xFraud: Explainable Fraud Transaction Detection

.....

About the paper

Authors: Susie Xi Rao, Shuai Zhang, Zhichao Han, Zitao Zhang, Wei Min, Zhiyao Chen, Yinan Shan, Yang Zhao, and Ce Zhang.

Affiliation: ETH Zurich, eBay China

- an explainable Fraud transaction prediction framework using Graph Neural Network (GNN)
- Proceedings of the VLDB Endowment, 2021

Intro to GNN

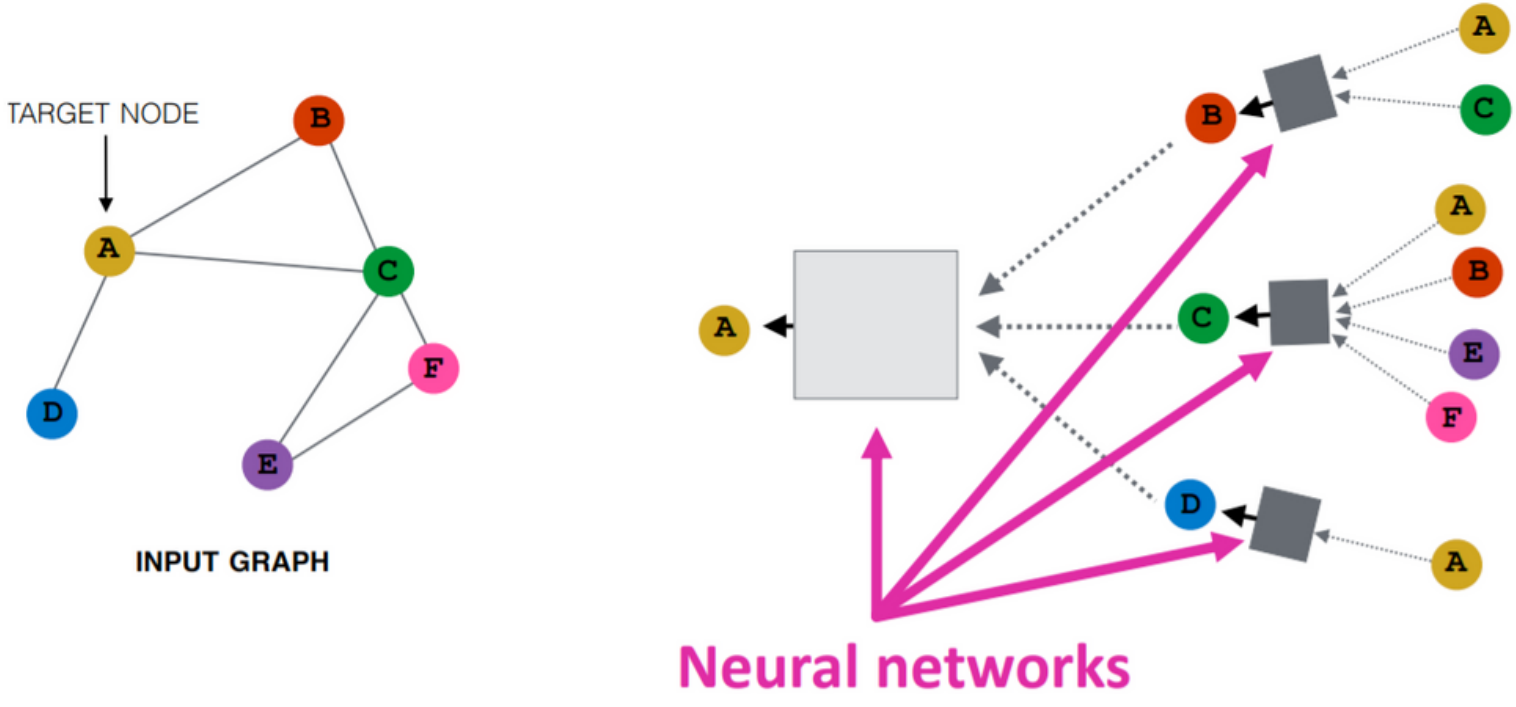

- GNNs are inspired by graph theory and convolutional neural networks.
- It takes the input graph comprising embeddings for edges, and nodes, and generates the output graph with transformed and updated embeddings by preserving the graph symmetry.

Intro to GNN (Cont'd)

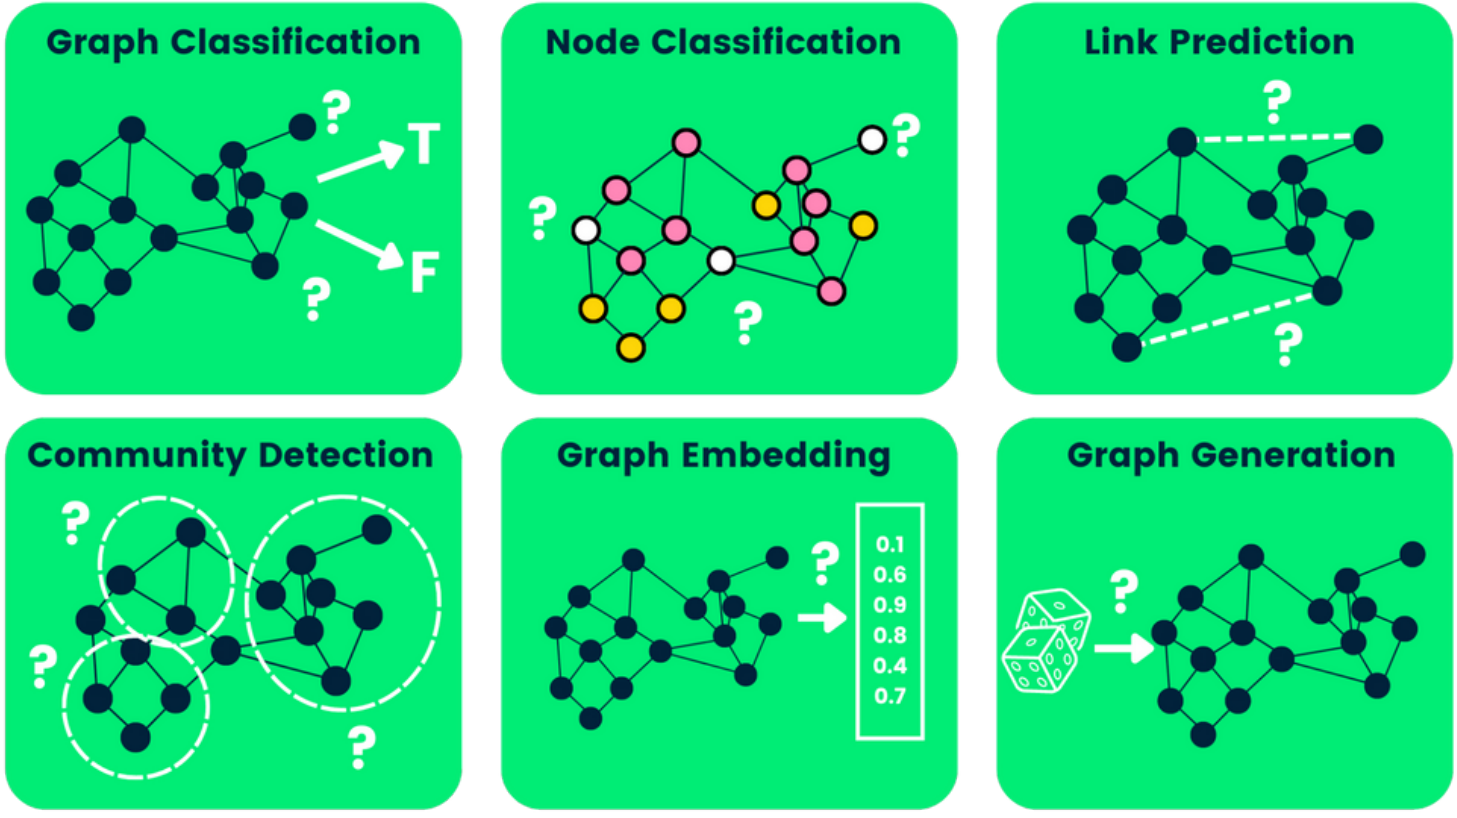

- **Types**: GCN, GAT, GraphSage etc.
- GNNs used in many state-of-art applications of recommender system, knowledge graph, and fraud detection.

Intro to GNN (Cont'd)

Methods	Interpretability	Human Efforts	Generalization	Training Efficiency
Rule-based	✔ ✔	✔ ✔	--	NA
Feature-based ML	✔	✔	✔	✔ ✔
GNN	-	✔ ✔	✔	✔

Goal in xFraud

- uncover fraudulent patterns hard to identify by humans
- assist human experts
- analyze GNN explainability in fraud detection

Contributions

- ideas borrowed from transformer (self-attention) on heterogeneous graphs
- industrial scale in a distributed setting
- quantitative + qualitative analyses in explaining GNN predictions

Previous Work

GEM: Graph Embeddings for Malicious Accounts (2018).

HGT: Heterogeneous Graph Transformer (2020).

Graph-based Fraud Detection Papers and Resources:
https://github.com/safe-graph/graph-fraud-detection-papers

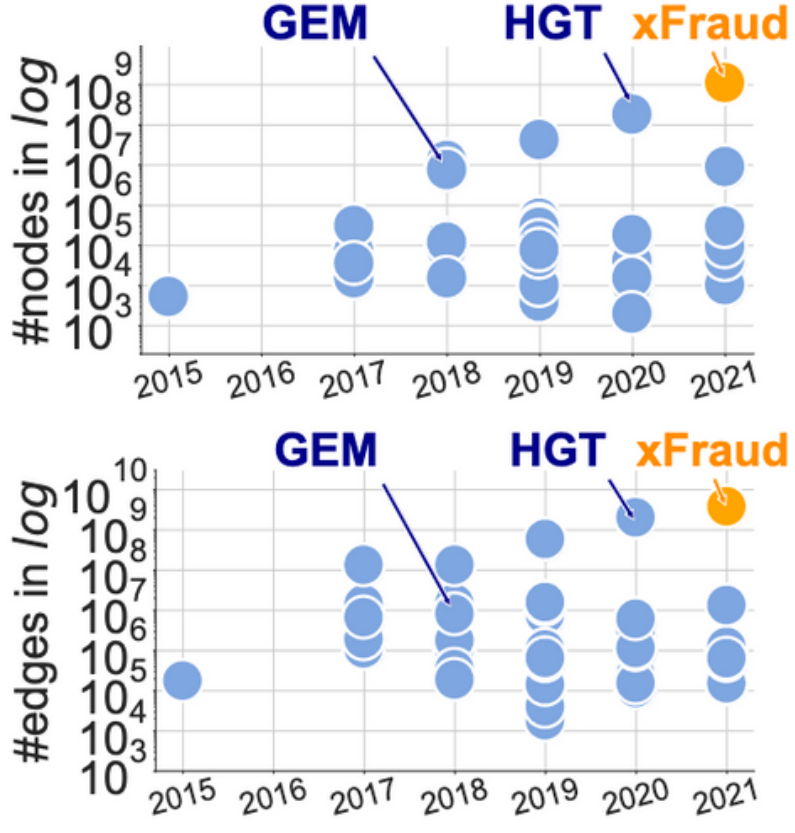

Node and edge numbers (log) of heterogeneous graph datasets in the literature.

Dataset

("B":billion;"M":million;"K":thousand)

*The ratio of frauds is only reported on the sampled datasets.

Dataset	Features	Graph type	#Nodes	#Edges	Fraud%*
eBay-xlarge	480	hetero	1.1B	3.7B	4.33%
eBay-small	114	hetero	289K	613K	4.30%
eBay-large	480	hetero	8.9M	13.2M	3.57%

Dataset	Node type	#Count	Node type%
eBay-xlarge	txn	857M	77%
	pmt	81M	7%
	email	72M	6%
	addr	62M	5%
	buyer	69M	5%
eBay-large	txn	3,752,225	42.40%
	pmt	1,180,114	13.30%
	email	1,307,179	14.80%
	addr	1,316,251	14.90%
	buyer	1,302,097	14.60%
eBay-small	txn	207,749	71.90%
	pmt	22,273	7.70%
	email	25,878	9.00%
	addr	7,138	2.40%
	buyer	25,815	9.00%

xFraud pipeline

Transaction	Buyer	Email	Payment token	Shipping address	Transaction features
1	1	john_eth@gmail.com	Credit card	Albert-Einstein-Strasse 1, Zurich	[..., ..., ...]
2	1	john_eth@gmail.com	Payment slip	Hauptstrasse 1, Zurich	[..., ..., ...]

Transactions → a heterogeneous graph

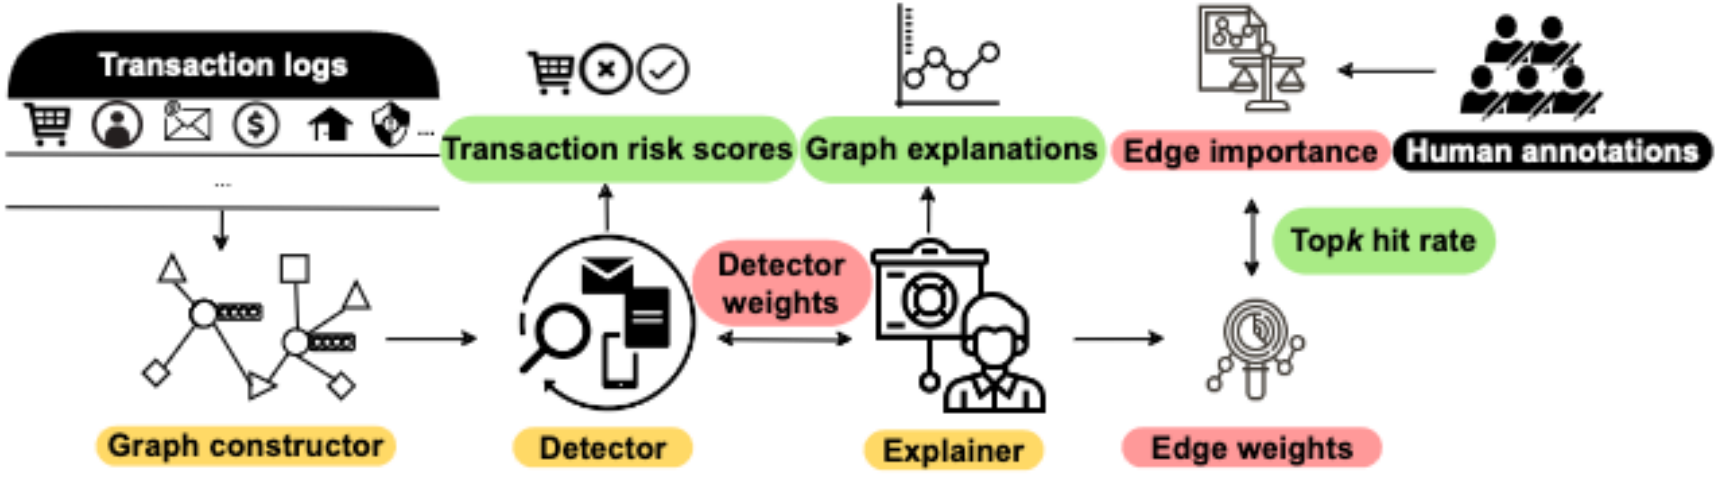

xFraud pipeline (detector + explainer)

Result

# machines	Model	AUC	Training time (s/epoch)	Inference time (s/batch)
8	GAT	0.8879	62.74	0.0557 ± 0.1966
	GEM	0.8961	**61.77**	**0.0167 ± 0.0054**
	xFraud detector+	**0.9074**	70.47	0.0799 ± 0.1868
16	GAT	0.8866	**33.11** (1.89×)	0.0557 ± 0.1966
	GEM	**0.8938**	33.56 (1.84×)	**0.0167 ± 0.0054**
	xFraud detector+	0.8892	38.72 (1.82×)	0.0799 ± 0.1868

- xFraud detector+ outperforms other models (machine #8)
- Accuracy decreases if resources increased (machine #16)
- Future work to develop better distributed algorithms for training heterogeneous graph models.

Result (Cont'd)

xFraud in real production at eBay

Precision (down-sampled)	Precision (original data)	Recall	BU checks flags	Real fraud
98%	32%	~10%	3	1
95%	16%	~20%	6	1

- eBay-xlarge (without downsampling): Fraud% <0.02%
- For 3 fraud candidates investigated by the business unit, 1 will be a real fraud.
- For every 6 fraud candidates investigated by the business unit, 1 will be a real fraud.

Result (Cont'd)

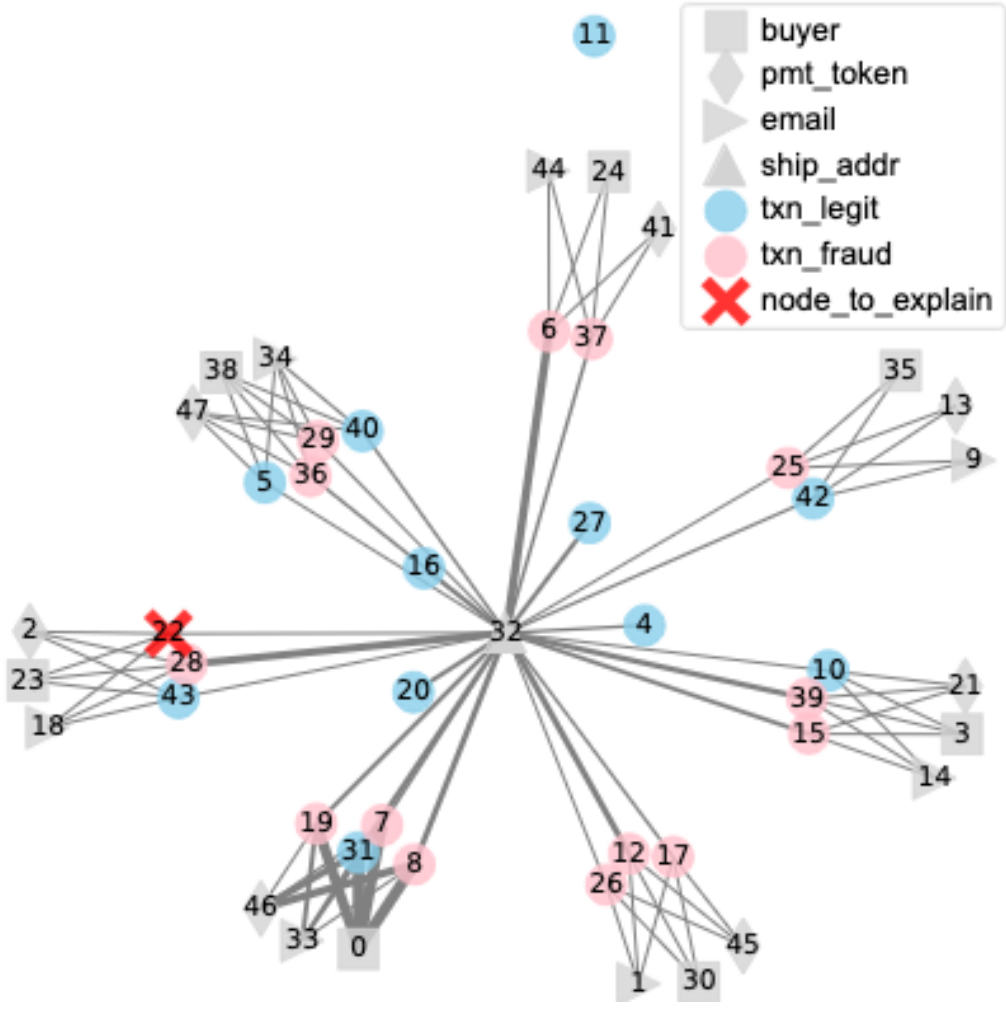

- Thicker ones are highly suspicious and should be under more detailed examinations.

Limitations

- Can't make accurate predictions if guest checkout allowed.
- May not work well for dynamic setting or real-time detection (Further study, BRIGHT, proposed complex Lambda architecture to address this).

Additional Resources

[1] Mingxuan Lu, Zhichao Han, Susie Xi Rao, Zitao Zhang, Yang Zhao, Yinan Shan, Ramesh Raghunathan, Ce Zhang, and Jiawei Jiang. 2022. **BRIGHT** - Graph Neural Networks in Real-time Fraud Detection. In Proceedings of the 31st ACM International Conference on Information & Knowledge Management (CIKM '22).

[2] Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. 2019. **Gnnexplainer**: Generating explanations for graph neural networks. In Advances in neural information processing systems. 9244–9255.

[3] Talk: https://www.youtube.com/watch?v=oHGxcG5Uo7Y

[4] Graph-based Fraud Detection Papers and Resources: https://github.com/safe-graph/graph-fraud-detection-papers