Building Agentic AI Systems with Open-Source Models

Course Duration: 48 Hours **Course code: BA-AI-SOSM**

1. Course Overview

The Building Agentic AI Systems with Open-Source Models course is designed to help learners understand, design, and implement autonomous and semi-autonomous AI agents using open-source large language models (LLMs) and modern AI frameworks. This course focuses on agentic workflows, tool-using agents, reasoning, memory, planning, and orchestration, enabling learners to build intelligent systems that can act, decide, and collaborate with minimal human intervention.

2. What you'll learn?

By the end of this course, learners will be able to:

- Understand agentic AI concepts and architectures
- Work with open-source LLMs (e.g., LLaMA, Mistral, Falcon)
- Design autonomous, tool-using AI agents
- Implement reasoning, planning, and memory mechanisms
- Build multi-agent systems and collaborative workflows
- Integrate agents with APIs, databases, and tools
- Apply retrieval-augmented generation (RAG) in agent systems
- Evaluate, monitor, and secure agentic AI applications
- Deploy agentic AI systems in real-world environments

3. Target Audience

This course is ideal for:

- AI/ML engineers and developers
- Data scientists transitioning to agentic AI
- Software engineers building intelligent applications
- Automation and DevOps professionals

- Product architects and technical leads
- Researchers and advanced students in AI
- Startups and innovation teams working with open-source AI

4. Pre-Requisites

To get the most out of this course, learners should have:

- Basic understanding of Python programming
- Familiarity with machine learning or NLP concepts
- Knowledge of APIs and RESTful services
- Basic experience with Linux and Git (recommended)

5. Course content

Module 1: Introduction to Agentic AI

- What is agentic AI?
- Autonomous vs reactive systems
- Agent architectures and design patterns
- Use cases and industry applications

Module 2: Open-Source LLM Ecosystem

- Overview of open-source language models
- Model selection and benchmarking
- Running models locally vs in the cloud
- Inference optimization techniques

Module 3: Foundations of AI Agents

- Agent components: perception, reasoning, action
- Prompt engineering for agents
- Tool calling and function execution
- Agent state and lifecycle management

Module 4: Reasoning, Planning, and Memory

- Chain-of-thought and structured reasoning
- Planning algorithms for agents
- Short-term and long-term memory design
- Vector databases and embeddings

Module 5: Tool-Using and API-Integrated Agents

- Integrating external tools and APIs
- Web search, code execution, and data tools
- Error handling and fallback strategies
- Secure tool access and permissions

Module 6: Retrieval-Augmented Agent Systems

- RAG architecture for agentic systems
- Document ingestion and indexing
- Context management and relevance ranking
- Improving accuracy and grounding

Module 7: Multi-Agent Systems

- Agent collaboration and communication
- Task decomposition and coordination
- Supervisor and worker agent patterns
- Conflict resolution and consensus

Module 8: Frameworks and Orchestration

- Agent frameworks overview (LangGraph, CrewAI, AutoGen, etc.)
- Workflow orchestration and state graphs
- Event-driven agent systems
- Scaling and performance considerations

Module 9: Evaluation, Safety, and Governance

- Evaluating agent performance and reliability
- Guardrails, safety, and alignment
- Observability and logging
- Ethical considerations and compliance

Module 10: Deployment and Real-World Applications

- Deploying agentic systems to production
- Cost optimization and monitoring
- Case studies: enterprise automation, research agents, copilots
- Capstone project: end-to-end agentic AI system