Adithya Ramesh

Email: adithyaramesh66@gmail.com

LinkedIn: linkedin.com/in/adithya6ramesh
GitHub: github.com/Adithya6ramesh
Portfolio: www.heyadithya.me

PROFESSIONAL SUMMARY

AI and backend engineer who turns LLMs into reliable, deployed products, with a focus on retrieval and voice technologies, from data pipelines to production APIs.

TECHNICAL SKILLS

- **AI/ML:** LLM implementation, Prompt Engineering, RAG pipelines, Multimodal AI, OpenCV, MediaPipe, OpenAI Whisper, OCR (Tesseract, Paddle OCR), Hugging Face Transformers, LangChain, Omni TTS, LoRA, MCP, Hybrid Search, RAG evaluation (Recall@k, MRR), reranking
- **Core Tech Stack:** Python, JavaScript, Flask, FastAPI, React.js, Firebase, REST APIs, MySQL, HTML, CSS
- **Tools & DevOps**: NLP, NumPy, Pandas, PyTorch, Vector Databases (ChromaDB), Docker, Hetzner Cloud Deployments, Git, GitHub, VS Code

EDUCATION

- **LBS Institute of Science and Technology for Women, Poojapura** — Trivandrum, Kerala
  Bachelor of Technology - Electronics and Computer Engineering — 2022- 2026

EXPERIENCE

- **Software Developer Intern | Bytsolv Kochi** — October 2025-November 2025
  - Built REST APIs for a restaurant application using Renostack and PostgreSQL.
  - Integrated APIs across backend services and debugged production issues end-to-end.

- **SEO Content Strategist Intern | Typito AI** — December 2025 - July 2026
  - Produced technical documentation and product-focused content that improved user understanding of AI-powered video editing features.

PROJECTS

- **Niyamadharshini | GitHub | Demo | Evals**
  - Built a **bilingual (Malayalam/English) RAG assistant over Kerala Niyamasabha records,** and cites the source page for every answer.
  - Designed a **three-stage text-extraction pipeline**: PyMuPDF for digital PDFs, falling back to Tesseract and then PaddleOCR for old scanned pages the earlier stages could not read. Indexed 50+ legal documents with BGE-M3 multilingual embeddings in ChromaDB; top-K semantic retrieval through LangChain feeds Gemini 2.5 Flash for grounded answers.
  - Added hybrid search (BGE-M3 + BM25, RRF) and reranking, raising rank-1 retrieval accuracy from 8/17 to 13/17 (MRR@10 0.58 → 0.81). Built a 20-question eval harness to compare retrieval setups; int8 quantization cut reranking time on CPU from 24.9 s to 10.5 s.
  - **Stack**: JavaScript, Python, Flask, ChromaDB, BGE-M3 Embedding Model, Gemini 2.5 Flash, Tesseract, Paddle OCR, LangChain, PyMuPDF, BM25
- **Murmur** | **GitHub** | **Live**
  - Built and deployed a **voice journaling app**: OpenAI Whisper transcribes spoken entries, and Gemini analyses their emotion and context.
  - Added fast search across past entries with Fuse.js, stored locally in the browser with Dexie.js (IndexedDB).
  - Grew to **50+ active users; ranked #8 on the Indie Hackers Build Board**.
  - **Stack**: React, Flask, Python, OpenAI Whisper, Google Gemini API, Tailwind CSS, Dexie JS, Fuse JS
- **Blinky** | **GitHub** | **Live**
  - Built an AI-powered web app that reviews an interface and flags usability, accessibility and visual-design issues, along with Generic AI content analysis and suggested fixes.
  - Grew to **20+ active users;** ranked **#4 on the Indie Hackers Build Board.**
  - **Stack**: Python Flask, React.js, HTML, CSS, JavaScript, Gemini API, Firebase

HONORS & ACHIEVEMENTS

- Kerala Police Cyberdome, **Hac'KP 2025 Finalist**
- **5+ Hackathon** Honours**.**
- **GHCI'24 and GHCI'25** Scholarship recipient
- **Smart India Hackathon** State-level Participant
- Tink-Her-Hack, TinkerHub – Top 100 Projects (2x)