

Complete Generative AI: From Basics to Expert Level

Duration: 80 hours (10 Days)

Course Outcomes

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

- Understand and apply core machine learning, deep learning, and transformer-based architecture
- Build, fine-tune, and deploy generative AI models using open-source and cloud-based tools
- Design effective prompts and preprocess unstructured data for retrieval-augmented generation (RAG) applications
- Develop intelligent AI agents, chatbots, and workflows using LangChain, LangGraph, and Microsoft Agent Framework
- Implement scalable, responsible, and production-ready generative AI solutions using modern LLM Ops best practices

Pre-requisites

- Basic Python or Object-Oriented programming knowledge
- Foundational ML knowledge
- Experience with Jupyter Notebooks or Python IDEs

Course Agenda

Chapter 01: Foundations of Machine Learning & Deep Learning

- Overview of Artificial Intelligence
- Overview of Machine Learning: Supervised, Unsupervised, Reinforcement Learning
- Deep Learning Concepts for Generative AI Models
- Introduction to GAN
- Overview of Generative AI
- Introduction to Transformers

Chapter 02: Open-Source Models with Hugging Face

- Explore Hugging Face Hub for open-source models
- Use the transformers library for text, image, and multimodal tasks
- Share apps via Gradio and Hugging Face Spaces
- Hands-on: Deploy your first Hugging Face AI app

Chapter 03: Prompt Engineering

- Understand Prompting Engineering
- Key Terminologies
- Elements of Prompting
- Prompt System Design
- Tips for Prompting
- Generative AI Tools for Prompting
- Lab: Prompting Techniques

Chapter 04: Preprocessing Unstructured Data for LLM Applications

- Extract data from PDFs, PowerPoints, Word, HTML, images, and tables

- Add metadata for enhanced RAG
- Apply layout detection, vision, and table transformers
- Hands-on: Preprocess enterprise data for RAG

Chapter 05: Finetuning Large Language Models

- What is Fine Tuning
- Pre-trained LLM vs Fine-Tuned LLM
- What is Instruction Fine-Tuning
- Data preparation for Fine-Tuning
- Approach to Fine-Tuning
- Hands-on: Fine-tune a Hugging Face model (QWEN)

Chapter 06: LangChain: Building Applications

- Introduction to LangChain
- Lab: Build a semantic search engine with LangChain
- Lab: Build a RAG agent with LangChain
- Lab: Build a SQL agent
- Lab: Build a Deep Research Agent

Chapter 07: The Evolution of Agentic Orchestration

- The Evolution of Agentic Orchestration (moving from LangChain to LangGraph)
- Lab: Build a Personal Assistant using LangChain and LangGraph (to understand difference between both)

Chapter 08: LangGraph for Agent Design

- Lab: Build a multi-source knowledge base with routing
- Lab: Build a custom RAG agent with LangGraph
- Lab: Workflows and Agents in LangGraph
- Lab: Build and run SQL Agent in LangGraph Studio

Chapter 09: Develop Agents in Microsoft Foundry

- Understanding AI agent development on Azure
- Develop an AI agent with Foundry Agent Service
- Integrate custom tools into your agent
- Develop a multi-agent solution with Foundry Agent Service
- Integrate MCP Tools with Foundry Agents

Chapter 10: Microsoft Agent Framework (Open Source) for AI Application Development

- Primitive to Microsoft Agent Framework (Semantic Kernel and Autogen)
- Lab: Semantic Kernel Basics
- Lab: Understanding Autogen Basics
- Overview of Microsoft Agent Framework
- Lab: Foundations of AI Agents with Microsoft Agent Framework and Ollama
- Lab: Advanced Multi-Agent Retrieval-Augmented Generation with Ollama