

RUDRA NARAYAN MEHER

AI Engineer | Generative AI & RAG | NLP Research | Applied ML
rudra.meher@rudrameher.in | LinkedIn | GitHub | Portfolio

PROFESSIONAL SUMMARY

AI engineer and product builder specializing in evidence-grounded document intelligence, retrieval-augmented generation (RAG), and low-resource NLP. Experience spans medical-legal AI systems, legal research platforms, and grammar-guided language modeling. Skilled in building end-to-end AI pipelines involving OCR, structured extraction, semantic retrieval, and citation verification. Focused on developing reliable, scalable AI applications through practical engineering, iterative experimentation, and human-centered evaluation.

TECHNICAL SKILLS

Languages: Python, Java, SQL **AI / ML:** LLMs, RAG, AI agents, NLP, scikit-learn, TensorFlow

Frameworks & data: LangChain, LangGraph, FastAPI, Flask, REST APIs, Next.js, PostgreSQL, MongoDB, Qdrant

Cloud & tooling: AWS, Azure, Google Cloud, Docker, Git, GitHub

EXPERIENCE

Co-Founder & AI Engineer | Shivru AI

2026 – Present

- Developed medical-legal document workflows combining OCR, structured extraction, retrieval, and source-grounded evidence citations.
- Built chronology generation, medical-billing reconciliation, missing-record checks, and citation-verification workflows for litigation review.

AI Engineer Intern | Auracle Labs

Feb 2026 – Jun 2026

- Contributed to technical course design, retrieval-augmented generation systems, and application memory functionality.

Technology Specialist | Oshmaal Pvt. Ltd.

Feb 2023 – Present

- Developed algorithmic trading logic and trading bots; worked on AI-based algorithms and predictive models.

SELECTED PROJECTS & RESEARCH

Sambalpuri NLP | Grammar-Guided Low-Resource Language Modeling

- Created an approximately 20K-example Odia-Sambalpuri-English resource using grammar-book material, lexical mappings, language tags, and rule-based transformations.
- Completed five experimental iterations with a roughly 200M-parameter model; tested language/category tags and 256- vs. 512-token inputs.
- Human review: 76% fully accurate; 89% usable or near-correct. Evaluation sample size and rubric are not yet documented.

Shivru AI | Medical-Legal Document Intelligence

- Implemented OCR-to-structured-data processing, evidence-grounded retrieval, medical chronology, and citation checking.
- Tested an end-to-end processing workload on a 500-page document set (not an accuracy benchmark).

Advotac | RAG-Based Indian Legal Research Platform

- Built an AWS-hosted legal research system covering approximately 860 Central Acts and legal judgments with Qdrant, OCR, and hierarchical JSON.
- Experimented with different document-chunking approaches to support legal retrieval workflows.

Mandi Price Forecasting | Applied Machine Learning

- Built an agricultural commodity price forecasting model using historical mandi data and seasonal patterns.

EDUCATION & ACHIEVEMENT

IIT Madras | BS in Data Science and Applications | 2024–2028 (Expected) **LogicLOOM — IIT Madras ML**

Hackathon: Top 2 and Top 3 placements in separate competitions.