EMINENCE AI ACADEMY

# Applied Generative AI

For developers moving into AI roles: build and ship LLM, RAG and agent systems in one month.

| FEE | DURATION | SESSIONS | DELIVERY |
| --- | --- | --- | --- |
| ₹75,000 | 1 month | 24 | Classroom & online |

## YOU CAN ALREADY CODE. THAT IS EXACTLY WHY THIS IS FRUSTRATING

You have read about RAG and agents, maybe wired up an API call, and you can tell the difference between a demo and a system. What you have not had is a month of supervised time to build the real thing — with the chunking that actually works, the evaluation that catches regressions, and the cost profile that survives contact with a finance team. That is the month this is.

## WHAT YOU WILL BE ABLE TO DO

- Work fluently with an AI coding assistant inside a real codebase
- Design and build a retrieval-augmented system against your own documents
- Build an agent that calls tools and completes a multi-step task
- Evaluate an AI feature properly instead of judging it by vibes
- Reason about cost, latency and failure modes before shipping
- Leave with a deployed project and the repository behind it

## IS THIS THE RIGHT PROGRAM FOR YOU?

**It is, if…**

- You write code today and are comfortable debugging your own work
- You want to move into an AI engineering role rather than improve your current one
- You can give weekday evenings plus Saturday labs for a month
- You want to finish with something deployed, not a notebook

**It is not, if…**

- You do not code. AI at Work is built for you instead, and costs less
- You want AI theory or research — this is an engineering program
- You cannot give the Saturday labs. That is where the project actually gets finished

## CURRICULUM — 20 EVENINGS AND 4 SATURDAY LABS

01 How these models actually behave — and prompting that survives production, not demos

02 AI-assisted development — Claude Code and Cursor in a real codebase: context files, working from a spec, reviewing and debugging generated code

03 Calling models from your own code — APIs, structured output, streaming, cost and latency

04 Retrieval — giving a model your own documents, end to end, and why naive RAG returns the wrong passage

05 Tools and agents — function calling, and an agent that completes a multi-step task

06 Making it trustworthy — evaluating an AI feature instead of judging it by vibes, guardrails, handling hallucination

07 Ship it — deploy, monitor and control cost, with the capstone built in the Saturday labs

This is the module outline, not a session-by-session syllabus. Order, depth and the specific tools used are adapted to each batch and to the trainer leading it.

## SCHEDULE AND FEE

**Monday to Friday evenings, scheduled around working hours**

Four full-day Saturday labs, ending in a Demo Day — in the Himayat Nagar classroom for the Hyderabad batch, on Google Meet for the online batch. Included in the 24 sessions, not additional to them.

You should already be comfortable writing code.

**₹75,000** in total — the same fee for the Hyderabad classroom batch and the live online batch. They run as separate batches with the same curriculum, the same number of sessions and the same trainer time.

## TALK TO US

**+91 76619 88777** · hello@eminenceai.io

409, 4th Floor, Amruta Street, Above Almond House, Himayat Nagar, Hyderabad, Telangana 500029, India

eminenceai.io/courses/applied-generative-ai

Eminence AI Academy does not offer placement services, placement assistance, recruiter or employer introductions, or any job-search service, and has no hiring tie-ups. We make no job or salary guarantee. Course fees are non-refundable and non-transferable. Any salary figures we publish are market data from named third-party research, not offers, projections, or our students' outcomes. Details correct at the date of download; the current version is always at https://eminenceai.io.