

# Ruchitha thota

+1(414)241-7475 | [ruchithota7777@gmail.com](mailto:ruchithota7777@gmail.com) | LinkedIn | North Carolina | open to relocate

## Summary

**AI / GenAI Engineer & Data Analyst** specializing in building production-ready ML, LLM, and RAG systems. Skilled in SQL, Python, Power BI, and cloud platforms to deliver scalable analytics and intelligent automation. Proven in driving business impact through predictive modeling, data pipelines, and AI-driven insights.

## Skills

- **Programming & Analytics**  
Python (Pandas, NumPy, Scikit-learn, TensorFlow, PyTorch), SQL (PostgreSQL, SQL Server, MySQL), R, Excel (Advanced Formulas, Power Query, VBA), DAX, MongoDB
- **Visualization & BI**  
Power BI (DAX, Star Schema, RLS), Tableau, Google Data Studio, Data Storytelling, KPI Dashboards, Business Intelligence Reporting.
- **Machine Learning & AI**  
Predictive Modeling, Forecasting, Gradient Boosting, Neural Networks, NLP, Explainable AI (SHAP, LIME), A/B Testing, Model Evaluation, Regression, classification, clustering (K-Means), feature engineering, model evaluation (F1, ROC-AUC, RMSE), Scikit-learn, TensorFlow, PyTorch, OpenCV, OCR, Computer vision tools
- **Cloud & Data Engineering**  
Azure (Data Factory, Synapse, Blob Storage), AWS (S3, Lambda, RDS), GCP (BigQuery), Snowflake, Databricks, Data Warehousing, Data Lake Design.
- **Agentic AI & Generative AI**  
LLMs (GPT, Llama, Azure OpenAI), LangChain (chains, tools, retrievers), LangGraph (state graphs, multi-agent workflows), Retrieval-Augmented Generation (RAG), semantic search, embeddings, ChromaDB, hybrid search, re-ranking, chunking strategies, embedding model selection
- **Model Context Protocol (MCP)**  
Context passing, agent-to-tool communication, multi-agent orchestration patterns
- **Backend & APIs**  
FastAPI, Flask, REST APIs, microservices, async processing, secure API design
- **Frontend**  
HTML, CSS, React, TypeScript, JavaScript

## PROFESSIONAL EXPERIENCE

### Data & BI Extern (Analytics & AI), Extern

Mar 2026 - Present | North Carolina, United states

- Built modular, AI-powered pipelines to process 200+ page mortgage blob files—combining OCR (Tesseract, PaddleOCR), PDF parsing (PyMuPDF), and RAG techniques for intelligent data extraction, classification, and search.
- Developed a document retrieval system using LlamaIndex and Retrieval-Augmented Generation (RAG), optimized for multi-document mortgage blobs. Enhanced precision through chunk tuning, metadata filtering, and evaluation of open-source LLMs like Mistral and Phi-2.
- Conducted end-to-end evaluation of the document intelligence system on 200+ page mortgage blobs—benchmarking OCR accuracy, RAG retrieval quality, and routing performance. Delivered a technical report outlining model trade-offs, optimization strategies, final deployment recommendations, and built a UI for demo purposes.

### Data & AI Intern, Brilliant

Sep 2025 - Dec 2025 | Michigan, United states

- Optimized SQL data models and analytical queries, reducing runtime from **9.2s to 6s** and enabling faster, data-driven decision-making across teams.
- Designed and deployed interactive Power BI dashboards, increasing executive adoption from **40% to 90%** and reducing manual reporting effort by **40%**.
- Built and evaluated machine learning models using Python (Scikit-learn), improving churn prediction **F1-score by 28%** and reducing false negatives by **22%**.
- Developed scalable ETL pipelines integrating data from Salesforce, HubSpot, and SQL systems, improving cross-functional reporting accuracy by **35%**.
- Applied feature engineering, hyperparameter tuning, and cross-validation to enhance model performance for customer segmentation and retention analytics.

- Automated data quality checks and validation workflows, improving dataset accuracy, consistency, and reliability for downstream analytics.
- Collaborated with Sales, Marketing, and Finance teams to translate business problems into data-driven and AI-powered solutions.
- Leveraged AI/ML techniques to enhance predictive analytics and support intelligent decision-making across business operations.

#### **Data Analyst / AI Intern, Food FIXR**

May 2025 - Jul 2025 | Wisconsin, United states

- Developed a Python-based OCR and text extraction pipeline to process food ingredient data, achieving **~95% accuracy** and converting unstructured data into structured, analyzable datasets.
- Applied NLP and data processing techniques to generate real-time ingredient insights, enabling data-driven, health-focused user recommendations.
- Designed and integrated data pipelines and APIs (including Stripe), ensuring seamless data flow for payment processing and subscription analytics.
- Built interactive dashboards and analytics workflows to monitor transactions, revenue trends, and user behavior, supporting business decision-making.
- Automated data validation and preprocessing steps to improve data quality, consistency, and reliability across application workflows.
- Collaborated with engineering teams to enhance system scalability, optimize APIs, and support data-driven product improvements.
- Leveraged AI-driven data processing and automation techniques to improve user experience and enable intelligent application features.

#### **Graduate Student Assistant (Data Analytics & AI), Concordia University Wisconsin**

Jan 2025 - Apr 2025 | Wisconsin, United states

- Developed privacy-preserving machine learning models for student retention analysis using federated learning, enabling secure analysis across distributed datasets.
- Applied advanced feature engineering, cross-validation, and hyperparameter tuning to improve model performance, increasing prediction accuracy by **85%**.
- Designed data pipelines and preprocessing workflows for sensitive datasets, ensuring **~99% data privacy compliance** while maintaining analytical integrity.
- Built predictive analytics solutions to identify at-risk students, contributing to a **15% improvement in retention outcomes** through data-driven interventions.
- Implemented synthetic data generation techniques to enable safe experimentation and model training on sensitive data.
- Collaborated with academic and IT stakeholders to translate business problems into scalable analytics and AI solutions.
- Contributed to a NAIRR pilot project with OpenMined, applying secure AI practices and advancing federated learning use cases.
- Applied AI/ML and emerging GenAI concepts to enhance data analysis workflows and support privacy-first intelligent systems.

#### **Data Analyst, Wipro Limited**

May 2021 - Feb 2024 | Hyderabad, India

- Automated SQL-based data pipelines and SLA reporting workflows, improving data reliability and reducing report failures from **11% to 2%**, enabling real-time analytics.
- Designed and optimized interactive dashboards (ServiceNow/BI tools), enhancing data accessibility and reducing query response time for business users.
- Built KPI-driven analytical reports to monitor incidents, workload trends, and operational performance, enabling faster, data-driven decision-making.
- Implemented automation scripts and rule-based workflows, reducing manual intervention and improving incident resolution time by **32%**.
- Collaborated with cross-functional teams to define data requirements, standardize metrics, and improve data governance and reporting consistency.
- Supported data integration and API-driven workflows, ensuring seamless data flow across systems and improving reporting accuracy.
- Applied foundational AI/automation concepts (rule-based logic, workflow optimization) to enhance operational analytics, laying groundwork for scalable AI-driven solutions.

## **PROJECTS**

---

### **Agentic AI System for Automated Data Analysis**

- Developed an AI-powered data analyst application that converts natural language queries into SQL, enabling dynamic analysis of user-uploaded datasets.
- Designed a multi-agent architecture (SQL generator, analysis agent) to automate data querying, insight generation, and decision support.
- Built an interactive dashboard using Streamlit and Plotly to visualize query results and support real-time data exploration.
- Implemented dynamic schema handling, allowing users to upload any CSV dataset and automatically generate accurate SQL queries.
- Integrated local LLMs (Ollama – Phi3/Mistral) to eliminate API costs while maintaining AI-driven insights.

#### **PDF Q&A Assistant Using Retrieval-Augmented Generation (RAG)**

- Developed an end-to-end PDF Question & Answer Assistant that allows users to upload documents and ask natural-language questions, with answers generated strictly from document content to reduce hallucinations.
- Implemented text extraction and intelligent chunking for large PDFs
- Generated semantic embeddings using Sentence Transformers
- Built vector similarity search with ChromaDB • Designed a RAG pipeline to retrieve relevant context before answer generation
- Integrated Groq (LLaMA-3) for fast, context-aware LLM responses • Built an interactive Streamlit UI for document upload and querying

#### **Snowflake AI Agent for Conversational Healthcare Analytics**

- Built an enterprise-grade AI agent using Snowflake Cortex to enable conversational analytics over both structured healthcare data and unstructured biomedical research.
- Developed a semantic data layer to translate natural language queries into SQL using Cortex Analyst, enabling efficient querying of large-scale clinical and claims datasets (1.4M+ patients, 65M+ encounters).
- Integrated Cortex Search to implement Retrieval-Augmented Generation (RAG) over PubMed biomedical research corpus for contextual insights.
- Designed and orchestrated an intelligent agent that dynamically routes user queries between structured and unstructured data sources, providing accurate, summarized, and cited responses.

#### **EDUCATION**

---

**Master of Science in Computer Science**, Concordia University Wisconsin (CGPA: 3.47)      Mar 2024 - Dec 2025 | Mequon, WI

#### **CERTIFICATION**

---

AWS Certified Cloud Practitioner – Amazon Web Services – Earned Apr 2025

