

Kameswara Sai Srikar Manda

(617) 212-7199 | srikarmks@gmail.com | srikarmk.dev | linkedin.com/in/srikarmk | github.com/Srikarmk

EDUCATION

Northeastern University

Boston, MA

Master of Science in Data Science, GPA **3.6/4.0**

Sep 2024 – Dec 2026

TECHNICAL SKILLS

Languages and Methods: Python, SQL, R, TypeScript, JavaScript; A/B testing, statistical analysis

LLMs and Agents: LangChain, LangGraph, CrewAI, MCP, RAG, Ollama, OpenAI and Anthropic APIs, Claude Code, Hugging Face, FAISS, LLM evals

ML and Statistics: PyTorch, TensorFlow, Keras, scikit-learn, XGBoost, NLTK, statsmodels, SciPy, NumPy, pandas

Data Engineering: Spark, Airflow, Snowflake, Hadoop, BeautifulSoup, Power BI, Tableau, Matplotlib, Streamlit

Databases: PostgreSQL, MySQL, MongoDB, DuckDB, Redis, Pinecone

Web, DevOps and Cloud: FastAPI, React, Next.js, Node.js, Docker, Kubernetes, Terraform, OpenShift, Podman, GitHub Actions, AWS, GCP

EXPERIENCE

Northeastern University

Sep 2026 – Present

Teaching Assistant, DS 3500: Advanced Programming with Data

Boston, MA

- Help a large undergraduate class learn Python for data work, under Prof. Deahan Yu
- Walk students through their first Git workflows, from installing and authenticating to resolving merge conflicts

Launch Plan

Jan 2026 – Jul 2026

Data Scientist

Boston, MA

- Built **LangGraph** workflows and **MCP server** integrations for a K-12 financial-literacy platform, letting LLM agents write and tailor lessons using internal tools and outside APIs
- Took content formatting and assessment assembly off the teaching team's plate by handing them to those agents
- Replaced a keyword search educators didn't trust with a **RAG pipeline** on Pinecone and custom embedding models
- Built the data layer from scratch with FastAPI, PostgreSQL and Airflow, pulling in content from several school districts

Soverin

Mar 2024 – Aug 2024

Data Engineer / Full Stack Developer

Remote

- Led full-stack work on an AI-tool marketplace (React, FastAPI, AWS, SQL) that grew to **10,000+ monthly users** and a catalog of 1,900+ tools
- Found where users struggled to search and reworked query flow and indexing: search **30% faster**, API latency **40% lower**
- Set up CI/CD with GitHub Actions and Postman tests, keeping the platform at **99.5% uptime** across releases

PROJECTS

RAG Poisoning: One Bad Passage in 10,000 | Ollama, Qwen3, Llama 3.1, Gemma 3, bge-small, Python

Oct 2026

- Hid one false or malicious passage among 10,000 in a RAG system and tested three open LLMs: one passage changed 21–42% of target answers, five copies 85–100%, and planted instructions were obeyed 2–56% of the time
- Filtering instructions out of retrieved text stopped every injection, at no accuracy cost on 783 clean questions

whygit | MCP server, Python MCP SDK, Claude Code, GitHub Actions, PyPI

Oct 2026

- An MCP server and CLI, published on PyPI, that tells coding agents and people what state a Git repo is in and how to fix it, covering 29 repo states and 33 common errors
- It never changes a repo, and tests prove it by fingerprinting every file in 38 test repos before and after each run (246 tests)

Silent RAG Failures | FAISS, sentence-transformers (MiniLM, bge-small), BEIR, LLM-as-judge

Sep 2026 – present

- A benchmark for spotting when a RAG system's retrieval quietly gets worse, with no relevance labels. It injects five real-world failures (half-finished embedding upgrades, stale documents, chunking changes, query drift, config mistakes) at 10–90% severity and checks which monitors catch them
- Built the retrieval pipeline, embedding cache, experiment runner and a re-embedding detector

Keystone | Python, NLP, LaTeX parsing, Next.js, pdf.js

Sep – Oct 2026

- A live site that shows what a research paper builds on, argues with and takes on faith, read straight from its LaTeX with no LLM and pinned to the page; 197 papers so far
- Publishes its own error rate on outside labels: 80% precision on SciCite (11,014 citations) and 69% on ACL-ARC

PAPERS AND CASE STUDIES

- Predictive Variance Decides: Region-Aware Gaussian Processes Should Adapt the Noise (preprint in preparation)
- PENMOT: Permutation-Equivariant Networks for Multi-Object Tracking (Northeastern, 2025)
- Experiment Forensics: a pre-registered re-analysis of 22,743 real A/B tests; all 13 planned checks replicated
- Smart Farming: A Machine Learning and IoT Approach (IRJET, 2023)