

Agentic AI Architect

Duration: 24 hours

Course Overview

This course delves into the architecture, orchestration, and governance of Agentic AI systems, equipping professionals with the strategic and technical competencies needed to build enterprise-grade autonomous solutions. It begins by contrasting agentic and traditional AI, introducing core principles of autonomy and decision-making loops before moving into a comparative study of open-source models such as LangChain, AutoGPT, and BabyAGI. Participants will design and orchestrate scalable multi-agent workflows, integrate them with cloud platforms like Azure and AWS, and implement real-time monitoring via modern DevOps practices. The course also emphasizes robust evaluation strategies, responsible AI implementation—including bias audits and ethical safeguards—and deployment best practices such as CI/CD, enterprise integration, and governance. Capstone workshops focus on practical use case architecture, enabling learners to deliver modular, production-ready agent systems while mapping strategic trends in open-source innovation and cross-domain agent collaboration.

Pre-requisites

- Good Understanding of AI/ML Principles
- Intermediate level proficiency in Python
- Familiarity with cloud platforms
- Exposure to tools like Jupyter notebooks, Git and Devops tools would be helpful.

Course Objectives

- Equip architects with deep understanding of Agentic AI principles and enterprise-grade design.
- Enable evaluation and selection of open-source models for different agentic use cases.
- Teach design of scalable, cloud-ready multi-agent systems and workflows.
- Emphasize responsible AI principles, lifecycle management, and risk mitigation.
- Provide practical experience architecting real-world multi-agent and autonomous systems.

Course contents

Module 1: Foundations of Agentic AI Systems

- Characteristics vs Traditional AI
- Principles of Autonomy & Decision Loops
- Key Enterprise Application Patterns

Module 2: Comparative Study of Open-Source Models

- Language, Vision, Multi-Modal, RL Model Categories
- LangChain, AutoGPT, BabyAGI Capabilities
- Model-Use Fit Analysis

Module 3: System Architecture & Orchestration

- Multi-Agent Workflow Design
- Memory, Tools, Inter-Agent Messaging
- LangChain vs AutoGPT vs Custom
- Architecture workshop

Module 4: Scalability & Cloud Integration

- Cloud Architectures (Azure, AWS, GCP)
- Container Orchestration (Docker, K8s)
- Real-Time Monitoring Pipelines
- Lab: Cloud Design

Module 5: Evaluation & Lifecycle Management

- Quantitative & Qualitative Evaluation
- Continuous Learning & Feedback Loops
- MLflow for Tracking
- Lab: Evaluation Strategy

Module 6: Responsible AI & Risk Mitigation

- Fairness, Explainability, Bias Audits
- Fail-Safes for Agentic Systems
- Ethical Design Patterns
- Design Audit Workshop

Module 7: Deployment Best Practices

- Environments, CI/CD, Blue/Green Deployments
- Enterprise Integration
- Data Governance & Access Control
- Deployment Planning

Module 8: Strategic Roadmapping & Future Outlook

- Cross-Domain Agent Collaboration
- Future Innovations
- Open-Source Trends & Partnerships

Module 9: Architecting Use Cases (Workshops)

- Customer Support Agent System Design
- Vision-Based Automation Architecture
- Modular Multi-Agent Pipelines