Agentic AI Developer

Duration: 24 hours

Course Overview

This comprehensive course offers a comprehensive journey into Agentic AI, guiding learners through the foundational concepts, architectural frameworks, and real-world applications of autonomous agents. Participants will explore cutting-edge tools such as Hugging Face Transformers, LangChain, and Llama 2 for language model orchestration; dive into agent planning with AutoGPT, BabyAGI, and reinforcement learning frameworks like Stable-Baselines3; and develop vision and multi-modal capabilities using YOLOv8, CLIP, and Whisper. The program emphasizes practical implementation through interactive modules on chatbot development, fine-tuning open-source models, and building intuitive UIs with Gradio and Streamlit—culminating in capstone projects that apply agentic intelligence to text-based customer support and multi-modal data entry. Integration with Azure pipelines, experiment tracking with MLflow, and Docker-based deployment ensure that learners build production-ready, scalable agent systems.

Pre-requisites

- Foundational Understanding of AI/ML
- Intermediate level proficiency in Python
- Exposure to tools like Jupyter notebooks, Git would be helpful.
- Understanding of Chatbot Workflows

Course Objectives

- Introduce participants to core Agentic AI concepts and architectures.
- Enable practical skills in working with open-source language models and orchestration frameworks.
- Build capabilities to design decision-making agents, vision/multimodal agents, and chatbots.
- Equip learners to evaluate, fine-tune, and deploy agentic AI models effectively.
- Empower participants to build and present complete agentic solutions via capstone projects.

Course contents

Module 1: Introduction to Agentic AI

- Overview of Agentic AI Concepts
- Key Use Cases and Application Areas
- Architecture and Component Overview

Module 2: Working with Open-Source Language Models

- Hugging Face Transformers (Bloom, Falcon)
- LangChain for LLM Orchestration
- Open-Assistant & Llama 2 Setup
- Hands-on : Model setup & orchestration

Module 3: Building Decision-Making Agents

- AutoGPT and BabyAGI Fundamentals
- Stable-Baselines3 & RLlib
- Planning with PDDL Solvers
- Hands-on: Agent planning & RL tasks

Module 4: Vision and Multi-Modal Agent Capabilities

- YOLOv8 and OpenCV/Mediapipe for Vision Tasks
- CLIP for Text-Image Understanding
- Whisper for Audio Processing
- Vision and Audio demos

Module 5: Interface Development & Integration

- Gradio and Streamlit for Agent UI
- Dockerizing Agentic AI Applications
- Azure Deployment Pipelines
- UI & deployment setup

Module 6: Agentic Chatbots

- Conversational Workflow with LangChain
- Multi-turn Dialogue Design
- Hands-On: Build a Use-Case Specific Chatbot
- Hands-on: Chatbot build

Module 7: Model Evaluation & Optimization

- Using MLflow for Experiment Tracking
- Hyperparameter Tuning & Model Versioning
- Integrating Evaluation into Agent Workflow
- Hands-on: Tracking & tuning

Module 8: Fine-Tuning Open-Source Models

- Hugging Face Trainer, PEFT, and Dataset Prep
- Custom Agent Behaviors Through Model Tuning
- Model fine-tuning

Module 9: Capstone Projects

- Text-Based Autonomous Customer Support Agent
- Multi-Modal Data Entry Agent