

Alexander Kranias

U.S. Citizen | 727-457-4433 | alexander.kranias@gatech.edu
alexkranias.com | linkedin.com/in/alexanderkranias | github.com/alexkranias

Education

Georgia Institute of Technology | Atlanta, GA

Bachelor of Science in Computer Science, GPA 4.00

August 2022 – Present
Expected Graduation, May 2026

Concentrations: Intelligence (AI/ML), Systems and Architecture

Organizations: GT Create-X Incubator Program, GT Start Up Exchange, Georgia Tech Varsity Rowing, GT Student Tour Guides

Coursework: Data Structures, Algorithms, Computer Architecture, Systems and Networks, Intro to AI, Robotics, Linear Algebra

Skills

Languages: Python, C, C++, CUDA, Java, React, HTML/CSS, JavaScript, Node.js

Tools/Frameworks: Docker, Kubernetes, CMake, Jupyter Notebook, Azure, Flask, PostgreSQL, SQL, ROS2, git, MATLAB, Simulink

Experience

Embedded Software Engineer Intern

Vermeer Corporation

May 2024– August 2024

Pella, Iowa

- Built C++ simulation of autonomous pile driver occupancy grids for Azure IoT cloud environment; integrated with Simulink/ROS2.
- Architected MQTT/DDS hybrid network with autonomous P2P MANETs to scale up to 50+ machines in remote 1000-acre sites.
- Developed a distributed VIN-data timestamp table that minimizes MQTT data transfer between Azure cloud and P2P networks.
- Created protobuf serialization benchmarks, nanogui UI, and custom quadtree occupancy grids (reduced protobuf size by 96%).

Software Engineer Intern

Raymond James Financial

May 2023 – July 2023

St. Petersburg, Florida

- Created a full-stack LLM copilot web app for financial advisors and investment bankers that deploys fine-tuned GPT-3 models.
- Developed HTML/CSS frontend and our backend with Azure OpenAI API, Outlook API, and internal CRM software.
- Designed project requirements; interviewed 20+ stakeholders; onboarded team to continue project; created roadmap with COO.

Research

Computer Architecture Research Assistant

Georgia Tech HPArch Lab (High Performance Architecture Lab)

August 2023 – Present

Atlanta, Georgia

- Developed materials for GPU Hardware/Software course; designed Tiled Matrix Multiplication project in CUDA; ran on cluster.
- Formulated cache replacement algorithm for Row Hammer mitigation that de-prioritizes cache blocks originating from hot rows.
- Built cache block memory access trackers in a remote Linux environment integrated into open-source ChampSim repo (C++).
- Benchmarked cache block accesses and DRAM hot-row utilization for 40+ traces; resulted in new IPCs of 80%-200% of initial.
- Developed memory paging and multithreaded round-robin process scheduling projects in a Linux docker environment.

Machine Learning Research Assistant

DuckAI Research Group

October 2022 – October 2023

Atlanta, Georgia

- Built DuckTrack: Accurate Computer Activity Tracking (tool to build multimodal computer agent datasets, 100+ downloads).
- Created ARB: Advanced Reasoning Benchmark for Large Language Models (accepted to MATH-AI Workshop at NeurIPS '23).

Extracurriculars

Co-Founder, SidelineAI (RippleX Finalist) | PyTorch, Flask, LangChain, React, Pinecone, PostgreSQL

July 2023 – Present

- AI search for sport recruiting; demoed to 17 colleges: 7 D1, 3 D2, 5 D3, and 2 NAIA; met with NBA VP of Strategy.
- Developed multimodal vector search of gameplay using LangChain, Pinecone VectorDB, PostgreSQL, and Flask CRUD requests.
- Designed vector embeddings to enable player and video embeddings search using action-timestamp and biography annotations.
- Created a Two-Stream CNN-TSM (Temporal Shift Module) architecture for real-time foul recognition in game footage.
- Built custom foul-labeling software to efficiently annotate over 44,000 televised NBA game clips (97 hours of footage).

Projects and Awards

3dReal (Stanford TreeHacks 2024 Winner) | Swift, CUDA, Python, Firebase, instant-ngp

February 2024 - Present

- Developed in 48 hours an iOS app that captures social NeRFs using NVIDIA instant-ngp, Firebase, Swift, and CUDA.
- Synchronously generates Neural Radiance Fields (NeRFs) using the cameras of multiple iOS devices, creating a shared 3D selfie.

Bite (HackMIT 2023 Winner - \$2K prize) | Python, Flask, Node.js, React, LogMeal API, IMUs

September 2023

- Built in 24 hours an intelligent wearable watch that uses computer vision to track your daily food and nutritional intake.