Lu Lin

Assistant Professor, College of Information Sciences and Technology
The Pennsylvania State University
E373 Westgate Building, University Park, PA 16802
lxl5598@psu.edu
(814) 867-4211
https://louise-lulin.github.io

RESEARCH

My research focuses on algorithmic foundations to enable machine learning models working reliably in the open world with noisy, biased and out-of-distribution inputs. I am fascinated by understanding the behavior of advanced AI diagrams (e.g., large language models, self-supervised learning, etc) and building trustworthiness in them.

EDUCATION

University of Virginia (UVa), Charlottesville, VA 09/2017 - 08/2022
Ph.D. in Computer Science

Beihang University, Beijing, China 08/2014 - 04/2017
M.S. in Computer Science

Beihang University, Beijing, China 09/2010 - 06/2014
B.S. in Computer Science

EMPLOYMENT

The Pennsylvania State University - University Park 09/2022 - present
College of Information Sciences and Technology
Tenure-track Assistant Professor

University of Virginia 08/2017 - 08/2022
Department of Computer Science
Research Assistant

Pinterest Labs, San Francisco, CA 06/2021 - 09/2021
Machine Learning Research Intern

LinkedIn, Sunnyvale, CA 06/2020 - 09/2020
Machine Learning Intern

DiDi AI Lab 09/2015 - 05/2016
Machine Learning Intern

GRANTS

[1] $660,920, **Co-PI** (PI: Jinghui Chen). *"Evaluating and Improving Trustworthiness in Large Language Models"*. Department of Homeland Security (Subaward from Arizona State University). 01/2024 - 12/2026.

[2] $32,244, **Sole PI**. *"Towards Adversarially Robust Anomaly Detection Through Diffusion Model"*. College of Information Sciences and Technology, The Pennsylvania State University. 07/2023 - 06/2024.

PUBLICATIONS

[1] Yi Nian, Yurui Chang, Wei Jin, **Lu Lin**, *Globally Interpretable Graph Learning via Distribution Matching*, in Proceedings of the ACM Web Conference (**WWW**), 2024.

[2] Zengyi Wo, Minglai Shao, Wenjun Wang, Xuan Guo, **Lu Lin**, *Graph Contrastive Learning via Interventional View Generation*, in Proceedings of the ACM Web Conference (**WWW**), 2024.

[3] Weiyu Sun, Xinyu Zhang, Hao Lu, Ying-Cong Chen, Ting Wang, Jinghui Chen and **Lu Lin**, *Backdoor Contrastive Learning via Bi-level Trigger Optimization*, in Proceedings of the 25th International Conference on Learning Representations (**ICLR**), 2024.

[4] Hangfan Zhang, Jinyuan Jia, Jinghui Chen, **Lu Lin** and Dinghao Wu, *A3FL: Adversarially Adaptive Backdoor Attacks to Federated Learning*, in Proceedings of the 37th Conference on Neural Information Processing Systems (**NeurIPS**), 2023.

[5] Songtao Liu, Zhengkai Tu, Minkai Xu, Zuobai Zhang, Peilin Zhao, Jian Tang, Zhitao Ying, **Lu Lin**, Dinghao Wu, *FusionRetro: Molecule Representation Fusion via Reaction Graph for Retrosynthetic Planning*, in Proceedings of the 40th International Conference on Machine Learning (**ICML**), 2023.

[6] Hangfan Zhang, Jinghui Chen, **Lu Lin**, Jinyuan Jia, Dinghao Wu, *Graph Contrastive Backdoor Attacks*, in Proceedings of the 40th International Conference on Machine Learning (**ICML**), 2023.

[7] **Lu Lin**, Jinghui Chen and Hongning Wang, *Spectral Augmentation for Self-Supervised Learning on Graphs*, International Conference on Learning Representations (**ICLR**), 2023.

[8] **Lu Lin**, Ethan Blaser and Hongning Wang, *Graph Structural Attack by Perturbing Spectral Distance*, in Proceedings of the 28th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (**SIGKDD**), Washington DC, USA, 2022.

[9] Yujia Wang, **Lu Lin** and Jinghui Chen, *Communication-Efficient Adaptive Federated Learning*, in Proc. of the 39th International Conference on Machine Learning (**ICML**), Baltimore, Maryland, USA, 2022

[10] Nan Wang*, **Lu Lin***, Jundong Li and Hongning Wang, *Unbiased Graph Embedding with Biased Graph Observations*, in Proceedings of the 31st conference in the International World Wide Web Conference (**WWW**), 2022.

[11] **Lu Lin**, Ethan Blaser and Hongning Wang, *Graph Embedding with Hierarchical Attentive Membership*, in Proceedings of the 15th International Conference on Web Search and Data Mining (**WSDM**), 2022.

[12] Yujia Wang, **Lu Lin** and Jinghui Chen, *Communication-efficient Adaptive Gradient Method for Distributed Nonconvex Optimization*, in Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (**AISTATS**), 2022.
– A short version of this paper also appears on International Workshop on Trustable, Verifiable and Auditable Federated Learning in **AAAI**, 2022.

[13] **Lu Lin**, and Hongning Wang, *Graph Attention Networks over Edge Content-Based Channels*, in Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (**SIGKDD**), 2020.

[14] Gong, Lin, **Lu Lin**, Weihao Song and Hongning Wang, *JNET: Learning User Representations via Joint Network Embedding and Topic Embedding*, in Proceedings of the 13th International Conference on Web Search and Data Mining (**WSDM**), 2020.

[15] **Lu Lin**, Zhen Luo, Dezhi Hong and Hongning Wang, *Sequential learning with active partial labeling for building metadata*, in Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (**BuildSys**), 2019.

[16] **Lu Lin**, Lin Gong, and Hongning Wang, *Learning Personalized Topical Compositions with Item Response Theory*, in Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining (**WSDM**), 2019.

[17] **Lu Lin**, Jianxin Li, Feng Chen, Jieping Ye and Jinpeng Huai, *Road traffic speed prediction: a probabilistic model fusing multi-source data*, in IEEE Transactions on Knowledge and Data Engineering (**TKDE**) 30.7 (2017): 1310-1323.

[18] **Lu Lin**, Jianxin Li, Richong Zhang and Chenggen Sun, *Opinion mining and sentiment analysis in social networks: a retweeting structure-aware approach*, in 2014 IEEE/ACM 7th International Conference on Utility and Cloud Computing (**UCC**), IEEE, 2014.

PRE-PRINTS

[1] Hangfan Zhang, Zhimeng Guo, Huaisheng Zhu, Bochuan Cao, **Lu Lin**, Jinyuan Jia, Jinghui Chen, Dinghao Wu, *On the Safety of Open-Sourced Large Language Models: Does Alignment Really Prevent Them From Being Misused?*, arXiv preprint arXiv:2310.01581, 2023.

[2] Bochuan Cao, Yuanpu Cao, **Lu Lin**, Jinghui Chen, *Defending against alignment-breaking attacks via robustly aligned LLM*, arXiv preprint arXiv:2309.14348, 2023.

[3] Songtao Liu, Rex Ying, Hanze Dong, **Lu Lin**, Jinghui Chen, Dinghao Wu, *How Powerful is Implicit Denoising in Graph Neural Networks*, arXiv preprint arXiv:2209.14514, 2022.

TEACHING

Instructor Spring 2024
PSU DS 340W: Applied Data Sciences (Undergrad)

Instructor Fall 2023
PSU IST 557: Data Mining: Techniques and Applications (Grad)

Instructor Spring 2023
PSU DS 310: Machine Learning for Data Science (Undergrad)

Teaching Assistant Spring 2020
UVa CS 6501: Graph Mining (Grad)

Head Teaching Assistant Fall 2019
UVa CS 6501: Natural Language Processing (Grad)

Head Teaching Assistant Spring 2019
UVa CS 6501: Text Mining (Grad)

AWARDS

• UVa CS John A. Stankovic Graduate Research Award	05/2022
• UVa Computer Science Fellowship	09/2017
• Outstanding Student Graduation Fellowship of Beijing City	06/2014
• Best Bachelor Thesis Award	06/2014

INVITED TALKS

Fairness in Machine Learning
Guest lecture at CS@Yale 03/2023

Graph Structure as A Double-Edged Sword in Machine Learning
Job talk at CS@WM, DS@NJIT, IST@PSU 02/2022

Shopping Journey Modeling via Long-term Checkout Reward
Intern project presentation at Pinterest 08/2021

Representation Learning on Heterogeneous Graphs at LinkedIn
Intern project presentation at LinkedIn 08/2020

Graph Attention Networks over Edge Content-based Channels
Oral presentation at Mining Text, Web & Social Media Session, KDD 08/2020

Graph Attention Networks over Edge Channels
Deep dive talk at LinkedIn 07/2020

Learning User Representation via Joint Network Embedding and Topic Embedding
Paper lightning talk, WSDM 02/2020

Learning Personalized Topical Compositions with Item Response Theory
Oral presentation, WSDM 02/2019

Opinion Mining and Sentiment Analysis in Social Networks
Oral presentation, UCC 02/2015

ACADEMIC SERVICES

Program Committee

- AAAI: The AAAI Conference on Artificial Intelligence 2021, 2022, 2023
- IJCAI: International Joint Conference on Artificial Intelligence 2021, 2022

Conference Reviewer

- NeurIPS: Neural Information Processing Systems 2022
- ICLR: International Conference on Learning Representations 2023
- AISTATS: International Conference on Artificial Intelligence & Statistics 2021, 2022
- KDD: ACM Conference on Knowledge Discovery & Data Mining 2021, 2022
- SIGIR: Conference on Research & Development in Information Retrieval 2021
- WSDM: ACM International Conference on Web Search & Data Mining 2021, 2023
- LoG: Learning on Graphs Conference 2022
- BigData: IEEE International Conference on Big Data 2021

Journal Reviewer

- TKDE: IEEE Transactions on Knowledge and Data Engineering 2021
- JAIS: Journal of Aerospace Information Systems 2022
- TIST: Transactions on Intelligent Systems and Technology 2022

DIVERSITY SERVICES

Graduate Society of Women Engineers (GradSWE) 08/2021 - present
Member and Undergrad Mentor University of Virginia
This committee hosts events to foster a more diverse and inclusive engineering community. I am serving as a mentor for undergrads.

Pinterest Employee Diversity Community 05/2021 - 08/2021
Member Pinterest
I joined the community as a member during my internship at Pinterest and attended fireside chats to learn and share experiences.

COMPETITIONS

- Winner of Yelp Dataset Challenge 2020 (Worldwide) $5000
  - The only winner of the year 2020
  - Technique is based on our WSDM 2020 paper "JNET: Learning User Representations via Joint Network Embedding and Topic Embedding"