

ADVANCED AI ENGINEERING PROGRAMME

AI & **Agentic AI** Engineer Programme

A sixteen-week, engineering-focused path from LLM fundamentals to production multi-agent systems
- built around what companies are actually hiring AI Engineers to design, ship and operate in 2026.

Build Real Agents

Every learner ships tool-using, multi-step agents – not chatbot demos.

Multi-Agent & Protocols

MCP, A2A and orchestration frameworks are core curriculum, not an afterthought.

Live, Mentor-Led

Taught personally by a practising Senior Data Scientist and AI/ML Programme Leader.

Six Capstone Tracks

Choose a capstone in your own domain, from a coding agent to enterprise RAG.

16

Weeks

80

Sessions

1.5 hrs

Per session

Mon-Fri

Live schedule

₹30-60L+

Experienced Agentic AI Engineers in India can earn this much annually, with entry-level roles starting around ₹8-15 LPA – **one of the fastest-growing, best-paid specialisations in technology right now.**

Prerequisite: working knowledge of Python (functions, loops, basic object-oriented programming). No prior machine learning or LLM experience is required.

Learn directly from someone who has **done** the work



Aravind Reddy Dudipala

Founder, Unicus Learning
Senior Data Scientist & AI/ML
Programme Leader

11 Years Experience

Berkeley Haas ML/AI Instructor

**Imperial College Business School
Instructor**

Aravind founded Unicus Learning and personally leads its training programmes in data analytics, data science, Power BI and agentic AI for working professionals. He brings **11 years of experience** as a Senior Data Scientist and AI/ML Programme Leader, and works with international universities, teaching Professional Certification courses at **Berkeley Haas** and the **Imperial College Business School**. Alongside this, he mentors working professionals across the globe in AI, machine learning, data science and data analytics.

"Agentic AI engineering sits closer to systems engineering than to prompt experimentation. I teach it that way - architecture, evaluation and production judgment first, not just which framework is trending this month."

That philosophy is why this programme treats multi-agent orchestration, evaluation and deployment as core curriculum from week one, and why the capstone requires a working, deployed agent rather than a notebook demo.

Every batch - up to 50 learners - is taught personally and live, session by session. No handoff to a teaching assistant.

Select sessions are also supported by visiting industry faculty with hands-on experience shipping production AI agent systems.

Engineering-first, **not** prompt-first

Real 2026 AI Engineer job postings read like a systems-engineering brief, not a prompting tutorial: LangChain, LangGraph, RAG pipelines, vector databases, multi-agent orchestration, FastAPI microservices, cloud deployment, and production monitoring for cost and latency. **This programme is built directly from that brief.**

That's why deployment, evaluation and observability get dedicated weeks instead of a closing lecture, and why multi-agent protocols like MCP and A2A are taught as core skills – they're already appearing as requirements in current job listings, not just nice-to-haves.

"Agentic AI engineering requires the mindset of a developer, the caution of a reliability engineer, and the judgment to know when a system needs a human in the loop."

The result mirrors how a working AI Engineer's project actually goes: LLM and prompting fundamentals first, then retrieval and single-agent tool use, then multi-agent systems, and finally the evaluation, safety and deployment skills that turn a prototype into something a business can run.

Who this is for

Software developers, data analysts and data scientists, and technical professionals ready to specialise in building and shipping production AI agent systems.

PREREQUISITE

Working knowledge of Python – functions, loops, basic object-oriented programming – is required. No prior machine learning or LLM experience is necessary; this programme builds those foundations from week one.

WHAT YOU WILL BE ABLE TO DO

- Engineer reliable prompts with **function calling** and structured outputs
- Build **RAG pipelines** with vector databases and hybrid retrieval
- Build tool-using agents with **LangChain and LangGraph**
- Design **multi-agent systems** using MCP and A2A communication patterns
- Evaluate agent quality with proper **eval and observability** tooling
- Add safety guardrails and human-in-the-loop checkpoints to autonomous systems
- Deploy and monitor an agent as a **live, production API**
- Show a deployed, end-to-end **capstone agent** – chosen from six real-world tracks – in interviews

Sixteen weeks, **nine** phases

01

Python & LLM Foundations

WEEKS 1-2

Advanced Python for AI systems, working with LLM APIs from OpenAI, Anthropic and open models, and the fundamentals of how LLMs actually behave.

Python

LLM APIs

FastAPI

OUTCOME Call, control and reason about an LLM's behaviour like an engineer, not a hobbyist

02

Prompt Engineering & Structured Outputs

WEEK 3

Advanced prompting patterns, function calling, structured outputs, prompt caching, and the basics of cost and latency tradeoffs.

Function Calling

Structured Outputs

OUTCOME Get an LLM to return reliable, machine-usable output every time

03

RAG & Retrieval Systems

WEEKS 4-5

Embeddings, vector databases, RAG architecture, hybrid retrieval and reranking – the retrieval layer nearly every agentic system depends on.

Vector Databases

RAG

Reranking

OUTCOME Ground an LLM's answers in your own data, with citations

04

Agent Frameworks & Tool Use

WEEKS 6-8

Building agents that plan, call tools, and maintain memory using LangChain and LangGraph, with LlamaIndex for retrieval-heavy workflows.

LangChain

LangGraph

LlamaIndex

OUTCOME Build a single agent that reliably completes a multi-step task

05

Multi-Agent Systems & Protocols

WEEKS 9-10

Orchestrating teams of agents that coordinate on a shared goal, using the Model Context Protocol (MCP) and Agent2Agent (A2A) communication patterns now appearing in job listings.

MCP

A2A

Orchestration

OUTCOME Design a system where multiple agents coordinate without stepping on each other

Phases **six** to nine

06

Evaluation, Observability & Safety

WEEKS 11-12

Building eval pipelines that track accuracy, latency and cost, observability tooling, and safety layers – guardrails, prompt-injection defence, human-in-the-loop checkpoints.

Evaluation

Observability

Guardrails

OUTCOME Prove an agent works – and catch it when it doesn't

07

Deployment & AgentOps

WEEKS 13-14

Sandboxed execution, containerising and deploying an agent as a live API, and monitoring cost, latency and behaviour once it's running in production.

Docker

Deployment

Monitoring

OUTCOME Ship an agent as a live, monitored service instead of a local demo

08

AI Engineering Communication

WEEK 15

Presenting agentic systems to technical and non-technical stakeholders, documentation practices, and running a convincing technical demo.

Documentation

Presentation

OUTCOME Explain what your agent does, and doesn't do, with confidence

09

Capstone Project

WEEK 16

One agentic system, start to finish, from a choice of six real-world tracks: design, build, evaluate and deploy – the centrepiece of your portfolio.

End-to-end project

OUTCOME Walk into interviews with a deployed, working agent you can demo live

Built for the **systems**-engineering side of AI

185

CODING ASSIGNMENTS

25

MODULE-END
ASSIGNMENTS

7

MINI PROJECTS

1

DEPLOYED CAPSTONE
AGENT

CAREER SUPPORT, BUILT IN

One-on-One Mentoring

Direct, personal guidance from your mentor throughout the programme – not a shared help desk.

Mock Interviews

Practice rounds structured around how AI Engineer interviews actually run, with real feedback.

Resume Preparation

A resume built around your actual agent projects and deployments, not a generic template.

Interview Preparation

Structured practice for both technical and HR rounds, so nothing on interview day is unfamiliar.


+ DSA

+ Aptitude

ADD-ON

Dedicated DSA (Data Structures & Algorithms) and Aptitude sessions are provided as part of interview preparation – useful for learners also targeting broader technical interview rounds.

Eleven skills, **one** stack



WHAT YOU WALK AWAY WITH

- A **deployed, portfolio-ready capstone agent** across RAG, tools and orchestration
- Confidence building **multi-agent systems** with MCP and A2A
- The ability to evaluate and explain agent behaviour to a non-technical stakeholder
- A skill set matched to **what AI Engineer job postings actually require**

CAPSTONE

Every learner finishes with a deployed, monitored agent – not just a notebook. **Six real-world tracks** to choose from, overleaf.

A project at **every** milestone

MILESTONE PROJECTS ALONG THE WAY

WEEK 3	Prompt engineering & structured output project Build a reliable function-calling pipeline that returns clean, structured output.
WEEK 5	RAG retrieval system project Build a grounded retrieval pipeline over a real document set, with citations.
WEEK 8	Single-agent tool-use project Build an agent that plans, calls tools, and completes a multi-step task.
WEEK 10	Multi-agent orchestration project Coordinate a small team of agents on a shared goal using MCP.
WEEK 12	Agent evaluation & observability project Build an eval suite that scores an agent on accuracy, latency and cost.
WEEK 14	Deployed agent with guardrails Package, containerise and deploy an agent as a live, monitored API.

CHOOSE YOUR CAPSTONE TRACK - WEEK 16

At Unicus, every learner has **the freedom to choose a capstone agent** aligned to their own domain and interests, and **we actively encourage learners to pursue it** - the tracks below are a starting point, not a limit.

TRACK 01

Autonomous coding agent

An agent that reads a codebase, proposes changes, and validates them by running tests.

Python · LangGraph · Sandboxed Execution

TRACK 02

Enterprise RAG knowledge assistant

A grounded Q&A agent over a private document set, with citations and confidence scoring.

RAG · Vector DB · Reranking

TRACK 03

Multi-agent workflow orchestrator

A team of agents that coordinate to complete a multi-step business process end to end.

MCP · A2A · Orchestration

TRACK 04

Browser automation agent

An agent that navigates real websites and completes multi-step tasks autonomously.

Browser Agents · Tool Use · Python

TRACK 05

Customer support agent with escalation

A support agent that resolves common issues and knows when to hand off to a human.

LangChain · Guardrails · APIs

TRACK 06

Research & report assistant agent

An agent that researches a topic across sources and produces a structured report.

RAG · Tool Use · Structured Output

PROGRAMME FORMAT

Live & Mentor-Led

16 weeks - Mon to Fri - 1.5 hrs/day

Enrol now

7013265897 | enquiry@unicuslearning.com
Hyderabad, India | www.unicuslearning.com