Onkar Jadhav

AI / Full Stack Software Engineer
Pune, Maharashtra
Bachelor of Technology in Computer Science

+91-7058604679
onkarj012@gmail.com
github.com/Onkarj012
linkedin.com/in/onkarj012

Summary

Final-year Computer Science student at PCCOE (9.08 CGPA) with hands-on experience building end-to-end AI systems, including an open-source LLM governance benchmark evaluated across the US, India, and EU legal scenarios, and a deep learning equity prediction system covering Nifty100 NSE equities. Proficient in Python, React.js, Node.js, FastAPI, and LangGraph.

Education

•Bachelor of Technology in Computer Science *Nov 2022 – Jun 2026*
Pimpri Chinchwad College of Engineering, Savitribai Phule Pune University CGPA: 9.08 (up to 7th semester)
– Relevant Coursework: Machine Learning, Data Structures & Algorithms, Database Management Systems, Operating Systems, Computer Networks, Object-Oriented Programming

•Higher Secondary Certificate *Jun 2022*
Shri Mhalsakant Junior College, Maharashtra State Board Score: 89.17%

Engineering Projects

•StockXpert – AI-Powered Equity Prediction & Stock Recommendation Platform *Nov 2025 – Present*
Python, TensorFlow, Next.js, FastAPI, TypeScript, Tailwind CSS, Cloudflare R2, AWS
– Architected a hierarchical deep learning system combining ResNet, GRU, Bi-LSTM, and multi-head attention to capture cross-horizon price dependencies across **Nifty100 NSE equities**, outperforming single-architecture baselines.
– Constructed a leakage-free feature pipeline from raw OHLCV data incorporating RSI, MACD, ADX, Bollinger Bands, SMA/EMA distance, and volatility indicators across strict train/validation/test splits.
– Implemented a multi-objective Huber loss with directional regularization and tail-event oversampling, achieving a **12% improvement in short-horizon return correlation** over baseline and reducing directional error on high-volatility equities.
– Designed a snapshot-first serving layer with Cloudflare R2 storage and market-aware TTL caching, delivering sub-second API responses across **all 100 Nifty stocks** during market hours while eliminating redundant inference costs.
– Built a full-stack trading dashboard (Next.js + FastAPI) featuring a live stock screener with real-time alerts, OHLCV candlestick charts, technical indicator overlays, and forecast tables (1D–10D) with calibrated confidence scores.

•GovBench – Open-Source LLM Governance Benchmark *Mar 2026 – Present*
Python, LLM Evaluation, AI Safety, Sentiment Analysis, Hallucination Detection
– Designed and open-sourced GovBench on GitHub, an LLM evaluation framework assessing model readiness for judicial and governmental deployment across **6 pillars**: demographic bias, procedural integrity, corruption resistance, jurisdictional awareness, transparency, and minority protection.
– Implemented controlled demographic isolation testing across **12 identity variants** per scenario with naturalistic prompting, measuring inherent bias in bail, sentencing, welfare, and immigration decisions across the US, India, and EU legal systems.
– Evaluated 4 production LLMs (Claude Sonnet 4.6, Gemini Flash Lite, DeepSeek V4, GPT-OSS 120B) across baseline, pressure, and adversarial modes, revealing top performers scored **90%+ overall** while identifying failures in minority protection (GPT-OSS: 63%) and jurisdictional awareness (DeekSeek V4: 56.7%).
– Built an automated scoring pipeline combining sentiment variance analysis, hallucination detection, and position drift tracking to produce composite deployment-readiness grades across evaluated models.

•PromptForge – Agentic Prompt Engineering Platform *Jan 2026 – Feb 2026*
Node.js, Express.js, BullMQ, Convex, OpenRouter, React.js
– Developed a platform for designing, testing, and optimizing LLM prompts, enabling developers and AI product teams to generate, refine, and structure prompts for faster delivery and more reliable outputs.
– Architected a creator-critic-evaluator pipeline that iteratively improves prompts through automated feedback, scoring clarity, completeness, constraints, and token efficiency across multiple versions.
– Implemented cross-model routing and benchmarking across OpenAI, Anthropic, Gemini, and open models, enabling direct comparison of quality, latency, and cost per prompt for model selection decisions.
– Built a fault-tolerant async backend using BullMQ workers and Convex persistence, supporting idempotent job retries and end-to-end logging for reliable prompt regression testing at scale.

Technical Skills

Languages: Python, JavaScript, SQL, C++, Java
Frontend: React.js, HTML, CSS
Backend: Node.js, Express.js, FastAPI, REST APIs, BullMQ
AI / ML: TensorFlow, XGBoost, LangGraph, OpenRouter, LLM Evaluation, RAG Pipelines, Agentic Workflows
Databases: MongoDB, PostgreSQL, Convex
Cloud & DevOps: AWS, Git, GitHub