

Implement Generative AI engineering with Azure Databricks (DP-3028)

Description

AI engineering is no longer reserved for large research teams. With Azure Databricks, professionals can now design, test, and deploy generative AI solutions at scale. This course guides you step-by-step through the implementation of powerful language models, while integrating advanced techniques such as Retrieval-Augmented Generation (RAG) and custom model optimization.

A hands-on approach to generative AI engineering

In practical terms, you'll learn how to work with AI models, adapt them to your business data, and evaluate their performance. Each training module is designed to combine theory with hands-on practice. You'll build complete and efficient workflows, leveraging the collaborative and scalable environment of Azure Databricks.

Course Content

Module 1: Getting started with language models in Azure Databricks

- Understand generative AI
- Understand large language models (LLMs)
- Identify the key components of LLM applications
- Use LLMs for natural language processing (NLP) tasks

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

- Explore the main concepts of a RAG workflow
- Prepare your data for RAG
- Retrieve relevant data using vector search
- Rerank your retrieved results

Module 3: Implementing multi-step reasoning in Azure Databricks

- What are multi-step reasoning systems?
- Explore LangChain
- Discover LlamaIndex
- Discover Haystack
- Explore the DSPy framework

Module 4: Fine-tuning language models with Azure Databricks

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

Module 5: Evaluating language models with Azure Databricks

- Compare LLM and traditional ML evaluation
- Evaluate LLMs and AI systems

- Evaluate LLMs using standard metrics
- Describe LLM-as-a-judge for evaluation

Module 6: Reviewing responsible AI principles for language models in Azure Databricks

- What is responsible AI?
- Identify risks
- Mitigate issues
- Use key safety tools to protect your AI systems

Module 7: Implementing LLMOps in Azure Databricks

- Transition from traditional MLOps to LLMOps
- Understand deployment patterns
- Describe MLflow deployment capabilities
- Use Unity Catalog to manage models

Lab / Exercises

- This course provides you with exclusive access to the official Microsoft lab, enabling you to practice your skills in a professional environment.

Documentation

- Access to Microsoft Learn, Microsoft's online learning platform, offering interactive resources and educational content to deepen your knowledge and develop your technical skills.

Participant profiles

- Data scientists
- AI engineers
- Machine learning developers
- Cloud solution architects

Prerequisites

- Understand the basics of artificial intelligence and machine learning
- Master the fundamentals of Azure Databricks
- Have some knowledge of natural language processing (NLP)

Objectives

- Implement generative language models in Azure Databricks
- Design RAG workflows to enhance LLM responses
- Fine-tune and evaluate the performance of Azure OpenAI models
- Apply responsible AI principles in your deployments
- Deploy and orchestrate models using LLMOps

Description

Implement Generative AI engineering with Azure Databricks (DP-3028)

Niveau

Intermédiaire

Classroom Registration Price (CHF)

900

Virtual Classroom Registration Price (CHF)

850
Duration (in Days)
1
Reference
DP-3028