Aneesh Dantuluri

669-251-7530 | adantuluri9@gatech.edu | linkedin.com/in/aneesh-dantuluri | San Jose, CA

Education

Georgia Institute of Technology Atlanta, GA

B.S. in Computer Science, Threads: AI/ML & Info Systems (Databases) — GPA: 4.00/4.00 *Expected May 2028*

- **Coursework:** Data Structures & Algorithms, Linear Algebra, Objects & Design, Intro to AI, Multivariable Calc
- **Honors:** Y Combinator Startup School 26 ($25k grant); 2nd @ Y Combinator Hackathon; Top 4.6% debater (HS)

Technical Skills

Languages: Java, Python, SQL, JavaScript, TypeScript, HTML/CSS,
Frameworks: React, Next.js, Flask, Vue.js, Node.js, PySpark, Scikit-learn, Tailwind CSS
Developer Tools: GitHub, Docker, AWS, GCP, Azure, Databricks, Postman, Supabase, Firebase, PostgreSQL, MCP
Libraries: PyTorch, TensorFlow, pandas, NumPy, Scikit-learn, spaCy, OpenCV, BeautifulSoup, ApacheSpark

Experience

Stanford AI Lab (SAIL) September 2026 – Present

Undergraduate Researcher *Palo Alto, CA*

- Optimizing agent workflows w/ GEPA evolutionary prompt improvement + reinforcement learning w/ Dilara Soylu
- Designed multi-objective feedback pipeline to recursively self-imrpove agent reasoning traces (token usage + latency)
- Scaled evolutionary search across Pareto frontiers to benchmark agent performance on complex reasoning tasks
- Achieved up to 35x higher sample efficiency than standard reinforcement learning

Stanford University (AI Measurement Science Lab) May 2026 – August 2026

Summer Research Intern *Palo Alto, CA*

- Built PyTorch-native demand accessor MLP model to analyze subject latent capabilities under Sang Truong
- Reached 81% accuracy on MMLU performance benchmark and 0.641 AUROC value
- Reduced AI benchmark size 2x (50 vs. 100 items) with no loss in reliability using a data-driven item-selection method, validated on held-out models to rule out overfitting

Capplica June 2023 – April 2026

Co-Founder *San Jose, CA*

- Led development of full-stack Vue.js & Firebase web app, scaling to 150K+ users and 300K peak monthly visits
- Automated data extraction w/ Selenium, spaCy NER, LLM RAG pipeline, compiling a database of 2,700+
- Introduced a database caching layer to reduce Firebase reads, cutting backend costs by 45%

DataNimbus June 2024 – August 2024

Software Engineer Intern *Plano, TX*

- Developed fault-tolerant ETL pipeline to ingest/analyze client ad campaign data using Python & Apache Spark on Databricks SDK; generated metric visualizations with campaign performance insights
- Engineered data ingestion/aggregation for Bing Ads SOAP API; designed custom logic (UDFs) to parse, flatten, and structure nested XML data; implemented asynchronous I/O to manage network latency
- Optimized Spark job performance through distributed caching, partitioning, and lazy evaluation

Projects

Pam: Robotic AI Caregiver (HackMIT) | *Python, PyTorch, YOLOE, OpenCV, Elasticsearch* September 2026

- Built persistent memory layer for Dementia patients: multimodal YOLO detection + Claude VLM scene processing
- Developed tool-based Deepgram voice agent w/ Google Workspace integrations, robot arm control, & face detection
- Trained action-chunking policy from scratch on self-curated demos, deployed to 6-DoF robotic arm for physical aid

Cortex: Cursor for Robot Training (2nd @ YC Hackathon) | *Python, Next.js, InsForge, IsaacSim* October 2026

- Built agent-native robotics MCP & Platform enabling any agent to train + control physical robots
- Developed end-to-end pipeline using simulated LLM demonstrations to train Vision-Language-Action policies
- Engineered agnostic agent harness w/ recursive training improvement, telemetry, adversarial training, and RL envs.
- Used InsForge cloud infrastructure to host & orchestrate vision models, training data, policies, and workflows

BrainSpace: AI Cloud Storage Tool | *Next.js, TypeScript, GCP, AWS, PostgreSQL* February 2026 – Present

- Built app powered by agentic graph-based RAG pipeline w/ interactive querying using LangChain & OpenAI
- Implemented ingestion, vectorization, semantic search with custom REST APIs using GCP & Supabase Postgres
- Optimized response performance using autonomous, iterative, self-refining data retrieval and state caching
- Developed AWS S3/Lambda infrastrucutre for file storage and serverless file processing

Leadership & Community Engagement

Speech and Debate Club August 2024 – May 2026

President *San Jose, CA*

- Raised $1,200+ for club funds; mentored 35+ students; led collaborations with 2 schools; managed weekly meetings