

# CHENG QIAN

chengq9@illinois.edu — Google Scholar

## EDUCATION

---

### University of Illinois Urbana Champaign, Ph.D. in Computer Science

2024 – Present

Member of BLENDER Lab, advised by Prof. Heng Ji.

Research Interest: Natural Language Processing, LLM Tool Learning and LLM-driven AI Agent; LLM Efficient Training and Knowledge Acquisition; LLM Reasoning; LLM-User Interaction.

### Tsinghua University, B.Eng. in Computer Science

2020 – 2024

Member of THUNLP, advised by Prof. Zhiyuan Liu.

## PUBLICATION & TALKS

---

### 2025

23. C. Qian\*, E. C. Acikgoz\*, H. Wang, X. Chen, A. Sil, D. Hakkani-Tür, G. Tur, and H. Ji, “Smart: Self-aware agent for tool overuse mitigation,” *Findings of ACL 2025*, 2025
22. C. Qian, E. C. Acikgoz, Q. He, H. Wang, X. Chen, D. Hakkani-Tür, G. Tur, and H. Ji, “Toolrl: Reward is all tool learning needs,” in *NeurIPS 2025*, 2025
21. P. Han\*, C. Qian\*, X. Chen, Y. Zhang, D. Zhang, and H. Ji, “Safeswitch: Steering unsafe llm behavior via internal activation signals,” in *Findings of EMNLP 2025*, 2025
20. C. Qian, H. Du, H. Wang, X. Chen, Y. Zhang, A. Sil, C. Zhai, K. McKeown, and H. Ji, “Modelingagent: Bridging llms and mathematical modeling for real-world challenges,” *Findings of EMNLP 2025*, 2025
19. R. Yang\*, H. Chen\*, J. Zhang\*, M. Zhao\*, C. Qian, K. Wang, Q. Wang, T. V. Koripella, M. Movahedi, M. Li, et al., “Embodiedbench: Comprehensive benchmarking multi-modal large language models for vision-driven embodied agents,” *ICML 2025*, 2025
18. K. Zhu, H. Du, Z. Hong, X. Yang, S. Guo, Z. Wang, Z. Wang, C. Qian, X. Tang, H. Ji, et al., “Multiagentbench: Evaluating the collaboration and competition of llm agents,” in *ACL 2025 Main Conference*, 2025
17. B. He, W. Zhang, J. Song, C. Qian, Z. Fu, B. Sun, N. Ding, H. Hong, L. Huang, H. Xue, et al., “Air: A systematic analysis of annotations, instructions, and response pairs in preference dataset,” *COLM 2025*, 2025
16. Y. Zhang, S. Li, C. Qian, J. Liu, P. Yu, C. Han, Y. R. Fung, K. McKeown, C. Zhai, M. Li, et al., “The law of knowledge overshadowing: Towards understanding, predicting, and preventing llm hallucination,” *Findings of ACL 2025*, 2025
15. B. Lyu, X. Cong, H. Yu, P. Yang, C. Qian, Z. Wang, Y. Qin, Y. Ye, Y. Lu, C. Qian, et al., “Enhancing open-domain task-solving capability of llms via autonomous tool integration from github,” in *ACL 2025 Main Conference*, 2025
14. H. Wang\*, C. Qian\*, M. Li, J. Qiu, B. Xue, M. Wang, H. Ji, and K.-F. Wong, “Toward a theory of agents as tool-use decision-makers,” *arXiv preprint arXiv:2506.00886*, 2025
13. H. Wang, C. Qian, W. Zhong, X. Chen, J. Qiu, S. Huang, B. Jin, M. Wang, K.-F. Wong, and H. Ji, “Acting less is reasoning more! teaching model to act efficiently,” *arXiv preprint arXiv:2504.14870*, 2025
12. C. Qian, Z. Liu, A. Prabhakar, Z. Liu, J. Zhang, H. Chen, H. Ji, W. Yao, S. Heinecke, S. Savarese, et al., “Userbench: An interactive gym environment for user-centric agents,” *arXiv preprint arXiv:2507.22034*, 2025
11. C. Qian, Z. Liu, A. Prabhakar, J. Qiu, Z. Liu, H. Chen, S. Kokane, H. Ji, W. Yao, S. Heinecke, et al., “Userll: Training interactive user-centric agent via reinforcement learning,” *arXiv preprint arXiv:2509.19736*, 2025
10. Q. He, C. Qian, X. Chen, B. He, H. Ji, et al., “Veri-r1: Toward precise and faithful claim verification via online reinforcement learning,” *arXiv preprint arXiv:2510.01932*, 2025
9. Y. Zhang, Q. Wang, C. Qian, J. Liu, C. Sun, D. Zhang, T. Abdelzaher, C. Zhai, P. Nakov, and H. Ji, “Atomic

reasoning for scientific table claim verification,” *arXiv preprint arXiv:2506.06972*, 2025

8. E. C. Acikgoz, C. Qian, H. Wang, V. Dongre, X. Chen, H. Ji, D. Hakkani-Tür, and G. Tur, “A desideratum for conversational agents: Capabilities, challenges, and future directions,” *arXiv preprint arXiv:2504.16939*, 2025
7. X. Chen, G. Li, Z. Wang, B. Jin, C. Qian, Y. Wang, H. Wang, Y. Zhang, D. Zhang, T. Zhang, et al., “Rm-r1: Reward modeling as reasoning,” *arXiv preprint arXiv:2505.02387*, 2025
6. X. Chen, S. Wang, C. Qian, H. Wang, P. Han, and H. Ji, “Decisionflow: Advancing large language model as principled decision maker,” in *Findings of EMNLP 2025*, 2025
5. S. Wu, C. Qian, Y. R. Fung, P. P. Liang, and H. Ji, “Alice: Proactive learning with teacher’s demonstrations for weak-to-strong generalization,” *arXiv preprint arXiv:2504.07316*, 2025
4. G. Zhang, Q. Su, J. Liu, C. Qian, Y. Pan, Y. Fu, and D. Zhang, “Isacl: Internal state analyzer for copyrighted training data leakage,” *arXiv preprint arXiv:2508.17767*, 2025
3. M. Cui, Y. Jiang, D. Zhou, C. Qian, Y. Zhang, and Q. Wang, “Shortagesim: Simulating drug shortages under information asymmetry,” *arXiv preprint arXiv:2509.01813*, 2025
2. R. G. Reddy, T. Dixit, J. Qin, C. Qian, D. Lee, J. Han, K. Small, X. Fan, R. Sarikaya, and H. Ji, “Winell: Wikipedia never-ending updating with llm agents,” *arXiv preprint arXiv:2508.03728*, 2025
1. H.-a. Gao, J. Geng, W. Hua, M. Hu, X. Juan, H. Liu, S. Liu, J. Qiu, X. Qi, Y. Wu, et al., “A survey of self-evolving agents: On path to artificial super intelligence,” *arXiv preprint arXiv:2507.21046*, 2025

## 2024

7. C. Qian\*, B. He\*, Z. Zhuang, J. Deng, Y. Qin, X. Cong, Y. Lin, Z. Zhang, Z. Liu, and M. Sun, “Tell me more! towards implicit user intention understanding of language model driven agents,” in *ACL 2024 Main Conference*, 2024
6. C. Qian, P. Han, Q. Luo, B. He, X. Chen, Y. Zhang, H. Du, J. Yao, X. Yang, D. Zhang, et al., “Escapebench: Towards advancing creative intelligence of language model agents,” in *ACL 2025 Main Conference*, 2024
5. C. Qian\*, S. Liang\*, Y. Qin, Y. Ye, X. Cong, Y. Lin, Y. Wu, Z. Liu, and M. Sun, “Investigate-consolidate-exploit: A general strategy for inter-task agent self-evolution,” *arXiv preprint arXiv:2401.13996*, 2024
4. B. He\*, N. Ding\*, C. Qian\*, J. Deng, G. Cui, L. Yuan, H. Hong, H.-a. Gao, L. Huang, H. Xue, et al., “The right time matters: Data arrangement affects zero-shot generalization in instruction tuning,” *ACL 2025 Findings*, 2024
3. S. Wu, M. Fung, C. Qian, J. Kim, D. Hakkani-Tur, and H. Ji, “Aligning llms with individual preferences via interaction,” *COLING 2025*, 2024
2. Y. Lu, S. Yang, C. Qian, G. Chen, Q. Luo, Y. Wu, H. Wang, X. Cong, Z. Zhang, Y. Lin, et al., “Proactive agent: Shifting llm agents from reactive responses to active assistance,” *ICLR 2025*, 2024
1. R. Tian, Y. Li, Y. Fu, S. Deng, Q. Luo, C. Qian, S. Wang, X. Cong, Z. Zhang, Y. Wu, et al., “Distance between relevant information pieces causes bias in long-context llms,” *Findings of ACL 2025*, 2024

## 2023

6. C. Qian, C. Han, Y. Fung, Y. Qin, Z. Liu, and H. Ji, “Creator: Tool creation for disentangling abstract and concrete reasoning of large language models,” in *Findings of EMNLP 2023*, 2023
5. C. Qian, C. Xiong, Z. Liu, and Z. Liu, “Toolink: Linking toolkit creation and using through chain-of-solving on open-source model,” in *NAACL 2024 Main Conference*, 2023
4. C. Qian, X. Zhao, and S. T. Wu, “”merge conflicts!” exploring the impacts of external distractors to parametric knowledge graphs,” in *COLM 2024*, 2023
3. Y. Qin\*, C. Qian\*, X. Han, Y. Lin, H. Wang, R. Xie, Z. Liu, M. Sun, and J. Zhou, “Recyclable tuning for continual pre-training,” in *Findings of ACL 2023*, 2023
2. Y. Qin, S. Hu, Y. Lin, W. Chen, N. Ding, G. Cui, Z. Zeng, Y. Huang, C. Xiao, C. Han, et al., “Tool learning with foundation models,” *ACM Computing Survey*, 2023

1. X. Team, “Xagent: An autonomous agent for complex task solving,” *XAgent blog*, 2023

2022

1. Y. Qin\*, C. Qian\*, J. Yi\*, W. Chen, Y. Lin, X. Han, Z. Liu, M. Sun, and J. Zhou, “Exploring mode connectivity for pre-trained language models,” in *EMNLP 2022 Main Conference*, 2022

(\* indicates equal contribution)

#### Patents:

Name: Training Method, Device, Medium, and Program Product for Predictive Model. Type: Invention. Inventors: Zhiyuan Liu, Maosong Sun, Cheng Qian (Technical Lead), Yujia Qin, Yankai Lin, Ruobing Xie, Jie Zhou, Tsinghua University & Tencent Inc. Ref. Num: CN202211242596.7.

#### Talks:

- Towards Reasoning as Action. Invited guest lecture at Northwestern University (CS 496 Agent AI). Host: Manling Li, 2025.5.
- Towards Reasoning as Action. Invited talk at IBM Research. Host: Avirup Sil, 2025.5.
- ToolRL: Reward is all Tool Learning Needs. Invited talk at NICE (NLP Academic Exchange Platform). Host: Hongru Wang, 2025.5.
- Implicit User Intention Understanding of LLM Agent. Invited talk at DEFIRST Reading Group. Host: Rebecca Salganik, 2025.5.
- Tool Learning and Tool Creation of LLM. Invited talk at SenseTime Smart City Group. Host: Hangyu Mao, 2023.6.

#### RESEARCH & WORK

---

**Summer Research Intern, Salesforce AI Research. Manager: Huan Wang** 2025  
Research Direction: Reinforcement learning of LLM agent; Agent’s user alignment; LLM Router

**Visiting Student, WInE Lab, CMU. Advisor: Prof. Sherry Wu** 2023  
Research Direction: LLM knowledge conflict.

**Visiting Student, BLENDER Lab, UIUC. Advisor: Prof. Heng Ji** 2023  
Research Direction: LLM agent tool creation.

**Undergrad Research Intern, CMU. Advisor: Prof. Chenyan Xiong** 2023  
Research Direction: Open-source adaptation of LLM tool use ability.

#### AWARDS

---

- Capital One ASKS Research Fellowship, UIUC. 2025
- Comprehensive Excellence Scholarship, Tsinghua University (1/180). 2021, 2023
- Volunteering & Social Excellence Scholarship, Tsinghua University (1/180). 2022
- Awards of Excellent Student Cadre, Tsinghua University. 2021
- Honorable Mention in 2023 Mathematical Contest in Modeling (MCM). 2023

#### SKILLS & SERVICE

---

##### English Skills

TOEFL (Best) 115/120 (Reading 30, Listening 30, Speaking 26, Writing 29).  
GRE Verbal Reasoning 162/170, Quantitative Reasoning 170/170, Analytical Writing 4/6.

##### Technical Skills

Proficient in C/C++, Python (PyTorch), LaTeX, Linux.  
Familiar with various neural networks and state-of-the-art deep learning techniques.

##### Service

Area Chair: ACL 2025, EMNLP 2025  
Reviewer: ARR 2023/2024/2025, ACL 2023/2024, EMNLP 2023/2024, NAACL 2023/2024, COLM 2024/2025, COLING 2025, NeurIPS 2025, AAAI 2025