

Aneesh Maganti

aneesh.maganti@nyu.edu • github.com/aminoa • linkedin.com/in/aneesh-maganti

EDUCATION

New York University, Tandon School of Engineering, Brooklyn, NY Fall 2024

Bachelor of Science, Computer Science

GPA: **3.80**

Relevant Courses: Algorithmic Machine Learning, Offensive Security, ML Visualization, Computer Architecture

SKILLS

Languages C++, Python, Javascript, C#, Java, Bash

Technologies PyTorch, Next.js, SDL, QT, React, Node, PostgreSQL, Sklearn, Docker, Linux

EXPERIENCE

Qualcomm, San Diego, CA, *Software Engineer Intern, AI Software/ML Group* May 2024 - Aug 2024

- Implemented various memory optimizations across the Qualcomm AI Runtime (QAIRT), reducing peak memory usage by 20% across benchmarked ResNet50 and LLAMA2 and LLAMA3 LLM models
- Optimized AIMET quantization pipeline via C++ PyBind API to reference static graph tensors within Python subprocess, reducing in-memory allocations by 50%
- Refactored QAIRT/AIMET toolchain via PyTorch 'meta' device to prevent unnecessary weight initialization and modified model export to use in-memory IR graph, reducing disk space and improving runtime

New York University, Brooklyn, NY, *Teaching Assistant (Machine Learning)* Sep 2023 - Dec 2023

- Instructed students weekly for machine learning topics of written and programming tasks through office hours
- Graded weekly assignments on the basis of proper algorithm implementation, code correctness and style

NYU Algorithms and Foundations Group, Brooklyn, NY, *ML Academic Researcher* Feb 2023 - Sep 2023

- Designed and implemented *Deltagonalshift*, a diagonal estimator for a dynamic matrix under Professor Christopher Musco to improve neural network optimization based on DeltaShift trace estimation algorithm
- Demonstrated Deltagonalshift was more effective than repeatedly running Hutchinson's diagonal estimator

MarketFusion, Los Altos, CA, *Software Engineering Intern* July 2022 - Sep 2022

- Developed client-side React.js web application registration and shopping pages for online food delivery service
- Facilitated user account creation by sending server requests to internal MySQL database via the Axios library

Corelink, Brooklyn, NY, *Software Engineer Intern* Sep 2021 - May 2022

- Implemented a C++ UDP network packet splitter to enable researchers to bypass Corelink's MTU limit from 20,000 to 64,000 bytes, increasing maximum throughput by 220%
- Designed Next.js/React interview scheduling platform using Auth0 for authentication and MongoDB backend

PROJECTS/ACTIVITIES

BUGS Open Source Club President Sep 2022 - Dec 2023

- Started and led 50+ member club by coordinating biweekly workshop and project coding events to discuss software engineering skills, foster contributions to open source, and create a fun, inclusive CS community.
- Led multiple workshops including discussion of open-source licenses, JavaScript Playwright automation, and an overview of emulation and the internals of C++ Game Boy emulator.
- Developed Next.js-React website NYU Syllabi with Netlify hosting and Docusaurus-based NYU CS Wiki websites

Dot Matrix - Game Boy Emulator August 2023

- Designed x86 C++ emulator for the Game Boy platform by implementing 255 standard + 240 cb instructions
- Simulated hardware features including registers, graphics (SDL), memory, timers, interrupts, and input handling

SentiTweet March 2023 - April 2023

- Created sentimental tweet generator via modified PyTorch PPLM library with GPT-2 to simulate conversations between Twitter users using natural language generation
- Employed D3.js visualization library to create graph of tweets, their sentiments, and relationships