Earn
Learn

GEN AI FOR
TECH PROFESSIONALS

AI — The Flight to your Dreams

DURATION	MODE	LOCATION
3 Months	Online & Offline	Pune + Online

Agentic AI · RAG · MCP · Multi-Agent Orchestration

Contents

01

The AI Industry Needs You

The AI industry needs professionals who can build efficient, ethical, and scalable systems while applying AI to real-world problems.

$665.7B
Big data & analytics market by 2030

27.6%
CAGR — global data analytics 2023–2030

$4.4T
Value added by GenAI to global economy

$667.9B
GenAI market 2030 (47.5% CAGR)

Rs. 9.57L+
Avg AI salary in India

10M+
AI jobs globally by 2027

Why Upskill Now?

The window to differentiate as an AI-capable engineer is **now**. RAG, Agentic AI, MCP, and Multi-Agent Systems are the four capabilities defining the next generation of AI infrastructure — and demand far exceeds supply.

Organisations worldwide are racing to build AI-powered products and automate workflows. Tech professionals who can architect and deploy these systems are the most sought-after talent in the global market.

02 Course Summary

Go beyond the basics and become a proficient Generative AI practitioner. This hands-on program is designed for developers, engineers, and tech professionals.

Course Duration	3 Months · 2-hour sessions, twice weekly
Delivery Mode	Online Live Sessions · Offline Face-to-Face (Pune)
Instructor Support	Live sessions, Q&A, doubt-solving & placement assistance
Format	Hands-on projects, real datasets, industry tools

03 Programme Highlights

1. **RAG — Retrieval-Augmented Generation**

   Build knowledge-grounded AI systems from your own documents — semantic chunking, vector stores, and context compression.

2. **MCP — Model Context Protocol**

   Connect AI agents to databases, APIs, and tools via standardised FastMCP servers — giving your agents the ability to act in the real world.

3. **Agentic AI — Autonomous AI Agents**

   Design agents that reason, plan, remember, and self-correct using ReAct loops and stateful memory architectures.

4. **Multi-Agent Orchestration**

   Build teams of specialised agents collaborating via LangGraph Directed Cyclic Graphs for complex parallel tasks.

04 What Will You Learn?

Module 1–2 Conversational & Multi-Agent AI

- **Build multi-turn conversational agents**

  Design production-ready chatbots with persistent context using the OpenAI SDK. Handle complex dialogue flows, manage state across turns and support dynamic user interactions at scale.

- **Function calling, structured outputs & tool use**

  Master function calling to connect LLMs to real-world tools and APIs. Enforce structured JSON outputs to ensure predictable, reliable data for downstream processing and system integration.

- **MCP tool servers with FastMCP**

  Set up Model Context Protocol (MCP) servers using FastMCP to expose tools, resources and prompts in a standardised, interoperable way — enabling any compliant AI client to use your services.

- **Connect agents to external APIs & databases**

  Integrate agents with REST APIs, SQL and NoSQL databases and third-party services. Build data-aware, action-taking systems that query live sources and write results back — end to end.

Module 3–4 Agentic AI & Model Context Protocol

- **ReAct reasoning for autonomous decision-making**

  Implement the Reasoning + Acting (ReAct) pattern so agents plan, act, observe tool outputs and self-correct — all without human intervention. Build agents that tackle multi-step tasks reliably.

- **Stateful short-term & long-term memory**

  Architect agents with working memory (conversation buffers) and persistent long-term memory backed by vector stores. Agents remember context across sessions, enabling continuity and personalisation.

- **Self-correction loops & output validation**

  Add production-grade guard-rails: automated evaluation chains, retry logic and structured validators that detect and fix hallucinations and schema violations in real time before they reach users.

- **Expose agent capabilities via MCP**

  Publish your agent's tools and resources as MCP endpoints — making them accessible to Claude, GPT and any MCP-compatible client. Enables seamless interoperability across the entire AI ecosystem.

What Will You Learn? (continued)

Module 5–6

Advanced RAG Systems

- **Semantic chunking & intelligent indexing**

  Move beyond fixed-size chunking to semantic boundary detection and hierarchical indexing. Preserve context and document structure to dramatically improve retrieval precision and answer quality.

- **Context compression (token-efficient retrieval)**

  Apply LLMLingua and similar compression techniques to shrink retrieved passages by up to 80% with minimal information loss — cutting inference costs while maintaining response accuracy.

- **Multi-source retrieval: FAISS & PostgreSQL**

  Build hybrid retrieval pipelines that query dense FAISS vector stores and structured PostgreSQL databases simultaneously, fusing results for richer, more comprehensive and factually grounded answers.

- **Hybrid BM25 + semantic similarity search**

  Combine BM25 keyword ranking with dense semantic embeddings to capture both exact-match and conceptual relevance — consistently outperforming either technique in recall, precision and F1 score.

Module 7–8

Multi-Agent Orchestration

- **Multi-agent architectures with LangGraph**

  Design supervisor–worker agent graphs in LangGraph with typed state, conditional routing and shared tool registries. Coordinate specialist agents to tackle complex, long-horizon tasks collaboratively.

- **Directed Cyclic Graphs (DCGs) for parallelism**

  Use DCG structures to run multiple agent sub-tasks concurrently, dramatically reducing end-to-end latency for compute-intensive workflows that decompose naturally into parallel branches.

- **Agent handoffs, routing & consensus**

  Implement intelligent intent-based routing so the right specialist agent handles each sub-task. Add consensus and voting mechanisms for high-stakes decisions that require multi-agent agreement.

- **Production deployment of agent systems**

  Package, containerise and deploy multi-agent systems to cloud environments. Instrument with structured logging, distributed tracing and graceful error recovery for reliable, observable production services.

05 Prerequisites & Who Should Attend

Programming	Intermediate Python — functions, classes, async/await
ML Awareness	Basic ML familiarity helpful but not required
APIs	REST API or SDK experience an advantage
Availability	8–10 hours per week commitment

06 Faculty & Advisors

At Aimllearn, we deliver in-depth AI and ML education, guided by experienced educators and business leaders with 30+ years of expertise.

SHRIRANG KARANDIKAR

PhD (USA) · Educator · Researcher · Industry Veteran

30+	7	40+	3
Years in advanced computational technologies	Patents granted	Peer-reviewed publications	Global industry giants

Intel · IBM · Shell

Shrirang is a doctorate (PhD.) from USA — a passionate educator and team-builder with over 30 years of expertise at Intel, IBM, and Shell. He holds 7 patents and has contributed to 40 peer-reviewed publications spanning ML, AI systems architecture, and large-scale computing.

> *"Bridge theory with real-world practice — every concept taught has a production use case behind it."*
>
> — Shrirang Karandikar, PhD

"I went to meet Dr. Karandikar sir — 35 years of experience is not a small deal. If you're serious about AI, there's no better place than AIML."
— Venkatesh Shaligram, B.Tech MIT Pune

Earn

Learn

START YOUR AI JOURNEY

AI — The Flight to your Dreams

Don't just use AI tools — build the next generation of them. Join a cohort of ambitious tech professionals and emerge as an expert in Agentic AI, RAG, MCP, and Multi-Agent Systems.

Agentic AI · RAG · MCP · Multi-Agent Orchestration

Industry-Ready Skills

Real agentic AI, RAG pipelines & multi-agent platforms.

Expert Guidance

PhD trainer, 30+ yrs at Intel, IBM & Shell. 7 patents, 40+ papers.

Flexible Learning

Online live + Offline Pune. Weekend & weekday batches.