Ayush Vardhan Jha

Bengaluru · ayushvjha@gmail.com · ayushvjha.com
linkedin.com/in/ayushvjha · github.com/ayushvjha · pramaanam.ayushvjha.com

SUMMARY

I'm an AI/ML engineer with about a year of production experience building LLM and classical-ML systems for legal and IP workflows: classification, LLM extraction, multi-agent drafting and citation verification. I ship end to end and evaluate rigorously, with leakage-free splits, calibration and held-out test sets. I built **Pramaanam**, a citation verifier for Indian Supreme Court judgments, and I'm now building **Factget**, a research and data agent, as founding engineer at TurtleLogix.

EXPERIENCE

AI/ML Engineer · TurtleLogix LLP, Bengaluru — Jul 2025 – Present

Founding Engineer, Factget — a research and data agent, built at TurtleLogix — Sep 2026 – Present

- Building a research and data agent that finds the data behind a question and shows where every number came from.

Client: Lakshmikumaran & Sridharan, Indian intellectual property law firm

IPR Docket: email triage platform — Go, Vue 3, PostgreSQL, Redis

- Built the platform that classifies incoming emails into **357 IP-law event types** and routes them to the responsible docketing teams.
- Microsoft Graph delta-sync ingestion, a PostgreSQL job queue with retry and recovery, and WebSocket real-time updates.
- Human-in-the-loop review flow that records reviewer decisions as labelled data for retraining.
- Keyset pagination and trigram indexing for search; clustered **8,745** reply templates into **58** standard ones to automate reply drafting.

Email classification and extraction service — Python, scikit-learn, PyTorch, FastAPI, vLLM

- Trained a two-stage classifier on **92,700** labelled emails with severe class imbalance, for a team handling **1,000–1,200 emails a day**.
- Compared nine models, incl. a fine-tuned transformer, with matter-level splits to prevent leakage; best **0.83 macro F1** on held-out data (4-class triage).
- Probability calibration plus an LLM verification step (Gemma on vLLM) for low-confidence predictions, with a rule-based fallback.
- LLM extraction of **24 case fields** with source references and OCR for PDF, image and Office files; shipped as a containerised FastAPI service.

FER drafting platform (team project) — Node/TypeScript, Django, React, Python, LangGraph, Postgres, Qdrant

- Five-service system turning an Indian First Examination Report and Form 3 into a drafted, objection-by-objection response.
- Document acquisition from InPASS, USPTO Global Dossier, EPO, JPO, KIPO and WIPO (Selenium workers, job queue with retry and orphan recovery); split discovery from download so attorneys pick the **10–20 relevant documents out of 200+** per patent family.
- OCR (PaddleOCR, Azure) and content-based LLM classification into 10 document types using head and tail excerpts, since Indian specifications put claims at the end; Form 3 extraction with a row-count self-check.
- **10-state FER lifecycle** in the orchestrator (Prisma, Postgres) with two-phase drafting: strategy and amended claims first, then attorney confirmation, then the final response; queue and boot-time recovery for in-flight drafts.
- React FER workspace: stepper, document picker, manifest editor, chat against the drafting engine, DOCX/PDF export.
- Hardened LLM calls: retries cut structured-output failures from about 30% to **under 1%**.

AI Research Intern · Manipal Institute of Technology — Jan 2025 – May 2025

Research on AI for legal documents.

Data Science Intern · Maiti Labs — Feb 2024 – Jul 2024

PROJECTS

Pramaanam · pramaanam.ayushvjha.com — citation verifier for Indian Supreme Court judgments — 2026

- Checks each citation in a legal draft against **38,397** Supreme Court judgments: case exists, name matches, quoted words appear at the cited paragraph.
- Deterministic checks, later-overruled flags from headnotes, and an opt-in LLM reading agent (custom loop) with a code-level guard so it never marks a citation VERIFIED.
- Held-out test split: **96.1%** right-case resolution on 667 items; **272/275** real quotes verified; **1/1,166** altered quotes passed.

Adverse Drug Event Detection — BioClinicalBERT, PyTorch, spaCy · GitHub — Aug – Sep 2025

- Fine-tuned BioClinicalBERT on **23,516** medical sentences to detect adverse drug events and drug–symptom relations, with a weighted loss for class imbalance.

PUBLICATIONS

- **Enhancing Potato Crop Health: A CNN-Based System for Early Detection of Early and Late Blight** — first author, Springer · Paper
- **Soft Actor-Critic Algorithm in High-Dimensional Continuous Control Tasks** — IEEE, 2025 · Paper

EDUCATION

B.Tech, Computer Science (AI & ML) · Manipal Institute of Technology — Aug 2021 – Jul 2025

SKILLS

Languages: Python, Go, SQL, JavaScript
LLMs and agents: LangGraph, vLLM, Hugging Face Transformers, sentence-transformers, Qdrant, LLM extraction and verification, agent guardrails, OCR
ML and evaluation: PyTorch, scikit-learn, fine-tuning, probability calibration, class imbalance, leakage-free splits, benchmark design
Backend and infra: FastAPI, Django, PostgreSQL, Redis, Docker, WebSockets, Microsoft Graph API, GitHub Actions, Oracle Cloud
Frontend: Vue 3, React