

NITISH CHOWDARY

 +1 (408) 581-5324  nitsancs@gmail.com  linkedin-nitish-chowdary  github-nitishsjsucs  portfolio-website

Education

San Jose State University	Aug 2025 – May 2027
Masters of Science – Software Engineering	GPA - 3.9
VIT University	2021 – 2025
Bachelor of Technology – Computer Science	CGPA - 3.84/4
Relevant Coursework	
Machine Learning, Gen AI, Reinforcement Learning, Recommender Systems, Deep Learning, Data Mining, Enterprise Software	

Work Experience

DoorDash – Machine Learning Intern, Ads Economics	May 2026 – Aug 2026
- Shipped Ads Econ Health, a production health-monitoring system for the ads marketplace that computes 90 metrics from 28 queries across Snowflake, Databricks, and Prometheus, writes them idempotently to a daily snapshot table, and delivers a digest to Slack - Designed a fail-closed LLM root-cause agent that collapses an alert storm into a single incident, using change-point onset detection and a Beta-Binomial posterior over root causes to localize failures across 6 upstream subsystems - Cut root-cause time on real alerts from ~1 hour to minutes by scoring evidence across 20 diagnostic checks (change correlation, blast radius, drift) and emitting only typed claims an executable predicate could verify against a warehouse row	
Broadridge Financial Solutions – Software Engineer	March 2025 – July 2025
- Optimized SQL (window functions, CTEs, partitioned aggregations) and Python processing that joined 8 normalized tables into model-ready financial datasets, cutting query runtime 40% - Raised automated test coverage 35% and resolved concurrency defects in a multi-threaded test module	
Harman Connected Services – Software Engineer Intern	Feb 2025 – April 2025
- Engineered a streaming data ingestion and analytics system for TEPCO’s IoT fleet, processing 10 million data points daily from 200,000+ field devices through Python parsers that aligned timestamps, imputed nulls, and normalized units across heterogeneous device telemetry - Detected device degradation and process drift with rolling-window statistics (7-day moving average, z-score thresholds) surfaced in Plotly dashboards tracking fleet-wide KPIs, cutting fault response time 30%	
Daira Edtech – Co-Founder & Technical Lead	Feb 2024 – Dec 2024
- Architected the LLM stack (RAG retrieval and agentic workflows) in Python and LangChain behind two learning platforms piloted at 15+ institutions and 4,200+ students, one of which (Jiveesha) won a \$100K grant - Owned the stack end to end, from ingestion pipelines (deduplication, schema validation, encoding normalization) feeding downstream ML models through to release, and set the engineering team’s code review and sprint standards	
Kona AI – Data Science Intern	Oct – Dec 2023 & Jun – Jul 2024
- Built distributed Dask pipelines processing millions of transaction records for AutoML fraud detection, with programmatic feature engineering (categorical encoding, scaling, interaction terms), model selection, and evaluation, cutting model iteration time 60% - Developed a natural-language-to-SQL assistant that grounded prompts in the live database schema, benchmarking generated-query accuracy across joins, filters, and aggregations	

Leadership & Achievements

- Felicis AI Fellow 2026**; won **\$25K** at the Felicis Fellows hackathon creating a robotics safety layer that uses a world model to predict a robot’s next actions and intervene before any harm
- Led 15-developer team; collaborated with **15+ institutions** for enterprise-scale AI software deployment
- Won **10+ hackathons**, including **1st at AWS sLLM Agents**, 2nd at AWS Hack Day 2025 and YC Hack Halloween, 3rd at AGI House small LLMs and Modern Stack, and 4th at CodeRabbit AI Hackathon (47 teams)

Projects

DoorDash – Agentic Causal Bandit Platform <i>Python, FastAPI, Next.js, Causal ML, Docker</i>  GitHub
- Extended a DoorDash-style Thompson Sampling incentive baseline with Networked LinTS, shape-constrained Bayesian optimization, and a delayed AIPW reward correction to handle interference between zones - Simulated 50,400 incentive decisions across 10 Bay Area zones, cutting cumulative regret 67.4% versus vanilla Thompson Sampling and measuring a 114% spillover ratio where cross-zone effects exceeded direct ones, showing independent per-zone experiments overstate true impact - Delivered a FastAPI and Next.js application with 8 analytical dashboards, a GPT-4o agent that drives experiment configuration behind human approval on writes, and pluggable Dasher swarm simulators (built-in, OASIS, MiroFish)
Predictive Maintenance – Classification & Anomaly Detection <i>Python, Scikit-learn, XGBoost, SHAP</i>  GitHub
- Benchmarked Random Forest, Gradient Boosting, Logistic Regression, and SVM classifiers on 10,000+ sensor records with engineered ratio, interaction, and rolling-window features, reaching 0.99 AUC-ROC across 5 failure classes - Shipped per-prediction SHAP and LIME explanations plus an unsupervised anomaly ensemble (Isolation Forest, LOF, One-Class SVM) for unlabeled monitoring, wired end to end into a Streamlit app running raw sensor input through to an explained prediction
Cortex – ML Research & Serving Platform <i>Rust, Python, MLflow, TensorFlow Serving, ClickHouse, Qdrant</i>  GitHub
- Built an ML platform on Kubernetes with MLflow tracking and A/B testing pipelines, TensorFlow Serving inference at <100ms P95, and a RAG path over Qdrant and Ollama with cost-aware routing across model providers - Wrote an event-sourced CQRS backend in Rust serving a GraphQL API with WebSocket subscriptions over ClickHouse streaming analytics, deployed on Kubernetes with Istio and autoscaling sized for 50K+ concurrent users
NanoAgent – Autonomous AI Agent Runtime <i>Zig, SSE Streaming, BLE/Serial, Embedded Systems</i>  GitHub
- Implemented an AI agent runtime from scratch in Zig: ReAct loop, multi-provider LLM backends, SSE streaming, and 3 pluggable transports (HTTP, BLE, Serial) - Held all allocation in a fixed-size arena, cross-compiling the runtime to a single ~450 KB binary across 50+ embedded targets

Technical Skills

Languages: Python, C++, Java, Go, Rust, Zig, SQL, TypeScript, JavaScript, Bash
AI/ML: PyTorch, TensorFlow, Scikit-learn, Transformers, LLMs, RAG, Agentic AI, Fine-Tuning, Evaluation, LLM-as-a-Judge, Prompt Engineering, Gradient Boosting (XGBoost, LightGBM), CNNs, Reinforcement Learning
ML Systems: Distributed Training, Model Serving (TensorFlow Serving), Kubeflow, MLflow, ONNX, Quantization, Model Context Protocol (MCP), LangChain, Observability, Service Reliability
Full Stack: React, Next.js, FastAPI, Node.js, REST APIs, GraphQL, gRPC, Microservices, API Design, Streamlit
Cloud/DevOps: AWS, GCP, Docker, Kubernetes, CI/CD, GitHub Actions, Jenkins, Prometheus, Linux/Unix, Git
Data: Pandas, NumPy, Dask, Kafka, Airflow, Snowflake, Databricks, ClickHouse, PostgreSQL, Qdrant, ETL Pipelines, Experimentation, A/B Testing, Multi-Armed Bandits, Causal Inference, Anomaly Detection, Product Metrics