

## EDUCATION

---

**Georgia Institute of Technology, Atlanta, GA:** *B.S. in Computer Science* Aug. 2023 – May 2026

**GPA:** 4.0/4.0

**Courses:** DL, ML, Systems for ML, Intro to AI, Perception and Robotics, Algorithms, OS, Undergraduate Research

## EXPERIENCE

---

**Georgia Tech, Zhang Lab:** *Undergraduate Researcher, Artificial Intelligence* Aug. 2024 – Present

- Built PRM-guided tree-search reasoning systems with high-throughput inference using SGLang and vLLM caching to generate synthetic trajectories with error correction and exploration.
- Trained multi-turn LLM agents with GRPO on SWE-bench, SWE-Smith, and SWE-Gym, adding reward-hacking defenses, tool-use checks, and rollout instrumentation.
- Generalized the system to Docker/MCP tool environments and scaled asynchronous PPO/GRPO/RLVR training using AReaL on Ray with a Kubernetes-managed sandbox server.

**DuckAI:** *AI Researcher & Engineer* Aug. 2023 – Present

- Engineered **DuckTrack**, a cross-platform computer-use capture tool that records screen, keyboard, and mouse with sub-frame alignment via OS input APIs and OBS synchronization.
- Enabled high-quality computer-use datasets that were later adopted in OSWorld and OpenCUA for benchmarking and foundation model training.
- Designed and implemented **PRMBench**'s full-stack annotation platform (React, FastAPI, SQLite), prototyping UI workflows with annotators to improve labeling speed and inter-annotator agreement for process reward model datasets.
- Built large-scale TTS data pipelines (YouTube and podcast ingestion, upscaling, VAD, diarization) and contributed optimizations to insanely-fast-whisper to accelerate transcription for speech datasets.

## SELECTED PROJECTS

---

**Tool-Integrated RL Agent Training System** May 2025 – Present

- Engineered a general tool-using agent framework with Docker sandboxes and an MCP interface (browser, terminal, code execution), supporting multi-turn rollouts across math, search, SWE-bench, and puzzle tasks.
- Integrated AReaL for asynchronous PPO/GRPO/RLVR training to maximize throughput under high tool latency, coordinating rollout workers with sandboxes hosted on a Kubernetes cluster.

**DuckTrack** Nov. 2023 – Jan. 2024

- Extended and maintained precise, cross-platform capture of screen, keyboard, and mouse events by benchmarking native input APIs and synchronizing with OBS to achieve millisecond-level, sub-frame alignment for computer-use datasets.

**PRM Annotation Platform** Nov. 2024 – Jun. 2025

- Developed a web app to label process rewards on multi-step CoT traces with side-by-side comparisons and QC flows, enabling open-source PRM dataset generation across math, logic puzzle, and science domains.

**RecipeGPT (AIATL hackathon)** Nov. 2023

- Led a 4-person team at the 36-hour AIATL hackathon to design and ship RecipeGPT by decomposing work, scoping features, and reviewing code across the stack.
- Implemented a GPT-4V and RAG pipeline over scraped recipes (FastAPI backend, React front end, Pinecone retrieval) with nutrition-aware filtering and dietary constraints.

**Photo Archive with ImageBind** Dec. 2023 – Jan. 2024

- Implemented a self-hosted search system over 200,000+ photos using multimodal embeddings from ImageBind.
- Tuned high-precision face clustering and daily iCloud sync to approximate Google Photos-style identity-based search in a private deployment.

**AlphaEvolve Extensions** May 2025 – Sept. 2025

- Replicated circle-packing results and extended OpenEvolve with new combinatorics and compiler tasks, upstreaming improvements and exploring scientific application domains.

## SKILLS

---

**Languages:** Python, Rust, C, Java, Go, JavaScript, Bash

**ML:** PyTorch, vLLM, SGLang, Ray, AReaL, veRL, HuggingFace

**Systems:** Linux, Docker, Kubernetes, SLURM, AWS, FastAPI, SkyPilot, uv, Git