Yilang Liu

yilang.liu@yale.edu yilangliu.github.io linkedin.com/in/yilangliu Google Scholar

EDUCATION

Yale University — New Haven, CT
Ph.D. Candidate in Mechanical Engineering — *Expected Graduation May 2027*
M.S. & M.Phil in Mechanical Engineering (en route to Ph.D.)

Carnegie Mellon University — Pittsburgh, PA
M.S. in Mechanical Engineering – Research Program — *May 2022*

Chongqing University — Chongqing, China
B.E. in Mechanical Design, Manufacture and Automation — *April 2020*

University of Cincinnati — Cincinnati, OH
B.S. in Mechanical Engineering (Cum Laude) — *April 2020*

PUBLICATIONS

Haoxiang You, **Yilang Liu**, Ian Abraham. "Accelerating Visual-Policy Learning through Parallel Differentiable Simulation". Advances in Neural Information Processing Systems (**Spotlight** NeurIPS 2025).

Yilang Liu, Haoxiang You, Ian Abraham. "Sample-Based Hybrid Mode Control: Asymptotically Optimal Switching of Algorithmic and Non-Differentiable Control Modes". IEEE International Conference on Robotics and Automation (ICRA) 2026.

Christian Hughes, **Yilang Liu**, Julia Engdahl, Houston Warren, Darrick Lee, Fabio Ramos, Travis Miles, Ian Abraham. "Asymptotically Optimal Ergodic Coverage on Generalized Motion Fields". Robotics: Science and Systems (RSS) 2026.

Haoxiang You*, **Yilang Liu***, Ian Abraham. "Efficient On-Policy Visual RL via Stochastic Decoupled Policy Gradient". CVPR 2026 Workshop on Bridging Vision, Language, and Action. (*Equal contribution.)

Yilang Liu, Haoxiang You. "Hybrid Mode Sample-Based Control". RSS 2025 Workshop on Whole-Body Control and Bimanual Manipulation: Applications in Humanoids and Beyond, 2025.

Yilang Liu, Amir Barati Farimani. "An Energy-Saving Snake Locomotion Pattern Learned in a Physically Constrained Environment With Online Model-Based Policy Gradient Method". ASME. J. Mechanisms Robotics, 2023.

Yilang Liu, Y. Chen, H. Li. "Design an Augmentation Exoskeleton to Enhance Lifting Strength". ASME IMECE2020-24658, 2020.

RESEARCH EXPERIENCE

Efficient On-Policy Visual RL via Stochastic Decoupled Policy Gradient — New Haven, CT
Advisor: Prof. Ian Abraham — *01/2026–Present*

- Proposed **Stochastic Decoupled Policy Gradient (SDPG)**, a lightweight visual RL method that trains visuomotor control policies end-to-end in **a few hours on a single NVIDIA RTX 4080 GPU**.
- Estimated policy gradients via random perturbations of trajectory rollouts, eliminating full trajectory differentiation and requiring **orders of magnitude fewer** batch-rendered environments to reduce compute and memory overhead.
- Designed an adaptive exploration strategy and reward-invariant normalization for numerically stable updates, and developed a unified perspective connecting classical policy-gradient methods.
- Outperformed baselines (DrQv2, TD-MPC2, DreamerV3, PPO) on visual MuJoCo benchmarks in training time, memory usage, and reward.
- Released a suite of realistic egocentric visual RL benchmarks spanning dexterous manipulation and challenging locomotion, and demonstrated sim-to-real transfer on a Unitree Go2 robot.

Asymptotically Optimal Ergodic Coverage on Generalized Motion Fields — New Haven, CT
Advisor: Prof. Ian Abraham — *06/2025–01/2026*

- Developed a flow-adaptive ergodic trajectory optimization framework for robotic exploration in time-varying environments using Maximum Mean Discrepancy (MMD) metrics.
- Derived a theoretical formulation of ergodic coverage on evolving domains, enabling coverage guarantees under dynamic flow fields.
- Designed and implemented a sampling-based trajectory optimization algorithm capable of planning exploration paths in non-convex and flow-constrained environments.
- Demonstrated the approach on multiple robotic platforms, including aerial drones (crazyflies) and Unitree Go2 quadruped robots, validating coverage performance in both simulation and real-world experiments.
- Achieved up to **42% improvement** in exploration performance over information-theoretic baselines in dynamic flow

environments.

Accelerating Visual-Policy Learning through Parallel Differentiable Simulation New Haven, CT

Advisor: Prof. Ian Abraham *01/2025–06/2025*

- Proposed **D.Va (Decoupled Visual-Based Analytical Policy Gradient)**, a novel method for accelerating visual policies by **10x wall-clock time** using parallel differentiable simulation.
- Proposed to decouple visual observations from computation graphs, eliminating differentiating through rendering and achieving training a humanoid running policy in **4 hours** on a single GPU.
- Implemented visual policy learning method from differentiable rendering as baseline comparison using PyTorch3D and proving smoother first-order policy gradient without differentiating through render.
- Benchmarked model-free and model-based visual policy learning approaches and showed **5x higher rewards** in humanoid running tasks against existing algorithms.
- Provided an empirical analysis of analytical policy gradients and highlighted the new opportunities for integrating policy learning with trajectory optimization techniques.

Sample-Based Hybrid Mode Control New Haven, CT

Advisor: Prof. Ian Abraham *12/2024–04/2025*

- Designed a sample-based control strategy for high-dimensional long-horizon tasks using hybrid control theory, achieving complex mode switching behaviors.
- Analyzed theoretical correctness and proved performance guarantees on optimizing mode sequencing.
- Achieved **85% higher reward** in manipulation and **41% in humanoid locomotion** vs. baselines, enabling contact-rich in-hand manipulation.
- Implemented state estimation module for real-world Unitree Go2 deployment based on an extended Kalman filter using proprioception information.
- Deployed **real-time torque-level** agile locomotion of a Unitree Go2 quadruped capable of switching from footstand to handstand within sample-based predictive control framework.

PROFESSIONAL EXPERIENCE

Dexmate Inc. Santa Clara, CA

Robotics Intern *10/2025–01/2026*

- Deployed the Vision-Language-Action (VLA) model **GR00T N1.6** on Dexmate Vega, a dual-arm mobile manipulator, for an Amazon-designated project, executing a long-horizon (**2-minute**) bimanual task—grasping a PVC box by its handle with one gripper while expanding it with the other, then placing it aside—achieving a **100% success rate**.
- Implemented the VLA model $\pi_{0.5}$ for a client project, deploying it on the physical robot for a bimanual task placing a condenser onto a conveyor; smoothed action-chunk transitions with Real-Time Action Chunking (RTC) in PyTorch for stable execution at ˜20 Hz, achieving a **100% success rate** validated in factory testing.
- Deployed the Action Chunking with Transformers (ACT) imitation-learning algorithm on the physical robot using a bimanual teleoperated dataset, achieving reliable bottle grasping and placing from vision.
- Trained a reinforcement learning policy with Proximal Policy Optimization (PPO) in simulation (Genesis and Isaac Sim), guided by RRT motion planning, achieving a **100% success rate** on cube grasping with a 6-DoF dexterous hand.

TEACHING

Teaching Fellow – ENAS 773: Introduction to Robotics, Control, and Learning Yale University

- Led weekly recitations and office hours, advised student final projects, and assessed coursework on robotics, control, and learning.

Teaching Assistant – 18-461: Introduction to Machine Learning for Engineers Carnegie Mellon University

- Guided students through machine learning fundamentals via recitations and office hours, mentored course projects, and graded assignments and exams.

ACADEMIC SERVICE

Conference Reviewer: ICRA (2022–2026), IROS (2022–2026), RSS (2025), ECCV (2024).

Journal Reviewer: IEEE Robotics and Automation Letters (RA-L).

SKILLS

Programming & ML: Python, C++, PyTorch, PyTorch3D.

Robotics & Tools: MuJoCo, Isaac Gym, Isaac Sim, Genesis, Brax; ROS 1/2; Unitree A1, Unitree Go2; LeRobot, Weights & Biases.