

DEEPESH AHLAWAT

New Delhi, India | deepeshahlawat001@gmail.com | [LinkedIn](#) | [GitHub](#)

PROFESSIONAL SUMMARY

AI Intern with hands-on experience building production-ready AI systems in Python, PyTorch, LangChain, and RAG pipelines, deployed on AWS and Oracle Cloud. Proven problem-solving track record across fintech and healthcare domains — including real-time anomaly detection, agentic LLM workflows, and complex document extraction. International First Place winner (Baidu, 2026) and National First Place winner (ISRO BAH'25). Looking to apply agentic AI, context engineering, and full-stack AI skills to impactful real-world products.

TECHNICAL SKILLS

Languages & APIs	Python (advanced), Bash, REST APIs, JSON, Regex, vector DB APIs (Pinecone, FAISS)
AI / ML Frameworks	PyTorch, TensorFlow, scikit-learn — model training, data preprocessing & evaluation
Agentic AI & LLMs	LangChain, LangGraph, RAG, GraphRAG, Context Engineering, LLM Tool-Use, Multi-Agent Orchestration
Databases & Storage	Vector databases (Pinecone, FAISS, Chroma), SQL (fundamentals), NoSQL concepts
Cloud & DevOps	AWS, GCP (familiar), Oracle Cloud, Docker, FastAPI, Git/GitHub, Weights & Biases, ONNX
Frontend & Backend	FastAPI (Python), Node.js, React / Next.js, REST API integration
OCR & Parsing	Baidu PaddleOCR (fine-tuning), Vision Transformers, OpenCV, PDF & document extraction

EXPERIENCE

Agentic AI Intern | ConnectHEOR, London, UK (Remote)

Jan 2026 – June 2026

- Architect multi-step autonomous AI agents using LangGraph and advanced LLM tool-use to automate complex reasoning and data extraction workflows across healthcare evidence synthesis use cases.
- Benchmark new agentic AI frameworks against existing solutions, evaluating performance on real-world document parsing and multi-step reasoning tasks.
- Design robust audit-trail logging systems for full transparency from initial prompt to final LLM output, ensuring production-grade reliability and reproducibility.

Research Intern (ML) | Physical Research Laboratory (PRL), ISRO, Ahmedabad, India

Sep 2025 – Dec 2025

- Designed, trained, and evaluated Physics-Informed Neural Networks (PINNs) on complex time-series data, covering the full ML pipeline from data preprocessing and feature engineering to model validation against established scientific catalogues. • Applied symbolic regression to translate complex physical dynamics into mathematical representations — demonstrating advanced model design and rigorous scientific evaluation methodology.
- Authored a research manuscript detailing deep-learning methodology and system optimisation, targeted for publication in an international scientific journal.

AI Intern | DigitalPaani, Gurugram, India

Apr 2025 – Sep 2025

- Engineered NLP modules using LLMs for deep document parsing, automating information extraction and converting unstructured text into structured, rule-compliant formats for fintech and operations use cases.
- Built production RAG pipelines using LangChain and Pinecone vector database to query and synthesise insights from dense technical documentation.
- Deployed an ensemble-based anomaly detection ML model monitoring plant signals in real-time, managing the complete pipeline: data preprocessing, model training, evaluation, and production deployment on Oracle Cloud.
- Integrated a full-stack conversational AI assistant using OpenAI, FastAPI (Python backend), and Oracle Cloud — shipped to production.

- Developed Python evaluation tools for error analysis, cost estimation, and failure detection, ensuring high reliability of LLM-generated outputs.

KEY PROJECTS

Chronos-VL: VLM-Based Complex Document Extraction | *PaddleOCR, PyTorch, Vision Transformers — International 1st Prize, Baidu 2026*

- Engineered an end-to-end Vision-Language Model pipeline using Baidu's 0.9B NaViT architecture to extract and digitise highly complex degraded documents — reducing Character Error Rate from 19.82% to 1.64% (76x improvement).
- Built output filtering systems to eliminate LLM hallucinations and a Two-Stage Pipeline (Visual Perception + Semantic Normalisation) integrating linguistic rules, demonstrating production-quality AI engineering.

Physics-Aware Neural Network for CME Detection | *Custom NN, Time-Series ML, Python — National 1st Prize, ISRO BAH'25*

- Designed a custom neural network outperforming state-of-the-art ARCANE systems by modelling continuous plasma dynamics with physics-informed features, validated against the Richardson/Cane ICME Catalogues (94.5% True Negative, 63.1% Detection Rate).
- Awarded National Rank 1 by the Minister of Space for building an explainable, production-ready AI system for critical space weather early-warning.

RagFin AI — Agentic Document Assistant | *LangChain, OpenAI GPT, Pinecone, FastAPI*

- Built an Agentic RAG-powered assistant for multi-step reasoning, dense document extraction, and structured report generation — directly applicable to fintech document automation workflows.
- Designed LLM workflows parsing unstructured data into strictly formatted structured outputs, demonstrating context engineering, vector database integration, and RAG pipeline proficiency.

EDUCATION

B.Tech — Dr. Akhilesh Das Gupta Institute of Professional Studies, GGSIPU

2023 – 2027

CGPA: 9.412 | University Rank 3

Relevant coursework: Machine Learning, Deep Learning, Data Structures & Algorithms, Linear Algebra, Probability & Statistics, Engineering Mathematics.

HONORS & AWARDS

- International First Prize, Baidu PaddleOCR Fine-Tuning Competition (2026) — \$1,000 prize; reduced Character Error Rate to 1.64% from 19.82% on complex historical manuscripts.
- National First Place, ISRO Bharatiya Antariksh Hackathon BAH'25 — Awarded by the Minister of Space, Government of India; ranked 1st nationwide out of thousands of teams.
- National Rank 70, Amazon ML Challenge (2025) — Top 70 nationwide in a large-scale multi-modal ML competition.
- 1st Place, TechGenesis 2.0 Hackathon (2025) — Built and shipped an innovative AI solution under strict time constraints.