Philip (Yizhou) Huang

Pittsburgh, PA | philiphuang@cmu.edu | philip-huang.com | linkedin.com/in/philip-yizhou-huang
github.com/philip-huang | scholar.google.com

Education

Carnegie Mellon University, Ph.D. in Robotics — Aug 2023 – Present
GPA: 4.12/4.33. Research: Multi-robot-arm coordination and planning (Advisor: Jiaoyang Li)

University of Toronto, M.Sc. in Computer Science — Sept 2021 – Aug 2023
GPA: 4.00/4.00. Research: Task and motion planning (Advisors: Florian Shkurti, Tim Barfoot)

University of Toronto, BASc. in Engineering Science — Sept 2016 – Jun 2021
GPA: 3.88/4.00 (Major: Machine Intelligence)

Selected Publications

- cuArms: GPU-Accelerated Conflict-Based Planning for Multi-Arm Manipulation, arXiv 2026 [pdf] [video]
- Prompt-to-Product: Generative Assembly via Bimanual Manipulation, IEEE RAM 2026 [pdf] [video]
- VAMP-MR: Vector-Accelerated Motion Planning and Execution for Multi-Robot-Arms, IEEE IROS 2026 [pdf]
- Benchmarking Shortcutting Techniques for Multi-Robot Arm Motion Planning, IEEE IROS 2025 [pdf][video]
- APEX-MR: Multi-Robot Asynchronous Planning and Execution for Cooperative Assembly, RSS 2025 [pdf][video]
- Field Testing of a Stochastic Planner for ASV Navigation Using Satellite Images, IEEE T-FR 2024 [pdf][video]
- Stochastic Planning for ASV Navigation Using Satellite Images, IEEE ICRA 2023 [pdf][video]
- Continual Model-Based Reinforcement Learning with Hypernetworks, IEEE ICRA 2021 [pdf][video]

Industry Experience

Research Scientist Intern, NVIDIA, Seattle, WA — May 2026 – Oct 2026
- Developed a GPU-accelerated multi-arm manipulation planner that generates plans for **8 arms across 40 tasks in under 1 s**, replacing days of RL training and improving precision from centimeter to **submillimeter** scale. Results were featured in NVIDIA's internal "Top-5 Things" email from the Chief Scientist and SVP of Research.
- Developed bimanual and multi-arm systems for assembly and manufacturing, including Blackwell GPU assembly.

Academic Experience

Robotics Researcher, AI for Robot Coordination at Scale Lab, CMU — Aug 2023 – Present
- Developed VAMP-MR, accelerating multi-arm motion planning by **two orders of magnitude** using CPU SIMD.
- Developed an asynchronous multi-robot planning framework that achieved **48% faster execution** compared to sequential planning and 36% compared to synchronous planning on average.
- Designed a multi-level reasoning framework for **the first dual-arm robotic system** for customized LEGO assembly using commercial bricks, integrating a physics engine, task and motion planning, and real-time control. Received **media coverage** in Modern Machine Shop.
- Developed a generalizable skill graph and ontology for autonomous robot assembly with the ARM Institute.

Robotics Researcher, Robot Learning & Vision Lab, University of Toronto — Jan 2020 – Aug 2023
- Conducted **km-scale, fully autonomous field tests** of ASVs equipped with GPS, vision, and sonar.
- Proposed a novel stochastic planning algorithm using satellite images, achieving a **15% reduction** in expected travel time across 1,000+ unique lake environments.
- Developed a **continual model-based reinforcement learning** approach with hypernetworks and demonstrated state-of-the-art performance across diverse simulated manipulation tasks.

Skills

Languages and Tools: Python, C++, HTML, Julia, Bash, PyTorch, MATLAB,Docker, Linux, Git, AI Agents
Specialties: Multi-Robot Systems, Task and Motion Planning, Manipulation, Field Robotics, Perception Systems