

Jarod Bloch

Pittsburgh, PA • 484-550-5079 • jarodb@andrew.cmu.edu • [linkedin.com/in/jarodbloch](https://www.linkedin.com/in/jarodbloch) • jarodblo.ch

Education

Carnegie Mellon University, School of Computer Science

May 2027

Bachelor of Science, Artificial Intelligence

Pittsburgh, PA

GPA: 3.3, Dean's List: Fall 2024 (High Honors), Fall 2025 (High Honors)

Awards: 3rd Place, CASI × Gray Swan AI Jailbreaking Hackathon (Fall 2025); 6th Place, top solo vs. teams of 2–3 (Spring 2026)

Member, **Plaid Parliament of Pwning (PPP)**, top-ranked competitive hacking (CTF) team; multi-time DEF CON champion

Coursework: ML for Cybersecurity (grad), Generative AI, Autonomous Agents, Natural Language Processing (grad)

Publications

Cheng, Y. F.*; **Bloch, J.***, et al. "Auditorily Embodied Conversational Agents: Effects of Spatialization and Situated Audio Cues on Presence and Social Perception." *Published in Proceedings of ACM CHI 2026*. <https://doi.org/10.1145/3772318.3791794>

*(Co-first author; co-presented oral talk at ACM CHI '26)

Experience

Connected Vehicle Software Engineering Intern | Python, LangGraph, RAG, BigQuery, Pandas May 2026 - August 2026
Ford Motor Company Dearborn, MI

- Building an agentic framework (Python, LangGraph) that orchestrates Ford's internal LLM platform (FordLLM) to predict severity classifications for vehicle diagnostic trouble codes (DTCs), helping the triage team prioritize EV issues.
- Designing the RAG pipeline (diagnostic documentation and live EV telemetry via BigQuery, Pandas) so the agent reasons from objective evidence rather than prior labels, establishing a consistent ground-truth for a previously subjective process.

Research Assistant | Python, Docker, vLLM, C#, Unity May 2025 - May 2026
Carnegie Mellon University Pittsburgh, PA

- Evaluated LLM robustness and recall under 100+ semantics-preserving JavaScript perturbations for Prof. Lujo Bauer (CyLab), identifying systematic failure modes in code understanding models.
- Boosted self-hosted LLM inference throughput 8.6× on a containerized vulnerability-benchmarking framework via quantization and GPU-utilization tuning, making large-scale adversarial evaluation tractable.
- Built a Unity-based experimental platform for the Augmented Perception Lab (Prof. David Lindlbauer) to study user perception of spatialized conversational AI agents in a 24-participant controlled study.

Teaching Assistant | C, Google Apps Script May 2025 - August 2025
Carnegie Mellon University Pittsburgh, PA

- Served as head of code reviews for 220+ primarily graduate-level students in Computer Systems (15-213/15-503), designing rubrics and coordinating sessions to evaluate C-based memory allocators, proxy servers, and shell programs.
- Built an automated, scalable TA management pipeline in Google Apps Script for student roster imports, assignment tracking, progress monitoring, and automatic statistics generation, adopted as infrastructure for future semesters.

Projects

Repo-Level Type Inference | Python, CodeQL, NetworkX, vLLM, Slurm February 2026 - May 2026

- Developed a language-agnostic neuro-symbolic type inference framework using Qwen3 models to annotate repository-scale codebases, achieving a 0.84 median type-match score (0-1) across 50 repositories while matching 96% of a model 13× larger.
- Reduced cross-file AI type errors by 37% by mapping dependency DAGs to inject explicit type signals into context.
- Guarded LLM output with a CodeQL layer that overrides hallucinated types via majority-voted points-to facts, omitting low-confidence guesses to prevent cross-file cascades.
- Optimized annotation quality via large-scale Slurm GPU ablation studies on 50 real-world repositories across Qwen3 model sizes (0.6B–30B), evaluating seeded root nodes, model types (Coder vs. Instruct), and context augmentation strategies.

Terrabot | Python, ROS, PyTorch, ONNX, OR-Tools August 2025 - December 2025

- Developed and deployed an autonomous greenhouse control system with a 4-person team on memory-constrained hardware (Raspberry Pi), integrating computer vision, SAT-based hardware fault diagnosis, and constraint-based adaptive scheduling.
- Architected a three-layer autonomous agent (behavioral FSMs, executive monitors, CP-SAT planner) that replanned daily schedules from live plant status by growth stage.
- Built a color-calibrated foliage segmentation pipeline (~97% pixel accuracy under variable lighting), selecting a fine-tuned UNet over a higher-accuracy ViT baseline that exceeded the Pi's memory limit to enable continuous on-device inference.

Skills

Machine Learning & AI: PyTorch, Hugging Face, vLLM, LoRA/PEFT, DPO, RL, LangGraph, RAG, Agentic Systems, OpenCV

Data & Tooling: scikit-learn, Pandas, NumPy, NetworkX, Weights & Biases (W&B), BigQuery

Programming Languages: Python, C, C++, Java, JavaScript, C#, SML, SQL

Infrastructure & MLOps: Docker, Slurm/HPC, AWS, Oracle Cloud, ONNX, FastAPI, CI/CD, Linux/Bash, Git