

Sadiksha Chitrakar

(203) 927-6288 | sadiksha.chitrakar101@gmail.com | linkedin.com/in/sadikshachitrakar | github.com/sadikshachitrakar

SUMMARY

AI Engineer with an M.S. in Data Science and 2+ years of software engineering experience building end-to-end LLM and machine learning systems, from RAG and agentic workflows to evaluation, APIs, testing, and containerized deployment. Proficient in Python, SQL, PyTorch, LangChain, and LangGraph.

EDUCATION

Master of Science in Data Science University of New Haven	August 2024 - May 2026 GPA: 3.97 West Haven, Connecticut, United States
Bachelor of Science (with Honors) in Computer Science Coventry University	October 2019 - November 2022 GPA: 3.94 Coventry, England, United Kingdom

EXPERIENCE

Graduate Research Assistant University of New Haven	August 2025 - December 2025 West Haven, Connecticut, United States
- Built and evaluated end-to-end machine learning pipelines for biological age prediction on 75,000+ NHANES records, achieving a 3.5-year MAE ($R^2 = 0.95$) under 5-fold cross-validation, with XGBoost reducing MAE by 18% versus Random Forest. - Engineered reproducible data pipelines across multiple NHANES survey cycles using MICE imputation, data standardization, feature engineering, cohort construction, and feature-importance analysis to identify predictive biomarkers.	
Software Engineer Xuno	January 2024 - January 2025 Ellensburg, Washington, United States
- Developed and shipped the Xuno cross-border payments app for iOS and Android, integrating REST APIs and driving architecture decisions that improved application scalability and maintainability, contributing to a 4.7-star rating and 500+ installs. - Reduced post-deployment bugs by 30% through unit, widget, and integration testing guided by production logs and runtime metrics.	
Flutter Developer Code Himalaya	November 2022 - January 2024 Lalitpur District, Nepal
Cut release-cycle time by 70% and published 30+ apps to the App Store and Google Play by automating build and deployment with Fastlane CI/CD pipelines.	

PROJECTS

CampusGuide: Agentic RAG University Assistant	GitHub
- Built a Retrieval-Augmented Generation (RAG) system over multi-format university documents (PDF, DOCX, HTML) using LangChain, all-mpnet-base-v2 embeddings, and ChromaDB, with top-5 semantic retrieval over 1,000-token chunks with 250-token overlap. - Benchmarked three 4-bit-quantized LLMs (Llama-3-8B, Mistral-7B, Phi-3-Mini) using an automated LLM-as-a-Judge harness (PatronusAI Lynx-8B), identifying Llama-3-8B as the highest-performing model with 53.2% factual accuracy and a grounding score of 0.54, compared with factual accuracies of 42.7% for Mistral-7B and 43.6% for Phi-3-Mini. - Reduced hallucinations on out-of-scope queries by engineering an agentic LangGraph workflow (query routing, tool calling, context grading, rewrite-and-retry, verification, escalation) that enforced context-grounded responses and correct abstention where weaker models fabricated content. - Exposed campus tools (document search, structured office/program/deadline lookup, deadline calculator) through an MCP server and a FastAPI endpoint, containerized with Docker and Docker Compose and validated with a pytest suite.	
AWS Serverless Data Pipeline for Demand Forecasting	GitHub
- Designed and deployed a serverless AWS pipeline (S3, Lambda, Glue, Athena, RDS/PostgreSQL, SageMaker, API Gateway) that automated ingestion, ETL, and SQL-based analytics over 100,000 multi-store retail records. - Delivered sales and inventory insights to non-technical stakeholders through a REST API built with API Gateway and Lambda and interactive Power BI dashboards.	
Keypoint Detection for Building Components	GitHub
- Fine-tuned a Keypoint R-CNN (ResNet-50/101-FPN) in PyTorch with transfer learning to detect building components and regress semantic keypoints from a custom COCO-format dataset (200 images augmented to 1,000 via Albumentations). - Achieved 0.89 mAP and 0.82 keypoint mAP (Object Keypoint Similarity), reducing validation loss from 10.05 to 3.73 while evaluating ResNet-50 and ResNet-101 backbones to optimize keypoint precision.	

SKILLS

Languages: Python, SQL, JavaScript, Dart, R
Machine Learning & AI: PyTorch, TensorFlow, scikit-learn, NLP, Computer Vision, Feature Engineering, LLMs, RAG, Agentic AI, Prompt Engineering, LangChain, LangGraph, Hugging Face Transformers, ChromaDB, MCP
Data, Cloud & MLOps: AWS (S3, Lambda, Glue, Athena, SageMaker, RDS, API Gateway), Docker, FastAPI, REST APIs, CI/CD, Pandas, NumPy, ETL, Data Modeling, Apache Spark, PostgreSQL, MySQL, MongoDB, Firebase
Tools: Git, GitHub, Power BI, Tableau, Excel, Jira

CERTIFICATIONS

Power BI (Data Modeling & Visualization) | Apache Spark Programming with Databricks | MLOps Essentials: Model Development and Integration