Aditya Chavan

Arlington, VA | +1 929-343-4232 | https://adityarchavan.com/ | ac8586@nyu.edu | LinkedIn | GitHub | CV | Scholar | Medium

EDUCATION

New York University, New York, USA
Master of Science, Computer Science ***Awarded a merit scholarship for the graduate program***
Coursework: Algorithms, Cloud Computing, Big Data & ML, Software Engineering, FI Algo Trading, MS Thesis: Blockchain Scaling

University of Mumbai, Mumbai, India
Bachelor of Engineering, Computer Science Engineering ***Awarded a merit scholarship for the undergraduate program***
Coursework: Data Structures, Algorithms, Artificial Intelligence, Web Development, Machine Learning & NLP, Database Systems

SKILLS

Programming: Python, Java, JavaScript, R, TypeScript, C, Kotlin, Node.js, Django, Spring, Flask, Express, HTML, CSS, Angular
Software and Methodologies: Hadoop, Dask, PySpark, Databricks, Tableau, Docker, Git, GraphQL, Agile, Scrum, Architecture Design
Machine Learning & Generative AI: scikit-learn, PyTorch, SciPy, Hugging Face, LangChain, LangGraph, Stable Diffusion, Llama
Other: MySQL, PostgreSQL, BigQuery, Redis, Memcached, Kafka, ActiveMQ, AWS: Kinesis, DynamoDB, Lex, SQS, Rekognition, Elasticsearch, Lambda, SNS, Splunk, OpenObserve, Grafana, Prometheus, Jenkins, Kubernetes, Dynatrace, JUnit, Cucumber

WORK EXPERIENCE

Software Engineer III, **Walmart Global Tech** *(Received the Bravo Award for Engineering Excellence)* February 2024 – Present
- Develop core services for Sidekick, Walmart's AI-powered workforce-management platform serving **1.6 million associates**, enabling task generation, assignment, and prioritization based on logistics data, historical activity, and business priorities.
- Designed and deployed **production-scale multi-agent systems** using LangGraph, A2A, MCP, and A2UI, enabling agent collaboration and tool-augmented workflows; built MCP servers exposing workforce task-generation, assignment, and prioritization capabilities to AI agents.
- Developed a **Semantic Data Access Platform** using MCP and LangGraph, enabling business teams to query BigQuery, Azure SQL, PostgreSQL, Cassandra, and graph databases through natural language.
- Designed Kafka-based ingestion pipelines and asynchronous processing services handling **150 million daily events** for Tier 0 services across **10,000+ stores**.

Research Software Engineer, **New York University Center for Neural Science** August 2023 – February 2024
- Developed **highly optimized** Python pipelines for functional neuroimaging (fMRI) analysis on NYU's high-performance supercomputing cluster, using a general linear model for trend fitting and compressive spatial summation.
- Used specialized neuroimaging libraries to process and map population receptive fields from fMRI data, revealing visual-stimulus processing; **optimized code** for a **5-fold increase in execution speed**.

Graduate Teaching Adjunct, **New York University** August 2021 – May 2023
- **Taught mathematical finance, calculus, and optimization** at NYU Courant through recitations, office hours, grading, and proctoring. Led data-driven decision-making recitations (Tableau, JMP, SQL, R) for 120+ NYU Stern MBA students.

Software Development Engineer, **Aleria Tech** *(Part of the AWS CTO Fellowship through the Techstars NYC accelerator)* Jun–Dec 2022
- Engineered a SaaS tool providing data-driven insights into partner organizations' diversity, equity, and inclusion initiatives.
- Implemented Twilio integration to send and receive messages through WhatsApp and Messages.
- Migrated the stack to AWS and containerized the Django app with Docker, using Lambda, S3, RDS, ECS, and Elastic Beanstalk to improve scalability, resilience, and speed and **increase uptime to 99%**.

Data Scientist (Graduate Researcher), **New York University Stern School of Business** June 2022 – September 2022
- Developed APIs with Python, Google Cloud Platform (Cloud Functions, Cloud Run), and AWS (Lambda, CloudWatch).
- Implemented generative AI solutions with APIs from Hugging Face, DeepFace, DALL·E, Truss, and Google Image Search.
- Led service migration to AWS, improving system compatibility and significantly increasing the user base.

Software Engineer, **Tata Consultancy Services** *(Received two spot awards for exceptional performance)* August 2020 – July 2021
- Worked on a communications platform handling **15 million daily requests** across SMS, email, push notifications, and IVR using Python, Java, Spring Boot, PostgreSQL, and Redis. Led Microsoft Teams integration as a new channel.
- Enhanced user privacy with a notification management dashboard used by **500K+ users**.
- Improved system reliability through Azure cloud migration and CI/CD pipeline implementation.

Software Development Engineer, **Sathvachan** December 2018 – July 2020
- Developed a full-stack wedding planning platform (Django, Python, MySQL) on DigitalOcean with advanced hotel booking, a vendor marketplace (photographers, caterers, musicians, decorators, florists, and more), and Google Ads integration.

PROJECTS

Walmart Global Tech Hackathon Platform *(Recognized as AI Champion by an EVP for distinguished contributions)* April–July 2026
- Built a Python-based registration and judging platform for Walmart's US Tech Hackathon 2026, used by 25,000+ Walmart Global Tech employees. Exposed 42 MCP tools and 37 REST endpoints via a Model Context Protocol server (FastMCP, FastAPI, Azure Cosmos DB) for team registration, project submissions, mentor/judge management, and skill-based team matching.
- Designed a 3-stage automated judging pipeline with three parallel LLM sub-agents (Opus, Sonnet, Haiku), consensus-based scoring on a 5-criterion rubric, and gates to advance, flag, or disqualify submissions before human judge rounds.
- Used LLM-based evaluations to assess complete submission packets, including code repositories, demo videos, live demos, and other artifacts, against predetermined rubrics aligned with business KPIs.

Generative AI-Powered Conversation Co-pilot *(Best LLM Solution award at Cornell LLaMahack)* December 2023 – February 2024
- Engineered a real-time video-conferencing co-pilot with LLMs (LLaMA2, GPT-4, PaLM2), Hugging Face, LangChain, and Gradio for context-aware interactions, plus OpenAI's Whisper for precise speech-to-text conversion.