

Xiaoyang Cao

Cambridge, MA · (+1) 617-216-8184 · xycao@mit.edu · Personal Website

EDUCATION

Massachusetts Institute of Technology

Cambridge, MA

M.S. in Technology and Policy (TPP)

Sep 2025 – May 2027 (expected)

M.S. in Computational Science and Engineering (CSE)

Laboratory for Information and Decision Systems (LIDS)

Research advisor: Prof. Michiel Bakker

Tsinghua University

Beijing, China

B.S. in Mathematics and Physics

Sep 2021 – Jun 2025

- GPA: 3.95 / 4.00 – ranked 3 / 60 (top 5%)

University of California, Berkeley

Berkeley, CA

Exchange Student – GPA: 3.93 / 4.00

Jan 2024 – Aug 2024

RESEARCH INTERESTS

Technical foundations for reliable and auditable AI; compound and multi-agent systems; reinforcement learning; safe autonomous systems.

PUBLICATIONS

- [1] **Xiaoyang Cao**, Siddarth Srinivasan, Michiel A. Bakker. “Do Modules Stay in Their Lane? Role Drift in Compound LLM Systems.” *Preprint*, 2026.
- [2] **Xiaoyang Cao**, Zelai Xu, Mo Guang, Kaiwen Long, Michiel Bakker, Yu Wang, Chao Yu. “RE-PO: Robust Enhanced Policy Optimization as a General Framework for LLM Alignment.” *International Conference on Learning Representations (ICLR)*, 2026.
- [3] **Xiaoyang Cao**, Zhe Fu, Alexandre M. Bayen. “Pareto Control Barrier Function for Inner Safe Set Maximization Under Input Constraints.” *American Control Conference (ACC)*, 2025.
- [4] **Xiaoyang Cao**, Dingyi Zhuang, Shenhao Wang, Jinhua Zhao. “Virtual Nodes Improve Long-term Traffic Prediction.” *Transportation Research Board Annual Meeting (TRB)*, 2025.

RESEARCH EXPERIENCE

Role Drift in Compound LLM Systems [1] [\[arXiv\]](#) [\[VentureBeat\]](#)

Apr 2026 – Present

Advised by Prof. Michiel Bakker, MIT

- Identified *role drift*: end-to-end reinforcement learning can improve terminal performance while causing specialized modules to abandon their intended functions.
- Developed role-level evaluations showing that 86% of an apparent 31-percentage-point accuracy gain disappeared once the decomposer’s assigned function was measured, demonstrating that end-to-end metrics can conceal internal failure.
- Proposed Role Anchor, a training objective that preserves the influence of role specifications while allowing modules to improve under end-to-end optimization.

RE-PO: Robust Preference Alignment for LLMs [2] [\[arXiv\]](#) [\[Project\]](#) Mar 2025 – Sep 2025

Advised by Prof. Yu Wang, Tsinghua University

- Proposed an EM-based framework that models preference-label correctness as a latent variable, jointly estimates label and annotator reliability, and systematically robustifies DPO, IPO, SimPO, and CPO.
- Proved convergence to the true noise level under calibration and improved AlpacaEval 2 win rates by up to 7.0 percentage points on Mistral-7B and Llama-3-8B, with additional gains on Arena-Hard and MultiPref.

Safe Control via Pareto Control Barrier Functions [3] [\[arXiv\]](#) May 2024 – Oct 2024

Advised by Prof. Alexandre Bayen, UC Berkeley

- Developed the Pareto Control Barrier Function algorithm by integrating Pareto multi-task learning into neural control barrier functions, maximizing the inner safe set while guaranteeing safety under input constraints.
- Matched the Hamilton–Jacobi viability kernel on a 2D inverted pendulum and achieved a 700% larger inner safe set than the benchmark on a 12-D quadrotor obstacle-avoidance task.

Virtual Nodes for Long-term Traffic Prediction [4] [\[arXiv\]](#) May 2024 – Sep 2024

Advised by Prof. Jinhua Zhao, MIT

- Introduced virtual nodes into spatio-temporal graph neural networks to mitigate over-squashing in long-horizon traffic forecasting, using a semi-adaptive adjacency matrix learned from node embeddings.
- Reduced RMSE by 6.27% and MAPE by 5.04% on a five-year, 716-sensor San Diego dataset; the learned adjacency emphasized traffic-intensive intersections, improving interpretability.

AWARDS & HONORS

Outstanding Winner & INFORMS Award, Mathematical Contest in Modeling – top 1% May 2023

Outstanding Innovation in Science & Technology Scholarship, Tsinghua University 2023, 2024

Comprehensive Excellence Scholarship, Tsinghua University Oct 2023

Outstanding Academic Scholarship, Tsinghua University Oct 2022

SKILLS & LANGUAGES

Programming: Python (PyTorch, scikit-learn, NumPy, Pandas), R, MATLAB

Optimization & Control: Gurobi, CVXPY, SciPy, Simulink

Languages: English (fluent), Chinese (native)