Agentic Intelligence in Practice: Building Autonomous AI Systems

Duration: 40 Hours

Course Overview

This hands-on course offers a practical deep dive into Generative and Agentic AI, guiding participants through the design and deployment of autonomous, reasoning-capable systems using both proprietary models like GPT-4 and open-source tools such as LLaMA 2, CLIP, and LangChain. Learners will explore multi-modal integration across text, vision, and audio, build intelligent chatbots, orchestrate agents with LangGraph, and manage experimentation with MLflow. With a strong focus on Azure OpenAI and scalable cloud deployment, the course equips professionals to create real-world AI solutions that reason, adapt, and operate autonomously across diverse domains.

Pre-requisites: Functional Knowledge of Python

Note: Students would be required to bring their own Azure Subscription to perform labs

---

Course Contents

Module 1. Introduction to Generative & Agentic AI

- Generative AI vs Agentic AI: Definitions and Key Differences
- Capabilities: Reasoning, Autonomy, Multi-Modal Understanding
- Ethical considerations & responsible AI practices

Module 2. Azure OpenAI & OpenAI Models

- Overview of Azure OpenAI service architecture & integration
- GPT-3.5, GPT-4, GPT-4 Turbo, GPT-5 (nano, mini, standard)
- Model selection criteria: latency, cost, accuracy, and context length
- Azure-specific deployment and governance features

Module 3. Open-Source Model Landscape

- Language Models: LLaMA 2, Mistral, Falcon, BLOOM
- Vision Models: YOLOv8, OpenCV, Mediapipe
- Multi-Modal Models: CLIP, Whisper, DALL-E Mini, Perceiver
- Comparative analysis: Proprietary vs Open-Source models
- Use-case mapping: Text, vision, audio, structured data

Module 4. Reasoning-Capable Models

- Chain-of-Thought (CoT) prompting

- Tool-augmented reasoning frameworks (ReAct, MRKL)
- Using reasoning in autonomous workflows
- Integration of reasoning with LangChain & LangGraph

Module 5. Multi-Modal AI

- Processing and integrating text, vision, audio, video, and structured data
- Cross-modal understanding and retrieval
- OpenAI multi-modal capabilities (e.g., GPT-4o, GPT-4.1)
- Case Study: Cross-modal data analysis for business insights

Module 6. LangChain Framework

- Core concepts: Chains, Agents, Memory, Tools
- Building autonomous workflows with LangChain
- LangChain for reasoning & decision-making tasks

Module 7. LangGraph for Agent Orchestration

- LangGraph architecture: Nodes, edges, state machines
- Designing and orchestrating multi-agent workflows
- Integrating LangGraph with LangChain for complex systems

Module 8. Building Chatbots with Agentic AI

- Using GPT & open-source models (e.g., LLaMA 2, Open-Assistant) for conversational agents
- Multi-turn conversations and workflow automation with LangChain
- Adding reasoning and tool use in chatbots via LangGraph
- UI integration with Gradio or Streamlit
- **Hands-On:** Build a task-specific chatbot with Azure OpenAI + LangChain

Module 9. Experimentation & Model Management with MLflow

- Introduction to MLflow for model tracking and experimentation
- Logging metrics, parameters, and artifacts for AI systems
- Fine-tuning, evaluation, and versioning with MLflow

Module 10. Azure Cloud Deployment

- Deploying GPT & open-source models on Azure
- Integration with Azure Functions, AKS, and Container Apps

- Securing endpoints & scaling deployments
- Monitoring and maintaining deployed AI systems

Module 11. Hands-On Project – Case Study

- Text-based autonomous customer support agent
- Multi-modal autonomous data entry assistant
- Vision-based inventory management system
- End-to-end implementation workflow

Module 12. Future Trends in Generative & Agentic AI

- Advances in GPT-5 reasoning & multi-modal capabilities
- Evolution of open-source ecosystems
- AI agent collaboration & self-improving systems
- Integration with IoT, robotics, and edge computing
