Reng Zheng

rengzheng@mit.edu — +1 (917) 558-6626 — U.S. Citizen
icarus.diy — linkedin.com/in/rengzheng — github.com/rengzhengcodes

Education

Massachusetts Institute of Technology — Cambridge, MA
Master of Engineering - Computer Systems — *Expected Jun 2027*

Massachusetts Institute of Technology — Cambridge, MA
Double Bachelor's of Science - Artificial Intelligence and Decision Making + Philosophy — *Expected Jun 2026*
Double Minors: Brain and Cognitive Science + Ancient Medieval Studies — *GPA: 4.9 / 5.0*

Courses: *Digital Systems Lab, Robotics Systems & Sciences, Graduate NLP, Computer Vision, Bayesian Inference, Geometric Computing*

Skills Summary

- **Languages**: Python, C++, Triton, SystemVerilog, Julia, RISC-V Asm, Java, JavaScript, Matlab, R, SQL
- **Frameworks**: Numba, PyTorch, Keras, ROS 2, PyBind11, FreeRTOS, Pandas, SciPy, Gurobi, Integer Set Library
- **Dev Tools**: Docker, Git, GitHub, SQLite, CMake, Shell Script, SSH, cocotb
- **EE Tools**: Logic Analyzer, Oscilloscope, KiCAD, Altium, LTSpice
- **Platforms**: Jetson Orin, Unix, Windows, Web, AWS, GCP, Modal, Arduino, STM32, Xilinx

Experience

MIT Motorsports — Cambridge, MA
Perception Lead — *Jun 2024 - Present*
Software Engineer — *Sep 2022 - Jun 2024*
- **Ground-Up Model Training**: Building a 10-hour dataset with 3x BlackFly cameras and an Ouster OS1-128 LiDAR. Training PointPainting on the dataset to produce a model targetting 86% of inferences having IoU=0.5 or better.
- **Built Perception Team**: Built and led a 4-person perception subteam within a 18-member autonomous-vehicle team.
- **Specced Sensor Requirements**: Deduced compute/sensor specs to do inference at 10Hz with 100ms latency.

MIT Energy-Efficient Multimedia Systems Group — Cambridge, MA
Undergraduate Researcher — *Oct 2022 - Present*
Supervisor: Michael Gilbert — *PIs: Joel Emer and Vivienne Sze*
- **Extended NVIDIA Timeloop**: Added 50+ Python bindings enabling a 40+ student-class to use in their capstone.
- **Researched Novel Network Model**: Modeled data distribution in distributed-buffer architectures in ISL.
- **Extending New Analytical Architecture Simulator**: Integrating distributed-buffer model into AccelForge.

MIT Hardware Architecture for Deep Learning Class — Cambridge, MA
Teaching Assistant — *Dec 2024 - Jun 2025; Dec 2025 - Present*
- **Teaching Duties**: Counseled 40+ students how DL implementations manifest in hardware, resulting in a 7/7 rating.
- **Assisted Final Projects**: Advised 24 project teams with their hardware-software co-optimization design capstone.
- **Revised Labs**: Edited 5 labs for the course to clarify wording, correct content errors, and suggest problem modifications.

MIT Digital Systems Laboratory Class — Cambridge, MA
Laboratory Assistant — *Sep 2025 - Dec 2025*
- **Debugged Student Issues**: Reviewed SystemVerilog 10 hours per week for code quality and errors for 150+ students.
- **Taught Test-Driven Development**: Coached generating testbenches and waveforms in cocotb for SystemVerilog.

Science of Intelligence Cluster — Berlin, Germany
Research Intern — *Jun 2024 - Aug 2024*
Supervisor: Yating Zheng — *PI: Pawel Romanczuk*
- **Algorithm Development**: Created 3 swarm shepherding methods that used convex hulls to estimate global details.
- **Simulated Results**: Verified simulation results across 1,000 runs, showing 98%+ task completion for all 3 methods.
- **Proved Robustness**: Demonstrated that multi-agent coordination is possible with limited local info for shepherding.

Projects

MIT Robotics Systems and Sciences Class Capstone — Cambridge, MA
Tech: Python, ROS 2, NVIDIA Jetson Orin Nano, ZED 2 Camera, Hokuyo LiDAR — *Mar 2024 - May 2024*
Used parallax to determine the racing line on an indoor track and a modified pure-pursuit algorithm to achieve robust motion planning. Used a finite state machine to navigate through a maze to detect and retrieve an objective with YOLOv5.

Midas: A One-Shot Greedy Computer Architecture Optimizer — Cambridge, MA
Tech: Python, Timeloop, Accelergy — *Apr 2024 - May 2024*
A greedy accelerator optimization algorithm with 60% greater cost reduction than coordinate descent. We achieve an additional 16% cost reduction by combining coordinate descent and Midas. Our method is only 10% worse on average than a full parameter-space grid search, using ∼30% of the runtime.

AudioNet: Generating Images from Audio with ControlNet — Cambridge, MA
Tech: Python, PyTorch, Sound2Scene dataset — *Mar 2024 - May 2024*
An attempt to generate images from audio inputs using stable diffusion and a ControlNet-inspired architecture. We produced 1,000+ images from audio on the Sound2Scene dataset and ∼5% of samples were reasonable inferences.