

[Lightning Talk @AML D EPFL 2022]

xFraud: Explainable Fraud Transaction Detection

Susie Xi RAO

srao@ethz.ch

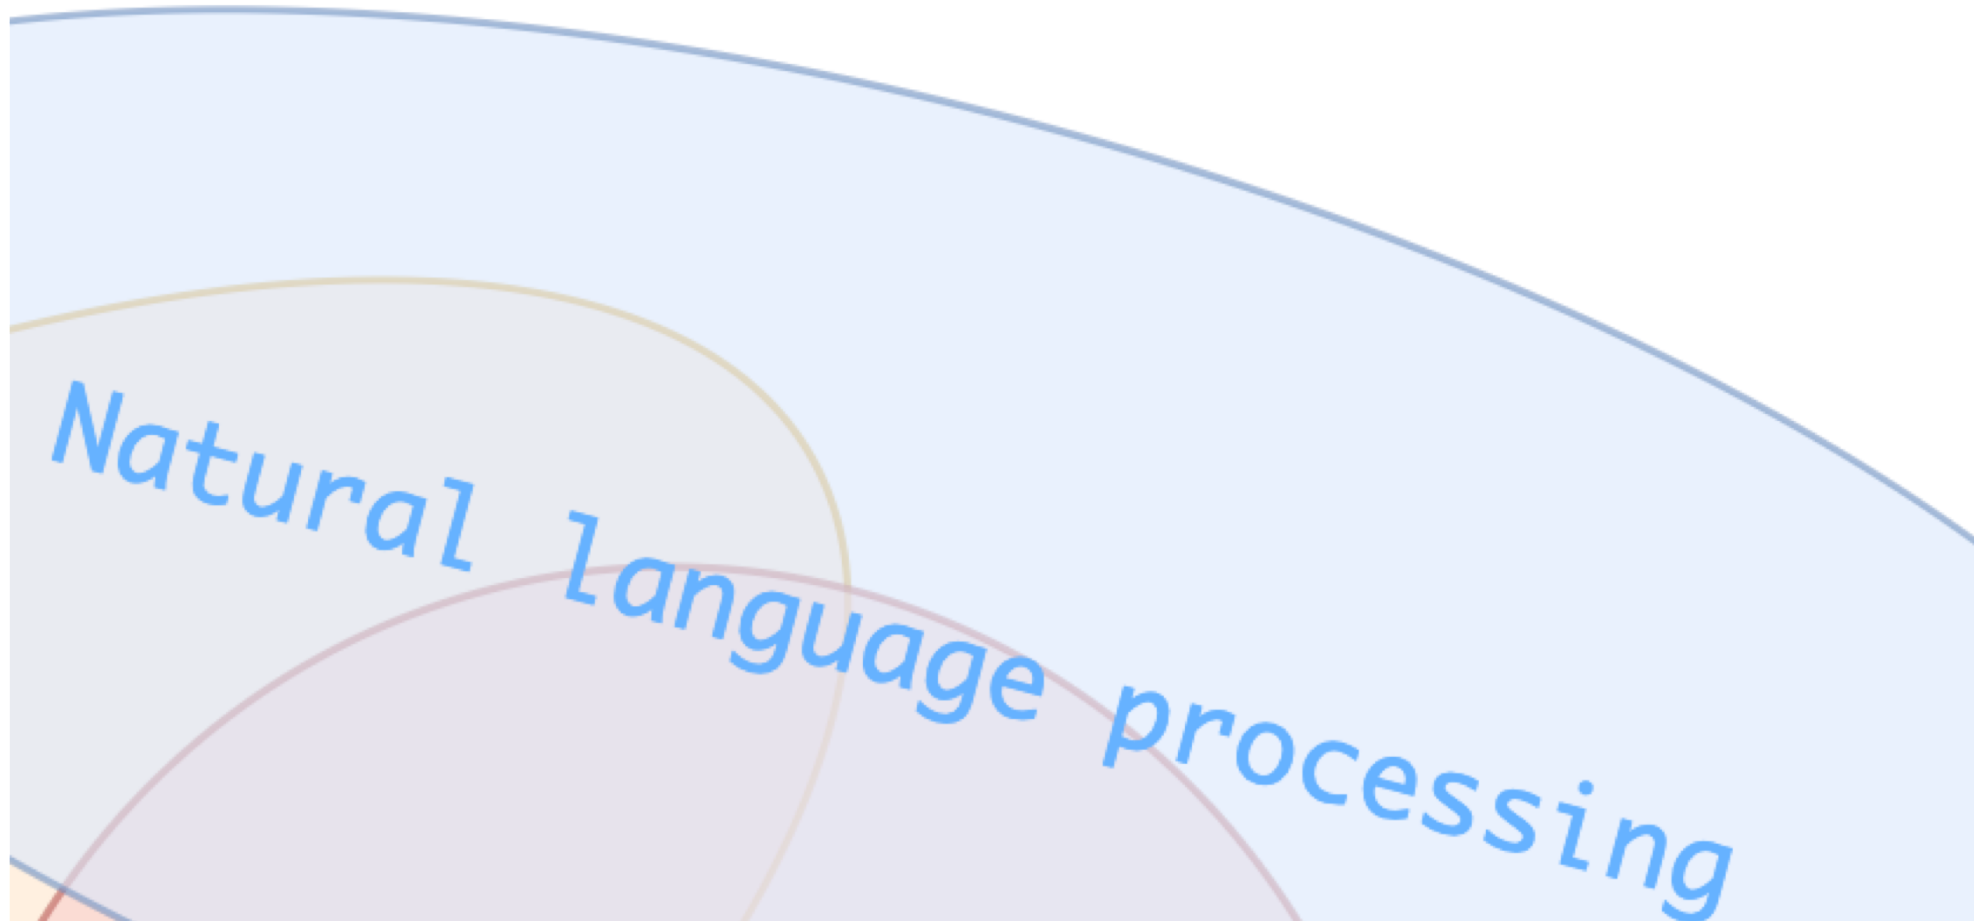
<https://susierao.github.io>

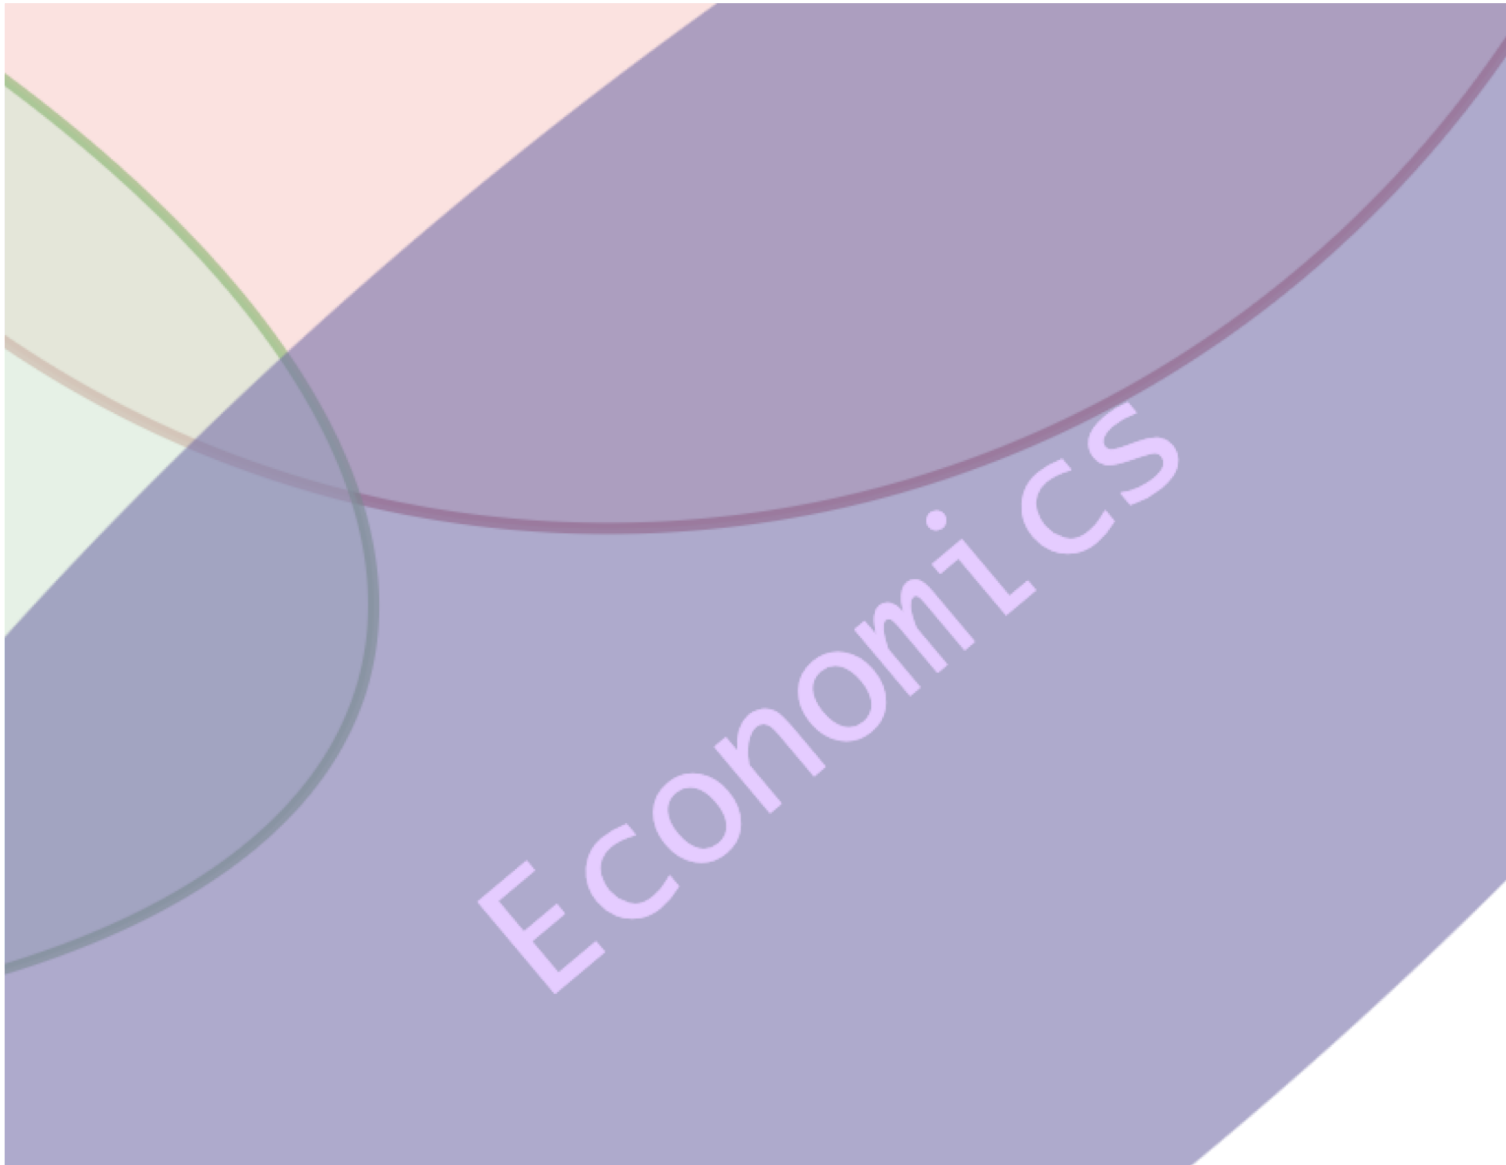
PhD Researcher, ETH Zurich (MTEC/INFK)

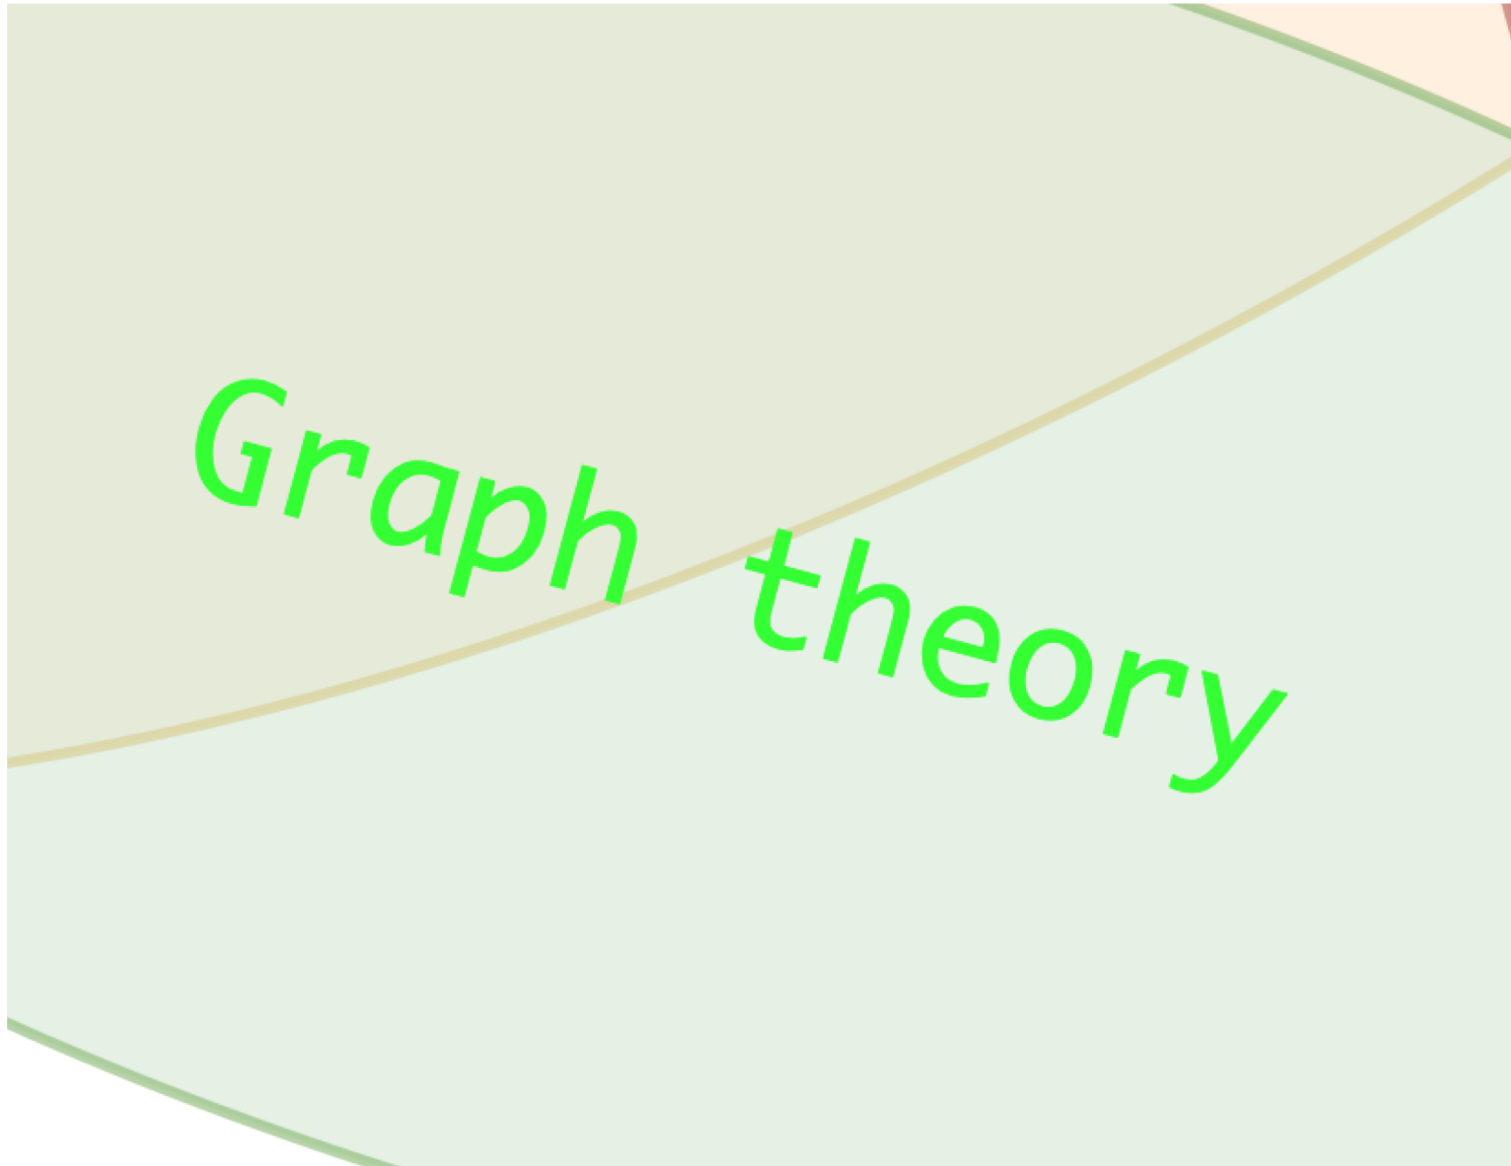
March 28, 2022

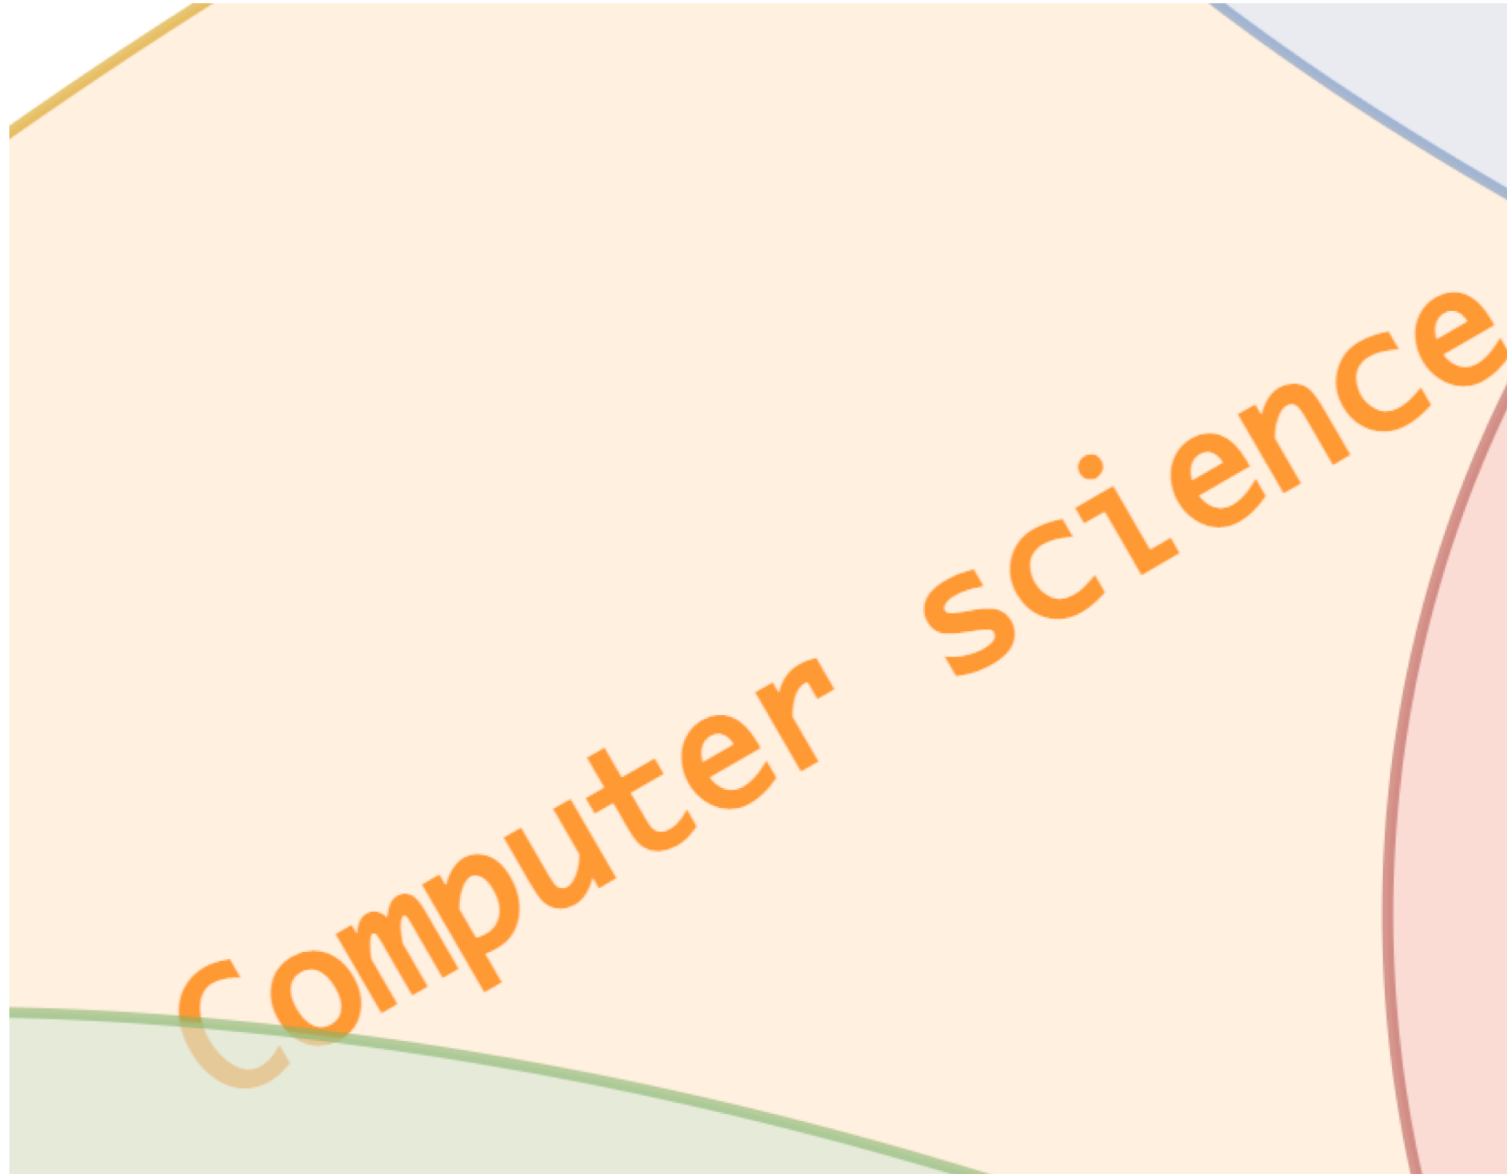


Self-introduction



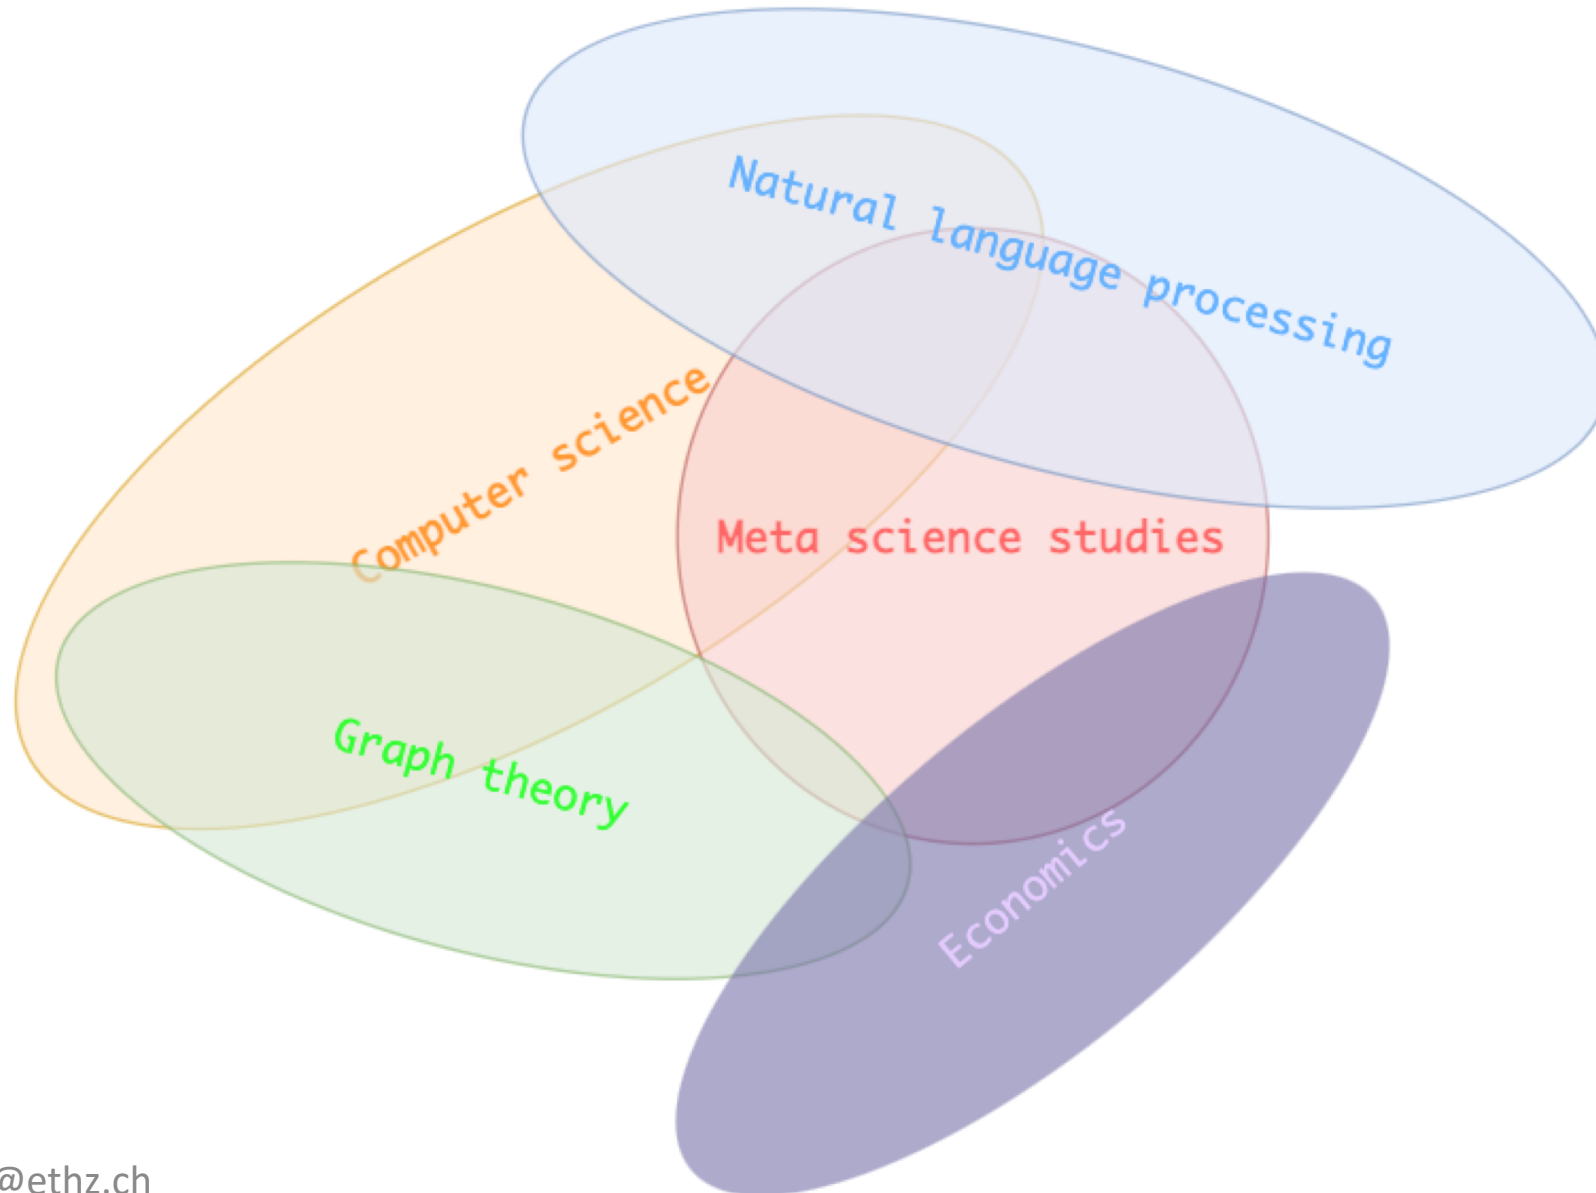








Meta science studies



Pattern mining in graph structured data in scholarly networks, citations networks, transaction networks in e-commerce, and firm networks.

Research focus

xFraud

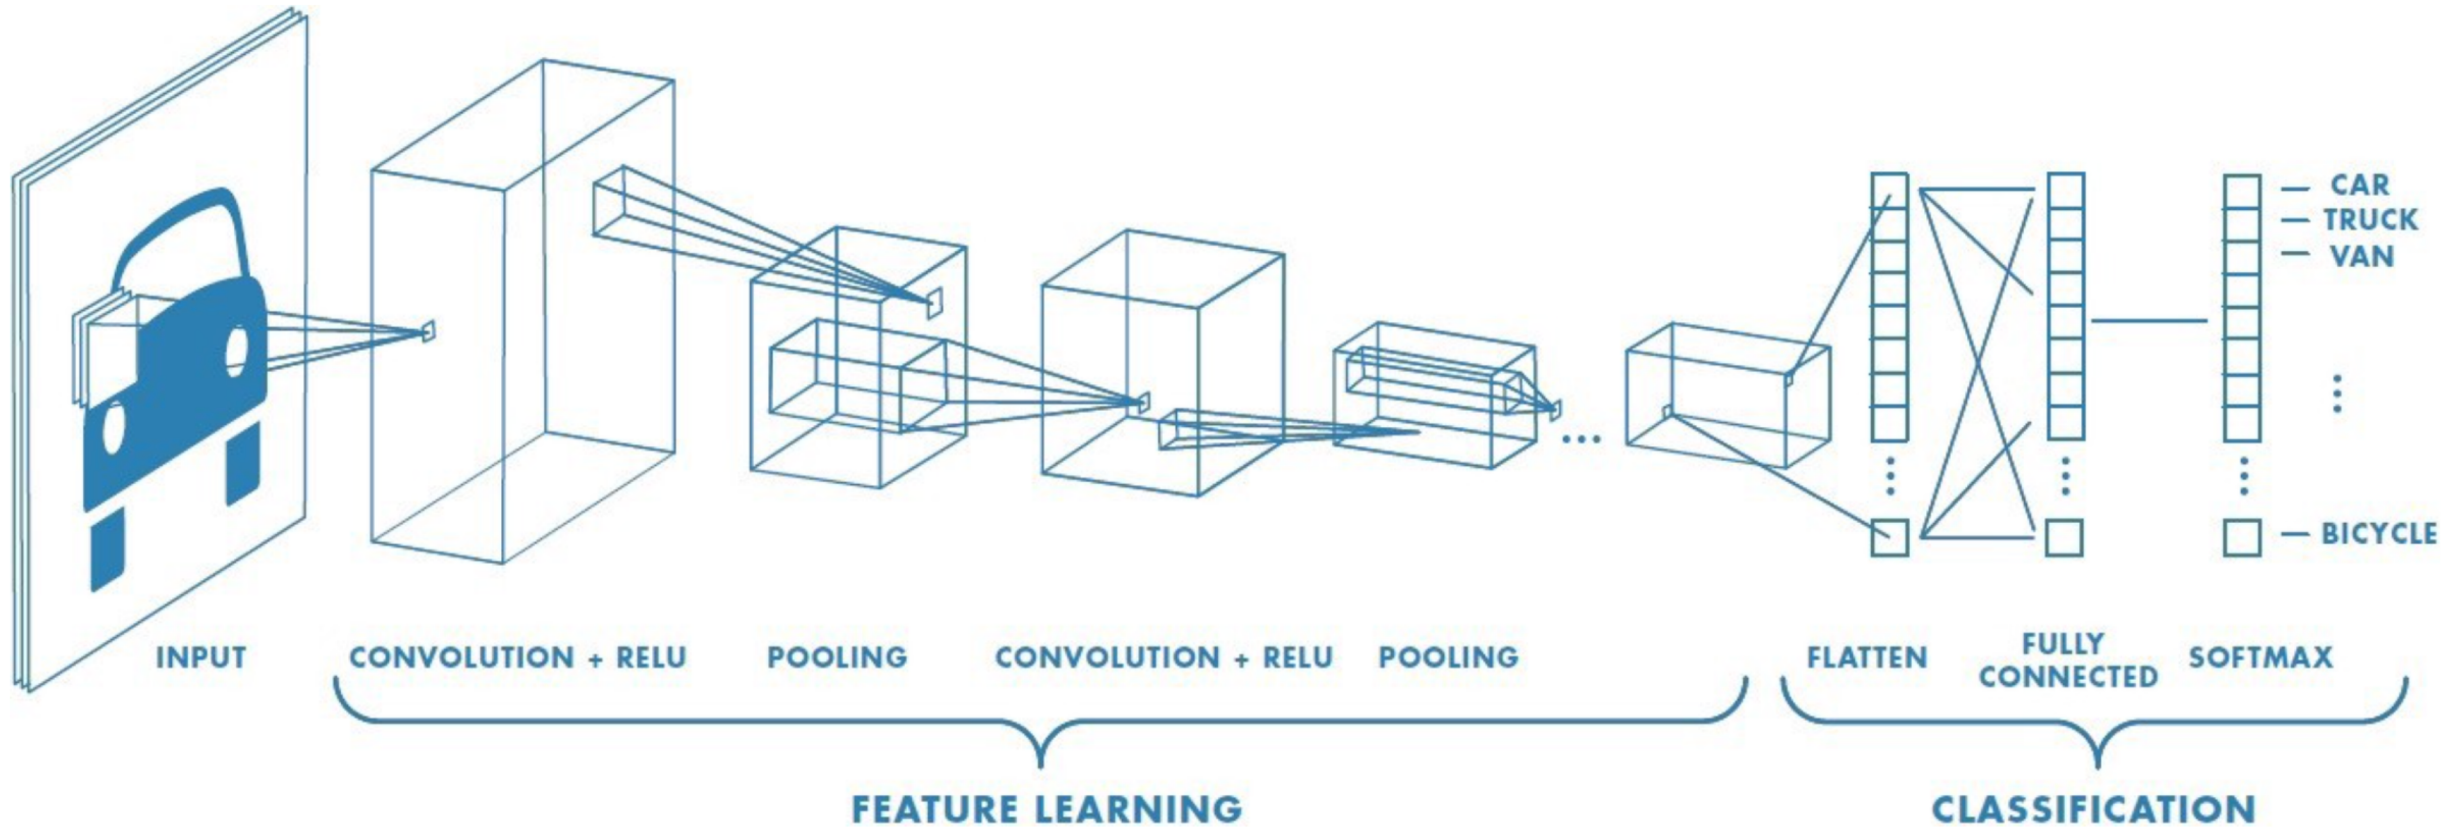
- an explainable *Fraud* transaction prediction framework using Graph Neural Network (GNN)
- accepted and forthcoming at *VLDB 2022*



A (very) brief introduction to GNN


- GCN, GAT, GraphSage
- The use of GNNs in many state-of-the-art applications of recommender system, knowledge graph, and fraud detection.

What we remember from convolutions on images





Source: <https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53>

Graph Convolutional Network (GCN) Kipf & Welling (2016)



Source: <https://tkipf.github.io/graph-convolutional-networks/>


Graph Attention Network (GAT) Veličković et al. (2017)




Multi-headed (self-)attention

Source: <https://arxiv.org/pdf/1710.10903v3.pdf>
<https://arxiv.org/pdf/1706.03762v5.pdf>


GraphSage Hamilton, Ying, and Leskovec (2017)



1. Sample neighborhood



2. Aggregate feature information from neighbors



3. Predict graph context and label using aggregated information

Source: <https://arxiv.org/pdf/1706.02216.pdf>

Susie Rao, srao@ethz.ch

GNN vs. rule vs. simple feature-based ML

	interpretability	human efforts	generalization	training efficiency
rule	++	++	--	NA
feature-based ML	+	+	+	++
GNN	-	++	+	+

GNN vs. rule vs. simple feature-based ML

	interpretability	human efforts	generalization	training efficiency
rule	++	++	--	NA
feature-based ML	+	+	+	++
GNN	-	++	+	+


Our GOAL in xFraud

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

xFraud Contributions

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

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



- ▶ **GEM: Graph Embeddings for Malicious accounts.** Liu, Ziqi, et al. "Heterogeneous graph neural networks for malicious account detection." *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 2018.
- ▶ **HGT: Heterogeneous Graph Transformer.** Hu, Ziniu, et al. "Heterogeneous graph transformer." *Proceedings of The Web Conference 2020*. 2020.

Ours: One of the largest heterogeneous graphs in the literature

Dataset summary

(“B”:billion;“M”:million;“K”:thousand)

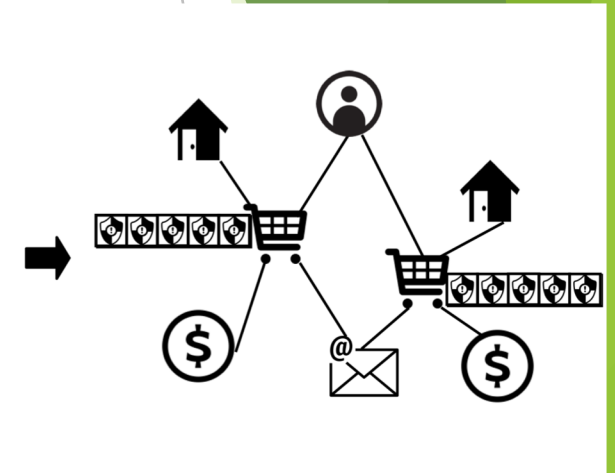
*The ratio of frauds is only reported on the sampled datasets.

Dataset	Features	Graph type	#Nodes	#Edges	Fraud%*
<i>eBay-xlarge</i>	480	hetero	1.1B	3.7B	4.33%

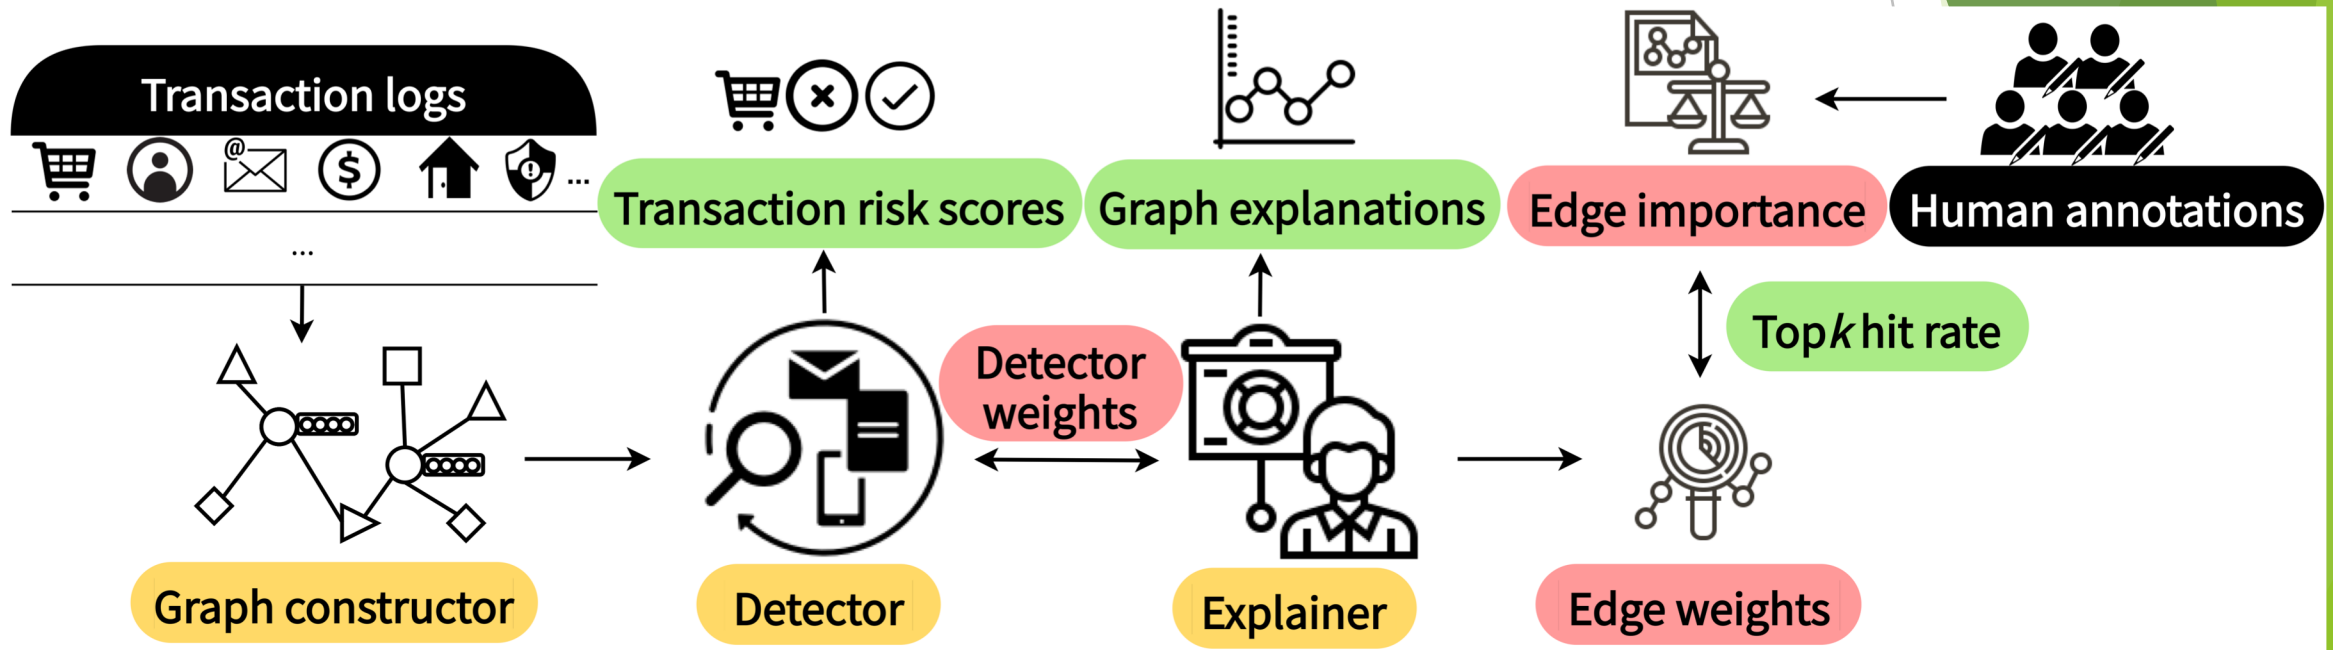
Note: eBay-xlarge (without down-sampling the benign class): Fraud% < 0.02%.

Transaction logs → a transaction graph

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	...



xFraud pipeline (detector and explainer)



End-to-end performance on the dataset *eBay-xlarge* (epochs: 128). We report the average scores over two different seeds (A and B).

# machines	Model	AUC	Training time (s/epoch)	Inference time (s/batch)
8	GAT	0.8830	62.74	0.0557 ± 0.1966
	GEM	0.8938	61.77	0.0167 ± 0.0054
	xFraud detector+	0.9038	70.47	0.0799 ± 0.1868

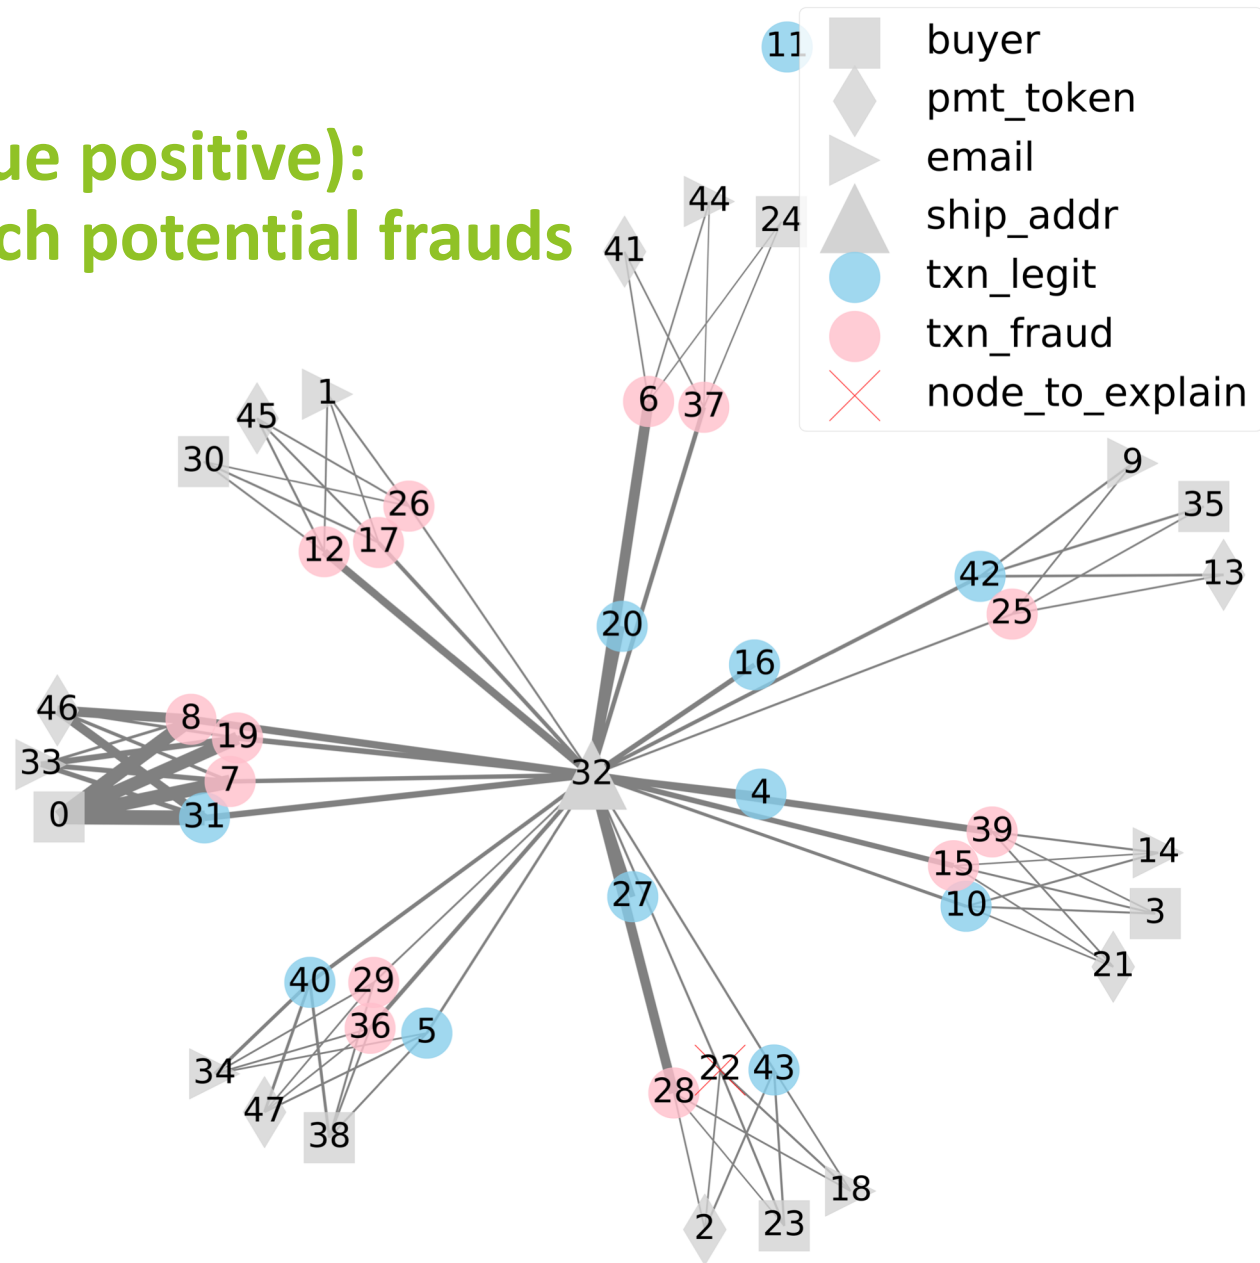
GEM: GCN with a heterogenous graph as input.

xFraud in real production at eBay

- ▶ *eBay-xlarge* (without downsampling the benign class)
 - ▶ Fraud% < 0.02%
- ▶ xFraud's good performance can save time for Business Unit (BU)!

Threshold	Precision (down-sampled)	Precision (original data stream)	Recall	BU checks N fraud candidates	k will be real fraud
0.983	98%	32%	~10%	3	1
0.977	95%	16%	~20%	6	1

Case study (true positive): xFraud helps to catch potential frauds



Rao, Susie Xi, et al. "xFraud: explainable fraud transaction detection."
Proceedings of the VLDB Endowment 15.3 (2021): 427-436.

<https://github.com/eBay/xFraud>

SCAN ME



Susie Xi RAO ^{1,2}
PhD Researcher

¹ Department of Computer Science

² Department of Management, Technology, and Economics

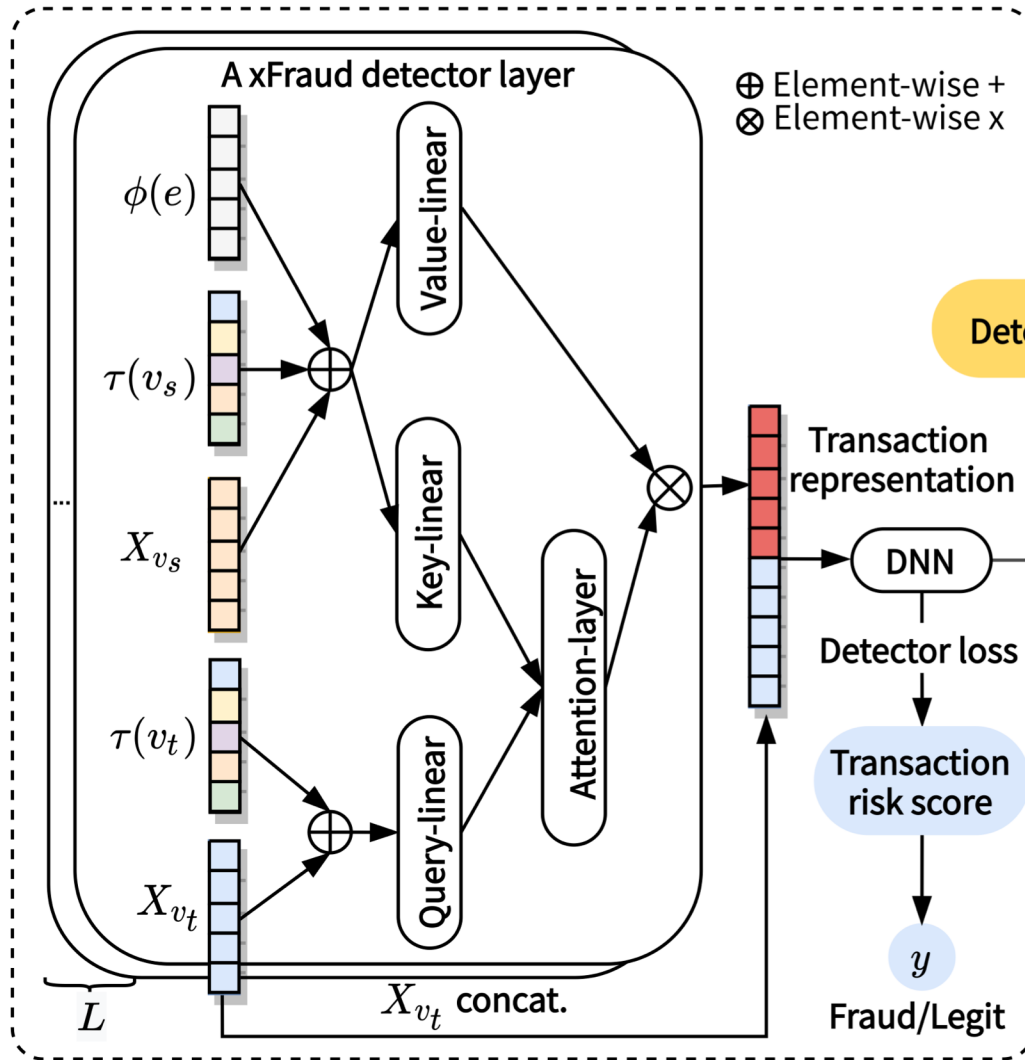
ETH Zurich

srao@ethz.ch


<https://susierao.github.io>

Appendix

xFraud detector and explainer



Detector



Hybrid Explainer

X : transaction features
 s : source node
 t : target node
 $\tau(v)$: node type embeddings
 $\phi(e)$: edge type embeddings

Distributed training of detector

