

# Amogh

AI/ML Engineer specializing in Agentic AI, RAG Systems, LLM Applications, and Backend Engineering

Agents · Available

Delhi, India · [github.com/ShriAmogh](https://github.com/ShriAmogh) · [www.linkedin.com/in/amogh-arora-b68056259](https://www.linkedin.com/in/amogh-arora-b68056259)

## SKILLS

---

Python, C/C++, SQL, FastAPI, Django, WebSockets, REST APIs, Selenium, Puppeteer, SQLite, PostgreSQL, MongoDB, Git, GitHub, VS Code, Docker, PyTorch, TensorFlow, Keras, Scikit-learn, Pydantic, LangChain, LangGraph, NLP, MCP, RAGAS, APO, Agentic Frameworks, Transformers, LLM APIs, Headless Browsers, DOM Extraction, OAuth Handshake, Data Structures & Algorithms, Machine Learning, Deep Learning, DBMS, Operating Systems, OOP, RandomForest

## ABOUT

---

Artificial Intelligence and Machine Learning undergraduate with hands-on experience building production-grade AI systems, agentic workflows, and scalable backend architectures. Experienced in Retrieval-Augmented Generation (RAG), LangGraph, MCP integrations, multi-agent systems, and LLM evaluation frameworks. Previously interned at Fractics, AlgorithmX, and Arya.ag, where I developed AI-powered automation, real-time systems, web scraping pipelines, and machine learning solutions for enterprise applications.

## EXPERIENCE

---

### AI Engineer Intern · Fractics

Feb 2026 – May 2026

internship

Engineered RAG and ingestion pipelines by integrating MCP servers into an existing RAG system with bidirectional WebSocket streaming; built an end-to-end RSS ingestion pipeline to enrich startup funding, acquisitions, and founders data. Developed advanced prompting and evaluation: Chain-of-Thought RAG with real-time UI streaming; APO system for prompt refinement using golden datasets; evaluated pipeline using RAGAS with custom domain metrics. Built authenticated scraping architecture using a headless Puppeteer pipeline with OAuth login handshakes for token refresh and gated data sources.

### Agentic AI Engineer Intern · AlgorithmX

Nov 2025 – Dec 2025

internship

Architected a scalable FastAPI backend with clean service-repository design, async ORM, and optimized DB interactions; built a bidirectional WebSocket system enabling session pooling, real-time state sync, and multi-user event pipelines. Developed a JSON-to-graph compiler converting JSON into a structured DAG. Engineered a multi-agent workflow using Microsoft Autogen, creating custom routing logic and orchestrated agent collaboration for automated code generation and validation.

### Software Developer Intern · Duco Consultancy

Jul 2025 – Aug 2025

internship

Designed and implemented MongoDB schema for scalable data handling. Built and deployed Express.js backend services, fully integrated with frontend. Streamlined workflows by loading large datasets into MongoDB for scalability.

### Software Engineer Intern · Arya.ag

Jul 2024 – Oct 2024

internship

Automated multi-page web data extraction with Selenium and CSV export. Achieved 96% accuracy with RandomForest on multilingual dataset via custom preprocessing and Google Translate API. Applied image processing for crop type classification.

## PROJECTS

---

### Multi-Agent Clinical AI System

Python · LangGraph · LangChain · Pydantic · FastAPI · MongoDB

- Architected a multi-agent clinical AI system that synthesizes messy longitudinal health records (PDF/JSON) into structured EHR-ready discharge summaries.
- Enforced zero-fabrication guardrails with a source-citation verification module; the agent triggers `escalate()` to alert clinicians of contradictions and unmanaged events instead of hallucinating.
- Designed an In-Context Correction Memory Loop with a simulated critic to inject localized corrections upstream and prevent recurrent token-matching errors.
- Optimized safety metrics to 1.00 safety recall on held-out patient sets, trapping critical allergen contraindications, missing rationales, and truncations.
- Tech: Python, LangGraph, LangChain, Pydantic, FastAPI, MongoDB

### Agentic AI RAG System with Caching

Python · LangChain · LangGraph · BM25 · ChromaDB · Pydantic · Cross-Encoder

- Hybrid agentic RAG pipeline combining BM25 lexical retrieval, dense vector search (all-MiniLM-L6-v2), and cross-encoder re-ranking to select the most relevant academic context with high precision.
- LangGraph-based execution graph with conditional routing, retry logic, and state propagation for deterministic orchestration of retrieval, generation, and evaluation.
- End-to-end agentic JSON generation and validation loop: LLM generates structured outputs, a Pydantic validator enforces schema correctness, feeds errors back, and ensures production-ready responses.
- Tech: Python, LangChain, LangGraph, BM25, ChromaDB, Pydantic, Cross-Encoder

### Startup Intelligence & Funding Analytics Pipeline

Python · FastAPI · Azure OpenAI · Selenium · Puppeteer · RSS · LinkedIn API

- Built an end-to-end startup intelligence pipeline that ingests articles from multiple RSS sources including TechCrunch, Inc42, The Hindu, and YourStory to identify funding, acquisition, and merger events.
- Developed an LLM-powered extraction framework to classify article relevance and transform unstructured news content into structured datasets containing startup, funding round, investor, and acquisition details.
- Engineered a multi-stage enrichment workflow integrating LinkedIn, CoreSignal, Explorium, and Tavily to gather founder profiles, company intelligence, investor information, and historical funding data.
- Implemented stealth browser automation using Puppeteer and Selenium for authenticated data collection, dynamic content extraction, and company website discovery across gated and JavaScript-heavy sources.
- Created a scalable async processing pipeline that generated analysis-ready datasets with startup metadata, funding amounts, investor details, and founder insights for market intelligence and research applications.
- Tech: Python, FastAPI, Azure OpenAI, Selenium, Puppeteer, RSS, LinkedIn API