Aarush Agarwal

aarusha@andrew.cmu.edu | (408)-594-8710 | LinkedIn | GitHub: AgarwalAarush | Website: aarushagarwal.dev

EDUCATION

Carnegie Mellon University *Expected Graduation, May 2027*

Bachelor of Science in AI & CS, School of Computer Science

- QPA: 3.65, CMU School of Computer Science Dean's List

EXPERIENCES

Shopify | *Machine Learning Engineer Intern* *May 2025 - Aug 2025*

- Built a buyer-fraud detection system with VertexAI, optimized BigQuery/Dataflow pipelines to increase predictive accuracy, and applied targeted feature selection to reduce training iteration time by 70%.
- Authored a patent for a novel AI framework where distributed, specialized agents use Neo4j graph traversal to collaboratively decompose and execute tasks, enhancing output quality.

Carnegie Mellon Physics Lab | *CUDA Research Assistant* *Aug 2024 - Oct 2025*

- Co-authored paper focused on a binning-based parallelized K-Nearest Neighbors algorithm, achieving up to a 200x speedup over FAISS (Facebook AI Similarity), Annoy (Spotify), and SCANN (Google) in low-dimensional spaces.
- Transitioned Python autograd and gradient functions to C++ and CUDA extension implementations and integrated with PyTorch JIT serialization, achieving a 10% decrease in KNN runtime.

PROJECTS

Yumi | *Backend Lead* *Oct 2025*

- **Awarded 3rd Place (3/150+ teams) at HackHarvard 2025** for an agentic social network that learns food profiles to eliminate dining friction, whether eating alone or coordinating with friends.
- Built a FastAPI backend with taste profile extraction and intelligence learning from user interactions, natural language reviews, sentiment analysis, multi-user preference merging, spatial restaurant search using PostGIS, and Twilio voice integration for automated reservations.

Medicly | *ML Lead* *Sep 2025*

- **Awarded Grand Prize (1/250+ teams) at HackCMU 2025** for a platform that slashes physical therapy evaluation time from 14 days to 10 minutes. Sponsors include Anthropic, Citadel, Stripe, Jane Street, & HRT.
- Engineered a multimodal pipeline with pose estimation, structured kinematics, joint-angle prediction, 3D mesh rendering, and LLM-powered recovery reports.

Nova | *Personal Exploration* *Jan 2025 - Aug 2024*

- Architected a multi-modal agentic AI voice assistant, integrating Whisper ASR, LangChain orchestration, dual TTS engines, wake-word detection, and intent classification to enable real-time system automation and conversational voice interaction.
- Built a macOS application in SwiftUI, featuring real-time voice visualization, seamless text/voice mode switching, CoreData-backed persistence, Keychain API key security, and sub-100ms latency TTS systems.

MV Test Tracker | *Backend Lead* *Feb 2023 – Jan 2024*

- Designed and built a robust backend using Firebase and Firestore, implementing real-time data synchronization, secure user authentication, and a custom scheduling optimization algorithm that reduced server call overhead by 40% and ensured even test distribution across multiple classes.
- Gained overwhelming support from district administration and 50+ teachers. Impacted 1200+ students at high school.

SKILLS

Languages C++, Python, SQL, C, Java, JavaScript, Swift, Lua | **Familiar:** Ruby, R, Bash, SML

Software Tensorflow, PyTorch, CUDA, Vue.js, React, Firebase/Firestore, Git, DBT, GCP, Neo4j, PostgreSQL, Supabase, OpenCV, Pandas, AWS

AWARDS

USACO Gold