

ANOUSHKA KANCHERLA

www.kancherla.net / Phone: (940) 999-4278 | Email: anoushka.kancherla@gmail.com

Data scientist and full-stack developer with experience building AI-powered systems, React applications, and production-ready data pipelines. Driven to deliver data-backed solutions that support experimentation and real-world decision-making.

EDUCATION

University of Texas at Austin Class of 2028 | Informatics, B.S., Human-Centered Data Science Track | GPA: 3.72

Certifications: Azure AI Fundamentals, Deep Learning Specialization, Machine Learning Specialization

New Tech High @ Coppell Class of 2024 | GPA: 4.126 (Weighted) | SAT: 1560 (Composite)

SKILLS

Programming Languages and IDEs: Python, JavaScript, SQL, R, Java, Visual Studio, Cursor

Databases & Data Systems: PostgreSQL, Supabase, MySQL, BigQuery, SQLAlchemy, ETL Pipelines

AI & ML: TensorFlow, scikit-learn, Neo4j (Cypher), LLM Orchestration, OpenCV

Web & Backend: React, Flask APIs, REST APIs, Figma Analytics & Visualization: Tableau, Looker Studio, Matplotlib

PROJECTS

Codebase Q&A

August 2026

- Architected a CLI tool that grounds Claude's codebase answers in real files via a **custom MCP server**, then **validates every claimed file:line citation against the live filesystem**, flagging invalid paths or ranges instead of trusting output by default
- Designed a **confidence and uncertainty framework** with code-level checks that override the model's self-reported confidence
- Adapted **explanation-stability methodology** from SHAP and feature-attribution research to LLM tool use, measuring answer reproducibility across repeated runs via Jaccard similarity

Personal Research Assistant *personal-research-assistant.vercel.app*

May 2026

- Built and deployed a full-stack AI research tool using the **Claude API**, streaming structured briefs from live web search via real-time NDJSON event routing
- Designed a **prompt-enforced source-confidence rating system** with per-finding reasoning transparency
- Integrated **Google OAuth and Drive API** for automatic brief persistence, cross-session recall, and audit logging

PulseAI Hook 'Em Hacks Hackathon

April 2026

- Solo-built and deployed a **multi-agent Gemini pipeline** to classify health datasets, run statistical hypothesis testing, and generate plain-English patient briefs
- Engineered a health equity module, **responsive visualizations**, and **K-Means clustering** across 20 sex-stratified biomarkers

GraphBeat UT Austin Women in CS Hackathon

February 2026

- Built an end-to-end **Graph-RAG app** using **Neo4j**, **Llama 3.3 (Groq)**, and **React** to deliver explainable music recommendations via song-relationship traversal
- Designed an optimized Neo4j schema via **neomodel OGM**, enforcing data integrity across a multi-node graph
- Developed a **responsive React (Vite) interface** with a multi-stage status stepper to visualize the Graph-RAG pipeline

Polarus HackTX 2025 Hackathon

October 2025

- Engineered a relational database using **MySQL + SQLAlchemy** to improve data storage for a civic accessibility platform
- Implemented backend data processing workflows, reducing manual data handling and ensuring consistent quality

RESEARCH & PROFESSIONAL EXPERIENCES

UT Austin School of Information Research Collaborator, Prof. Nathan TeBlunthuis

June 2026 – Present

- **Designing an interrupted time-series study** testing whether five major AI coding tool releases shifted reciprocity norms in open-source contribution
- Built and validated an **excess-word-frequency pipeline** across GH Archive issue, PR, and commit text

MIRA Safety IT / AI Development Intern

January 2026 – May 2026

- Migrated backend data layers from a third-party BaaS to **PostgreSQL**, centralizing access logic and enforcing clean frontend/backend separation
- Realigned **REST API** routes across the full stack and resolved broken authentication flows to restore functionality
- Implemented a missing admin service with **structured error handling** and hardened data integrity constraints

WIZE Computing Academy Summer Data Intern

June – August 2025

- Designed and deployed an **automated ETL pipeline** using **Google Apps Script + BigQuery**, eliminating manual data updates
- Analyzed datasets in **SQL** to surface insights informing operational improvements
- Built **real-time Looker Studio dashboards**, expanding data visibility for non-technical stakeholders

University of Texas at Dallas Applied-Logic, Programming Languages & Systems Research Internship

June – July 2023

- Ran experiments using **FOLD-R++** and **s(CASP)**, generating reproducible research findings on explainable AI
- Authored a research paper on explainable logic models, strengthening academic documentation of methods

PROFESSIONAL CERTIFICATIONS

Agentic AI Certification *DeepLearning.AI*

December 2025

- Built agentic AI workflows in Python using **reflection, tool use, planning, and multi-agent design patterns**
- Integrated LLMs with external tools (databases, APIs, web search, code execution) for end-to-end automation

Deep Learning Specialization *DeepLearning.AI*

July 2025

- Built neural networks (CNNs, RNNs, LSTMs) and applied **regularization, optimization, and hyperparameter tuning** to improve performance
- Implemented sequence models for **NLP and time-series analysis**, enhancing prediction accuracy and model robustness

Microsoft Certified: Azure AI Fundamentals *Microsoft*

June 2025

Machine Learning Specialization *Stanford Online*

August 2024

- Developed supervised models using **NumPy, scikit-learn, and TensorFlow**, improving predictive accuracy across tasks