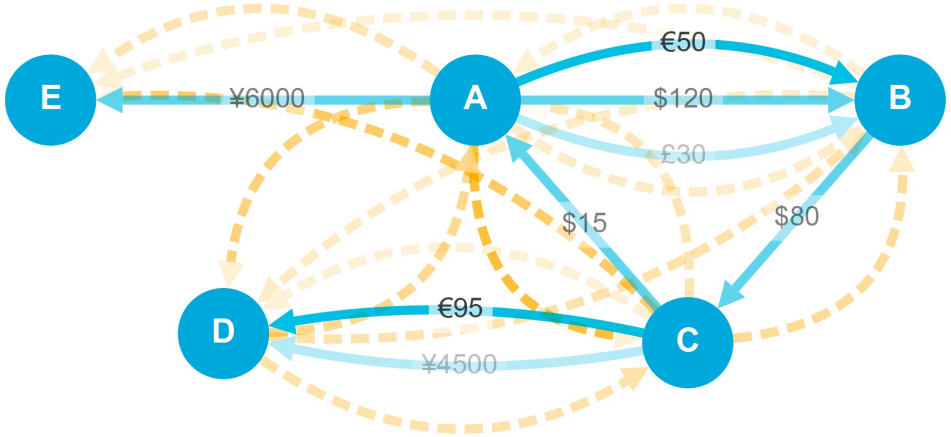
Can we build financial crime models that are robust against adversarial attacks?

Are standard adversarial training techniques applicable to multigraph GNNs?

Adversarial Robustness



Adversarial Attacks on Images

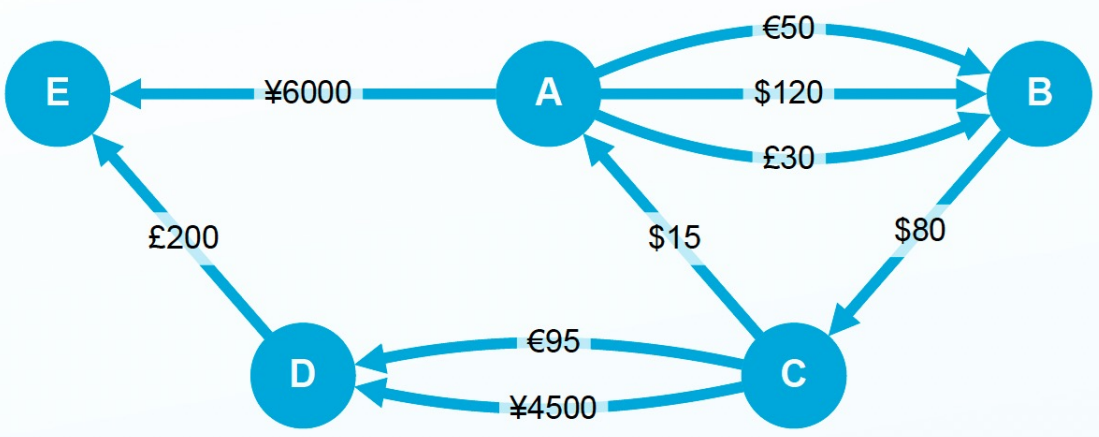


Soft-perturbed multigraph $\tilde{\mathcal{G}} = (\tilde{\boldsymbol{B}}, \boldsymbol{X}, \boldsymbol{Z})$

Can we do the same for multigraphs?

Dennis Heijmans et al. “Towards Adversarially Robust Multigraph Neural Networks for Financial Crime Analysis”, LoG 2026.

Structural Perturbations on Multigraphs (I)



Attributed multigraph

	e_0	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8
A	−1	−1	−1	0	0	0	0	−1	1
B	1	1	1	−1	0	0	0	0	0
C	0	0	0	1	−1	−1	0	0	−1
D	0	0	0	0	1	1	−1	0	0
E	0	0	0	0	0	0	1	1	0

Incidence matrix

Key Contribution: Incidence-Matrix Based Multi-GNN Formulation

Dennis Heijmans et al. “Towards Adversarially Robust Multigraph Neural Networks for Financial Crime Analysis”, LoG 2026.

Structural Perturbations on Multigraphs (II)

	e_0	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8
A	−1	−1	−1	0	0	0	0	−1	1
B	1	1	1	−1	0	0	0	0	0
C	0	0	0	1	−1	−1	0	0	−1
D	0	0	0	0	1	1	−1	0	0
E	0	0	0	0	0	0	1	1	0

Incidence matrix

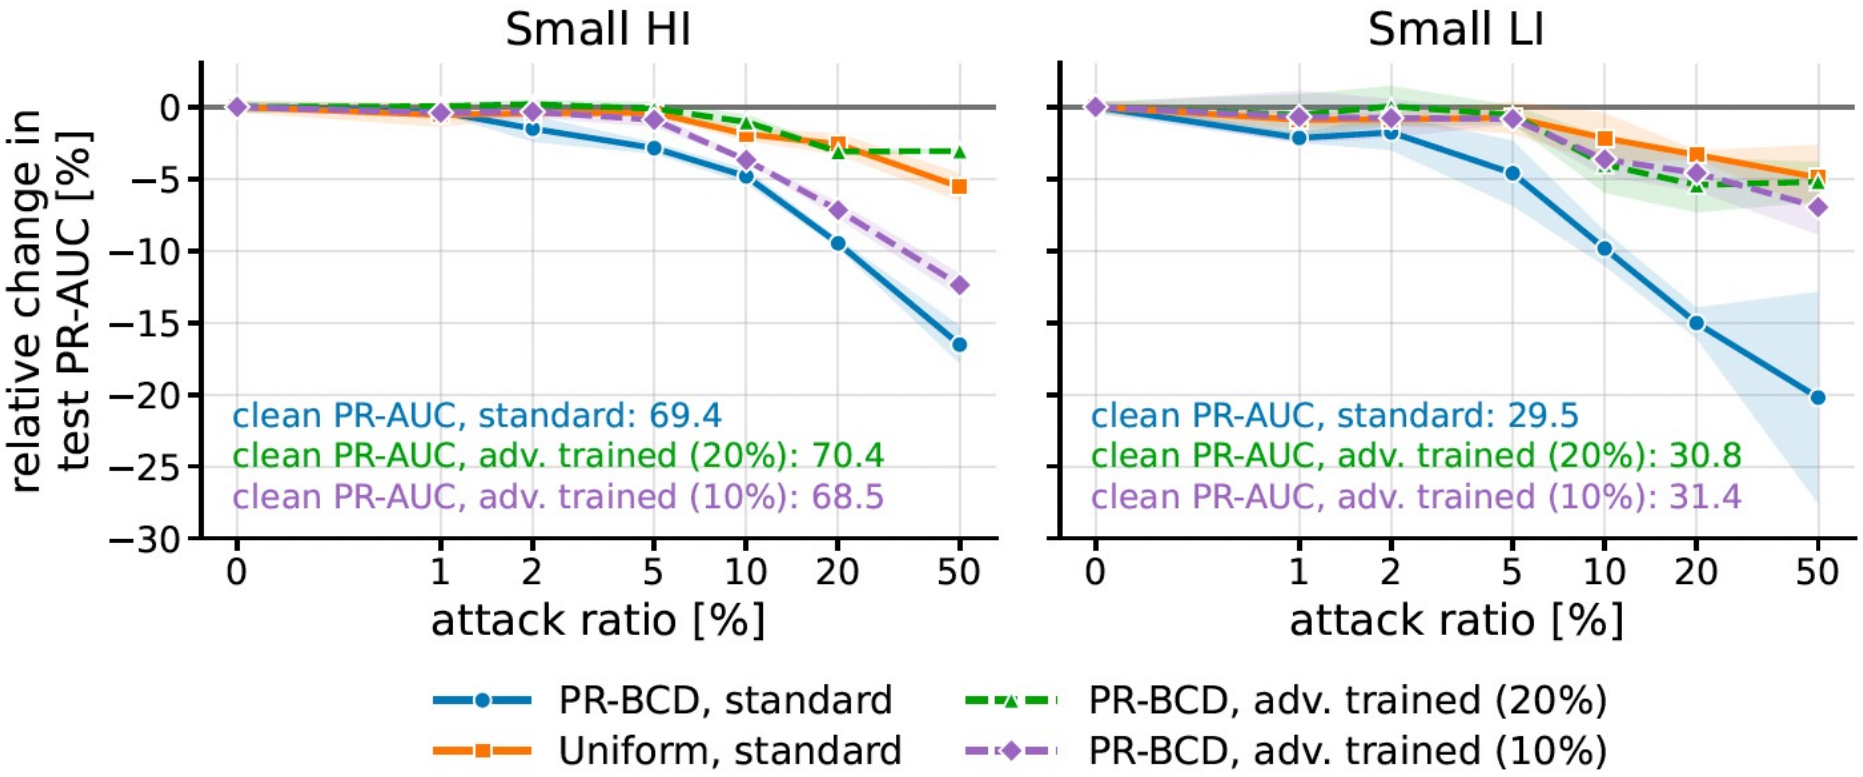
	e_0	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8
A	−1	−1	−1	0.15	0.22	0.05	0.40	−1	0.50
B	0.75	0.50	0.25	−1	0.15	0.07	0.30	0.20	0.20
C	0.15	0.25	0.38	0.50	−1	−1	0.30	0.15	−1
D	0.05	0.15	0.22	0.20	0.25	0.75	−1	0.15	0.15
E	0.05	0.10	0.15	0.15	0.38	0.13	0	0.50	0.15

Soft-perturbed incidence matrix

White Box Edge-rewiring attacks (PR-BCD to exploit the sparsity)

Dennis Heijmans et al. “Towards Adversarially Robust Multigraph Neural Networks for Financial Crime Analysis”, LoG 2026.

Results on AMLWorld Datasets



Edge-rewiring only between criminal accounts

Dennis Heijmans et al. “Towards Adversarially Robust Multigraph Neural Networks for Financial Crime Analysis”, LoG 2026.

What have we ignored so far? (III)

Criminal networks can cross institutional and jurisdictional boundaries!

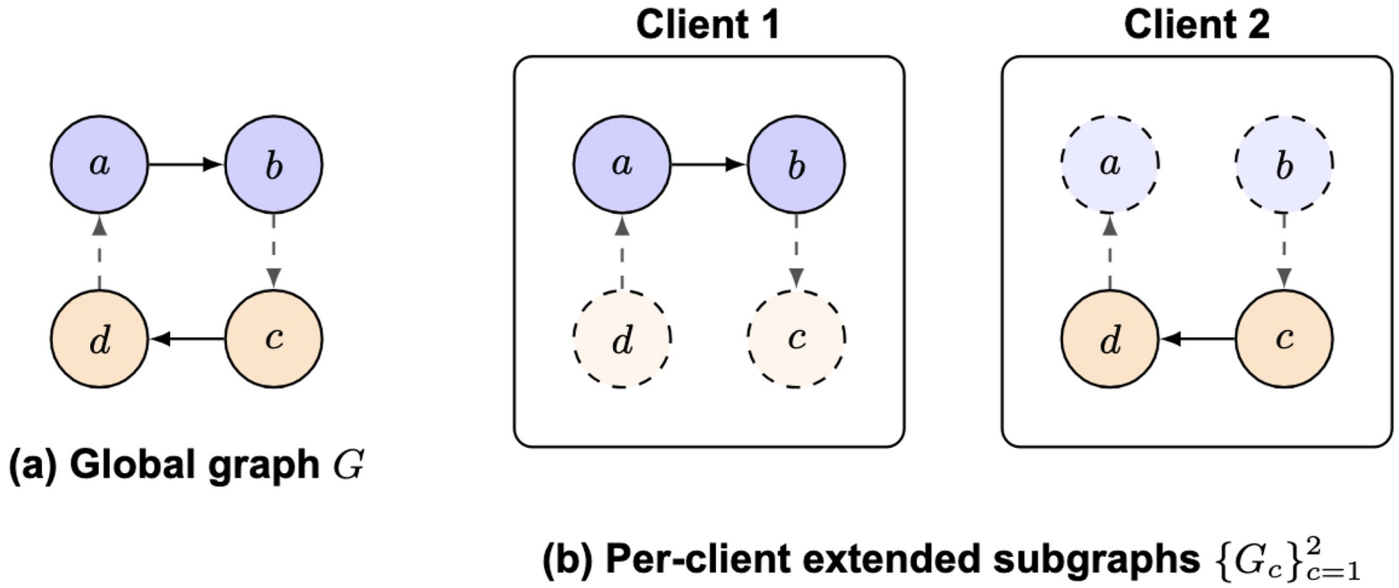
Financial transaction data is distributed across institutions and jurisdictions

Can we train decentralized GNN models as an effective countermeasure?

Can we detect subgraph patterns that cross-institutional boundaries?

Decentralization: Structural Observability Problem

Suppose each node has access to its immediate neighborhood, including the remote neighbors.

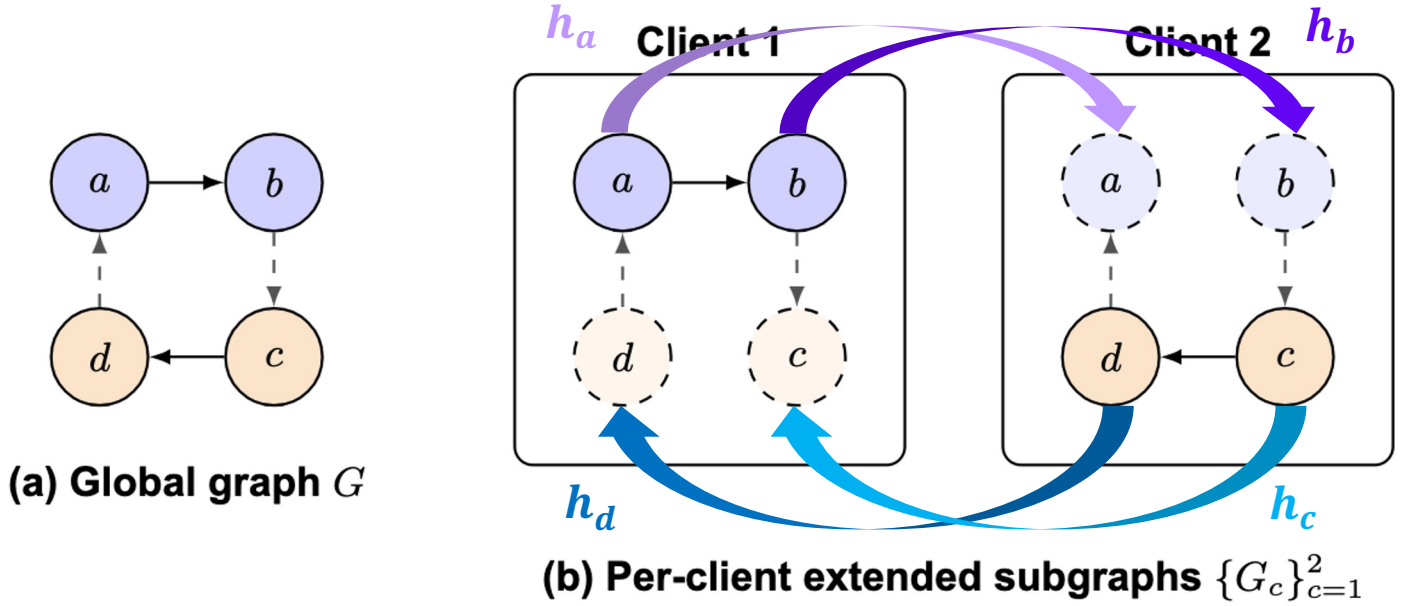


Clients observe the boundary structure but cannot reproduce the information that flows across it.

Two consequences: (i) incorrect node representations during the forward pass
and (ii) incomplete learning signals during the backward pass.

Recovering the Missing Information

Our Solution: Layer-wise Exchange Across Cross-Client Edges.



(a) Global graph G

(b) Per-client extended subgraphs $\{G_c\}_{c=1}^{2}$

What remains private:
- Graph topology
- Node/edge features
- Node/edge labels

Exchanges (i) **embeddings** during forward pass

(ii) **gradients** during backward pass.

Selin Ceydeli, “Layer-wise Exchange for Subgraph Federated Learning”, TU Delft MSc Thesis, June 2026.
Paper version on arXiv:2605.06433 since May 2026.
Financial AI Security, Privacy and Safety Workshop @ ICAIF 2026.

.

What have we ignored so far? (IV)

Criminal networks can cross institutional and jurisdictional boundaries!

Raw transaction data cannot be shared (privacy concerns and regulations)

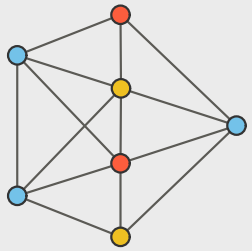
How can we preserve the privacy of during model training and inference?

Can we execute financial crime models using fully homomorphic encryption?

Privacy-Preserving GNN Inference

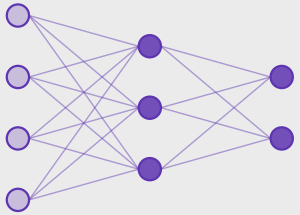
We consider a trusted third-party (e.g., AMLA) that performs the GNN **inference**.

What the server has?



Global graph structure

- Connectivity
- Who->Whom

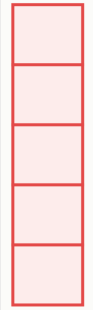


GNN model

- Pre-trained

Unencrypted

Remains private!



Transaction and account features

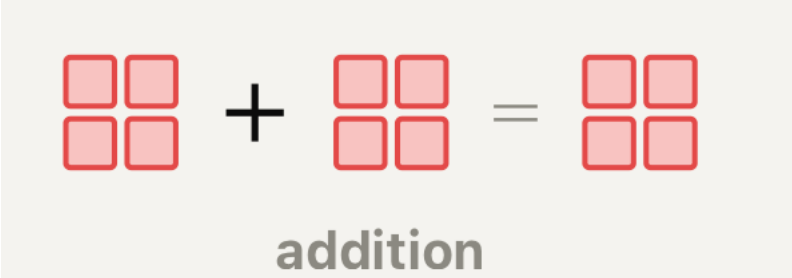
- Amount
- Timestamp
- Currency
- ...

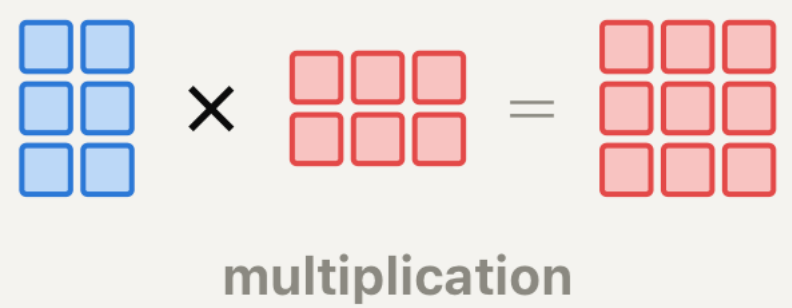
Encrypted

H. Çağrı Bilgi et al. “Privacy-Preserving GNN Inference on Financial Transaction Graphs via Homomorphic Encryption”, arXiv:2610.18054.
Financial AI Security, Privacy and Safety Workshop @ ICAIF 2026..

What Fully Homomorphic Encryption Allows?

Homomorphic encryption allows computation directly on **encrypted** values.





FHE CONSTRAINTS

No direct random access. Use vectorized matrix operations!

No comparison. ReLU computation must be approximated!

This approximation can be done without any loss of utility!

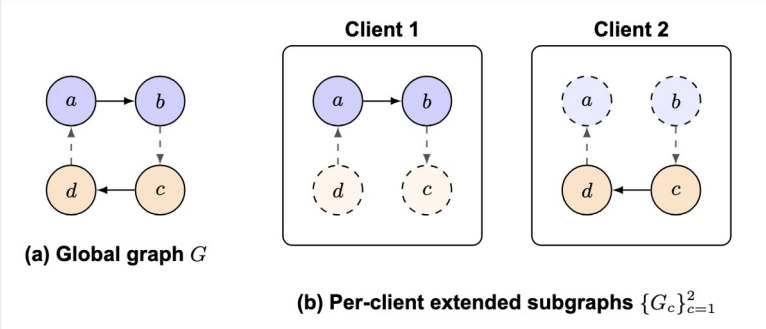
1.65 s
per transaction (batch size 16)

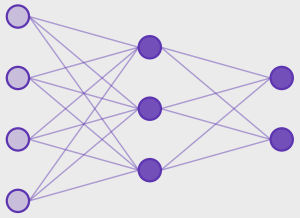
NVIDIA RTX 5090 GPU

H. Çağrı Bilgi et al. “Privacy-Preserving GNN Inference on Financial Transaction Graphs via Homomorphic Encryption”, arXiv:2610.18054.
Financial AI Security, Privacy and Safety Workshop @ ICAIF 2026..

Can we put things together? Yes!

Decentralization: No global graph

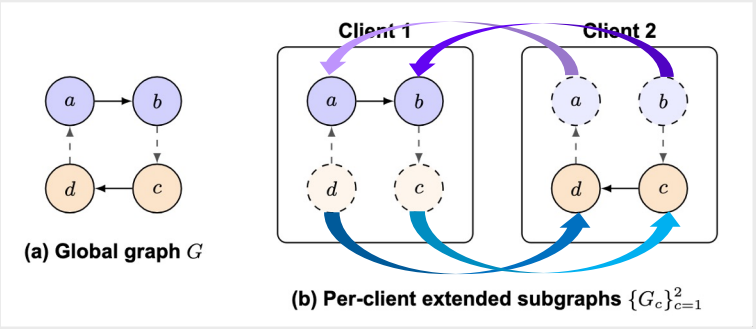




GNN model
- Pre-trained

✓ Topology not shared

Input Features and Activations



- Amount
- Timestamp
- Currency

✗ Encrypted

Decentralized inference on encrypted features & activations!

Conclusions

- Model utility has significantly improved in the past few years
  - New GNN and graph transformer architectures, LLMs
- Adversarially robust models are not out of reach
  - Adversarial training helps limit the utility loss
- Decentralised model training is possible
  - Layer-wise exchange prevents the utility loss
  - Communication cost? Privacy leakage?
- Preserving the privacy without a loss of utility
  - Fully homomorphic GNN inference possible today
  - Differential privacy? How much utility will be lost?

Financial Crime Detection:

Utility, Robustness, Decentralization & Privacy

November 15, 2026
Financial AI Security, Privacy & Safety
An ICAIF 2026 Workshop

Kubilay Atasu
Associate Professor, Software Technology Dept.

TU Delft Delft University of Technology

