

ROHAN DHANRAJ YADAV

AI/ML Engineer | Generative AI | Agentic AI Systems

✉ rohan.dhanraj.y@gmail.com | 📞 +91-7008958143 | 📍 Bengaluru, India

🌐 LinkedIn | 📁 Portfolio | 🐙 GitHub

PROFESSIONAL SUMMARY

AI/ML Engineer specializing in **Generative AI, Multi-Agent Systems, and AI Platform Development**, with 7+ years of experience including 4+ years in production AI systems. Proven expertise in building scalable **LLM-powered applications, RAG pipelines, and Complex Agentic architectures** that handle 10K+ daily queries and process 1M+ records.

Led cross-functional teams to design and deploy **Multi-Agent Systems**, autonomous workflows, and enterprise-grade AI platforms. Strong experience across **end-to-end AI lifecycle** — from data pipelines and model fine-tuning to API development and cloud deployment.

Delivered high-impact AI solutions for **Fortune 500 clients**, achieving up to **50% performance improvements** and building intelligent systems for real-world applications.

CORE COMPETENCIES

- **Languages/Frameworks:** Python, Django, Flask, FastAPI, Streamlit
- **Cloud/DevOps:** Git, GitHub, Docker, CI/CD, GCP, AWS, MLOps
- **GenAI & Agentic Systems:** LLMs (OpenAI, Anthropic, Vertex AI, Google Gen AI Platform, Cohere, AI21 Labs), LangChain, LangGraph, RAG, Prompt Chaining, Chain-of-Thought, Agent Orchestration, Multi-Agent Collaboration, Deep Research, Autonomous Agents
- **ML & NLP:** Scikit-Learn, TensorFlow, Keras, PyTorch, Optuna, Transformers (Encoder/Decoder), BERT/RoBERTa, Sentiment Analysis, NER, Text Classification, Hyperparameter Optimization, PyCaret
- **Data Engineering & Databases:** Apache Airflow, ETL, SQL (MySQL, PostgreSQL, BigQuery), NoSQL (MongoDB, Cassandra), Vector DBs (Pinecone, Qdrant, LanceDB, Chroma, Pgvector, Milvus, Weaviate)
- **Analytics:** NumPy, Pandas, Matplotlib, Seaborn, Plotly, SciPy, Statistical Analysis, Predictive Analytics
- **Web Scraping & Tools:** BeautifulSoup, Scrapy, Selenium, Playwright, Pydantic Structured Output

PROFESSIONAL EXPERIENCE

AI Tech Solutions Ltd. | AI/ML Engineer — Generative AI & Agentic Systems

Nov 2024 – Oct 2025

- Led a **15-member team** building scalable AI workflows using LangGraph, natively achieving a **38% accuracy improvement** via hybrid RAG.
- Architected multi-modal autonomous assistants executing GPT-4 context orchestration, rapid image search, and LLM-based tool decisions.
- Developed **AI-Tech-Crawler**, an open-source web scraping platform natively utilizing LLMs and Pydantic schemas for data extraction.
- Scaled complex containerized orchestration models deploying **Docker** against Vector DBs for robust enterprise semantic search.

Tech: Python, LangChain, LangGraph, FastAPI, LLMs, RAG, Agent Orchestration, Docker, Vector DBs

Aimlytics Technology | Software Engineer — Generative AI & ML

Sep 2022 – Oct 2024

- Led an **8-member team** resolving **50%+** of critical deployment issues and delivering custom AI workflows for Fortune 500 clients (Cisco, ITC).
- Architected a robust RAG-based virtual sales pipeline (**SalesMoji**) structurally scaling to consistently process **4,500+ daily API requests**.
- Fine-tuned **BERT & RoBERTa** NLP models driving **18% performance improvements** across a 15-language sentiment tracking stack.
- Automated rigorous ETL architectures utilizing **Apache Airflow** to seamlessly ingest and process **100K+ daily records**.

Tech: Python, RAG, LangChain, Airflow, BERT, RoBERTa, FastAPI, Flask, TensorFlow, PyTorch, Vector DBs, GCP

iNeuron.ai | Machine Learning Engineer Intern

Aug 2021 – Sep 2022

- Led **10 engineers** building automated ML pipeline architectures, natively deploying scalable classification networks via **Flask and Docker**.
- Reduced backend model training times by natively **35%** employing **Optuna** for strictly automated hyperparameter optimizations.
- Built a multi-class bioinformatics processing system successfully parsing **77 protein expression markers** accurately across 8 functional classes.

Tech: Python, Flask, Scikit-Learn, Optuna, MySQL, Cassandra, Docker, AWS

Bureau Veritas Group | Metallurgical Quality Engineer

Aug 2016 – Nov 2019

- Directed quality assurance and supervised metallurgical standards compliance across major industry frameworks (Jindal, Tata Steels).
- Targeted data-driven quality enhancement leveraging **Python (NumPy, Pandas, SciPy)** for rigorous technical statistical and predictive tracking.

Tech: Data Analysis, Python, NumPy, Pandas, SciPy, Statistical Analysis, Quality Control

KEY PROJECTS & ARCHITECTURES

- **Omodore:** No-code AI agent platform with autonomous reasoning, hybrid RAG, and multi-agent orchestration via LangGraph.
- **Rbetraj:** Automated dropshipping platform powered by AI-Crawler with LLM-based structured data extraction.
- **GALAMBO:** Multi-modal agentic AI assistant with GPT-4, real-time image search, and autonomous tool usage.
- **SalesMoji:** AI-powered sales intelligence system reliably processing 100K+ records/day and 4,500+ daily API requests.
- **AISERA:** Enterprise AI system for Fortune 500 clients with BERT fine-tuning and production issue resolution.
- **evalmyBRAND:** Aspect-based sentiment analysis system (RoBERTa, 18% improvement) with 15+ language support.
- **InventtEd:** AI-powered educational platform with CLIP, Flan-T5, automated essay grading, and virtual tutoring.

EDUCATION

B.Tech. Mechanical Engineering

Biju Patnaik University of Technology | 2016

GATE Qualified (ID: A2WR3)

CERTIFICATIONS & ACHIEVEMENTS

- 🏆 **ML Zoom Camp:** DataTalksClub (8F163B)
- 🏆 **HackerRank Gold:** Python, Problem Solving, Statistics
- 🏆 **Kaggle 30 Days ML:** Rank 69 / 7,573 Teams