Vishruth Anand

vanand@gatech.edu | (480) 548-8705 | linkedin.com/in/vishruth-anand | vishruthanand.com

EDUCATION

Georgia Institute of Technology GPA: *4.0*

BS/MS in Computer Science (Fintech + Machine Learning), MSQCF Candidate *Expected Graduation: December 2027*

Coursework: Design/Analysis of Algorithms, Computer Vision, Database Systems, Knowledge-Based AI, Probability/Statistics, Computer Systems/Networks, Deep Learning, Data Structures/Algorithms, Computational Modeling & Algorithms, ML in Computational Biology

WORK EXPERIENCE

Software Engineering Intern *May 2026 - August 2026*

Databricks

- Re-architected Databricks' production lineage pipeline from a table-specific model into a generic typed-node graph, adding Unity Catalog Volume support across runtime events, Kafka ingestion, and dual control-plane stores with a zero-regression rollout
- Delivered the #1 requested lineage feature, addressing a pain point raised by 27 of 30 customers and reducing common Volume-lineage analysis from 25 lines of SQL to ~5

Software Development Engineer Intern *May 2025 - August 2025*

Amazon Web Services

- Re-architected Amazon DCV's web video pipeline by eliminating server-side encode/decode hops and rendering webcam streams locally via WebRTC and canvas overlays, achieving 120 FPS and 0.15 ms render latency (266× improvement)
- Built a fault-tolerant, real-time coordination layer across a Chrome extension, proxy browser, and remote client components, using service workers and native messaging APIs to maintain 99.9% synchronization reliability across dynamic edge cases including tab detachment, DPI changes, and multi-display movement

Undergraduate Researcher - Software & Machine Learning *August 2023 - Dec 2024*

Robotic Musicianship Lab - Georgia Tech Vertically Integrated Project

- Designed an AI-powered MuseScore plugin using QML, enabling real-time melody generation, sequence prediction, and style transfer across 15+ musical genres, integrating transformer-based models directly into composers' workflows
- Constructed a large-scale symbolic music data pipeline, scraping and sanitizing 650,000+ MIDI files using Spotify and OpenAI APIs, Python, and SQL, and preparing model-specific datasets for architectures such as Coconet and SymphonyNet
- Deployed a Flask-based inference API connecting the plugin to transformer-based music generation models, supporting sub-second inference latency for 50+ professional musicians

Software Engineer Intern *May 2024 - August 2024*

PyPs.In

- Shipped a production Chrome extension for intelligent LinkedIn engagement, combining DOM traversal and asynchronous extension messaging with LLM-driven generation to double reply rates and reduce manual outreach time by ~80%
- Implemented a Node.js + Express backend integrating the Gemini API and a ranking heuristic that scored posts by engagement signals, follower metrics, keyword relevance, and content context to prioritize high-value interactions
- Integrated a local small language model using Transformer.js to pre-filter reply-worthy posts and generate lightweight responses on-device, reducing Gemini API usage by ~30% without adding perceptible latency

Neuroimaging Research Intern *June 2022 - August 2022*

University of Arizona Memory Development and Disorders Lab

- Analyzed 350+ MRI scans using ITK-SNAP and R to identify hippocampal subfield differences between neurotypical and Down syndrome youth; co-authored a forthcoming publication and presented findings to 1000+ attendees at the AZBio Awards

OTHER

Quantitative Research Analyst @ Georgia Tech Investments Committee (Quant Sector)

- Built a walk-forward ETF allocation framework using Hidden Markov Models, PCA-based feature reduction, and regime-conditioned mean-variance optimization, evaluating 84,000+ configurations across 14 years of out-of-sample data
- Achieved 0.982 OOS Sharpe, a +0.387 lift over equal weight, and stronger downside protection; pitched the strategy to Georgia Tech's Investments Decision Group and secured funding for implementation

WaitWhat (Hackathon Winner @ NexHacks 2026) | waitwhat.tech

- Built a low-latency lecture assistant for live transcription, LLM Q&A, quizzes, and instructor analytics, designing transcript chunking and context compression that cut prompt usage ~70% while preserving real-time responsiveness

SKILLS

Programming Languages: Python, Javascript/Typescript, Scala, Java, C++, Bash/Shell, C, SQL, Rust, Assembly, R

Frameworks & Tools: AWS, React.js, Node.js/Express.js, PyTorch/TensorFlow, Git, Linux, Docker, Android Studio, ITK-Snap, GDB