DP-3028 : Implement Generative AI Engineering with Azure Databricks

Course Description:

This learning path explores how to design, build, fine-tune, evaluate, and operationalize generative AI solutions using Azure Databricks. By combining Apache Spark's scalability with collaborative capabilities, you will develop robust AI systems integrating large language models (LLMs), Retrieval-Augmented Generation (RAG), multi-stage reasoning, and more.

Duration: 8 Hours

Prerequisites:

Familiarity with fundamental AI concepts and Azure Databricks (*Recommended: "Get started with artificial intelligence" and "Explore Azure Databricks" modules*)

Learning Objectives:

- Understand LLM fundamentals and their NLP applications
- Implement Retrieval-Augmented Generation (RAG) workflows
- Apply multi-stage reasoning frameworks
- Fine-tune Azure OpenAI language models
- Evaluate LLM performance with standard metrics
- Adopt responsible AI practices in LLM workflows
- Manage LLM operations using LLMOps with Azure Databricks

Content Coverage :

Module 1: Get Started with Language Models in Azure Databricks

- Introduction
- Understand Generative AI
- Understand Large Language Models (LLMs)
- Identify key components of LLM applications
- Use LLMs for Natural Language Processing (NLP) tasks
- Exercise: Explore language models

Module 2: Implement Retrieval-Augmented Generation (RAG) with Azure Databricks

- Introduction
- Explore the main concepts of a RAG workflow
- Prepare your data for RAG
- Find relevant data with vector search
- Rerank your retrieved results
- Exercise: Set up RAG

Module 3: Implement Multi-Stage Reasoning in Azure Databricks

- Introduction
- What are multi-stage reasoning systems?
- Explore LangChain
- Explore LlamaIndex
- Explore Haystack
- Explore the DSPy framework
- Exercise: Implement multi-stage reasoning with LangChain

Module 4: Fine-Tune Language Models with Azure Databricks

- Introduction
- What is fine-tuning?
- Prepare your data for fine-tuning
- Fine-tune an Azure OpenAI model
- Exercise: Fine-tune an Azure OpenAI model

Module 5: Evaluate Language Models with Azure Databricks

- Introduction
- Compare LLM and traditional ML evaluations
- Evaluate LLMs and AI systems
- Evaluate LLMs with standard metrics
- Describe LLM-as-a-judge for evaluation
- Exercise: Evaluate an Azure OpenAI model

Module 6: Review Responsible AI Principles for Language Models

- Introduction
- What is responsible AI?
- Identify risks
- Mitigate issues
- Use key security tooling to protect your AI systems
- Exercise: Implement responsible AI

Module 7: Implement LLMOps in Azure Databricks

- Introduction
- Transition from traditional MLOps to LLMOps
- Understand model deployments
- Describe MLflow deployment capabilities
- Use Unity Catalog to manage models
- Exercise: Implement LLMOps