

Sudeeksha Vandrangi

sudeekshav.proj@gmail.com | [LinkedIn](#) | [GitHub](#) | Tempe, AZ (Open to relocation or Remote)

EDUCATION

Master of Science in Data Science, Analytics, and Engineering Arizona State University, Tempe, AZ, USA	May 2025 4.0/4.0
Bachelor of Technology in Metallurgical and Materials Engineering National Institute of Technology, Rourkela, India	Jun. 2023 8.91/10.0

TECHNICAL SKILLS

- **Programming Languages:** Python, SQL (Hive, SparkSQL), Scala, Bash, C
- **Database and DB tools:** MySQL, NoSQL, PostgreSQL, Amazon Redshift, MongoDB, DynamoDB, Redis
- **Distributed Systems and Cloud:** AWS (EC2, S3, SageMaker), Azure, GCP Vertex AI, Spark, Kubernetes, Docker, Airflow, Databricks
- **Frameworks:** J2EE, Spring Boot, Hibernate, Spring MVC, Django, Anaconda, RESTful APIs, Power BI, Tableau, Looker
- **Methodologies and Tools:** Agile, DevOps, Lean Development, CI/CD, SDLC, Linux, Git/GitHub, JIRA, Docker, Kubernetes
- **AI Frameworks and Tools:** OpenAI API, Claude, Pinecone, ElevenLabs, Gradio, Cursor, Windsurf, Loveable, Prompt Engineering, Retrieval Augmented Generation, Model Context Protocol, LangChain, LangFlow, CrewAI
- **ML Frameworks:** A/B Testing, Causal Inference, Time Series Forecasting, Natural Language Processing, Deep Learning

EXPERIENCE

Data Scientist Intern | Lennox International - Richardson, TX Jun. 2024 – Aug. 2025

- **Improved forecast accuracy by 10%, as measured by MAPE, on over 9 million records** by designing and deploying a scalable **Temporal Fusion Transformer** (state-of-the-art model) by integrating **Apache Spark** and **Azure ML Studio**.
- **Reduced training and pipeline execution time by 35%**, cutting iteration cycles from 72 hours to under 24 hours by optimising distributed compute resources in **Databricks**.
- **Enabled automated real-time business alerts**, accelerating stakeholder response time by **30%**, by integrating notification services via **Azure Communication APIs**.
- **Increased system reliability by 25%** by diagnosing and resolving edge-device data synchronisation issues across **15K+** field devices and backend services.
- **Delivered a production-ready forecasting and monitoring platform on schedule for Q3 2025 rollout** by driving cross-functional collaboration across engineering and product teams.

Technical Trainee | Tata Steel Pvt. Ltd. - Jamshedpur, IND May 2022 – Jul. 2022

- **Reduced defective slab output by 15%** by engineering data validation and anomaly detection scripts in **Python** and **SQL** to monitor real-time production logs.
- **Built a logistic regression model (90% accuracy), cross-validated to predict slab defects**, supporting early interventions and reducing defect rate by identifying potential causal factors.
- **Automated reporting workflows** using **Python** to generate monthly summaries and dashboards, **reducing manual analytics workload by 40%** and enabling faster insights.

PROJECTS

MCP-Powered Agentic RAG with Firecrawl and Qdrant | Python, FastMCP, Qdrant, Firecrawl, Docker Jun. 2025

- Improved retrieval accuracy by **~30% over standard RAG pipelines** by designing an extensible **Retrieval-Augmented Generation** system that combines **Qdrant vector storage** and **Firecrawl web search** for low-latency semantic retrieval.
- Reduced maintenance overhead by **~40%** by implementing a modular architecture with clearly defined interfaces, enabling teams to add new data sources, retrieval strategies, or tools without modifying core logic..
- Enabled seamless integration with developer tools by **exposing FastMCP-powered APIs** for clients like **Cursor IDE**, outperforming conventional static pipelines in flexibility and modularity.

Modular LLM Orchestration and Feedback Pipeline | Python, Streamlit, OpenAI API, Cursor May 2025

- **Increased qualitative feedback review efficiency by ~40%** by developing a modular, local-first orchestration framework inspired by workflows at Listen Labs, designed to manage task routing, prompt versioning, and structured audit trails.
- **Enhanced traceability and compliance**, by logging of all LLM interactions in append-only records, improving transparency over conventional single-model pipelines.
- Improved extensibility and scalability, enabling research teams to rapidly add new tasks, prompt templates, and feedback workflows **without re-architecting the system**.

CodeNudge: AI-Powered Technical Interviewer | Python, Langflow, WhisperAPI Sept. 2024

- Built an MVP in under **24 hours** to validate real-time technical interview simulation, enabling rapid iteration and early user feedback.
- Boosted candidate preparation effectiveness over static interview guides by implementing adaptive question chaining and live hinting workflows.
- Engineered seamless multi-turn dialogue by integrating **Langflow orchestration** with **Whisper API** speech-to-text transcription, delivering dynamic performance insights in real time.

CLTV Prediction & Segmentation for Personalized Insurance | Python, Gradient Boosting, MiniBatch K-Means May 2024

- **Improved marketing efficiency by ~20%** by predicting Customer Lifetime Value with **92.5% R² accuracy**, enabling tailored premium policy offerings and retention strategies.
- **Segmented over 100,000 customers into 3 actionable clusters**, reducing resource allocation inefficiencies by **~15%** and supporting data-driven campaign optimization.