

ALOK KUMAR TRIPATHI

✉ aloktripathe@gmail.com [in linkedin.com/in/aloktripathi1](https://www.linkedin.com/in/aloktripathi1) github.com/aloktripathi1

Education

Indian Institute of Technology, Madras

Sep 2023 – Sep 2027

Bachelor of Science in Data Science and Applications

TECHNICAL SKILLS

Languages: Python, TypeScript, JavaScript, SQL

LLM / Agentic: RAG, LangChain, FAISS, Qdrant, Multi-agent Orchestration, OpenAI API, Prompt Engineering, Structured Output

Frameworks/Tech: FastAPI, React, Next.js, PyTorch, Scikit-learn, HuggingFace Transformers, Tailwind CSS

Databases: PostgreSQL, Redis, Qdrant, FAISS, Supabase

Developer Tools: Git, Docker, Vercel, AWS, GitHub Actions, WandB, Jupyter

Projects

SachCheck | Production Multi-Agent Fact-Verification Pipeline

[GitHub](#) | [Live](#)

- Built and shipped a multi-agent pipeline end to end: one model extracts atomic claims from any article or image, parallel evidence agents query 4 independent sources, an 11-signal heuristic scorer aggregates evidence, and a second model synthesizes the final verdict. Image mode adds OCR with a vision-model fallback and a 9-signal manipulation detector.
- Enforced strict schema validation at every agent handoff with retry-and-repair logic for malformed outputs; cost-aware model routing cut per-check cost from \$0.04 to \$0.006 with no accuracy loss. Served via FastAPI with SSE streaming so verdicts render claim by claim as they arrive.

Lumen | AI Research Paper Intelligence System

[GitHub](#)

- Building a RAG system for ML research: users upload any paper or ArXiv link, ask technical questions, and get citation-grounded answers with exact page references, with no hallucinated citations. Supports multi-paper comparison and related work discovery.
- Qdrant vector DB with local sentence-transformer embeddings for retrieval; Celery and Redis manage async paper ingestion; outputs validated through structured schemas.

Music Genre Classification | Noisy Audio Mashup Classification

[GitHub](#) | [Live](#)

- Classified 3020 noisy audio mashups into 10 genres under a severe train-test distribution gap; identified the root cause via EDA using ANOVA F-scores per stem before choosing a model architecture.
- Ran 8 systematic experiments tracked via WandB, from a scratch CNN (LB 0.52) to an AudioSet-pretrained Audio Spectrogram Transformer (0.93); final weighted ensemble of AST, ResNet-50, and CNN reached **0.9614 Macro F1**, ranking **top 2.5% of 1200+ teams**.

Achievements

- **Anthropic x IIT Madras Club Builder Hackathon:** 1st Runner Up for SachCheck, a production multi-agent AI pipeline.
- **Kaggle:** Top 2.5% of 1200+ teams in Audio Genre Classification (0.9614 F1); Top 4% of 800+ in Thoracic Pathology Detection; Top 1.4% (12th of 866) in Plant Leaves Super-Resolution.
- **Open Source:** Merged [PR #5415](#) and [PR #5609](#) on mem0ai/mem0, adding native `embed_batch()` support across 6 embedder classes (Ollama, LMStudio, Together, HuggingFace, VertexAI, GoogleGenAI), eliminating N sequential HTTP calls in favor of 1 batched call.