

Medhansh Kapoor

AI/ML Engineer | Full-Stack Developer

Email: medhansh541@gmail.com | Phone: +91 8368680865 | GitHub: [Medhansh-741](https://github.com/Medhansh-741) | LinkedIn: [medhansh-kapoor](https://www.linkedin.com/in/medhansh-kapoor)

Skills

Languages: Python, TypeScript, JavaScript, SQL, C, C++

Backend / Web: FastAPI, REST APIs, Next.js, Node.js, WebSockets, Celery, Pydantic, React, Supabase, Git

Applied AI: PyTorch, LangGraph, LangChain, LlamaIndex, RAG, AI Agents, Multi-agent Systems, YOLOv8, ONNX, OpenCV, Sentence Transformers, Embeddings, Semantic Search, Prompt Engineering, Function Calling, Multimodal AI, NLP

Data / Infra / MLOps: ETL Pipelines, PostgreSQL, PostGIS, Vector/Graph Databases (Qdrant, Neo4j), Redis, SQLite, GDAL, GeoServer, Martin Tile Server, OpenLayers, AWS S3, GCP, Docker/Podman, RHEL/Ubuntu

Experience

IndiaAI Mission (MeitY) – AI/ML Intern | [\[Offer Letter\]](#)

8 June 2026 – July 2026

Building a standalone dataset-quality evaluation toolkit for AIKosh (under MeitY's IndiaAI Mission) – bridging the ICMR MIDAS 2.0 framework gap, slated for production integration by internship end.

Next.js, FastAPI, Celery, Redis, Supabase (PostgreSQL), AWS S3, Pydantic, JWT Auth, Jinja2, WeasyPrint

- **Metadata Intake:** Engineered an 8-step wizard with pre-signed URLs for direct-to-storage uploads, bypassing the backend to enhance security and reduce server load.
- **Assessment Engine:** Architected an asynchronous, parallelized scoring engine to evaluate datasets across 15 government-defined quality domains, generating Composite Quality and Privacy Risk scores for automated release eligibility (Open/Controlled/Restricted).
- **Integration:** Built automated multi-format reporting (JSON/HTML/PDF) and webhook APIs, enabling AIKosh to auto-ingest verified dataset metadata upon assessment completion.

ISSA – DRDO – Student Trainee, Ministry of Defence | [\[Offer Letter\]](#)

18 May 2026 – July 2026

Building a self-contained, fully offline GIS platform enabling classified defence environments with no internet access to securely upload, process, and visualize geospatial map data.

FastAPI, PostGIS, GDAL/ogr2ogr, GeoServer, Martin Tile Server, OpenLayers, Docker/Podman, RHEL

- **Ingestion Pipelines:** Built automated ETL pipelines for vector (Shapefiles to PostGIS with reprojection/indexing) and raster data (GeoTIFFs auto-published via REST), eliminating manual GIS server setups.
- **Serving Layer & Infra:** Enabled real-time map-tile delivery to browser clients via a containerized 4-service microservices backend, securely deployed on a firewalled RHEL environment for fully offline, classified operations.

Geminid Systems – Software Development Intern | [\[Offer Letter\]](#) | [\[Completion\]](#)

May 2026 – Jun 2026

Evaluated enterprise AI toolchains and shipped production integration tests on live Salesforce CRM infrastructure.

Vanna.ai, LlamaIndex, LangChain, Librosa, PyDub, Torchaudio, Salesforce Apex, Einstein AI, Agentforce

- **Benchmarking:** Evaluated **Vanna.ai**, **LlamaIndex**, **LangChain** for NL-to-SQL, and audio frameworks (Librosa, PyDub) for feature extraction; identified that **agent architecture outweighs model choice** for multi-table reasoning.
- **Salesforce AI Platform:** Built Apex REST services and SOAP integrations; conducted prompt engineering experiments on live CRM data using Einstein AI and Agentforce.

Projects

JanSamadhan – Autonomous Civic Surveillance Platform | [\[Live\]](#) | [\[GitHub\]](#) | [\[Demo Video\]](#)

Mar 2026

Detects civic issues via CCTV/dashcam, auto-generates complaint tickets, verifies repairs, and retrains itself.

YOLOv8, ONNX, OpenCV, FastAPI, PostgreSQL/PostGIS, Gemini 2.5

70.3% Precision | 57.7% Recall | 0.68 mAP50

- **Model & Inference:** Trained YOLOv8 on **4,783 images** with 30% engineered negatives and FN-bucketing for recall diagnosis; achieved **20ms ONNX inference** on a FastAPI microservice deployed via **GCP Cloud Run**.
- **Auto-Ticketing & Reliability:** Engineered a burst-frame extraction pipeline feeding a **4-tier reliability engine** and **DIGIPIN geospatial deduplication** (4m² grid), processing 256 end-to-end complaints in **0.36s** per ticket.
- **Verification & Active Learning:** Enforced automated post-repair rescans to verify fixes before ticket closure; bucketed field data into labeled classes for continuous baseline-gated retraining.
- **LLM Routing (Seva):** Integrated Gemini 2.5 Flash using a 4-step Chain-of-Thought to autonomously route complaints across **42** civic categories, auto-routing **125/256 tickets** with zero manual input.

NyayaAI – Multi-Agent Legal Intelligence Platform | [\[Live\]](#) | [\[GitHub\]](#) | [\[Demo Video\]](#)

Mar 2026 (36-hr Sprint)

5-agent pipeline: legal case in – research, strategy, drafted documents, and reasoned explainability out.

LangGraph, FastAPI, Qdrant, Neo4j, Redis, WebSockets, Celery, Groq, Gemini

5-Agent Workflow | 11 Legal Domains

- **Orchestration & Intake:** Built a stateful 5-stage LangGraph pipeline (Intake → Research → Strategy → Drafting → Explainability) with OCR-aware processing and **SHA-256 caching** to mitigate LLM hallucinations.
- **RAG & Determinism:** Indexed **4,582 chunks** (7 legal acts) in Qdrant via Sentence Transformers; engineered a diagnostic harness to catch retrieval drift and OCR instability across N-run tests.
- **Reliability & Graph:** Implemented a triple-engine fallback (Groq → Gemini → offline rules) for zero-downtime generation; built a Neo4j legal knowledge graph (**1,410 nodes, 1,837 relationships**) for explainable document drafting.
- **Realtime Infrastructure:** Developed an event-driven custom Redis Pub/Sub, WebSockets, and Celery backend for live citizen-lawyer negotiation chat, bypassing third-party APIs.

Achievements

India Innovates '26 – National Finalist

[\[Certificate\]](#)

Winner in Digital Democracy track; presented **JanSamadhan** live before senior policy leaders at Bharat Mandapam; project forwarded to central ministries.

Prayatna 3.0 Hackathon – Finalist

[\[Certificate\]](#)

Built NyayaAI's **5-agent legal AI backend** during a **36-hour** sprint.

Education

Manipal University Jaipur – Bachelor of Technology in Computer Science & Engineering

Aug 2025 – 2029

VVDAV Public School, New Delhi – CBSE Class XII, PCM + Computer Science

2024