DATAMARCOS

LIVE ONLINE · MENTOR-LED PROGRAM

LLMOps & AI Evaluation in Production

8 weeks · 8–10 hours a week · Part-time, live online

An eight-week, part-time live program for engineers who have already built a prototype with large language models and now need to evaluate, observe, deploy and harden it. A natural follow-on if you have already shipped a RAG or agent prototype — including graduates of our GenAI program.

Fees and EMI options shared on a counsellor call.

Program guide · Curriculum, projects and format. No prices are listed in this brochure.

Noida One, 303-304, Rasoolpur Nawada, Industrial Area, Sector 62, Noida, Uttar Pradesh 201309, India

www.datamarcos.com · hello@datamarcos.com · +91 88519 30644

About this program

Who it is for

- Software, data and ML engineers who need to take an LLM app beyond a prototype
- People who have completed a GenAI or similar program and want a production follow-on
- Engineers in the US, UK and other regions working on evaluation, observability or deployment

Who it is not for

- Beginners who have not yet built a first LLM or RAG project — start with a foundations program such as Applied Agentic AI or GenAI & Agentic AI Engineer
- Leaders who want strategy without writing code — see AI Product & Strategy for Leaders

Starting point

- Comfortable writing Python
- You have built, or can bring, an LLM, RAG or agent prototype — a personal project is enough
- Familiarity with Git and a local development workflow

DATAMARCOS

What you will be able to do

These are skills and work products from the program — not placement, salary or promotion promises.

- 1. Design a golden set, a regression check and a human-review sample for an LLM feature
- 2. Instrument traces and basic quality, latency and cost visibility
- 3. Version prompts and application code so a change can be reviewed before release
- 4. Evaluate retrieval quality on a RAG-style path, not only final answers
- 5. Expose the system behind a simple FastAPI service with a repeatable deploy path
- 6. Document failure modes, guardrails and a rollback plan for a capstone service

Projects

Project 1: Evaluation harness

A golden set, a scoring script and a short report on your starting prototype. This is the backbone of the later weeks.

Project 2: Instrumented service

Traces plus a simple view of quality, latency and cost, with prompts and code under version control.

Project 3: Capstone: a hardened LLM service

A FastAPI-wrapped service with an evaluation report, a release checklist, a safety note and a short runbook. Depth matters more than a new model demo.

DATAMARCOS

Curriculum · 8 weeks

8–10 hours a week, live and mentor-led. Open each week in the live sessions for labs and feedback.

Week 1: Production failure modes

- What breaks when a notebook becomes a service: quality drift, latency, cost, prompt changes
- A reference architecture for an LLM application
- Choosing the prototype you will harden during the program
- Writing down current behaviour before you change it

Lab: Run your prototype against a first slice of representative tasks and record failures.
You ship: A short baseline report of quality issues.

Week 2: Evaluation design

- Golden sets, task types and labels a team can maintain
- Automatic checks versus human review, and when each is enough
- Regression tests so a prompt change cannot silently land
- Reporting results so a product partner can read them

Lab: Build a small golden set and a script that scores your prototype.
You ship: An evaluation harness and a first score report.

Week 3: Observability: traces, latency, cost and quality

- Tracing prompts, retrieval and tool calls
- Basic dashboards for latency, cost and quality over time
- Sampling production traffic for review
- What to alert on — and what is just noise

Lab: Add tracing to your prototype and review a handful of traces in class.
You ship: Instrumented runs with a simple quality and latency snapshot.

Week 4: Prompt and release management

- Versioning prompts and application code together
- Environments and a simple release bar
- Reviewing a change before it reaches users
- Rollbacks when quality drops

Lab: Put your prompts in version control and run a pre-release evaluation.
You ship: A documented release checklist.

Week 5: RAG evaluation and retrieval quality

- Measuring retrieval separately from the final answer
- Chunking, indexing and query changes as hypotheses you can test
- Common retrieval failures in production data
- A small experiment log

Lab: Score retrieval on a labelled subset and try one improvement.
You ship: A retrieval evaluation note with before-and-after scores.

Curriculum (continued)

Week 6: Safety, guardrails and red-teaming basics

- Prompt injection, data leakage and over-confident answers, in practical terms
- Lightweight guardrails and fallback behaviour
- A short red-team pass on your own service
- What still needs a human review step

Lab: Run a short red-team script and record issues.
You ship: A safety note and a list of residual risks.

Week 7: Serving with FastAPI, environments and CI

- Wrapping the application in a simple FastAPI service
- Configuration, secrets and environment separation at a practical level
- A basic GitHub Actions path for tests and a build
- Docker as a packaging option

Lab: Expose an endpoint and run your evaluation suite in CI.
You ship: A served prototype with a passing test job.

Week 8: Capstone: instrument, evaluate and harden

- Assembling the evaluation report, traces, safety note and runbook
- Live capstone walk-throughs and mentor feedback
- What you would do next in a real team

Lab: Present the capstone service and its evaluation evidence.
You ship: Capstone service, evaluation report and a short operations note.

Format, fees and refunds

Part-time and live online, in a small mentor-led cohort. Weekend-morning batches are offered for learners in the US. Evening batches (India time) are offered for learners in the Gulf and Africa. Europe-oriented groups typically meet after work and on Saturdays. Recordings are shared if you miss a live session. Exact days, times and the next cohort date are confirmed on a counsellor call.

Designed with the US and the UK in mind; join live from anywhere.

A Datamarcos certificate of completion is awarded when you finish the program work. Cohort start dates are shared on a counsellor call — we do not publish a date until then.

Fees and EMI options shared on a counsellor call.

7-day refund policy: if the program is not right for you, ask within 7 days of enrollment.

Mentor

Live sessions are mentor-led by Deepesh Mishra, founder of Datamarcos, a BITS Pilani alumnus, and a solution architect and corporate trainer in data and AI who has trained teams at 200+ organizations.

DATAMARCOS

Questions, answered

Do I need prior LLM experience?

Yes. You should already be comfortable in Python and have built, or be able to bring, an LLM, RAG or agent prototype. This is a production follow-on, not a first course in generative AI. Applied Agentic AI or GenAI & Agentic AI Engineer are better starting points if you are new to this work.

How long is the program?

Eight weeks, part-time and live online, at about 8–10 hours a week including live sessions, labs and your capstone work.

What will I build?

You will harden a prototype: an evaluation harness, tracing, a release checklist, retrieval and safety notes, and a FastAPI-wrapped capstone service with a short runbook.

Are the sessions live?

Yes. Sessions are live and mentor-led, with recordings if you miss a class. Weekend-morning batches are offered for learners in the US, with other regional times available. Exact times are confirmed on a counsellor call.

Which tools are covered?

Python, pytest, FastAPI, Docker, GitHub Actions, and an observability stack such as LangSmith or Langfuse, plus OpenTelemetry. You work on a prototype you bring or rebuild, rather than a single vendor platform.

What about fees?

Fees and EMI options shared on a counsellor call.

Is there a refund policy?

Yes. Datamarcos has a 7-day refund policy. If the program is not right for you, ask within 7 days of enrollment.

DATAMARCOS

Talk to a counsellor

Share a few details on datamarcos.com and a counsellor will talk through program fit, the next cohort, fees and enrollment.

Fees and EMI options shared on a counsellor call.

7-day refund policy: if the program is not right for you, ask within 7 days of enrollment.

Noida One, 303-304, Rasoolpur Nawada, Industrial Area, Sector 62, Noida, Uttar Pradesh 201309, India

www.datamarcos.com · hello@datamarcos.com · +91 88519 30644

This brochure does not promise placement, salary or promotion outcomes. It is not legal, immigration or official-government certification advice.