Cheng Tang

chengt6@illinois.edu | +1 (217) 600-1859
tach1018.github.io | github.com/tach1018

Research Interests

My research focuses on reinforcement learning and optimization for sequential decision-making under uncertainty, with an emphasis on theoretical guarantees and real applications in post training or RLHF. I study policy-gradient convergence, robust offline learning, and stochastic networks, and am interested in learning-guided optimization and resource allocation.

Education

University of Illinois Urbana-Champaign Aug 2025 – Present
Ph.D. in Computer Science | GPA: 4.0/4.0
Advisor: Prof. Yunzong Xu

Tsinghua University Sep 2021 – Jun 2025
B.S. in Mathematics; B.Ec. in Economics and Finance
Graduation Commencement Speaker for the School of Economics and Management

Research Experience

Microsoft Research Asia May 2026 – Aug 2026
Research Intern, Machine Learning Group
- Led a first-author project on adaptive policy-gradient methods for linear softmax bandits, spanning algorithm design, theoretical analysis, implementation, and experiments.
- Proved fast-escape and global-convergence results; manuscript submitted to AISTATS 2027.

University of Illinois Urbana-Champaign Aug 2025 – Present
Doctoral Research | Advisor: Prof. Yunzong Xu
- Study finite-time queue peaks in stochastic networks. Coauthored a manuscript on logarithmic peak growth after geometric thresholds under interior slack, with matching lower bounds.
- Studied linear contextual bandits through smoothed analysis, examining how random perturbations to contexts affect exploration and regret.

Duke University Mar 2024 – Nov 2024
Research Intern, Remote | Advisor: Prof. Pan Xu
- Studied robust offline RL with linear function approximation and f-divergence regularization; co-first-author paper at ICML 2025 with instance-dependent upper bounds and information-theoretic lower bounds.

Carnegie Mellon University Aug 2024 – Oct 2024
Research Intern | Advisor: Prof. Ding Zhao
- Developed offline safe RL methods for learning high-reward policies from imbalanced datasets under safety constraints, led the project from formulation through theoretical and empirical evaluation. Technical report: [link]

ModelBest Oct 2023 – Feb 2024
Research Intern | Advisor: Prof. Zhiyuan Liu
- Redesigned LLaMA generation and decoding logic; evaluated the methods on LongBench and LEval.

Publications/Manuscripts

** Equal contribution*

Robust Offline Reinforcement Learning with Linearly Structured f-Divergence Regularization
Cheng Tang*, Zhishuai Liu*, Pan Xu
ICML 2025; also accepted at the NeurIPS ORxML Workshop 2025. | Paper | Code
Introduces structured robust regularization and pessimistic value iteration, with instance-dependent upper bounds and information-theoretic lower bounds.

RL-Inf: Tracking Non-local Training Data Influence for Online Reinforcement Learning
Shixuan Liu*, **Cheng Tang***, Yuzheng Hu*, Fan Wu, Han Zhao, Jiaqi-W Ma
NeurIPS 2026, accepted; also accepted at the AIW Workshop at COLM 2026.
Studies non-local training-data influence in online reinforcement learning.

Adaptive Delightful Policy Gradient in Linear Softmax Bandits: Fast Escape and Global Convergence
Cheng Tang, Changsheng Chen, Zhengjun Huang, Siwei Wang
Under review at AISTATS 2027. Research internship project at Microsoft Research Asia.

Finite-Time Queue Peak Laws in Stochastic Networks: Logarithmic Scaling After Geometric Thresholds
Hao Liang, **Cheng Tang**, Yunzong Xu (alphabetical order)
Under review at Operations Research. Presented in INFORMS

Steering the Feedback Loop: Policy Gradients for Online Performative Optimization
Changsheng Chen, **Cheng Tang**, Lanran Fang, Grani A. Hanasusanto
Under review at AISTATS 2027.

CAMEL: Collaborative Agentic Memory via Evidence Linking across Speakers and Time
Zhengjun Huang, Zhoujin Tian, Junle Chen, **Cheng Tang**, Wenxuan Liu, Fangyuan Zhang, Qintian Guo, Liyi Zeng, Xiaofang Zhou, Zhaoquan Gu
Under review in ICLR 2027.

MedQPA-Gen: Medical Question Proposing and Answering for Report Generation
Weijie Liang, Xiyue Zhu, Ruike Zhu, Chenhao Li, **Cheng Tang**, Zhiyu Liu, Zhihua Gong, Shirui Luo, Yudu Li, Volodymyr Kindratenko
Findings of ACL 2026, accepted.

Honors and Awards

CS PhD Fellowship Addendum, UIUC	2025
Dean's Award of Zhili College, Tsinghua University	2024
Highest honor for students in Zhili College	
Excellent Comprehensive Scholarship, Tsinghua University	2024
Outstanding student in Top Open Program	2023
First Prize, China Undergraduate Mathematical Contest in Modeling	2023
Silver Medal, Chinese Mathematical Olympiad	2020

Other academic activities

Volunteer at ICML 2025, Vancouver, Canada
Host at QingLi academic festival, Tsinghua Univeristy, China.