Generative AI Course Syllabus

Section 1: Foundations (AI/ML Basics)- 8 Hours

- ➢ What is Artificial Intelligence and Machine Learning
- ➢ Supervised vs Unsupervised Learning
- ➢ Basic NLP Concepts and Applications
- ➢ AI vs GenAI: Key Differences and Use Cases
- ➢ Real-world ML Use Cases
- ➢ Lab Session: Sentiment Analysis with Classical NLP

Section 2: Python for AI- 6 Hours

- ➢ Numpy, Pandas – Essential Data Manipulation
- ➢ Regex for Text Processing
- ➢ Calling REST APIs and Handling JSON
- ➢ Web Scraping with BeautifulSoup/Requests
- ➢ Lab Session: Collecting and Analyzing Web Data for NLP Tasks

Section 3: Deep Learning & Transformers- 10 Hours

- ➢ Neural Network Basics
- ➢ Recurrent Neural Networks (RNNs)
- ➢ Transformers and Attention Mechanism
- ➢ BERT, GPT, and other Transformer Models
- ➢ Lab Session: Building a Transformer-based Text Classifier

Section 4: Generative AI Concepts- 6 Hours

- ➢ What is GenAI and Why Now?
- ➢ Architecture of LLMs (BERT, GPT, T5)
- ➢ Diffusion Models, GANs – High-level Overview
- ➢ Real-world Applications of GenAI (Text, Image, Audio)
- ➢ Lab Session: Generate Text with Pre-trained GPT Model

Section 5: Prompt Engineering – 4 Hours

- ➢ Prompt Types: Zero-shot, Few-shot, Chain-of-thought
- ➢ Prompt Templates for Various Use Cases
- ➢ System Instructions and Role-based Prompting
- ➢ Lab Session: Writing Effective Prompts using OpenAI Playground

Section 6: OpenAI & LLM APIs – 6 Hours

- ➢ Using OpenAI's GPT-4, DALL·E, Whisper APIs
- ➢ Function Calling and Output Structuring
- ➢ Handling Token Limits and Response Management
- ➢ Lab Session: Build a Q&A; Bot with ChatGPT API

Section 7: LangChain – 8 Hours

- ➢ LangChain Introduction and Setup
- ➢ LLM Wrappers, Chains, and Tools
- ➢ Agents and Retrieval Augmented Generation (RAG)
- ➢ LangChain with Vector Stores
- ➢ Lab Session: Build a Document-aware Chatbot using LangChain + FAISS

Section 8: Hugging Face Transformers – 6 Hours

- ➢ Loading Pretrained Models
- ➢ Inference and Tokenization
- ➢ Fine-tuning Basics
- ➢ Lab Session: Sentiment Classifier using Hugging Face

Section 9: Vector Databases – 4 Hours

- ➢ Introduction to Vector DBs: FAISS, Chroma, Pinecone
- ➢ Storing and Searching Embeddings
- ➢ Use Cases in Retrieval-Augmented Generation (RAG)
- ➢ Lab Session: Build a Simple Vector Search over PDF Documents

Section 10: Real-Time GenAI Projects – 12 Hours

- ➢ Capstone Project 1: Resume Screener with LLM
- ➢ Capstone Project 2: AI Chatbot for Course FAQs
- ➢ Capstone Project 3: Chat with PDF using RAG
- ➢ Bonus Mini Project: Text-to-Image App using DALL·E API

Section 11: MLOps & Deployment – 6 Hours

- ➢ Using Streamlit and Gradio for Demos
- ➢ Deploying LLM Apps to Hugging Face Spaces and Streamlit Cloud
- ➢ Containerization with Docker (Basics)
- ➢ Lab Session: Deploy GenAI App to Web

Section 12: Responsible AI and Ethics – 2 Hours

- ➢ Hallucination and Bias in LLMs
- ➢ Guardrails and Moderation APIs
- ➢ Watermarking, Copyright & Fair Use
- ➢ Lab Session: Analyzing Hallucination in Model Responses