Applied Generative AI with Python

Duration: 56 hrs

Course Overview

This course offers a robust foundation in Python programming, machine learning fundamentals, and cutting-edge generative AI systems. In Section I, learners begin by brushing up on Python essentials—from syntax and data structures to control flow, NumPy, Pandas, and data visualization—before progressing into core ML concepts, frameworks like Scikit-learn and TensorFlow, and a mini project focused on classification, regression, and clustering. Section II transitions into the world of Generative AI, covering foundational architectures, large language models (LLMs), prompt engineering, and Retrieval-Augmented Generation (RAG) systems. Participants will explore agentic AI frameworks such as LangChain, LangGraph, and CrewAI, alongside practical labs and demos using tools like Ollama and Azure AI Foundry. The course concludes with advanced fine-tuning techniques and a critical focus on ethical and responsible AI practices, preparing learners to build, evaluate, and deploy intelligent systems with confidence and accountability.

Course Prerequisites

Python, Machine learning and Deep learning (conceptual)

Course Contents

Section I: Brushing upon Python and ML Fundamentals

- **Module 1: Introduction to Python for ML**
    - Why Python?
    - Environment setup (Jupyter/Colab)
    - Basic syntax, variables, operations
- **Module 2: Python Data Structures**
    - Lists, Tuples, Dictionaries
    - Indexing, slicing, iteration
- **Module 3: Control Flow & Functions**
    - If-else, loops
    - Defining functions
- **Module 4: NumPy & Pandas**
    - Arrays, slicing
    - Multi-dimensional operations
    - DataFrames, Series
    - Missing values, grouping, merging
- **Module 5: Data Visualization**

- Matplotlib & Seaborn
- Line, bar, histogram plots

- **Module 6: ML Concepts**
  - Supervised, Unsupervised, Reinforcement
  - Scikit-learn basics
  - Deep Learning Neural Network Understanding
  - Intro to TensorFlow, Transformers, GANs, Diffusers
- **Module 7: Pydantic Framework**
  - Core Concepts like Base Model, Type Hinding, annotation, default values etc.

  (topics aligned with langgraph state creations

- **Module 8: Mini Project (homework)**
  - Classification, regression, clustering

Section II: Generative AI

- **Module 9: Introduction to GenAI**
  - Architecture & applications
  - GANs
- **Module 10: Large Language Models (LLMs)**
  - Architecture & types
  - Major LLMs: Llama, Qwen, Cohere, Falcon
  - Text AI tasks: translation, summarization, speech
  - Image AI tasks: detection, segmentation, captioning, VQA
  - Labs
- **Module 11: Prompt Engineering**
  - Techniques for text, image, code
  - Context Engineering
  - Standard templates
  - Using Meta's Llama

    - Generative AI tools helpful for Upgradation and learning
    - Labs
- **Module 12: Basic LLM Systems (RAG)**
    - RAG concepts
    - LangChain: Intro, prompt templates, chains, agents
    - Ollama for running the model locally
    - Labs: chatbot, RAG Implementation
- **Module 13: Advanced RAG and QnA Systems**
    - RAG vs QnA
    - Tabular & SQL QnA systems
    - Advanced RAG: Agentic RAG, Multimodal RAG
    - Labs: Hands-on labs and real-world implementation
- **Module 14: Agentic AI Systems (Lab based Topics)**
    - Agent Frameworks: LangChain, LangGraph, Autogen, CrewAI, Agno.
    - Multi-agent Systems
    - MCP – Introduction and Use cases
    - No code/low code agent creation

- **Module 15: Azure AI Foundry (only demo)**
    - Retrieval-Augmented Generation (RAG) Implementation
    - Fine-tuning Models on Azure
    - Implementing Content Filters
    - Model Evaluation and Monitoring
    - Agents
    - Demo
- **Module 16: Fine-Tuning Techniques**
    - Quantization & optimization
    - Fine-tuning Gemma
    - Labs
- **Module 17: Ethical & Responsible Generative AI**
    - Ethical & Responsible Generative AI principles

- Open-source ethical frameworks
- Cost Estimation: API pricing, Model pricing
- LLMOps: Discussion
- Labs