

Arjun S. Tomar

+91 96857 09169 | arjun16tomar@gmail.com | linkedin.com/in/arjunstomar | github.com/arjun16-t

Final year B.Tech student specialising in AI/ML, focused on building product-ready ML pipelines and scalable backend systems. Experienced in end-to-end ML development — from retrieval-augmented generation and model training to production deployment. Seeking internships in ML/AI Engineering.

EDUCATION

PCET's Pimpri Chinchwad University Pune, Maharashtra
B. Tech, Computer Science & Engineering — AI/ML Specialisation (CGPA: 9.1/10) 2023 – 2027
Certifications: RH104 Linux (RedHat), AI Engineering (IBM), Django (Meta), RAG (DeepLearning.AI), MCP (Hugging Face)

ENGINEERING PROJECTS

AthenaChat - Chatbot-As-A-Service | *Django REST, Jina Embeddings, Qdrant, Groq, Celery* 2026

- Designed a chatbot-as-a-service platform with a Django REST backend, Qdrant vector store, and Groq-powered LLM inference, architected for multi-client deployment with isolated knowledge bases per tenant
- Architected a **multi-tenant** RAG pipeline with per-client isolated vector search using Qdrant payload filtering and SHA256-based deduplication, preventing redundant ingestion across tenants
- Engineered a hybrid search system combining dense semantic embeddings (**Jina v5 Nano**) with sparse SPLADE vectors, achieving higher retrieval precision than keyword-only or semantic-only baselines

HireMailer | *Python, Django, Celery, Redis, PostgreSQL* 2025

- Architected a distributed three-service system (Web, Worker, Beat) using Celery task queues and Redis as the message broker, enabling non-blocking bulk email operations at scale.
- Implemented a full delivery reliability pipeline with pre-send MX record validation, 3x automated retries on failure, bounce handling, and one-click follow-up automation.
- Engineered an intentional throttling mechanism simulating human-like sending behaviour, processing **100+ emails per batch in 8-10 minutes** with zero spam flagging; deployed production-ready on Railway.

NSE Stock Prediction Engine with xAI | *PyTorch, Streamlit, FinBERT, SHAP, Git, Pandas-TA* 2026

- Built an end-to-end ML pipeline training three architectures (GRU, LSTM, custom Transformer encoder) on **55,000 sequences** across 50 Nifty 50 stocks to predict next-day closing price and market direction for NSE-listed stocks
- Engineered a **12-feature technical indicator pipeline** (RSI, EMA, Bollinger Bands, ATR) with per-stock MinMaxScaler normalization, enabling inference on any NSE-listed stock beyond the training set
- Implemented late fusion of price model outputs with FinBERT-based news sentiment (recency-weighted aggregation over a 7-day window from yFinance and Google RSS), improving direction accuracy from **62.3% to 74.6%**
- Deployed an interactive Streamlit app with live NSE data fetching, fuzzy search across **2,360 tickers**, and SHAP GradientExplainer integration for feature importance visualization and model explainability

Autonomous Vehicle Road Sign Detection | *Python, YOLOv8, OpenCV* 2024

- Fine-tuned YOLOv8 on a custom road sign dataset with hyperparameter optimization for real-time traffic sign detection at simulated high speeds
- Achieved **>80% mAP** on validation data, optimizing the precision-speed tradeoff for real-time inference applicability

TECHNICAL SKILLS

Languages: Python, C/C++, Java, SQL, JavaScript
ML Frameworks: TensorFlow, Scikit-Learn, PyTorch, LangChain, Hugging Face, YOLOv8, Sentence-Transformers, OpenCV, SHAP
Backend & Tools: Django, Flask, FastAPI, Celery, Redis, Git, Linux, Streamlit
Database: Postgres, SQLite, Qdrant, ChromaDB
Concepts: Computer Vision, Transformers, RAG, Agentic Workflows, Async Task Queues, MVC Architecture, Deep Learning, NLP