📘 Module 1: Introduction to Agentic AI

- What is Agentic AI?
- Evolution from Rule-Based Systems to Autonomous Agents
- Key Concepts: Agents, Goals, Planning, Autonomy
- LLMs vs. Agentic AI: What's the Difference?
- Real-world use cases and applications
- Demo: Simple Task Agent using GPT + Python

🧠 Module 2: Core Building Blocks of Agentic AI

- Agent Architectures: Reactive, Deliberative, Hybrid
- Intent Recognition and Goal Specification
- Memory and State Handling
- Tools & APIs Integration
- Planning Algorithms: Simple vs. LLM-based Planning
- Demo: Build an Intent-to-Action Mapping Agent

⚙️ Module 3: Tools and Ecosystem

- LangChain, AutoGen, ReAct, CrewAI, n8n
- Vector Databases: Pinecone, Weaviate, ChromaDB
- Plugin Ecosystem: APIs, Zapier, Webhooks
- Integrating Web Search, Files, Emails
- Demo: Using LangChain with a Search Tool

🧩 Module 4: Agent Memory and State Management

- Short-Term vs. Long-Term Memory
- Using ChromaDB / Weaviate for memory persistence
- Embeddings and Retrieval-Augmented Generation (RAG)
- Multi-turn Conversation Context Handling
- Demo: AI Chatbot with Searchable Memory

🤖 Module 5: Multi-Agent Systems and Agent Collaboration

- When and Why Multi-Agent Systems?
- Architectures: Parallel vs. Hierarchical vs. Role-based
- Task Distribution and Communication Between Agents
- CrewAI or AutoGen for Multi-Agent Flow
- Demo: Writer + Researcher + Critic Agent Flow

🔄 Module 6: Workflow Orchestration and Autonomous Loops

- Self-looping Agent: Monitor → Plan → Act → Reflect
- Self-healing and error-handling strategies
- Task Prioritization and Goal Re-evaluation
- Autonomy vs. Control Balance
- Demo: Build a Self-updating News Summary Bot