Parshvi Jain

+91-7452905310 | parshvijain1508@gmail.com | linkedin.com/in/jainparshvi | github.com/parshvi1508 | parshvi.tech

Education

ABES Engineering College — Ghaziabad, India
B.Tech in Computer Science, CGPA: 8.48/10 — *Aug. 2023 – May 2027*

Indian Institute of Technology Madras — Chennai, India
BS in Data Science and Applications (Online) — *Aug. 2024 – May 2028*

Experience

Trurism — Remote, India
Backend Developer Intern — *Oct. 2025 – Jan. 2026*

- Architected 20-module FastAPI backend with PostgreSQL, Alembic, and Redis; designed JWT/RBAC auth, rate limiting, and multi-tenant white-label architecture for isolated B2B2C operations across travel verticals.
- Built ML orchestration layer connecting model inference to booking and recommendation flows; integrated Razorpay HMAC-SHA256 webhooks and XML.Agency API for live flight and hotel data; deployed via Azure CI/CD with Docker.

MedEvidences AI — Remote, USA
Backend Developer Intern — *Apr. 2025 – Aug. 2025*

- Migrated Flask monolith to async FastAPI; rebuilt 30+ REST APIs with OAuth2 and NPI middleware, achieving 25% latency reduction by eliminating blocking I/O across database and auth calls.
- Led zero-downtime PostgreSQL migration from Azure to AWS RDS using async SQLAlchemy; implemented dual-layer session and chat memory architecture for stateful AI workflows.

Projects

XAI Forensics | *Python, PyTorch, FastAPI, HuggingFace, LIME, Next.js, Docker* | GitHub | Demo

- Built a 3-panel forensic XAI system: LIME token attribution (300 perturbation samples), greedy counterfactual removal, and dual-model divergence scoring across DistilBERT-SST2 and Twitter-RoBERTa; rejected attention weights as explanations per Jain & Wallace 2019.
- Dockerized FastAPI backend (4 endpoints, per-request latency tracking) with Next.js dashboard; chose LIME over SHAP to preserve local fidelity per prediction, avoiding SHAP's global approximation assumptions.

Pan-India Renewable Energy Forecasting | *Python, TensorFlow, LSTM, scikit-learn, AWS S3* | GitHub

- Achieved 2.2% MAPE jointly forecasting energy consumption and renewable generation across 40+ Indian states (96,360+ hourly records); outperforms ARIMA (4.8%), linear regression (6.1%), exponential smoothing (5.3%); RMSE = 0.15 GWh ($p < 0.01$). Published at IEEE Conference 2025.
- 4-layer LSTM (13K params), 60-day sliding windows, dropout-based uncertainty quantification (95% prediction intervals); sub-1s CPU inference; designed for existing government grid infrastructure without GPU requirements.

Data Analyst Agent API | *Python, FastAPI, Gemini 2.5, Docker, OpenRouter* | GitHub

- Automated analysis API supporting 7 file formats via a 3-stage Extract-Analyze-Format pipeline; semaphore-based concurrency (max 3 parallel requests) with UUID-isolated Docker sandboxes per execution.
- 4-attempt exponential backoff with Gemini 2.5-Flash primary and OpenRouter fallback; system continues serving when primary model endpoint fails.

Technical Skills

Languages: Python, SQL
AI/ML: PyTorch, TensorFlow, scikit-learn, HuggingFace Transformers, LSTM, LIME, NLP, Explainable AI (XAI), LLM integration, prompt engineering, retrieval-augmented generation (RAG)
Backend: FastAPI, Flask, REST API design, async architecture, JWT, OAuth2, RBAC, multi-tenant systems, WebSockets
Infrastructure: Docker, AWS (S3, RDS), Azure, PostgreSQL, Redis, MySQL, Firebase, CI/CD, Git