

NAVEEN ADHARSH C

Coimbatore, India | +91 63792 39548 | naveenadharshc@gmail.com
[LinkedIn](#) | [Github](#)

SUMMARY

Data Science graduate student and AI/ML practitioner who designs and ships production-grade AI systems and machine learning models across reinforcement learning, generative AI / RAG, and forecasting — with PyTorch, TensorFlow, FastAPI, and vector databases. Backed by 4 research manuscripts under review across RL and generative AI. Seeking an AI/ML internship to build models that run in real systems.

PROJECTS

[SmartQuant — Reinforcement Learning Trading Agent with Risk-Aware Position Sizing](#)

- Engineered a custom Gymnasium trading environment with a transaction-aware reward, training PPO / Recurrent-PPO agents to jointly control trade direction, risk sizing, and capital allocation.
- Built a full out-of-sample evaluation suite — Sharpe, Sortino, drawdown, Calmar, and a buy-and-hold benchmark — with LSTM price forecasting and an interactive Plotly performance dashboard.

[InterviewAI — RAG-Grounded Technical Interview Platform](#)

- Built a full-stack, RAG-grounded AI interview platform (Next.js 15, FastAPI, PostgreSQL) that runs adaptive, role-based technical interviews — generating questions and grading answers strictly from retrieved source material to eliminate hallucination.
- Designed a multi-role retrieval pipeline (PyMuPDF, LangChain text-splitting, sentence-transformers MiniLM, ChromaDB) ingesting 7 ML textbooks into 7,600+ vector embeddings across 3 role-specific collections, with full traceability — every question logs ranked similarity scores and source-page citations.
- Architected a pluggable LLM-provider layer (Gemini / local Ollama / mock) behind a clean interface for cost-free local inference and offline CI, producing structured hiring reports (score, recommendation, skill breakdown) — backed by 21 automated tests.

[Finalytics — Retrieval-Augmented Equity Research Engine](#)

- Built a RAG platform that extracts insights from annual reports and SEC-style filings, pairing SentenceTransformers embeddings and a ChromaDB vector store with a FastAPI backend.
- Shipped interactive Plotly dashboards (candlestick charts, live market analytics) and an LLM-driven bank risk-scoring module to support analyst peer comparison and trend intelligence.

[ECHO — Real-Time Meeting Copilot with Agenda Tracking](#)

- Built a platform-agnostic meeting copilot that captures live audio, transcribes speech in real time, and tracks the conversation against a predefined agenda.
- Engineered real-time agenda-completion logic comparing the live transcript to agenda items to surface unresolved topics, presented through a Streamlit dashboard with a live transcript stream and progress view.

RESEARCH EXPERIENCE

Graph-Constrained Multi-Agent RL for Decentralized IoT Self-Healing

Under review

- Proposed GC-MARL, integrating graph-spectral reward shaping, Byzantine robustness, and decentralized policy consensus for 30-node IoT networks under cascading failures.
- Achieved 50% faster algebraic-connectivity recovery, 39% lower topology oscillation, and 82% Byzantine resilience versus baselines using PPO-based MARL and convergence-bound analysis.

Multimodal CLIP Embedding Compression for RAG Systems

Manuscript complete — under review

- Benchmarked 12 compression methods across MS-COCO and Flickr30k using CLIP embeddings, FAISS indices, and Flan-T5 — among the first systematic studies of compression's effect on multimodal RAG generation quality.
- Introduced the Embedding Retrieval Utility (ERU) metric; showed retrieval metrics overstated downstream generation loss by up to 65x, with PCA preserving quality under aggressive compression.

Impact of Feature Composition on Hybrid RNN Performance & Reliability

Under review

- Ran a 20-run stability analysis of 5 RNN architectures (incl. Hybrid LSTM-GRU) for stock prediction, isolating the effect of OHLCV, FinBERT sentiment, and technical-indicator features.
- Demonstrated that feature composition drives model stability more than architectural complexity.

Regime-Adaptive Reinforcement Learning for Equity Trading

Manuscript complete

- Benchmarked PPO, Recurrent PPO, MAML, and context-conditioned policies across forced bull/bear/sideways regime-shift benchmarks on historical US equities.
- Found PPO strongest on overall Sharpe while context-conditioned policy best preserved bear-market performance; meta-learning underperformed in low signal-to-noise regimes.

TECHNICAL SKILLS

Languages: Python, SQL, C++

Machine Learning: Scikit-learn, XGBoost, feature engineering, model evaluation, hyperparameter tuning

Deep Learning: PyTorch, TensorFlow, CNNs, LSTM/GRU, Transformers, autoencoders, fine-tuning

Reinforcement Learning: PPO, Recurrent PPO, Multi-Agent RL, custom Gym environments, reward shaping

Generative AI / NLP / LLMs: RAG, LangChain, LlamaIndex, embeddings, vector search, prompt engineering, speech-to-text

Computer Vision: CLIP, image-text retrieval, embedding compression (PCA / quantization)

Tools & Deployment: FastAPI, Flask, Docker, FAISS, ChromaDB, PostgreSQL, MongoDB, Plotly, Git, Linux

EDUCATION

Integrated MSc, Data Science — Amrita Vishwa Vidyapeetham, Coimbatore

Expected 2027

GPA 8.14/10 · Focus: Machine Learning, Deep Learning, Generative AI

Higher Secondary (Computer Science) — ARB International School, Pollachi

2022 · 87.6%

CERTIFICATIONS

Deep Learning Specialization (DeepLearning.AI)

Python for Data Science, AI & Development (IBM)

Build Deep Learning Models with TensorFlow (Codecademy)

Build a Machine Learning Model (Codecademy)