

SIHONG XIE

Room 303, W4 Building, HKUST-GZ, Guangzhou, Guangdong, China

Email: xiesihong1@gmail.com

A. BIOGRAPHY

Education

University of Illinois at Chicago

PhD in Computer Science

Sep, 2010 - Aug, 2016

Sun Yat-Sen University, Guangzhou, China

M.E. in Software Engineering

Sep, 2008 - Jun, 2010

B.E. in Software Engineering

Sep, 2004 - Jun, 2008

Employment

HKUST (Guangzhou), Guangzhou, Guangdong, China

Jun, 2023 - present

Associate Professor (with tenure), AI Thrust, Information Hub

Lehigh University, Bethlehem, PA

Aug, 2016 - Jun 2023

Associate Professor (tenured in May 2023), Department of Computer Science and Engineering

Assistant Professor, Department of Computer Science and Engineering

University of Illinois at Chicago, Chicago, IL

Aug, 2012 - Aug, 2016

Research Assistant, Department of Computer Science

Supervisor: Philip S. Yu

LinkedIn Corporation, Mountain View, CA

May, 2013 - Aug, 2013

Applied Researcher, Search and Network Analysis Group

Supervisors: Anmol Bhasin, Baoshi Yan, Mohammad Amin

IBM Research, Hawthorne, NY

May, 2012 - Aug, 2012

Research Intern, Exploratory Stream Analytics Group

Supervisor: Wei Fan

B. PUBLICATIONS

Keys: Doctoral student^D; Master student^M; Undergraduate student^U; Doctoral advisor^{DA}; Industrial partner^I

Journals

- [1] Yue Sun, Chao Chen, Yuesheng Xu, **SiHong Xie**, Rick S. Blum, Parv Venkitasubramaniam On the Generalization Discrepancy of Spatiotemporal Dynamics-informed Graph Convolutional Networks. *Frontiers in Mechanical Engineering*. 2024
- [2] Mengzhu Sun, Xi Zhang, Jianqiang Ma, **SiHong Xie**, Yazheng Liu, and Philip S. Yu^{DA}. Inconsistent Matters: A Knowledge-guided Dual-consistency Network for Multi-modal Rumor Detection *IEEE Transactions on Knowledge and Data Engineering (TKDE)*. 2023. Impact factor: 9.235 (according to <https://www.computer.org/csdl/journal/tk/about/14395>).

- [3] Xi Zhang, Yuan Su, Siyu Qu, **Sihong Xie**, Binxing Fang and Philip S. Yu^{DA}. IAD: Interaction-Aware Diffusion Framework in Social Networks *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 2018. Impact factor: 9.235 (according to <https://www.computer.org/csdl/journal/tk/about/14395>).
- [4] **Sihong Xie** and Philip S. Yu^{DA}. Active zero-shot learning: A novel approach to extreme multi-labeled classification. *International Journal of Data Science and Analytics*, 2016. Impact Factor: 2.52 (according to <https://research.com/journal/international-journal-of-data-science-and-analytics>).
- [5] Senzhang Wang, **Sihong Xie**, Xiaoming Zhang, Zhoujun Li, Philip S. Yu^{DA}, and Yueying He. Coranking the future influence of multiobjects in bibliographic network through mutual reinforcement. *ACM Trans. Intell. Syst. Technol.*, 7(4):64:1–64:28, May 2016. Impact factor: 3.971 (according to <https://dl.acm.org/journal/tist/about>)
- [6] Guan Wang, **Sihong Xie**, Bing Liu, and Philip S. Yu^{DA}. Identify online store review spammers via social review graph. *ACM Trans. Intell. Syst. Technol.*, 3(4):61:1–61:21, September 2012. Impact factor: 3.971 (according to <https://dl.acm.org/journal/tist/about>)

Refereed Conferences

- [7] Rui Xu^D, Yi Chen, Jiujiu Chen^D, and **Sihong Xie***. Information Flow Reveals When to Trust Language Models. The 43rd International Conference on Machine Learning (ICML). 2026. (*Oral, acceptance rate 168/23918≈0.7%*).
- [8] Yazheng Liu^D, Xi Zhang, **Sihong Xie***, and Hui Xiong. Towards the Explainability of Temporal Graph Networks via Memory Backtracking and Topological Attribution. The 43rd International Conference on Machine Learning (ICML). 2026. (*Spotlight, acceptance rate 536/23918≈2.2%*).
- [9] Rufeng Chen^D, Yue Chang^D, Xiaqiang Tang^D, Hechang Chen, and **Sihong Xie***. PSG-Nav: Probabilistic Scene Graph Navigation via Multiverse Decision Making. The 43rd International Conference on Machine Learning (ICML). 2026. (*Poster, acceptance rate 26.6%*).
- [10] Chenghua Guo, **Sihong Xie**, and Xi Zhang. X-EviProbe: Post-hoc Parameter-free Evidential Uncertainty Quantification for Frozen Graph Neural Networks. The 43rd International Conference on Machine Learning (ICML). 2026. (*Poster, acceptance rate 26.6%*).
- [11] Zhaofan Zhang^D, **Sihong Xie**, and Hui Xiong. Adaptive-Critical: Identifying Critical States for Risk-Sensitive Navigation via Distributional Reinforcement Learning. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). 2026. (*Acceptance rate 36%*).
- [12] Zhaofan Zhang^D, Minghao Yang, **Sihong Xie**, and Hui Xiong. Perturbation-Resilient Autonomous Navigation with Distributionally Robust Reinforcement Learning. The 33rd International Joint Conference on Artificial Intelligence (IJCAI). 2026.
- [13] Yue Chang^D, Rufeng Chen^D, Zhaofan Zhang^D, Yi Chen, and **Sihong Xie**. RAG-3DSG: Enhancing 3D Scene Graphs with Re-Shot Guided Retrieval-Augmented Generation. The 19th European Conference on Computer Vision (ECCV). 2026. (*Acceptance rate 27.52% = 2883/10473*).
- [14] Yi Wang^M, Xiaqiang Tang^D, Keyu Hu^U, Haojie Lu, **Sihong Xie***. REFO: Reinforced Evolutionary Faithfulness Optimization for Large Language Models. The Fortieth AAAI Conference on Artificial Intelligence (AAAI-26). 2026. (*Poster paper, acceptance rate 17.52%*).

- [15] Yi Wang^M, Haojie Lu, CHEN Li, **Sihong Xie**^{*}. Deconstructing Spatial Complexity: Hierarchical Decomposition for LLM Spatial Reasoning. The 2026 Conference on Empirical Methods in Natural Language Processing (EMNLP). 2026.
- [16] Yazheng Liu^D, Xi Zhang, **Sihong Xie**^{*} and Hui Xiong. Robust Explanations of Graph Neural Networks via Graph Curvatures. The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS). 2025. (*Poster paper, acceptance rate 24.52%*).
- [17] Adithya Kulkarni, Mohana Chakraborty, **Sihong Xie**^{*} and Qi Li. Budget Allocation Exploiting Label Correlation between Instances. The 41st Conference on Uncertainty in Artificial Intelligence (UAI). 2025. (*Poster paper, acceptance rate 30.7%*).
- [18] Xiaqiang Tang^D, Qiang Gao, Jian Li, Nan Du, Qi Li, and **Sihong Xie**^{*}. MBA-RAG: a Bandit Approach for Adaptive Retrieval-Augmented Generation through Question Complexity. The 31st International Conference on Computational Linguistics (COLING). 2025. (*Short paper, acceptance rate 20.6%*).
- [19] Xiaqiang Tang^D, Jian Li, Keyu Hu^U, Nan Du, Xiaolong Li, Xi Zhang, Weigao Sun, and **Sihong Xie**^{*}. CogniBench: A Legal-inspired Framework and Dataset for Assessing Cognitive Faithfulness of Large Language Models. The 63rd Annual Meeting of the Association for Computational Linguistics (ACL). 2025. (*Poster paper, acceptance rate 20.3%*).
- [20] Xiaqiang Tang^D, Jian Li, Nan Du, and **Sihong Xie**^{*}. Adapting to Non-Stationary Environments: Multi-Armed Bandit Enhanced Retrieval-Augmented Generation on Knowledge Graphs. The 39th Annual AAAI Conference on Artificial Intelligence (AAAI). 2025. (*Poster paper, acceptance rate 23.4%*).
- [21] Yazheng Liu^D and **Sihong Xie**^{*}. Explanations of GNN on Evolving Graphs via Axiomatic Layer edges. The Thirteenth International Conference on Learning Representations (ICLR). 2025. (*Poster paper, acceptance rate 26.65%*).
- [22] Rui Xu^D, Chao Chen, Yue Sun, Parvathinathan Venkatasubramanian, and **Sihong Xie**^{*}. Wasserstein-Regularized Conformal Prediction under General Distribution Shift. The Thirteenth International Conference on Learning Representations (ICLR). 2025. (*Poster paper, acceptance rate 26.65%*).
- [23] Xingyuan Chen^M, Rui Xu^D, Yazheng Liu^D, Wenxing Huang, Minxuan Huang, Weiyan Chen, and **Sihong Xie**^{*}. Causality-Informed Machine Learning for ICU Length of Stay Prediction: A Healthcare Case Study. Springer Nature Computer Science. A conference version is published on CCF BigData Conference, 2025.
- [24] Chao Chen, Chenghua Guo, Rufeng Chen^D, Guixiang Ma, Ming Zeng, Xiangwen Liao, Xi Zhang, and **Sihong Xie**^{*}. Training for Stable Explanation for Free. The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS). 2024. (*Poster paper, acceptance rate 25.8%*).
- [25] Chenghua Guo, Han Yu, Jiaxin Liu, Chao Chen, Qi Li, **Sihong Xie**^{*}, and Xi Zhang^{*}. Linear Uncertainty Quantification of Graphical Model Inference. The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS). 2024. (*Poster paper, acceptance rate 25.8%*).
- [26] Yuefei Lyu, **Sihong Xie**, Chaozhuo Li, Xi Zhang^{*}. Enhancing Robustness of Graph Neural Networks on Social Media with Explainable Inverse Reinforcement Learning. The Thirty-eighth Annual

- Conference on Neural Information Processing Systems (NeurIPS). 2024. (*Poster paper, acceptance rate 25.8%*).
- [27] Ru-Bin Li, Man-Sheng Chen, Xin-Yu Ding, Chang-Dong Wang, **Sihong Xie**, Shuangyin Liu, Min Chen, and Mohsen Guizani. Periodic Prompt on Dynamic Heterogeneous Graph for Next Basket Recommendation. The IEEE International Conference on Data Mining (ICDM) 2024. (*Short paper, acceptance rate 19.5%*).
 - [28] Yazheng Liu^D, Xi Zhang, and **Sihong Xie**. A Differential Geometric View and Explainability of GNN on Evolving Graphs. The International Conference on Learning Representations (ICLR). 2023. (*Poster paper, acceptance rate 31.8%*).
 - [29] Adithya Kulkarni, Mohna Chakraborty, **Sihong Xie**, Qi Li. Optimal Budget Allocation for Crowdsourcing Labels for Graphs. The Conference on Uncertainty in Artificial Intelligence (UAI). 2023. (*Acceptance rate 243/778*).
 - [30] Yuefei Lv, Xiaoyu Yang, Jiabin Liu^D, **Sihong Xie**, Philip S. Yu^{DA}, and Xi Zhang. Interpretable and Effective Reinforcement Learning for Attacking against Graph-based Rumor Detection. IJCNN 2023.
 - [31] Yazheng Liu^D, Xi Zhang, and **Sihong Xie**. Trade less Accuracy for Fairness and Trade-off Explanation for GNN. Second International Workshop on Data science for equality, inclusion and well-being challenges, December 17-20, 2022, Osaka Japan
 - [32] Eric Enouen^U, Katja Mathesius^U, Sean Wang^U, Arielle Carr, and **Sihong Xie**. Efficient first-order predictor-corrector multiple objective optimization for fair misinformation detection. IEEE BigData, 2022. (*Regular paper, acceptance rate 122/633=19.2%. Eric, Katja, and Sean are undergraduate authors*)
 - [33] Sean Wang^U, Arielle Carr, and **Sihong Xie**. A Predictor-Corrector Method for Multi-objective Optimization in Fair Machine Learning. REU Symposium at the BDCAT conference, 2022. (*Sean is an undergraduate author.*)
 - [34] Yifei Liu, Chao Chen^D, Yazheng Liu, Xi Zhang, and **Sihong Xie**. Multi-objective Explanations of GNN Predictions. ICDM, 2021. (*Regular paper, acceptance rate 10%. [Selected for publication on the journal KAIS](#)*).
 - [35] Chao Chen^D, Yifan Shen, Guixiang Ma^I, Xiangnan Kong, Srinivas Rangarajan, Xi Zhang, and **Sihong Xie**. Self-learn to Explain Siamese Networks Robustly ICDM, 2021. (*Short paper, acceptance rate 20%*).
 - [36] Nasim Sabetpour, Adithya Kulkarni, Sihong Xie, and **Qi Li**. Truth Discovery in Sequence Labels from Crowds ICDM, 2021. (*Regular paper, acceptance rate 10%. [Selected for publication on the journal KAIS](#)*).
 - [37] Kai Burkholder^U, Kenny Kwok^U, Yuesheng Xu^U, Jiabin Liu^D, Chao Chen^D, and **Sihong Xie**. Certification and Trade-off of Multiple Fairness Criteria in Graph-based Spam Detection. CIKM, 2021. (*Full paper, acceptance rate 21.7%*).

- [38] Hang Yin, John Lee, Thomas Hartvigsen, Xiangnan Kong, and **Sihong Xie**. Energy-Efficient Models for High-Dimensional Spike Train Classification using Sparse Spiking Neural Network. KDD, 2021. (*Full paper, acceptance rate 15.5%*).
- [39] Shengli Zhu^U, Jakob Coles^U, and **Sihong Xie**. Active Search using Meta-Bandits. CIKM, 2020. (*Poster paper*).
- [40] Yifei Liu, Chao Chen^D, Yazheng Liu, Xi Zhang, and **Sihong Xie**. Shapley Values and Meta-Explanations for Probabilistic Graphical Model Inference. CIKM, 2020. (*Full paper, acceptance rate 20%*).
- [41] Yingdong Dou, Guixiang Ma^I, Philip S. Yu^{DA}, and **Sihong Xie**. Robust Spammer Detection by Nash Reinforcement Learning. KDD, 2020. (*Full paper, acceptance rate 16.9%*).
- [42] Chao Chen^D, Yifei Liu, Xi Zhang, and **Sihong Xie**. Scalable Explanation of Inferences on Large Graphs. ICDM, 2019. (*Short paper, acceptance rate 18.5%*).
- [43] Youshan Zhang, Sihong Xie, and **Brian D. Davison**. Transductive Learning via Improved Geodesic Sampling. BMVC, 2019. (Regular paper, acceptance rate 28%).
- [44] Shuaijun Ge, Guixiang Ma^I, **Sihong Xie**, and Philip S. Yu^{DA}. Securing Behavior-based Opinion Spam Detection. BigData, 2018. (Regular paper, acceptance rate 19.7%).
- [45] Lei Zheng, Yixue Wang, Lifang He, **Sihong Xie**, Fengjiao Wang, and Philip S. Yu^{DA}. PER: A Probabilistic Attentional Model for Personalized Text Recommendations. BigData, 2018. (Regular paper, acceptance rate 19.7%).
- [46] Vahid Noroozi, Sara Bahaadini, Lei Zheng, **Sihong Xie**, Weixiang Shao, and Philip S. Yu^{DA}. Semi-supervised Deep Representation Learning for Multi-View Problems. BigData, 2018. (Regular paper, acceptance rate 19.7%).
- [47] **Sihong Xie** and Philip S. Yu^{DA}. Next Generation Trustworthy Fraud Detection. 2018 IEEE 4th International Conference on Collaboration and Internet Computing (CIC). (Short paper).
- [48] Hu Xu, **Sihong Xie**, Shu Lei, and Philip S. Yu^{DA}. Dual attention network for product compatibility and function satisfiability analysis. AAAI, 2018. (Regular paper, acceptance rate 25.8%).
- [49] Hu Xu, **Sihong Xie**, Lei Shu, and Philip S. Yu^{DA}. Product function need recognition via semi-supervised attention network. BigData, 2017. (Regular paper, acceptance rate 17.8%).
- [50] Xiaokai Wei, **Sihong Xie**, Bokai Cao, and Philip S. Yu^{DA}. *Rethinking Unsupervised Feature Selection: From Pseudo Labels to Pseudo Must-Links*. ECML. 2017. (Regular paper, acceptance rate 29%).
- [51] Vahid Noroozi, Lei Zheng, Sara Bahaadini, **Sihong Xie**, and Philip S. Yu^{DA}. Seven: Deep Semi-supervised Verification Networks. IJCAI, 2017. (Regular paper, acceptance rate 25.9%).
- [28] **Sihong Xie**, Qingbo Hu, Weixiang Shao, Jingyuan Zhang, Jing Gao, Wei Fan, and Philip S. Yu^{DA}. Effective crowd expertise modeling via cross domain sparsity and uncertainty reduction. SDM. SIAM, 2016.

- [29] **Sihong Xie**, Shaoxiong Wang, and Philip S. Yu^{DA}. Active zero-shot learning. CIKM, 2016.
- [30] Chun-Ta Lu, **Sihong Xie**, Weixiang Shao, Lifang He, and Philip S. Yu^{DA}. Item recommendation for emerging online businesses. IJCAI, 2016.
- [31] Yuan Yuan, **Sihong Xie**, ChunTa Lu, Jie Tang, and Philip S. Yu^{DA}. Interpretable and effective opinion spam detection via temporal patterns mining across websites. BigData. IEEE, 2016.
- [32] Chenwei Zhang, **Sihong Xie**, Yaliang Liu, Jing Gao, Wei Fan, and Philip S. Yu^{DA}. Multi-source hierarchical prediction consolidation. CIKM, 2016.
- [33] J. Zhang, C. T. Lu, M. Zhou, S. Xie, Y. Chang, and P. S. Yu. Heer: Heterogeneous graph embedding for emerging relation detection from news. In *2016 IEEE International Conference on Big Data (Big Data)*, 2016.
- [34] Jingyuan Zhang, Bokai Cao, **Sihong Xie**, Chun-Ta Lu, Philip S. Yu^{DA}, and Ann B. Ragin. Identifying connectivity patterns for brain diseases via multi-side-view guided deep architectures. SDM. SIAM, 2016.
- [35] Qingbo Hu, **Sihong Xie**, Jiawei Zhang, Qiang Zhu, Songtao Guo, and Philip S. Yu^{DA}. Heterosales: Utilizing heterogeneous social networks to identify the next enterprise customer. WWW. ACM, 2016.
- [36] Jingyuan Zhang, Luo Jie, Altaf Rahman, **Sihong Xie**, Yi Chang, and Philip S. Yu^{DA}. Learning entity types from query logs via graph-based modeling. CIKM. ACM, 2015.
- [37] Xiaokai Wei, **Sihong Xie**, and Philip S. Yu^{DA}. Efficient partial order-preserving unsupervised feature selection on networks. In *SDM*. SIAM, 2015.
- [38] **Sihong Xie**, Qingbo Hu, Jingyuan Zhang, Jing Gao, Wei Fan, and Philip S. Yu^{DA}. Robust crowd bias correction via dual knowledge transfer from multiple overlapping sources. In *BigData*. IEEE, 2015.
- [39] **Sihong Xie**, Qingbo Hu, Jingyuan Zhang, and Philip S. Yu^{DA}. An effective and economic bi-level approach to ranking and rating spam detection. In *DSAA*. IEEE, 2015.
- [40] **Sihong Xie**, Jing Wang, Mohammad S.Amin, Baoshi Yan, Anmol Bhasin, Clement Yu, and Philip S. Yu^{DA}. A context-aware approach to detection of short irrelevant texts. In *DSAA*. IEEE, 2015.
- [41] Bowen Dong, **Sihong Xie**, Jing Gao, Wei Fan, and Philip S. Yu^{DA}. Onlinecm: online consensus maximization with missing values. In *SDM*. SIAM, 2015.
- [42] Qingbo Hu, **Sihong Xie**, Shuyang Lin, Senzhang Wang, and Philip S. Yu^{DA}. Clustering embedded approaches for efficient information network inference. *Data Science and Engineering*, 2015.
- [43] Qiongbo Hu, **Sihong Xie**, Shuyang Lin, Wei Fan, and Philip S. Yu^{DA}. Frameworks to encode user preferences for inferring topic-sensitive information networks. In *SDM*. SIAM, 2015.

- [44] Qiongbao Hu, **Sihong Xie**, Shuyang Lin, Senzhang Wang, and Philip S. Yu^{DA}. CENI: a hybrid framework for efficiently inferring information networks. In *ICWSM*. AAAI, 2015.
- [45] **Sihong Xie**, Jing Gao, Deepak Turaga, Wei Fan, and Philip S. Yu^{DA}. Class-distribution regularized consensus maximization for alleviating overfitting in model combination. In *KDD*. ACM, 2014.
- [46] Senzhang Wang, **Sihong Xie**, Xiaoming Zhang, Zhoujun Li, Philip S. Yu^{DA}, and Xinyu Shu. Future influence ranking of scientific literature. In *SDM*. SIAM, 2014.
- [47] Guoqiong Liao, Philip S. Yu^{DA}, Qianhui Zhong, **Sihong Xie**, Zhen Shen, Changxuan Wan, and Dexi Liu. Trajectory event cleaning for mobile rfid objects. In *15th International Conference on Mobile Data Management*. IEEE, 2014.
- [48] Guoqiong Liao, Yuchen Zhao, **Sihong Xie**, and Philip S. Yu^{DA}. An effective latent networks fusion based model for event recommendation in offline ephemeral social networks. In *CIKM*. ACM, 2013.
- [49] Chun-Ta Lu, **Sihong Xie**, Xiangnan Kong, and Philip S. Yu^{DA}. Inferring the impacts of social media on crowdfunding. In *WSDM*. ACM, 2013.
- [50] **Sihong Xie**, Xiangnan Kong, Jing Gao, Wei Fan, and Philip S. Yu^{DA}. Multilabel consensus classification. In *ICDM*. IEEE, 2013.
- [51] Chuan Shi, Xiangnan Kong, Philip S. Yu^{DA}, **Sihong Xie**, and Bin Wu. Relevance search in heterogeneous networks. In *EDBT*. ACM, 2012.
- [52] **Sihong Xie**, Wei Fan, and Philip S. Yu^{DA}. An iterative and re-weighting framework for rejection and uncertainty resolution in crowdsourcing. In *SDM*. SIAM, 2012.
- [53] **Sihong Xie**, Guan Wang, Shuyang Lin, and Philip S. Yu^{DA}. Review spam detection via temporal pattern discovery. In *KDD*. ACM, 2012.
- [54] **Sihong Xie**, Guan Wang, Shuyang Lin, and Philip S. Yu^{DA}. Review spam detection via time series pattern discovery. In *WWW (short paper)*. ACM, 2012.
- [55] Guan Wang, **Sihong Xie**, Bing Liu, and Philip S. Yu^{DA}. Review graph based online store review spammer detection. In *ICDM*. IEEE, 2011.
- [56] **Sihong Xie**, Wei Fan, Olivier Verscheure, and Jiangtao Ren. Efficient and numerically stable sparse learning. In *ECML/PKDD*, 2010.
- [57] **Sihong Xie**, Wei Fan, Jing Peng, Olivier Verscheure, and Jiangtao Ren. Latent space domain transfer between high dimensional overlapping distributions. In *WWW*. ACM, 2009.
- [58] Erheng Zhong, **Sihong Xie**, Wei Fan, Jiangtao Ren, Jing Peng, and Kun Zhang. Graph-based iterative hybrid feature selection. In *ICDM*. IEEE, 2008.

C. HONORS AND AWARDS

IEEE Senior Member	2025.12-present
NSF CAREER Award	2022

Faculty Grants for International Connections, Lehigh University	2017
SIAM SDM 2012 Student Travel Award	May, 2012
ACM SIGKDD 2012 Student Travel Award	Aug, 2012
IEEE ICDM 2008 Data Mining Contest Crown Award (highest award)	2008
Outstanding Graduate Student Scholarship, Sun Yat-Sen University	2008 - 2009
Outstanding Student Scholarship First Prize, Sun Yat-Sen University	2006 - 2007

D. RESEARCH FUNDING

Extramural research grants (total funding amount: \$1,556,481, RMB 300,000 total amount awarded to HKUST-GZ: RMB 300,000, total amount awarded to Lehigh: \$1,056,481, total amount awarded to the faculty member as a PI: \$806,481, total amount awarded to the faculty member as a Co-PI: \$250,000.)

- **Tencent**, “*Motif-based explainable graph foundation model for novel suspicion detection*”. RMB 300,000. 2024-2025
- **NSF**, “*CAREER: Bilevel Optimization for Accountable Machine Learning on Graphs*”. \$556,481. 2022-2027
- **NSF**, “*III: Small: Collaborative Research: Algorithms, systems, and theories for exploiting data dependencies in crowdsourcing*”. \$500,000. Lehigh share: 50% . PI: Sihong Xie. Co-PI: Qi Li. 2020-2023
- **NSF**, “*SaTC: CORE: Small: Collaborative: Learning Dynamic and Robust Defenses Against Co-Adaptive Spammers*”, NSF. \$500,000. Lehigh share: 50%. PI: Philip S.Yu. Co-PI: Sihong Xie. 2019-2023

Intramural research grants (total funding amount: \$350,000, total amount awarded to the faculty member as a Co-PI: \$350,000.)

- **Lehigh Research Future**, “*Mathematical Optimization and Data Science*”. \$250,000. PI: Frank Curtis. Co-PI: Aida Khajavirad, Alberto Lamadrid, Ted Ralphs, Srinivas Rangarajan, Daniel Robinson, Karmel S. Shehadeh, Larry Snyder, Parv Venkitasubramaniam, Luis Nunes Vicente, **Sihong Xie**. 2021-present
- **Lehigh Accelerator** “*Efficient, Explainable and Robust Data Scientific Methods for Smart Engineering Systems*”. \$100,000. PI: Parv Venkitasubramaniam. Co-PI: Hector Munoz-Avila, Brian Davison, Jeff Heflin, Rick Blum, Ted Ralphs, Paolo Bocchini, **Sihong Xie**. 2019-present

E. EDITOR/EDITORIAL REVIEW

Frontiers Topic Editor: Frontiers in Big Data (Topic: Trustworthiness in Data Mining or Truth Discovery in Data Mining) May 2022-May 2023.

F. SCHOLARLY PRESENTATIONS

Invited presentations

- Explainable Graphical Models
Shandong University, Qingdao, China Aug, 2024
Anhui University, Hefei, China Jun, 2024
- Fairness in machine learning
Outreach to K6-12 students at LEAP Academy, Rutgers University, Camden Nov, 2021
- Robust misinformation detection for e-commerce.
Seminar, Meta (former Facebook) Nov, 2021
- Interpretable, Robust, and Fair Learning on Graphs.
Seminar, CS Department at Iowa State University Oct, 2021

Seminar, OptML group at Lehigh ISE department	Mar, 2021
[5] Secure Opinion Spam Detection	
The Second International Workshop on Truth Discovery and Fact Checking	Aug, 2020
[6] Next Generation Trustworthy Fraud Detection	
Worcester Polytechnic Institute, MA	Dec, 2018
[7] Crowdsourcing for Natural Language Processing applications.	
Beijing University of Posts and Telecommunications, Beijing, China.	Jun, 2017
Sun Yat-Sen University, Guangzhou, China.	Jul, 2017
[8] Enhancing predictive power via consensus models	
University of Illinois, Chicago, IL	Nov, 2014
[9] Context-aware spam detection	
LinkedIn Corporation, Mountain View, CA	Aug, 2013

Refereed presentations

[1] MARS: Memory Attention-Aware Recommender System	
DSAA 2019, Washington, DC	Sep, 2019
[2] Securing Behavior-based Opinion Spam Detection	
BigData 2018, Seattle, WA	Dec, 2018
[3] Next Generation Trustworthy Fraud Detection	
CIC 2018, Philadelphia, PA	Oct, 2018
[4] Exploiting multiple sources for trustworthy content managements	
BigData Conference 2015, Santa Clara, CA	Nov, 2015
[5] Overfitting and regularization in model combination	
KDD 2014, New York City, NY	Aug, 2014
[6] Multi-label consensus classification	
ICDM 2013, Dallas, TX	Dec, 2013

G. TEACHING AND RESEARCH ADVISING

Course taught (21 offerings of 5 courses from 2016 Fall - 2024 Fall.)

AIAA 5046 (HKUST-GZ) “Fundamentals of Machine Learning”, Spring 2024.

AIAA 5047 (HKUST-GZ) “Responsible Machine Learning”, Fall 2023, Fall 2024, Fall 2025.

CSE 326/426 (Lehigh) “Fundamentals of Machine Learning”, Spring 2019, Fall 2020, Fall 2021, Fall 2022.

CSE 325/425 (Lehigh) “Natural Language Processing”, Fall 2017, Fall 2018, Fall 2019, Spring 2021, Spring 2022, Spring 2023.

CSE 241 (Lehigh) “Database Systems & Apps”, Spring 2017, Spring 2018, Fall 2019, Spring 2020, Fall 2020, Fall 2021, Fall 2022.

CSE 398/498 (Lehigh) “Text Mining”, Fall 2016

Doctoral advising

Yue Chang, AI Thrust, HKUST-GZ (co-adviser: Hui Xiong). *Embodied AI*. 2025 Fall - Present.

Rufeng Chen, AI Thrust, HKUST-GZ. *Embodied AI*. 2024 Fall - Present.

Zhaofan Zhang, AI Thrust, HKUST-GZ (co-adviser: Hui Xiong). *Embodied AI*. 2024 Fall - Present.

Jiujiu Chen, AI Thrust, HKUST-GZ (co-adviser: Hui Xiong). *Embodied AI*. 2024 Fall - Present.

Xiaqiang Tang, AI Thrust, HKUST-GZ. *LLM and agent*. 2024 Fall - Present.

Rui Xu, AI Thrust, HKUST-GZ. *Conformal prediction; uncertainty quantification*. 2023 Fall - Present.

Yazheng Liu, AI Thrust, HKUST-GZ. *Explainable graphical models*. 2023 Fall - Present.

Chao Chen, CSE, Lehigh University. *Explainable graphical models* [34, 35, 37, 40, 42]. 2018 Summer - 2023 Spring.

- First job placement: Postdoc at HIT(SZ).
- Supported as RA (9 semesters via NSF grants and a Lehigh Accelerator fund) and TA (1 semester).
- Paper presentations on the corresponding conference(s): [42, 40, 34, 35].

Jiaxin Liu, CSE, Lehigh University. *Fair graphical models* [37]. 2019 Fall - 2023 Spring.

- Supported as RA (6 semesters via NSF grants) and TA (2 semesters).
- Paper presentations on the corresponding conference(s): [37].

Xuehan Chen, CSE, Lehigh University. *Explainable NLP*. 2021 Fall - 2022 Spring.

- Supported as RA (1 semester via an NSF grant) and TA (3 semesters).

Master advising (sorted by advising end time)

Wenzhan Yu, AI Thrust, HKUST-GZ. *VLM reasoning*. 2025 Fall - Present.

Qirong Guo, AI Thrust, HKUST-GZ. *Steering LLM*. 2025 Fall - Present.

Yi Wang, AI Thrust, HKUST-GZ. (co-adviser: Li Chen). *LLM and Embodied AI*. 2024 Fall - Present.

Xingyuan Chen, AI Thrust, HKUST-GZ. *XAI for medicine* [23]. 2023 Fall - 2025 Spring.

Xiangrong Liao, AI Thrust, HKUST-GZ. *Fair RL for EV car charging station placement*. 2023 Fall - 2025 Spring.

Tianchuang Bai, AI Thrust, HKUST-GZ. *Faithful LLM for medicine*. 2023 Fall - 2025 Spring.

Hui Wu, AI Thrust, HKUST-GZ. *Pose estimation from video*. 2023 Fall - 2025 Spring.

Wenqi Qiu, AI Thrust, HKUST-GZ. *An HCI study for EEG signal processing*. 2023 Fall - 2025 Spring.

Shujing Feng, CSE at Lehigh. 2022 Fall - 2023 Spring. Graduation expected in 2024 Spring. *Constrained optimization for NLP*.

Yuesheng Xu, CS at NYU. 2022 Summer - 2023 Fall. Graduation expected in 2023 Spring. *Physics-informed graph neural networks* [37].

Praveer Narwelkar, CSE. 2021 Spring - 2021 Fall. Graduated in 2019 Spring. Placement unknown. *Multi-modal misinformation detection*.

Yue Shu, ECE. 2018 Spring. Graduation date and placement unknown. *Multiple instance learning for sentimental information extraction*.

Kexin Ding, ECE. 2018 Spring. Graduation date and placement unknown. *Building a web interface for crowdsourcing a multi-modal (text-image) sentimental dataset*.

Undergraduate advising (sorted by advising end time)

Ningyuan Xiao. 2025 Summer - present. *Privacy and Security for Multi-Agent LLM System*

Keyu Hu. 2024 Summer. *Reliable LLM* [19].

Sean Wang, Cornell University. 2022 Summer - present. Graduation date and placement unknown. *Multi-objective fair machine learning* [32, 33].

Yifan She, Kevin Coxhead, Kamilla Muninova, and Alonso Cornejo Leon. Lehigh University. 2022 Spring - 2022 Fall. Graduation date and placement unknown. *Capstone project: explainable misinformation detection system*.

Katja Mathesius, Drake University. 2022 Summer. Graduation date and placement unknown. *Multi-objective fair machine learning* [32, 33].

Sherry Huang, Lehigh University. 2022 Spring. Graduated in 2022 Spring. Placement unknown. *Fair machine learning*.

Noah Backman, Chris Munoz, Andy Kelly, and Jinan Hong. Lehigh University. 2021 Spring - 2021 Fall. Graduation date and placement unknown. *Capstone project: explainable misinformation detection system*.

Eric Enouen, Ohio State University. 2021 Summer - 2021 Fall. Graduation date and placement unknown. *Constrained multi-objective optimization* [32].

Kenny Kwock, Lehigh University. 2020 Summer - 2021 Fall. Graduated in 2022 Spring. First job: Amazon. *Fair multi-objective spam detection* [37].

Kieran O'Connor, Lehigh University. 2021 Spring - 2021 Fall. Graduation date and placement unknown. *Explainable graphical models*.

Yuesheng Xu, Lehigh University. 2019 Summer - 2021 Spring. Graduated 2021 Spring. Admitted to Master at NYU. *Graphical bandits, fair spam detection* [37].

Kai Burkholder, Lehigh University. 2020 Summer. Graduation date and placement unknown. *Fair multi-objective spam detection* [37].

Anmol Shrestha, Mavin Buranslip, Dylan Schaschl, John Dutra. Lehigh University. 2020 Spring - 2020 Fall. Graduation date and placement unknown. *Capstone project: explainable misinformation detection system*.

Jakob Coles, Lehigh University. 2019 Summer - 2020 Spring (REU). Graduated in 2021 Spring. Admitted to Ph.D. at UW-Madison. *Graphical bandits* [39].

Sam Chebruch, Lehigh University. 2019 Summer - 2020 Spring. Graduated 2021 Spring. Placement unknown. *Graphical bandits* [39].

Shengli Zhu, Lehigh University. 2017 Summer - 2020 Spring. Graduated 2020 Spring. Admitted to Master at CMU. Now at Meta. *Graphical bandits, relation extraction, spam detection game* [39].

Yifan Zhang, Lehigh University. 2020 Summer - 2020 Fall. Graduation date and placement unknown. *Explainable spam detection*.

Kexin Shu, CS at De Anza College, Santa Clara, CA. 2019 Summer. Graduation date and placement unknown. *Spam detection*. Admitted by UC Berkeley.

Jacob Nocentino. 2019 Summer (REU). Graduation date and placement unknown. *Graphical bandits*.

Jeremy Ding, Lehigh University. 2019 Summer. Graduation date and placement unknown. *Graphical bandits*.

Jun-Lucas Pritsker, CSE. 2017 Winter - 2018 Spring. Graduation date and placement unknown. *Human-like sentimental mapping from images to texts*. co-adviser: Xiaolei Huang.

Basilio Garcia Castillo, CSE. 2016 Fall - 2017 Spring. Graduated 2017 Spring. First job: Google NYC. *Rule-based sentiment analysis*.

Tom Bousso, CS at UC Berkeley, CA. 2017 summer. Graduation date and placement unknown. *Active learning for implicit sentiment analysis*.

Hannah Lambert, Chijioke Umezina, Jun-Lucas Pritsker, CSE. 2017 Fall. Graduation date and placement unknown. *Senior Design: crowdsourcing for aspect-based sentiment analysis*.

Zhenyu Li, ISE. 2017 Fall. Graduation date and placement unknown. *Fine-grained sentimental information extraction using topic models*.

Yilin Zhu, CS at De Anza College, Santa Clara, CA. 2017 Summer. Graduation date and placement unknown. *Spam detection using graphical model*.

Dissertation committee

Xiasi Wang, DSA, HKUST-GZ. 2025 Fall. Thesis Defense Chairperson. *A Study on Robustness and Generalization for Visual and Multimodal Learning*.

Qianyi Cai, AI, HKUST-GZ. 2025 Summer. Thesis Defense Chairperson. *Continual Graph Learning via Prompt Tuning*.

Junyi Li, AI, HKUST-GZ. 2025 Summer. Thesis Defense Chairperson. *Causal Learning with Large Pretrained Models*.

Zichen Tang, DSA, HKUST-GZ. 2025 Summer. Thesis Defense Chairperson. *Communication Optimization for Federated Learning under Bandwidth Constraints*.

Jiashun Cheng, DSA, HKUST-GZ. 2025 Summer. Thesis Defense Chairperson. *Generalizable and Efficient Transfer Learning for Graph and Language Models*.

Botao Wang, DSA, HKUST-GZ. 2025 Summer. Thesis Defense Chairperson. *Exploring Information Flow through Graphic Causal Networks for Multivariate Time Series Anomaly Detection*.

Tailin Zhou, IoT, HKUST-GZ. 2025 Summer. Thesis Defense Chairperson. *Understanding, Improving, and Implementing Federated Learning on Heterogeneous Data*.

Griffin Kent, ISE, Lehigh University. 2023 Spring. Proposal and defense. *Stochastic Methods in MultiLevel and Multi-Objective Optimization*.

Mohammad Hesam Shaelaie, ISE, Lehigh University. 2023 Spring. Proposal and defense. *Explainable Graph Learning and Unified Framework of Reinforcement Learning and Branch and Bound*.

Weihang Yuan, CSE, Lehigh University. 2022 Spring. General exam. Graduation date and placement unknown. *Deep Goal Reasoning*.

Morgan Fine-Morris, CSE, Lehigh University. 2022 Spring. General exam. Graduation date and placement unknown. *Learning Planning Knowledge in Structured Domains with Numeric Fluents*.

Noah Reifsnyder, CSE, Lehigh University. 2022 Spring. Defense in 2022 Spring. First job: Parallax Advanced Research. *A Study of Goal Reasoning Agents' Expectations for Nondeterministic and Continuous domains*.

Suyun Liu, ISE, Lehigh University. 2021 Fall. Defense in 2022 Spring. First job: Amazon. *Stochastic Multi-Objective Optimization and Its Application to Fairness in Machine Learning*.

Zhiyu Chen, CSE, Lehigh University. 2021 Spring. Defense in 2022 Spring. First job: Amazon. *Dataset Search and Augmentation*.

Sicong Kuang, CSE, Lehigh University. 2020 Spring. Defense in 2020 Spring. First job: LinkedIn. *Improving Semantic Representation in Bias Detection*.

H. SERVICES

University, College, and Department

Red-Bird MPhil Program, HKUST-GZ: Selection & Interview Committee (S&I). Sep, 2024 - present
AI Thrust, HKUST-GZ: Search and Substantiation Committee. Jun, 2024 - present
HKUST-GZ: Teaching-& Education-track Appointment Committee (TEAC). Mar, 2024 - present
CSE department, Lehigh University: Master of Data Science curriculum committee. 2022 Spring
CSE department, Lehigh University: DE&I committee. 2021 Spring - 2022 Fall
CSE department, Lehigh University: Tenure-track faculty search committee. 2018-2019, 2020-2021
CSE department, Lehigh University: Ad hoc committee on Teaching Load and Planning. 2020-2021
Lehigh University: Faculty Senate (Faculty affair sub-committee). 2020 Fall
CSE department, Lehigh University: Curriculum committee. 2016 - 2020
CSE department, Lehigh University: Graduate Admissions Committee. 2016 - 2020

Workshop organization

Mis2-TrueFact 2022: Joint International Workshop on Misinformation and Misbehavior Mining on the Web & Making a Credible Web for Tomorrow, Workshop co-host with KDD 2022. Homepage: <http://claws.cc.gatech.edu/mis2-truefact-kdd2022.html>.

The Third International Workshop on Truth Discovery and Fact Checking: Making a Credible Web for Tomorrow, Virtual Workshop co-host with KDD 2021. Homepage: <https://www.microsoft.com/en-us/research/event/kdd-2021-truefact-workshop-making-a-credible-web-for-tomorrow/>.

The Second International Workshop on Truth Discovery and Fact Checking: Making a Credible Web for Tomorrow, Virtual Workshop co-host with KDD 2020. Homepage: <https://www.microsoft.com/en-us/research/event/kdd-2020-truefact-workshop-making-a-credible-web-for-tomorrow/>.

Conference and journal services

Area Chair: NeurIPS (2026), WWW (2025, 2026), CIKM (2026).

Senior PC member: AAAI (2022), SDM (2024).

PC member: ICLR (2022-2026), NeurIPS (2022-2025), ICML (2023-2026), COLM (2026), ECCV (2026), AISTATS (2024, 2026), UAI (2026), KDD (2018-2022, 2026), WWW (2024), SIGIR (2020-2025), ACL (2025-2026), AAAI (2019-2022, 2025-2026), IJCAI (2020, 2022), SDM (2018-2022), ICDM (2017-2021, 2026), CIKM (2017-2022).

External conference reviewer: SIGAPP SAC (2012), APWeb (2012), PRICAI (2012), ICDM (2012 - 2014), ASONAM (2013), WWW (2014), SDM (2014 - 2016), BigData (2015), AAAI (2016).

Journal reviewer: Neural Networks by Elsevier (2024), the Journal of AI Research (JAIR) (2021), Frontiers in Big Data (2022), Transactions on Knowledge Discovery from Data (TKDD) (2014, 2016)

- 2018), Transactions on Knowledge and Data Engineering (TKDE) (2017-2020, 2022, 2024), and the Transactions on the Web (TWeb) (2018, 2019, 2022), Springer Nature (2024).

Conference session chairs

SDM (2021), SIGIR (2021), ICDM (2021), UAI (2022).

Grant Panelist

NSF Panelist (2017, 2020, 2021, 2022, 2023)