Simin Liu

simin.liu.1314@gmail.com • siminl.com • U.S. Citizen (Work Authorized)

I'm a Robotics PhD at CMU specializing in dexterous manipulation at the intersection of learning and model-based planning. Previously, I worked on learned safe control for quadrotors and multitask RL for locomotion. I've shipped all my research on hardware (bimanual KUKA, quadrotors, legged millirobots), giving me a strong intuition for what makes algorithms work on real systems. My model-based background (trajectory optimization, MPC) gives me a distinctive angle on current RL bottlenecks: in recent work, I found that planner-generated demos can greatly accelerate RL training, and I have concrete ideas for pushing RL efficiency and robustness along these lines.

Education

Carnegie Mellon University **2020 – 2026**
PhD in Robotics (Qualcomm Graduate Fellow)
Thesis: Reachable Sets for Control and Planning: From Reactive Safety to Contact-Rich Manipulation
Advisors: Prof. Changliu Liu, Prof. John Dolan
Selected Coursework: Deep RL & Imitation Learning, Machine Learning Theory and Algorithms, Mathematical Statistics, Convex Optimization, Underactuated Robotics, Motion Planning

University of California, Berkeley **2015 – 2019**
B.S. Electrical Engineering & Computer Science (Research Honors, Academic High Honors, Dean's List)
Thesis: Adaptive Control for Legged Locomotion via Meta-Learning
Advisors: Prof. Sergey Levine

Experience

Research Intern **09/2024 – 10/2025**
Robotics and AI Institute (formerly Boston Dynamics AI Institute) *Cambridge, MA*

Built a planner that enables a bimanual system to move large, heavy objects using whole-arm contact. Unlike prior sampling-based approaches, which could produce whole-arm plans but at poor quality, this planner globally optimizes over grasp sequencing and in-grasp motion jointly: producing consistent, efficient plans suitable for hardware deployment and reinforcement learning.
Methods: Graph of Convex Sets (GCS), RRT, contact dynamics, trajectory optimization, motion primitives, reachable sets

- Demonstrated whole-arm reorientation of large, heavy objects on a bimanual KUKA system. Outperformed the state-of-the-art RRT-based baseline across several cost functions: 60% lower actuation effort, 68% lower object path length, 25% fewer regrasps. Returns feasible solutions on 90% of randomized queries. [Demo]
- Designed the planner as a graph of convex sets: nodes are object configurations reachable within a fixed grasp, edges are regrasps. Shortest-path jointly optimizes grasp sequencing and in-grasp motion, which are typically handled separately by discrete search and trajectory optimization.
- Using the planner to generate synthetic demonstrations for downstream RL; exploring multiple RL-from-demonstrations formulations (trajectory imitation, AMP-style discriminators, force-guided teacher-student) to improve sample efficiency.
- Shipped in a multi-author Python research codebase: Dockerized for reproducibility, standard Git/PR workflow with senior-researcher code review, component-level tests and benchmark-based validation.

Research Assistant **09/2020 – current**
Intelligent Control Lab, CMU *Pittsburgh, PA*

Built full-stack safe control systems for agile robotic platforms, combining model-based planning with learned safety filters.
Methods: machine learning, control barrier functions (CBF), MPC

- Synthesized robust-adaptive safety controllers for nonlinear systems with unknown model parameters, composable with online system identification for end-to-end safety. Generated a collision-avoidance filter for a quadrotor with unknown drag in minutes. [Demo]
- Developed "neural control barrier functions": a neural safe controller parameterization that scales synthesis to systems with high state dimension. Learned a pendulum-balancing filter for a 10D quadrotor-pendulum with thrust limits in under 2 hours.
- Designed a full-stack architecture for the quadrotor-pendulum: MPC at the planning layer, neural CBF safety filter, cascaded PID inner loop. Demonstrated safety enforcement (preventing pendulum fall) during trajectory tracking. [Demo]

Research Assistant **09/2017 – 05/2019**
Robotic AI & Learning Lab, UC Berkeley *Berkeley, CA*

Developed a model-based RL pipeline for multi-task locomotion with online adaptation to novel conditions.
Methods: multitask RL, sim2real, domain randomization, few-shot learning

- Demonstrated few-shot adaptation to unseen tasks (novel terrains, state-estimation error, payload variation) with a 3-8x increase in path-following reward over a no-adaptation baseline. [Demo]
- Trained an adaptable dynamics model on a few hours of data generated via domain randomization over leg loss, terrain variation, and payload changes. At test time, the model adapts few-shot and is rolled out inside a sampling-based MPC.
- Validated the pipeline on hardware (legged millirobot).

Skills

Robot learning: reinforcement learning (PPO, SAC); planner-generated synthetic demos; teacher-student training; multitask RL; sim2real and domain randomization; sim and hardware data pipelines; imitation learning
Manipulation: whole-arm contact; bimanual manipulation; contact modeling and simulation
Motion planning and controls: trajectory optimization (direct methods and DDP/iLQR); MPC; sampling-based MPC (CEM, MPPI); sampling-based planning (RRT, PRM); Graph of Convex Sets (GCS); planning with motion primitives; safe control (CBF)
Hardware: KUKA iiwa; quadrotors; legged millirobots
Optimization: convex programs (QP, SOCP); NLP (SQP, interior-point)
Software: Python, C++, ROS2, Drake, MuJoCo, PyTorch, IPOPT/SNOPT, Linux, Git, Docker

Publications

[1] "Approximately Optimal Global Planning for Contact-Rich SE (2) Manipulation on a Graph of Reachable Set." **S. Liu**, T. Zhao, B.P. Graesdal, P. Werner, J. Wang, J. Dolan, C. Liu, T. Pang.
Under review (IEEE T-RO).

[2] "Synthesis and Verification of Robust-Adaptive Safe Controllers." **S. Liu**, K. Yun, J. Dolan, and C. Liu. In 2024 European Controls Conference.

[3] "Certifying Robustness of Learning-Based Keypoint Detection and Pose Estimation Methods." X. Luo, T. Wei, **S. Liu**, Z. Wang, L. Mattei-Mendez, T. Loper, J. Neighbor, C. Hutchison, and C. Liu.

In 2025 ACM Transactions on Cyber-Physical Systems.

[4] "Safety Index Synthesis via Sum-of-Squares Programming." W. Zhao, T. He, T. Wei, **S. Liu**, and C. Liu. In 2023 American Controls Conference.

[5] "Learning the positions in countsketch." Y. Li, H. Lin, **S. Liu**, A. Vakilian, D.P. Woodruff. In 2023 International Conference in Learning Representations. **Oral presentation, top 25%.**

[6] "Safe Control Under Input Limits with Neural Control Barrier Functions." **S. Liu**, C. Liu, and J. Dolan. In 2022 Conference on Robot Learning.

[7] "On learned sketches for randomized numerical linear algebra." **S. Liu**, T. Liu, A. Vakilian, Y. Wan, DP Woodruff.
arXiv preprint arXiv:2007.09890_ (2020).

[8] "Learning to Adapt in Dynamic, Real-World Environments Through Meta-Reinforcement Learning." A. Nagabandi*, I. Clavera*, **S. Liu**, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn.
In 2018 International Conference on Learning Representations.

Honors and Awards

- Qualcomm Innovation Fellowship: 18 selected from 200 **2023**
- UC Berkeley Undergraduate Research Honors: 20 selected from 500 **2019**
- Computing Research Association GHC Research Scholars **2018**
- Microsoft Undergraduate GHC Scholarship **2017**
- Top 10% of TAs in UC Berkeley EECS Department **2017–2018**
- UC Berkeley College of Engineering Dean's List: top 10% highest GPAs **2016–19**
- Tau Beta Pi, Eta Kappa Nu, and Phi Beta Kappa: **2016**
  National engineering, computer science, liberal arts honor societies

Talks

- "Towards Higher-Quality Planning for Contact-Rich Manipulation", Robotics Seminar, Duke **2026**
- "Towards Higher-Quality Planning for Contact-Rich Manipulation", Manipulation Seminar, CMU **2026**
- "Reactive Safety for Self-Driving", Qualcomm Fellowship awardee talks **2024**
- "Synthesis of Control Safety Filters for Uncertain Systems", European Controls Conference (ECC) **2024**
- "Scalable Synthesis of Control Safety Filters", Conference on Robot Learning (CORL) **2021**
- "Scalable Synthesis of Control Safety Filters", Safe Autonomous Systems Lab, UCSD **2021**
- "Sample Inefficiency in Deep-Reinforcement Learning for Robotics", Women in Machine Learning Workshop, ICML **2020**

Teaching

- CMU 16-711 (Kinematics, Dynamics, and Control), taught by Hartmut Geyer **Fall 2022**
- CMU 16-811 (Math Fundamentals for Robotics), taught by Michael Erdmann **Fall 2021**
- UC Berkeley CS188 (Introduction to Artificial Intelligence), taught by Sergey Levine **Spring 2019**
- UC Berkeley CS170 (Efficient Algorithms and Intractable Problems), **Fall 2017-2018**
  Taught by Umesh Vazirani, Prasad Raghavendra, Sanjam Garge

Service

Mentoring

- Yogita Choudhary, CMU MS to CMU PhD
- Chase Dunaway, New Mexico Tech BS to MIT PhD
- Dvij Kalaria, CMU MS to UC Berkeley PhD
- Shivesh Khaitan, CMU MS to ML Engineer at Aurora
- Haoru Xue, CMU MS to UC Berkeley PhD
- Kai Yun, CMU MS to MIT PhD

Academic Reviewer

- Robotics: T-RO, R-AL, ICRA, IROS, RSS
- Controls: ACC
- Machine Learning: L4DC, ICLR