Chandan Rudrappa

chandanrudrappa00@gmail.com | +1-(682)-320-5138 | Linkedin | Portfolio | Github | Edison, NJ | Open to relocation

Summary

AI Engineer with **2.5+ years of experience** shipping LLM-powered systems: multi-agent workflows on LangGraph, RAG over vector search, and Claude API classification with prompt caching, built for explainable, human-in-the-loop use. Benchmarks every system against baselines instead of trusting prompt quality alone.

Experience

Research Analyst — Applied AI & LLM Systems Jul 2025 – Present
The University of Texas at Arlington, School of Social Work *Arlington, TX*

- Built a human-in-the-loop LLM classification pipeline on the Claude API that tags waiver sections with confidence scores and rationales for analyst review, piloted on 20 themes and cutting **analyst coding time by 60%**.
- Scaled that pilot to the full **57-theme** codebook with a RAG system embedding **1,300+** human-coded MAXQDA segments in LanceDB, retrieving nearest examples so every prediction is explainable and traceable to its source.
- Implemented a multi-label NLP classifier (TextCNN and fine-tuned BERT) benchmarked on micro/macro-F1 across the same 20 policy categories, establishing a reproducible offline baseline to validate the LLM approach.
- Automated section-by-section comparison of Medicaid 1915(c) waiver filings with a Streamlit tool that runs sentence and word-level text diffs and exports HTML reports, saving the team **40+ hours of manual review every month**.

Graduate Teaching Assistant Jan 2025 – May 2025
The University of Texas at Arlington *Arlington, TX*

- Cleaned, transformed, and validated real-world community datasets across two centers, increasing usability metrics for student projects and setting up rigorous baselines for statistical modeling.
- Mentored **100+** undergraduate students in Probability, Statistics, and Calculus, guiding hypothesis testing and experimental setups and boosting comprehensive test metrics by an average of **20%**.

Data Scientist Mar 2022 – Jul 2023
Analyttica Datalab Pvt. Ltd. Bengaluru, India

- Defined column-type-to-function recommendation logic for **1,400+** platform functions, powering a Neo4j-backed engine that surfaced context-aware suggestions (e.g., time-series functions for datetime columns).
- Designed and delivered **150+** reusable data science functions for the LEAPS analytics engine, migrating legacy R time-series libraries to Python and shipping releases through Docker, Kubernetes, and Jenkins CI/CD pipelines.
- Developed **60+** end-to-end machine learning workflows and notebooks integrated into the platform, covering regression, classification, clustering, and time-series models.
- Refactored **80+** static legacy plots into interactive Plotly dashboards, cutting front-end delivery cycles by **25%**; led ML demo sessions for Siemens and Aditya Birla Group (**30 participants** each).

Technical Skills

Generative AI & LLMs: Anthropic Claude API, Claude Code, RAG, vector search (LanceDB), text embeddings, prompt engineering & caching, GenAI evaluation, LangChain, LangGraph

Machine Learning & NLP: scikit-learn, PyTorch, TensorFlow, Hugging Face Transformers, BERT, TextCNN, multi-label classification, model evaluation

Programming & Databases: Python, SQL, R, DuckDB, Polars, MySQL, SQLite, BigQuery, Neo4j, C, C++

Cloud: Azure (Azure ML, Microsoft Foundry), AWS (Lambda, SQS, API Gateway, DynamoDB, ECR, S3, EC2)

Big Data & Visualization: Apache Spark, PySpark, Hadoop, ETL pipelines, data warehousing, Streamlit, Plotly, Tableau

Tools & MLOps: MLflow, model registry & versioning, drift monitoring, Terraform, ONNX Runtime, Docker, Kubernetes, Jenkins, GitHub Actions, CI/CD, pytest, moto, Git, Jupyter

Projects

NFL Scouting-Report Agents — Multi-Agent LLM System | GitHub | Live Demo Oct – Nov 2026

Orchestrates six Claude Haiku agents in LangGraph over 26 deterministic DuckDB stat tools to write NFL opponent scouting reports; a fact-checker re-runs every cited tool call and returns failed claims for revision, reaching 100% final claim accuracy across 120 reports. An evaluation harness that catches 99.9–100% of planted errors showed the multi-agent design beats a single agent by +0.72 on a 1–5 blind-graded scale (95% CI +0.33 to +1.10), at $0.03–0.05 per report.

Financial Complaint Classifier — Serverless Inference on AWS | GitHub Sep 2026

Classifies consumer-finance complaints into six financial product categories with fine-tuned DistilBERT, served on an event-driven AWS pipeline (Lambda, SQS, DynamoDB) built in Terraform. A TF-IDF + logistic regression baseline beat DistilBERT (0.893 vs 0.874 macro-F1); INT8 ONNX quantization cut model size 75% without observed degradation.

Education

Master of Science in Data Science, *The University of Texas at Arlington, **GPA: 4.0/4.0*** Aug 2023 - May 2025
Coursework: Neural Networks, Computer Vision, Big Data, Probability & Statistics, Data Mining, Machine Learning

Certifications: Microsoft Certified — Operationalizing Machine Learning and Generative AI Solutions (AI-300), Aug 2026