

Xujiang Zhao, Ph.D.

✉ zhaouxj32@gmail.com

🔍 Google Scholar

🌐 zxj32.github.io

📞 518-423-0658

Employment History

- 2022 - Present 📌 **Research Staff Member.** NEC Laboratories America.
- 2021 Summer 📌 **Research Intern.** NEC Laboratories America.
- 2019 Summer 📌 **Research Intern.** Alibaba Damo Academy.

Education

- 2019 – 2022 📌 **Ph.D. University of Texas at Dallas** in Computer Science.
Thesis title: *Multidimensional Uncertainty Quantification for Deep Neural Networks.*
- 2014 – 2017 📌 **M.Sc. University of Science and Technology of China** in Computer Science.
- 2010 – 2014 📌 **B.S. Chongqing University** in Civil Engineering.

Research Interests

I am generally interested in AI, machine learning, and natural language Processing, with a primary focus on:

- Agentic AI
- LLM Post Training
- Reinforcement Learning
- Generative AI
- Efficient AI
- AI Safety & Trustworthy AI

Applied Research Project

• Next-Generation Agentic Frameworks

Developing autonomous, self-evolving systems that leverage Reinforcement Learning and test-time scaling to enhance complex reasoning. These frameworks integrate long-term memory, multi-modal vision reasoning, and robust code generation to enable agents that refine their own strategies over time.

Related Publications: ICML2026, ACL 2026, ICLR 2025, NeurIPS 2025, NAACL 2025, ACL 2025, EACL 2026

• Domain-Specific Efficient Training of Foundation Models

Developing efficient domain-specific foundation models through SFT, LORA, model compression, knowledge distillation, knowledge unlearning, fine-tuning, and knowledge graph to enhance efficiency and adaptability.

Related Publications: ACL 2026, EACL 2026, NAACL 2025, NAACL 2024, EMNLP 2023, CSUR 2025, NeurIPS 2021

• Trust, Privacy, and Responsible LLM

Ensuring LLMs generate reliable, fair, and privacy-preserving outputs, particularly for high-stakes applications where transparency and trust are critical.

Related Publications: ACL 2026, NeurIPS 2025, EACL 2026, NAACL 2024, ICLRW 2024, EMNLP 2024

• AI Safety & Trustworthy AI

Advancing the foundations of reliable AI through robust safety frameworks and trustworthy modeling. Our work centers on building interpretable and resilient systems that maintain high-level performance while providing formal guarantees of stability, transparency, and ethical accountability

Related Publications: NeurIPS 2020 (Spotlight), NeurIPS 2020, AAI 2026, WWW 2026 SDM 2025, Frontiers in Big Data 2024, ICASSP 2023, CIKM 2022, ICDM 2022, ICASSP 2022, EMNLP 2021, Fusion 2019, AAI 2021, AAI 2019, IEEE Big Data 2019, IEEE Big Data 2018, ICDM 2018, MILCOM 2018

Research Publications

†: Denote Corresponding Author; *: Denote equal contribution; *Students* : Denote students I supervise.

Conference Proceedings

NeurIPS(5), ICML(1) ICLR(1), NAACL(4), EMNLP(3), ACL(4), EACL(3), AAAI(4), ICDM(2), CIKM(4), SDM(1), ICASSP(2), IJCAI(1), WWW(2), IEEE BigData(2), PAKDD(1), FUSION(1), MILCOM(1)

1. Yutong Cheng, Haifeng Chen, Wenchao Yu, **Xujiang Zhao**, Peng Gao, Wei Cheng. "Escaping Whack-a-Mole: Code Documentation Optimization via Dependency-Guided Bi-level Search.". Forty-Third International Conference on Machine Learning. **ICML 2026**. [Agentic AI] [Code Generation]
2. Peng Xia, Peng Xia, Jianwen Chen, Hanyang Wang, Jiaqi Liu, Kaide Zeng, Yu Wang, Siwei Han, Yiyang Zhou, **Xujiang Zhao**, Haifeng Chen, Zeyu Zheng, Cihang Xie, Huaxiu Yao. "SkillRL: Evolving Agents via Recursive Skill-Augmented Reinforcement Learning." **ICLR 2026 MemAgents Workshop Best Paper Award Runner-Up**. [Agentic AI] [Self Evolving]
3. Linlin Yu, **Xujiang Zhao**†, Dong Li, Yanchi Liu, Wei Cheng, Zhengzhang Chen, Chen Zhao, Feng Chen, Haifeng Chen. "Uncertainty-Aware Test-Time Search for Optimization Problem Solving." In Proceedings of the Association for Computational Linguistics: **ACL 2026**. [Agentic AI] [Test Time Scaling]
4. Binchi Zhang, **Xujiang Zhao**, Jundong Li, Haifeng Chen, Zhengzhang Chen. "Mind the Gap in Cultural Alignment: Task-Aware Culture Management for Large Language Models." In Proceedings of the Association for Computational Linguistics: **ACL 2026**. [Post Training]
5. Xuyuan Liu, Shengyu Chen, Xinshuai Dong, Yanchi Liu, **Xujiang Zhao**, Haoyu Wang, Yujun Yan, Haifeng Chen, Zhengzhang Chen. "Representation Interventions Enable Lifelong Unstructured Knowledge Control." In Proceedings of the Association for Computational Linguistics: **ACL 2026**. [Post Training] [Memory]
6. Minghao Guo, Qingcheng Zeng, **Xujiang Zhao**, Yanchi Liu, Wenchao Yu, Mengnan Du, Haifeng Chen, Wei Cheng. "DeepSieve: Information Sieving via LLM-as-a-Knowledge-Router." in Findings of the European Chapter of the Association for Computational Linguistics. **EACL, 2026**. [Agentic AI] [Workflow]
7. Wangyang Ying, Yanchi Liu, **Xujiang Zhao**, Wei Cheng, Zhengzhang Chen, Wenchao Yu, Yanjie Fu, Haifeng Chen. "Multi-Agent Procedural Graph Extraction with Structural and Logical Refinement." in Findings of the European Chapter of the Association for Computational Linguistics. **EACL, 2026**. [Agentic AI] [Workflow]
8. Minhua Lin, Zhengzhang Chen, Yanchi Liu, **Xujiang Zhao**, Zongyu Wu, Junxiang Wang, Xiang Zhang, Suhang Wang, Haifeng Chen. "Decoding Time Series with LLMs: A Multi-Agent Framework for Cross-Domain Annotation." in Findings of the European Chapter of the Association for Computational Linguistics. **EACL, 2026**. [Agentic AI] [Time Series]
9. Dong Li, Zhengzhang Chen, **Xujiang Zhao**, Linlin Yu, Zhong Chen, Yi He, Haifeng Chen, Chen Zhao. "MARLIN: Multi-Agent Reinforcement Learning for Incremental DAG Discovery." In Proceeding of the 40th AAAI Conference on Artificial Intelligence, **AAAI, 2026**. [Agentic AI] [RL]
10. Zhixia He, Chen Zhao, Minglai Shao, Xintao Wu, **Xujiang Zhao**, Dong Li, Qin Tian, Linlin Yu. "Out-of-Distribution Detection with Positive and Negative Prompt Supervision Using Large Language Models." In Proceeding of the 40th AAAI Conference on Artificial Intelligence, **AAAI, 2026**. [Trustworthy AI]
11. Qin Tian, Chen Zhao, Xintao Wu, Dong Li, Minglai Shao, **Xujiang Zhao**, Wenjun Wang. "Class-Domain Incremental Learning on Graphs via Disentangled Knowledge Distillation." In Proceeding of The Web Conference, **WWW, 2026**. [Trustworthy AI]
12. Dong Li, **Xujiang Zhao**†, Linlin Yu, Yanchi Liu, Wei Cheng, Zhengzhang Chen, Zhong Chen, Feng Chen, Chen Zhao, Haifeng Chen. "SolverLLM: Leveraging Test-Time Scaling for Optimization Problem via LLM-Guided Search." In Advances in neural information processing systems, **NeurIPS, 2025**. [Agentic AI] [Test Time Scaling]
13. Shengkun Tang, Cong Zeng, Yuanzhou Chen, Zhiqiang Shen, Wenchao Yu, **Xujiang Zhao**, Haifeng Chen, Wei Cheng, Zhiqiang Xu. "Human Texts Are Outliers: Detecting LLM-generated Texts via Out-of-distribution Detection." In Advances in neural information processing systems, **NeurIPS, 2025**. [Trustworthy AI]
14. Qiwei Zhao*, Dong Li*, Yanchi Liu, Wei Cheng, Yiyu Sun, Mika Oishi, Takao Osaki, Katsushi Matsuda, Huaxiu Yao, Chen Zhao, Haifeng Chen, **Xujiang Zhao**†. "Uncertainty Propagation on LLM Agent." In Proceedings of the Association for Computational Linguistics: **ACL 2025**. [Agentic AI] [Uncertainty]

15. Jonathan Light, Yue Wu, Yiyu Sun, Wenchao Yu, Yanchi Liu, **Xujiang Zhao**, Ziniu Hu, Haifeng Chen, Wei Cheng. "Scattered Forest Search: Smarter Code Space Exploration with LLMs." In The Thirteenth International Conference on Learning Representations. **ICLR**, 2025. [[Agentic AI](#)] [[Code Generation](#)]
16. Xinyuan Wang, Yanchi Liu, Wei Cheng, **Xujiang Zhao**, Zhengzhang Chen, Wenchao Yu, Yanjie Fu, Haifeng Chen. "MixLLM: Dynamic Routing in Mixed Large Language Models." In Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics. **NAACL**, 2025. [[Efficient AI](#)]
17. Xianjun Yang, Wei Cheng, **Xujiang Zhao**, Wenchao Yu, Linda Petzold, Haifeng Chen. "Position Really Matters: Towards a Holistic Approach for Prompt Tuning." in Findings of the Association for Computational Linguistics. **NAACL**, 2025. [[Post Training](#)]
18. Ruomeng Ding, **Xujiang Zhao**[†], Chen Zhao, Minglai Shao, Zhengzhang Chen, and Haifeng Chen. "Evidence-Based Out-of-Distribution Detection on Multi-Label Graphs." In Proceedings of the 2025 SIAM International Conference on Data Mining. **SDM**, 2025. [[Trustworthy AI](#)]
19. Chengyuan Deng, Zhengzhang Chen, **Xujiang Zhao**, Haoyu Wang, Junxiang Wang, Jie Gao, Haifeng Chen. "Correlation-aware Online Change Point Detection." In Proceedings of the 34nd ACM International Conference on Information and Knowledge Management, **CIKM**, 2025. [[Trustworthy AI](#)]
20. Chen Ling, **Xujiang Zhao**[†], Xuchao Zhang, Wei Cheng, Yanchi Liu, Yiyu Sun, Mika Oishi, Takao Osaki, Katsushi Matsuda, Jie Ji, Guangji Bai, Liang Zhao, Haifeng Chen. "Uncertainty Quantification for In-Context Learning of Large Language Models." In Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics. **NAACL**, 2024. [[Reliable LLM](#)]
21. Nan Zhang, Yanchi Liu, **Xujiang Zhao**, Wei Cheng, Runxue Bao, Rui Zhang, Prasenjit Mitra, Haifeng Chen. "Pruning as a Domain-Specific LLM Extractor," in Findings of the Association for Computational Linguistics. **NAACL**, 2024. [[Post Training](#)] [[Model Compression](#)]
22. Yijia Xiao, Yiqiao Jin, Yushi Bai, Yue Wu, Xianjun Yang, Xiao Luo, Wenchao Yu, **Xujiang Zhao**, Yanchi Liu, Quanquan Gu, Haifeng Chen, Wei Wang, Wei Cheng. "Large Language Models Can be Good Privacy Protection Learners." in Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing. **EMNLP**, 2024 [[Post Training](#)] [[Privacy](#)]
23. Yujie Lin, Chen Zhao, Minglai Shao, Baoluo Meng, **Xujiang Zhao**, and Haifeng Chen. "Towards counterfactual fairness-aware domain generalization in changing environments." **IJCAI**, 2024. [[Trustworthy AI](#)]
24. Yujie Lin, Chen Zhao, Minglai Shao, **Xujiang Zhao**, and Haifeng Chen. "Adaptation Speed Analysis for Fairness-aware Causal Models." In Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, **CIKM**, 2023. [[Trustworthy AI](#)]
25. Weili Shi, Xueying Yang, **Xujiang Zhao**, Haifeng Chen, Zhiqiang Tao, and Sheng Li. "Calibrate Graph Neural Networks under Out-of-Distribution Nodes via Deep Q-learning." In Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, **CIKM**, 2023. [[Trustworthy AI](#)] [[RL](#)]
26. Chen Ling, Xuchao Zhang, **Xujiang Zhao**[†], Yanchi Liu, Wei Cheng, Mika Oishi, Takao Osaki, Katsushi Matsuda, Haifeng Chen, Liang Zhao. "Open-Ended Commonsense Reasoning with Unrestricted Answer Candidates," in Findings of 2023 Conference on Empirical Methods in Natural Language Processing.: **EMNLP**, 2023 [[Application](#)]
27. **Xujiang Zhao**, Xuchao Zhang, Chen Zhao, Jin-hee Cho, Lance Kaplan, Dong Hyun Jeong, Audun Jøsang, Haifeng Chen, Feng Chen, "Multi-Label Temporal Evidential Neural Networks for Early Event Detection." In IEEE International Conference on Acoustics, Speech and Signal Processing, **ICASSP**, 2023. [[Application](#)]
28. **Xujiang Zhao**, Krishnateja Killamsetty, Rishabh Iyer, Feng Chen. "How Out-of-Distribution Data Hurts Semi-Supervised Learning." IEEE International Conference on Data Mining, **ICDM**, 2022. [[Trustworthy AI](#)]
29. Xueying Yang, Jiamian Wang, **Xujiang Zhao**, Zhiqiang Tao. "Calibrate Automated Graph Neural Network via Hyperparameter Uncertainty." In Proceedings of the 31st ACM International Conference on Information and Knowledge Management, **CIKM**, 2022. [[Trustworthy AI](#)]
30. **Xujiang Zhao**, Xuchao Zhang, Wei Cheng, Wenchao Yu, Yuncong Chen, Haifeng Chen, Feng Chen. "SEED: Sound Event Early Detection via Evidential Uncertainty". In IEEE International Conference on Acoustics, Speech and Signal Processing, **ICASSP**, 2022. [[Application](#)]

31. Haoliang Wang, Chen Zhao, **Xujiang Zhao**, Feng Chen. "Layer Adaptive Deep Neural Networks for Out-of-distribution Detection". In Pacific-Asia Conference on Knowledge Discovery and Data Mining, **PAKDD**, 2022. [[Trustworthy AI](#)]
32. Krishnateja Killamsetty, **Xujiang Zhao**, Feng Chen, Rishabh Iyer. "RETRIEVE: Corset Selection for Efficient and Robust Semi-Supervised Learning". In Advances in neural information processing systems, **NeurIPS**, 2021. [[Efficient AI](#)]
33. Liyan Xu, Xuchao Zhang, **Xujiang Zhao**, Haifeng Chen, Feng Chen, Jinho D. Choi. "Boosting Cross-Lingual Transfer via Self-Learning with Uncertainty Estimation". In Proceedings of 2021 Conference on Empirical Methods in Natural Language Processing, **EMNLP**, 2021. [[Application](#)]
34. Zhuoyi Wang, Chen Zhao, Yuqiao Chen, Hemeng Tao, Yu Lin, **Xujiang Zhao**, Yigong Wang, Latifur Khan. "CLEAR: Contrastive-Prototype Learning with Drift Estimation for Resource Constrained Stream Mining." In Proceeding of The Web Conference, **WWW**, 2021. [[Application](#)]
35. Yibo Hu, Yuzhe Ou, **Xujiang Zhao**, Feng Chen. "Multidimensional Uncertainty-Aware Evidential Neural Networks." In Proceeding of the Thirty-fifth AAAI Conference on Artificial Intelligence, **AAAI**, 2021. [[Trustworthy AI](#)]
36. **Xujiang Zhao**, Feng Chen, Shu Hu, Jin-Hee Cho. "Uncertainty Aware Semi-Supervised Learning on Graph Data." In Advances in Neural Information Processing Systems, **NeurIPS**, 2020, (**Spotlight; Acceptance rate: 4%**). [[Trustworthy AI](#)]
37. Weishi Shi, **Xujiang Zhao**, Qi Yu, Feng Chen. "Multifaceted Uncertainty Estimation for Label-Efficient Deep Learning." In Advances in neural information processing systems, **NeurIPS**, 2020. [[Efficient AI](#)]
38. Adil Alim, **Xujiang Zhao**, Jin-Hee Cho, Feng Chen. "Uncertainty-Aware Opinion Inference Under Adversarial Attacks." In 2019 IEEE International Conference on Big Data, **Big Data**, 2019. [[Trustworthy AI](#)]
39. **Xujiang Zhao**, Yuzhe. Ou, Lance. Kaplan, Feng. Chen, and Jin-Hee. Cho. "Quantifying Classification Uncertainty using Regularized Evidential Neural Networks." accepted to **AAAI 2019 Fall Symposium Series**, Artificial Intelligence in Government and Public Sector. [[Trustworthy AI](#)]
40. **Xujiang Zhao**, Shu Hu, Jin-Hee Cho, and Feng Chen. "Uncertainty-based Decision Making using Deep Reinforcement Learning." In 2019 22nd International Conference on Information Fusion , **FUSION**, 2019. [[RL](#)]
41. **Xujiang Zhao**, Feng Chen, and Jin-Hee Cho. "Deep Learning for Predicting Dynamic Uncertain Opinions in Network Data." In 2018 IEEE International Conference on Big Data, **Big Data** 2018. [[Trustworthy AI](#)]
42. **Xujiang Zhao**, Feng Chen, and Jin-Hee Cho. "Deep Learning based Scalable Inference of Uncertain Opinions." In 2018 IEEE International Conference on Data Mining, **ICDM** 2018. [[Trustworthy AI](#)]
43. **Xujiang Zhao**, Feng Chen, and Jin-Hee Cho. "Uncertainty-Based Opinion Inference on Network Data Using Graph Convolutional Neural Networks." IEEE Military Communications Conference , **MILCOM**, 2018. [[Trustworthy AI](#)]


Journal Article

1. Chen Ling, **Xujiang Zhao**[†], Jiaying Lu, Chengyuan Deng, Can Zheng, Junxiang Wang, Tanmoy Chowdhury, Yun Li, Hejie Cui, Xuchao Zhang, Amit P., Wei Chen, Haoyu Wang, Yanchi Liu, Zhengzhang Chen, Haifeng Chen, Chris White, Quanquan Gu, Jian Pei, Carl Yang, Liang Zhao., "Domain Specialization as The Key to Make Large Language Models Disruptive: A Comprehensive Survey", **ACM Computing Surveys**, 2025. [[Post Training](#)]
2. Ali Riahi Samani, **Xujiang Zhao**, Feng Chen. "Distribution Shift, Generalization and OOD Challenge in Offline Reinforcement Learning: A comprehensive survey." **Neural Computing and Applications** , 2026. [[Reinforcement Learning](#)]
3. Junji Jiang, Chen Ling, Hongyi Li, Guangji Bai, **Xujiang Zhao**, Liang Zhao. "Quantifying Uncertainty in Graph Neural Network Explanations." **Frontiers in big Data**, 2024. [[Trustworthy AI](#)]
4. Zhen Guo*, Zelin Wan*, Qisheng Zhang*, **Xujiang Zhao***, Feng Chen, Jin-hee Cho, Qi Zhang, Lance Kaplan, Dong Hyun Jeong, Audun Jøsang. "A Survey on Uncertainty Reasoning and Quantification for Decision Making: Belief Theory Meets Deep Learning." **Information Fusion** (Impact IF: 17.5), 2023. [[Trustworthy AI](#)]

Under Review



1. Peng Xia, Peng Xia, Jianwen Chen, Hanyang Wang, Jiaqi Liu, Kaide Zeng, Yu Wang, Siwei Han, Yiyang Zhou, **Xujiang Zhao**, Haifeng Chen, Zeyu Zheng, Cihang Xie, Huaxiu Yao. "SkillRL: Evolving Agents via Recursive Skill-Augmented Reinforcement Learning." **NeurIPS 2026** Under Review. [\[Agentic AI\]](#) [\[Self Evolving\]](#)
2. Dong Li, Yanchi Liu, **Xujiang Zhao**[†], Wei Cheng, Zhengzhang Chen, Xintao Wu, Zhong Chen, Chen Zhao, Haifeng Chen. "Memory-Induced Inference-Time Adaptation for Continual Learning in Small Language Models." **NeurIPS 2026** Under Review. [\[Agentic AI\]](#) [\[Memory\]](#)
3. Dong Li, Yanchi Liu, **Xujiang Zhao**[†], Wei Cheng, Zhengzhang Chen, Xintao Wu, Zhong Chen, Chen Zhao, Haifeng Chen. "HierFlow: Hierarchical Dual-Space Search for Automatic Agentic Workflow Generation." **NeurIPS 2026** Under Review. [\[Agentic AI\]](#) [\[Workflow\]](#)
4. Shang Ma, Haifeng Chen, Yuanzhou Chen, Yanchi Liu, **Xujiang Zhao**, Wenchao Yu, Yanfang Ye, Xusheng Xiao, Wei Cheng. "The Geometry of Fakeness: OOD Learning for AI-Image Detection with Information Bottleneck." **NeurIPS 2026** Under Review. [\[Trustworthy AI\]](#)
5. Jiaqi Liu, Shi Qiu, Mairui Li, Bingzhou Li, Peng Xia, Siwei Han, Haonian Ji, Letian Zhang, Hardy Chen, Haoqin Tu, Xinyu Yang, **Xujiang Zhao**, Haifeng Chen, Jiawei Zhou, Xiao Wang, Hongtu Zhu, Yun Li, Jiaheng Zhang, Yuyin Zhou, Sheng Wang, James Zou, Zeyu Zheng, Cihang Xie, Mingyu Ding, Huaxiu Yao. "AutoResearchClaw: End-to-end Autonomous Research From Idea to Paper." **NeurIPS 2026** Under Review. [\[Agentic AI\]](#) [\[Self Evolving\]](#)
6. Hongyu Cao, Yanchi Liu, Kunpeng Liu, **Xujiang Zhao**, Wei Cheng, Yanjie Fu, Haifeng Chen. "Data-centric Small LLM Learning: A Gradient Admission Perspective." **NeurIPS 2026** Under Review. [\[Agentic AI\]](#) [\[Memory\]](#)
7. Xinshuai Dong, Haifeng Chen, Xuyuan Liu, Shengyu Chen, Haoyu Wang, Yanchi Liu, **Xujiang Zhao**, Kun Zhang, Zhengzhang Che. "TaPE: Certified Robustness Against Table Permutations with Tabular Positional Encoding." **NeurIPS 2026** Under Review. [\[Post Training\]](#)
8. Bangwei Guo, **Xujiang Zhao**[†], Yanchi Liu, Wei Cheng, Shengyu Chen, Dongyue Li, Morimoto Masaharu, Takayuki Kuroda, Dimitris N. Metaxas, Haifeng Chen. "TopoAgent: A Structure-Aware Perception-to-Reasoning Framework for Diagram-to-Graph Topology Extraction with Large Vision-Language Models." **EMNLP 2026** Under Review. [\[Agentic AI\]](#) [\[Vision Reasoning\]](#)
9. Xinyu Wu, Yanchi Liu, Dong Li, **Xujiang Zhao**, Wei Cheng, Zhengzhang Chen, Chen Zhao. "FlexTag: Instruction Tagging for Multi-Perspective Understanding with Small Language Models." **EMNLP 2026** Under Review. [\[Post Training\]](#)
10. Kangyu Zhu, **Xujiang Zhao**[†], Hao Chen, Peng Xia, Yiming Liang, Siwei Han, Huaxiu Yao, Haifeng Chen. "PRISM: Progressive Reasoning with Iterative Structured Memory for Multi-Step Visual Understanding." **ECCV 2026** Under Review. [\[Agentic AI\]](#) [\[Vision Reasoning\]](#)

HONORS & Awards

- 2026  ICLR 2026 MemAgents Workshop Best Paper Award Runner-Up.
- 2025  The Excellent Invention Award at NEC Corporation.
- 2024  Our ACM Computing Surveys paper on Domain Specialization of LLMs is honorably mentioned by **The 2024 Economic Report of the President from the White House**.
- 2020  NeurIPS Spotlight Presentation (Top 4%).
- 2018  ICDM 2018 Student Travel Award.
- 2014  Outstanding Graduate Award of Chongqing University.
 First-class College Scholarship at Chongqing University.
- 2013  National Scholarship.
- 2012  First Prize in The National Drawing Skills and Advanced Technology.

Mentoring Experience






Thesis Committee

- Chen Ling  Ph.D. from *Emory University*, 2024
-  EMNLP 2023, NAACL 2024, ICLR-W 2024, Frontiers in Big Data 2024, ACM Computing Surveys 2025

Interns






- Dong Li  Ph.D. student from *Baylor University*, 2025 Spring - 2026 Summer
-  NeurIPS 2025, ACL 2025, NeurIPS 2026 (Submitted)
- Bangwei Guo  Ph.D. from *Rutgers University*, 2025 Fall - 2026 Spring
-  EMNLP 2026 (Submitted)
- Kangyu Zhu  Ph.D. from *University of North Carolina at Chapel Hill*, 2026 Spring
-  ECCV 2026 (Submitted)
- Peng Xia  Ph.D. from *University of North Carolina at Chapel Hill*, 2025 Fall
-  NeurIPS 2026 (Submitted)
- Linlin Yu  Ph.D. from *The University of Texas at Dallas*, 2025 Summer
-  NeurIPS 2025, ACL 2026
- Qiwei Zhao  Ph.D. student from *University of North Carolina at Chapel Hill*, 2024 Summer
-  ACL 2025
- Ruomeng Ding  Master student from *Georgia Institute of Technology*, 2023 Summer
-  AAAI-W 2023, SDM 2025

Invited Talks


- Sep 2025  Uncertainty Quantification and Reasoning for Large Language Models, **Seminar Talk** at Brigham Young University.
- May 2025  Uncertainty Quantification and Reasoning for Large Language Models, **Keynote** at SDM 2025 ARRML Workshop.
- Dec 2024  Uncertainty Quantification in LLMs, **Keynote** at IEEE BigData 2024 RobustMLDS Workshop.
- Dec 2020  Uncertainty Aware Semi-Supervised Learning on Graph Data, **Spotlight Presentation** at NeurIPS 2020.
- Nov 2018  Deep Learning-based Scalable Inference of Uncertain Opinions, **Keynote** at Institute of Information Engineering, Chinese Academy of Sciences (CAS).

Professional Activities

Program Chair

- 2026  5th KDD Workshop on Uncertainty Reasoning and Quantification in Decision Making.
- 2025  4th KDD Workshop on Uncertainty Reasoning and Quantification in Decision Making.
- 2024  3rd KDD Workshop on Uncertainty Reasoning and Quantification in Decision Making.
- 2023  2nd KDD Workshop on Uncertainty Reasoning and Quantification in Decision Making.
-  1st AAAI Workshop on Uncertainty Reasoning and Quantification in Decision Making.

Journal Reviewer

-  IEEE Transactions on Knowledge and Data Engineering

Professional Activities (continued)

- IEEE Transactions on Services Computing
- ACM Transactions on Knowledge Discovery from Data
- Frontiers in Big Data
- Neurocomputing

Program Committee Member / Reviewer

- 2022-2026 IEEE International Conference on Learning Representations (ICLR)
- 2022-2025 International Conference on Machine Learning (ICML)
- 2021-2025 Conference on Neural Information Processing Systems (NeurIPS)
- 2022-2025 AAAI Conference on Artificial Intelligence (AAAI)
- 2020-2025 ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)
- 2022-2025 Association for Computational Linguistics (ACL) Rolling Review (ARR)
- 2022 SIAM International Conference on Data Mining (SDM)
- 2022-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- 2022 ACM International Conference on Web Search and Data Mining (WSDM)