

Runyu (Cathy) Zhang

32 Vassar Street, Cambridge MA, 02139

✉ runyuzha@mit.edu • 🌐 runyu-cathy-zhang.github.io/

Education

Massachusetts Institute of Technology

School of Engineering Distinguished Postdoctoral Fellow

Laboratory for Information and Decision Systems (LIDS)

Department of Civil and Environmental Engineering (CEE)

Advisors: Asuman Ozdaglar and Gioele Zardini

Cambridge, MA

July 2025 – Present

Harvard University

Ph.D. in Applied Mathematics, School of Engineering and Applied Sciences

Advisors: Na Li Committee: Jeff Shamma, David Parkes, Bo Dai

Dissertation: Advancing Multi-Agent Systems with Scalable and Robust Learning and Control

Cambridge, MA

Sept. 2019 – May 2025

Peking University

B.Sc. in Scientific and Engineering Computing

School of Mathematical Sciences

Beijing, China

Sept. 2015 – July 2019

Other Academic Experiences

Salesforce Research

Research Intern

Advisor: Yu Bai

Palo Alto, CA

June 2022 – September 2022

University of California, Los Angeles

Student Researcher

Advisor: Deanna Needell

Los Angeles, CA

June 2018 – September 2018

Research Interests

Areas: Control and Optimization, Reinforcement Learning, Artificial Intelligence, Generative Models, Multi-Agent Systems, Game Theory

Topics: Scalable, robust, and strategic coordination of multi-agent systems; learning-based and distributed control; multi-agent reinforcement learning and equilibrium selection; diffusion models for constrained optimization and planning; risk-sensitive and robust reinforcement learning; online and adaptive control. Applications include robotics, autonomous mobility and transportation systems, smart cities, power and energy systems, large language models, and scientific machine learning.

For a detailed overview of my research agenda and future directions, please see my online research statement [here](#).

Awards, Honors, and Scholarships

MIT Postdoctoral Fellowship for Engineering Excellence

2025

University of Michigan President's Postdoctoral Fellowship (offered, declined)

2025

Featured in IEEE Control Systems Magazine, "PhDs in Control"

2025

Northeast Robotics Colloquium (NERC) Rising Star Cornell University, Ithaca, New York	2025
EECS Rising Star	2024
2023 American Control Conference (ACC) Student Travel Award	2023
Finalist of the Two Sigma Diversity PhD Fellowship (1 of 8)	2022
Certificates of Distinction and Excellence in Teaching Derek Bok Center for Teaching and Learning, Harvard University	2020
Excellent Graduate Peking University	2019
Second Prize, China Undergraduate Mathematical Contest in Modeling, Beijing District Peking University	2016
Elite Undergraduate Training Program School of Mathematical Sciences, Peking University	2016–2019
First Prize in Beijing National High School Student Mathematics Competition	2014

Teaching Experience

ES 202: Learning, Estimation and Control of Dynamical Systems *Spring 2024*
Engineering and Applied Sciences, Harvard University **Rating: 4.75/5.0**
Section Leader. Led office hours and provided feedback on weekly readings and in-class discussions. Contributed to the design and grading of assignments.

ES155: Systems and Control *Fall 2020*
Engineering and Applied Sciences, Harvard University **Rating: 4.83/5.0**
Section Leader. Led weekly sections and office hours. Contributed to the design and grading of weekly assignments and the midterm and final exams. Provided feedback on course projects for 20+ students.

Guest Lectures and Tutorials.....

Invited Tutorial Lecturer, UT Austin Game Theory Summer School *Aug. 13, 2026*

Guest Lecturer, MIT AI Accelerator and U.S. Air Force Test Pilot School Workshop *Aug. 7, 2026*
 Massachusetts Institute of Technology

Guest Lecturer, MIT System Design and Management Seminar *Feb. 18, 2026*
 Massachusetts Institute of Technology

Guest Lecturer, Advanced Topics in Machine Learning *Oct. 16, 2024*
 Professor Bo Dai's course, Georgia Institute of Technology; topic: distributed learning and control in networked multi-agent systems

Research Mentoring and Student Supervision

At MIT

- Jiarui Li, Ph.D. student, MIT Civil and Environmental Engineering, 2025–present.
- Xinling Li, Ph.D. student, MIT Transportation, 2025–present.
- Yujun Huang, Ph.D. student, MIT Civil and Environmental Engineering, 2025–present.
- Yubo Cai, Ph.D. student, MIT CEE / Computational Science and Engineering, 2025–present.
- Riccardo Fiorista, graduate student, MIT Transportation, 2025–present.
- Karan Mahesh, graduate student, MIT Mechanical Engineering, 2026–present.

- Zeina Rmaile, MIT Summer Research Program intern and Georgia Institute of Technology undergraduate, Summer 2026.
- Augusto Mondelli, visiting master's student from ETH Zürich, 2026.
- Florian Lercher, visiting student from Technical University of Munich, 2026.
- Raphael Karmalker, visiting master's student at MIT Energy Initiative, 2026.

At Harvard

- Phevos Paschalidis, Harvard College, 2023–2024.
- Aryan Naveen, Harvard College, 2023–2024.
- Johan Olsson, visiting master's student from Lund University, 2023.

Outside Harvard and MIT

- Mingjia He, Ph.D. student, ETH Zürich, 2025–present.
- Binghang Lu, Ph.D. student, Purdue University, 2025–present.
- Zheyuan Deng, master's student, Brown University, 2025–present.
- Yuze Ge, master's student, School of Data Science, Fudan University, 2026–present.

Invited Talks

1. Conference on Decision and Control (CDC), Workshop on “Control Architecture Theory”, Honolulu, HI, 2026.
2. INFORMS Annual Meeting, Invited Session on “Online Learning and Reinforcement Learning”, San Francisco, CA, 2026.
3. SIAM Great Lakes Section Annual Meeting, Purdue University, West Lafayette, IN, 2026.
4. Artificially Intelligent Cities Workshop, New York University, New York, NY, 2026.
5. American Control Conference (ACC), Invited Session on “Autonomous Risk-Aware Perception, Planning, and Control”, New Orleans, LA, 2026.
6. American Control Conference (ACC), Workshop on “Toward Safe and Scalable Multi-Agent Systems: Bridging Control, Reinforcement Learning, and Generative AI”, New Orleans, LA, 2026.
7. Applied Mathematics Seminar, University of California, Irvine, CA, 2026.
8. Northeast Robotics Colloquium (NERC), Rising Star Spotlight, Cornell University, Ithaca, NY, 2025.
9. Advances in Neural Information Processing Systems (NeurIPS), San Diego, CA, 2025.
10. Invited Seminar Talk, Duke University, Durham, NC, 2025.
11. Invited Seminar Talk, University of New Hampshire, Durham, NH, 2025.
12. Invited Seminar Talk, Nanyang Technological University, Singapore, 2025.
13. Invited Seminar Talk, ETH Zürich, Zürich, Switzerland, 2025.
14. Invited Seminar Talk, Peking University, Beijing, China, 2025.
15. Invited Seminar Talk, Tsinghua University, Beijing, China, 2025.
16. Modeling and Optimization: Theory and Applications (MOPTA), Invited Session on “Multi-Agent Reinforcement Learning for Dynamic Games”, Lehigh University, Bethlehem, PA, 2024.
17. Conference on Decision and Control (CDC), Invited Session on “Strategic Coordination in Multi-Agent Systems and Its Applications”, Singapore, 2023.
18. INFORMS Annual Meeting, Invited Session on “Decentralized Learning and Optimization”, Phoenix, AZ, 2023.

19. ONR Science of Autonomy Program Review, Arlington, VA, 2023.
20. American Control Conference (ACC), San Diego, CA, 2023.
21. International Symposium on Mathematical Theory of Networks and Systems (MTNS), University of Bayreuth, Bayreuth, Germany, 2022.
22. International Conference on Continuous Optimization (ICCOPT), Invited Session on “Multi-Agent Learning in Games”, Lehigh University, Bethlehem, PA, 2022.
23. American Control Conference (ACC), Invited Session on “Online Learning and Games for Control”, Virtual, 2021.

Professional Service

Conference Leadership and Workshop Organization.....

Publications Co-Chair, 2027 Conference on Learning for Dynamics and Control (L4DC) 2027
KTH Royal Institute of Technology, Stockholm, Sweden

Student Activities Chair, 2027 IEEE International Conference on Intelligent Transportation Systems (ITSC) 2027
Boston, Massachusetts

Workshop Organizer and Session Chair, “Generative AI Meets Control and Optimization: Theory, Algorithms, and Systems” 2026
IEEE Conference on Decision and Control, Honolulu; co-organized with Jiawei Zhang and Asuman Ozdaglar

Workshop Organizer and Session Chair, “Toward Safe and Scalable Multi-Agent Systems: Bridging Control, Reinforcement Learning, and GenAI” 2026
American Control Conference, New Orleans; co-organized with Gioele Zardini and Na Li

IEEE Control Systems Society NextCom.....

Activities Co-Chair, IEEE CSS Next Generation Committee 2026–present

Member, IEEE CSS NextCom General Activities Committee Spring 2025–present

Co-organizer, IEEE CSS NextCom Jazz Riverboat Cruise at ACC 2026 May 27, 2026
New Orleans, Louisiana; More than 200 people networking event

Seminar Organization and Hosting.....

Student Liaison and Host, MIT LIDS Seminar: May 4, Apr. 21, 2026 and Nov. 12, 2025
Lillian Ratliff, Jeff Shamma, and Florian Dörfler

Student Co-Organizer, Harvard Machine Learning Foundations Seminar Fall 2022–Fall 2023

Student Co-Organizer, Harvard-EEML Student Seminar Fall 2021–Spring 2022

Outreach and Public Engagement.....

Speaker, MIT Civil and Environmental Engineering Annual Research Day Mar. 9, 2026
Massachusetts Institute of Technology

Guest Speaker and Graduate-Student Panelist, Women in SEAS Research Apr. 13, 2024
Harvard University; co-hosted by the Harvard Computer Society, Harvard Undergraduate Engineering Society, and Harvard Women in Computer Science

Speaker, SEAS Lightning Talks: Ideas for Societal Impact Feb. 28, 2024
Harvard John A. Paulson School of Engineering and Applied Sciences

Research Group Representative, AM/CS/EE Ph.D. Recruitment Event Oct. 6, 2022

Harvard University; represented Professor Na Li's research group to prospective doctoral students

Speaker, Harvard SEAS Research Showcase
Harvard University

Mar. 8, 2022

Speaker, Electrical Engineering Open House Teaser Talk
Harvard University

Mar. 11, 2021

Technical Reviewer for International Journals.....

TAC	IEEE Transactions on Automatic Control
AUT	Automatica
L-CSS	IEEE Control Systems Letters
IEEE-IS	IEEE Intelligent Systems
SCL	Systems & Control Letters
DGA	Dynamic Games and Applications
JMLR	Journal of Machine Learning Research
JSAIT	IEEE Journal on Selected Areas in Information Theory

Technical Reviewer for International Conferences.....

NeurIPS	Advances in Neural Information Processing Systems
ICML	International Conference on Machine Learning
ICLR	International Conference on Learning Representations
AISTATS	International Conference on Artificial Intelligence and Statistics
L4DC	Learning for Dynamics and Control Conference
ACC	American Control Conference
CDC	IEEE Conference on Decision and Control

Proposal Writing

DARPA MATHBAC Proposal: COSMIC 2026
COMpositional Swarm Mechanisms for Instrumented Communication. Technical proposal contributor; revised the multi-agent architecture and technical narrative for the DARPA Mathematics of Boosting Agentic Communication (MATHBAC) program.

NSF ACCESS Discover Allocation (Principal Investigator) 2026
Spectral-Norm-Aware Continual Learning with the Muon Optimizer. Proposal reviewed and approved for 750,000 ACCESS Credits.

MIT-Amazon Science Hub / Amazon Robotics Research Proposal 2026
Technical proposal contributor; developed research directions and technical narrative for a prospective multi-year collaboration with Amazon Robotics.

Siemens Energy and MIT Research Collaboration 2026
Scalable and Adaptive Multi-Agent Reinforcement Learning for Topology Optimization in Power Grids. Research co-lead and proposal co-developer; co-developed the proposal and co-leads the design of scalable and robust multi-agent reinforcement-learning methods. \$500,000 over 24 months; currently in the final paperwork phase between Siemens Energy and the MIT Energy Initiative.

Languages

Chinese (native), English (fluent)

List of Publications

Preprints / in preparation

- [U1] Z. Deng, B. Lu, H. Feng, S. Huang, D. Wang, Y. Xu, Z. Zhang, Y. Sun, C. Mou, **R. Zhang**, Y. Hao, B. Poczos, X. Li, “CONTRAMEM: Learning Self-Evolving Procedural Memory from Contrasting Multi-Model Trajectories”, 2026.
- [U2] X. Li, **R. Zhang**, G. Zardini, “Implementation-Based Incentive Design for Autonomous Mobility-on-Demand and Transit Systems”, 2026.
- [U3] B. Lu, Z. Deng*, **R. Zhang***, B. Hu, Y. Zhao, Y. Tian, C. Mou, G. Lin, X. Li, “Muon-OGD: Muon-Based Spectral Orthogonal Gradient Projection for LLM Continual Learning”, 2026.
- [U4] B. Lu*, **R. Zhang***, C. Mou, N. Li, G. Lin, “AdamFLIP: Adaptive Momentum Feedback Linearization Optimization for Hard Constrained PINN Training”, 2026.
- [U5] **R. Zhang**, J. Zhang, G. Zardini, S. Amin, A. Ozdaglar, “Denoising as Projection: Constrained Optimization with Gradient-Guided Diffusion”, 2026.
- [U6] M. He, A. Censi, **R. Zhang**, E. Frazzoli, G. Zardini, “Co-Investment with Payoff-Sharing Mechanism for Cooperative Decision-Making in Network Design Games”, 2025.
- [U7] J. Li, A. Zanardi, F. Pecora, **R. Zhang**, G. Zardini, “FICO: Finite-Horizon Closed-Loop Factorization for Unified Multi-Agent Path Finding”, 2025.
- [U8] **R. Zhang**, G. Zardini, A. Ozdaglar, J. Shamma, N. Li, “Zeroth-Order Constrained Optimization from a Control Perspective via Feedback Linearization”, 2025.
- [U9] **R. Zhang**, G. Zardini, A. Ozdaglar, J. Shamma, N. Li, “Equilibrium Selection for Multi-Agent Reinforcement Learning: A Unified Framework”, 2024.

Journal Articles

- [J1] X. Han, A. Malkawi, Z. Li, **R. Zhang**, N. Li, “Optimizing Radiant Floor Heating Control with Night Setbacks Using Model-Free Reinforcement Learning and Transfer Learning”, *Building and Environment*, vol. 287, p. 113771, 2026.
- [J2] J. Li, F. Pecora, **R. Zhang**, G. Zardini, “Adaptive-Horizon Conflict-Based Search for Closed-Loop Multi-Agent Path Finding”, *IEEE Robotics and Automation Letters*, vol. 11, no. 9, pp. 10505–10512, 2026.
- [J3] P. Paschalidis, A. Naveen, **R. Zhang**, N. Li, “Cooperative Multi-Agent Graph Bandits: UCB Algorithm and Regret Analysis”, *IEEE Transactions on Automatic Control* (in press), 2026.
- [J4] **R. Zhang**, N. Li, A. Ozdaglar, J. Shamma, G. Zardini, “Optimism as Risk-Seeking in Multi-Agent Reinforcement Learning”, *IEEE Control Systems Letters*, vol. 10, pp. 1–6, 2026, Additionally selected for presentation at the American Control Conference 2026.
- [J5] **R. Zhang**, W. Li, N. Li, “Optimal Network Control of Spatially Exponential Decaying Linear Quadratic Regulator”, *Automatica*, vol. 173, p. 112061, 2025.
- [J6] J. Haddock, T. Will, J. Vendrow, **R. Zhang**, D. Molitor, D. Needell, M. Gao, E. Sadovnik, “Neural Nonnegative Matrix Factorization for Hierarchical Multilayer Topic Modeling”, *Sampling Theory, Signal Processing, and Data Analysis*, vol. 22, p. 4, 2024.

- [J7] **R. Zhang**, Z. Ren, N. Li, “Gradient Play in Stochastic Games: Stationary Points, Convergence, and Sample Complexity”, *IEEE Transactions on Automatic Control*, vol. 69, no. 10, pp. 6499–6514, 2024.
- [J8] E. X. Chen, X. Han, A. Malkawi, **R. Zhang**, N. Li, “Adaptive Model Predictive Control with Ensembled Multi-Time Scale Deep-Learning Models for Smart Control of Natural Ventilation”, *Building and Environment*, vol. 242, p. 110519, 2023.
- [J9] Y. Li, Y. Tang, **R. Zhang**, N. Li, “Distributed Reinforcement Learning for Decentralized Linear Quadratic Control: A Derivative-Free Policy Optimization Approach”, *IEEE Transactions on Automatic Control*, vol. 67, no. 12, pp. 6429–6444, 2022.

Conference Proceedings

- [C1] J. Li, **R. Zhang**, G. Zardini, “Certificate-Driven Closed-Loop Multi-Agent Path Finding with Inheritable Factorization”, *2026 IEEE 65th Conference on Decision and Control (CDC)* (in press), Honolulu, Hawaii, USA, 2026.
- [C2] **R. Zhang**, G. Zardini, “Random-Subspace Sequential Quadratic Programming for Constrained Zeroth-Order Optimization”, *2026 IEEE 65th Conference on Decision and Control (CDC)* (in press), Honolulu, Hawaii, USA, 2026.
- [C3] Z. Ren*, **R. Zhang***, B. Dai, N. Li, “Scalable Spectral Representations for Multiagent Reinforcement Learning in Network MDPs”, *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 258, pp. 550–558, 2025.
- [C4] **R. Zhang**, A. Raghunathan, J. Shamma, N. Li, “Constrained Optimization From a Control Perspective via Feedback Linearization”, *Advances in Neural Information Processing Systems*, vol. 38, pp. 62328–62364, 2025.
- [C5] J. Olsson, **R. Zhang**, E. Tegling, N. Li, “Scalable Reinforcement Learning for Linear-Quadratic Control of Networks”, *2024 American Control Conference (ACC)*, pp. 1813–1818, 2024.
- [C6] P. Paschalidis, **R. Zhang**, N. Li, “Cooperative Multi-Agent Graph Bandits: UCB Algorithm and Regret Analysis”, *2024 American Control Conference (ACC)*, pp. 4044–4049, 2024.
- [C7] **R. Zhang**, Y. Hu, N. Li, “Soft Robust MDPs and Risk-Sensitive MDPs: Equivalence, Policy Gradient, and Sample Complexity”, *International Conference on Learning Representations (ICLR)*, 2024.
- [C8] **R. Zhang**, H. Ma, N. Li, “Multi-Agent Coverage Control with Transient Behavior Consideration”, *Proceedings of the 6th Annual Learning for Dynamics and Control Conference (L4DC)*, vol. 242, pp. 1464–1476, 2024.
- [C9] **R. Zhang**, W. Li, N. Li, “On the Optimal Control of Network LQR with Spatially-Exponential Decaying Structure”, *2023 American Control Conference (ACC)*, pp. 1775–1780, 2023.
- [C10] **R. Zhang**, Y. Zhang, R. Konda, B. Ferguson, J. Marden, N. Li, “Markov Games with Decoupled Dynamics: Price of Anarchy and Sample Complexity”, *2023 IEEE 62nd Conference on Decision and Control (CDC)*, pp. 6573–6580, 2023.
- [C11] **R. Zhang**, Y. Zheng, W. Li, N. Li, “On the Relationship of Optimal State Feedback and Disturbance Response Controllers”, *22nd IFAC World Congress*, vol. 56, no. 2, pp. 7503–7508, 2023.
- [C12] Y. Zhang, **R. Zhang**, Y. Gu, N. Li, “Multi-Agent Reinforcement Learning with Reward Delays”, *Proceedings of the 5th Annual Learning for Dynamics and Control Conference (L4DC)*, vol. 211, pp. 692–704, 2023.
- [C13] **R. Zhang***, Q. Liu*, H. Wang, C. Xiong, N. Li, Y. Bai, “Policy Optimization for Markov Games: Unified Framework and Faster Convergence”, *Advances in Neural Information Processing Systems*, vol. 35, pp. 21886–21899, 2022.

- [C14] **R. Zhang**, J. Mei, B. Dai, D. Schuurmans, N. Li, “On the Global Convergence Rates of Decentralized Softmax Gradient Play in Markov Potential Games”, *Advances in Neural Information Processing Systems*, vol. 35, pp. 1923–1935, 2022.
- [C15] **R. Zhang**, Z. Ren, N. Li, “Gradient Play in Stochastic Games: Stationary Points and Local Geometry”, *25th International Symposium on Mathematical Theory of Networks and Systems (MTNS)*, vol. 55, no. 30, pp. 73–78, 2022.
- [C16] **R. Zhang**, Y. Li, N. Li, “On the Regret Analysis of Online LQR Control with Predictions”, *2021 American Control Conference (ACC)*, pp. 697–703, 2021.
- [C17] Y. Li, Y. Tang, **R. Zhang**, N. Li, “Distributed Reinforcement Learning for Decentralized Linear Quadratic Control: A Derivative-Free Policy Optimization Approach”, *Proceedings of the 2nd Conference on Learning for Dynamics and Control (L4DC)*, vol. 120, pp. 814–814, 2020.
- [C18] M. Gao, J. Haddock, D. Molitor, D. Needell, E. Sadochnik, T. Will, **R. Zhang**, “Neural Nonnegative Matrix Factorization for Hierarchical Multilayer Topic Modeling”, *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pp. 6–10, 2019.

Refereed Workshop Papers

- [W1] **R. Zhang**, J. Zhang, G. Zardini, S. Amin, A. Ozdaglar, “Denoising-Corrected Gradient Guidance for Feasible Diffusion Planning: A Constrained-Optimization Perspective”, *Workshop on Hybrid Architectures for Embodied Autonomy: Bridging Learning, Planning, and Control at the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Pittsburgh, Pennsylvania, USA, 2026, **Spotlight presentation**.