



MFE-IT

Reference: DP-3028

Generative AI with Azure Databricks DP-3028 Training Course

Design, Train and Deploy LLM Solutions

Duration: 1 Day | Hours: 6 h

Remote · Sessions guaranteed from 1 registrant · 60% hands-on practice

DESCRIPTION

This training course teaches you to implement generative AI solutions on the Azure Databricks platform. You will learn to work with large language models (LLMs), build RAG pipelines, fine-tune models and deploy production-ready AI applications using the tools of the Databricks ecosystem: MLflow, Unity Catalog and Delta Lake.

The training prepares you for the official Microsoft DP-3028 certification and combines theory with intensive hands-on practice on the Azure Databricks platform. You will build end-to-end generative AI pipelines, from data ingestion to RAG retrieval, fine-tuning, evaluation and deployment — applying the same patterns used by enterprise AI teams in production.

LEARNING OBJECTIVES

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

- Work with large language models on Azure Databricks
- Build RAG (Retrieval-Augmented Generation) pipelines
- Fine-tune and customise LLMs
- Use MLflow to track and manage AI models
- Deploy generative AI solutions in production
- Apply responsible AI principles

PREREQUISITES

- Python programming experience
- Basic understanding of machine learning
- Familiarity with Azure cloud services
- Experience with Databricks recommended

Because each participant is unique, a personalised interview is systematically organised in advance with our expert to design a training programme perfectly aligned with their objectives, level and professional challenges.

TARGET AUDIENCE

- Data scientists and ML engineers

- Azure architects
- Developers working on AI and data projects
- Profiles preparing for the DP-3028 certification

DETAILED PROGRAMME

The training alternates between theoretical input and hands-on practice (approximately 60% of the time). Modules are built around practical exercises based on real-world business use cases.

Module 1 – Introduction to LLMs and Generative AI

- Understanding LLMs and their key components
- Identifying use cases for generative AI in NLP and business applications
- Positioning of Azure Databricks in the modern AI stack

Module 2 – Designing Generative AI Applications

- Designing applications that leverage generative models
- Using Spark to explore and integrate these models in a collaborative environment
- Architecture patterns for enterprise GenAI workloads

Module 3 – Building RAG Pipelines

- Understanding RAG workflows and preparing data
- Setting up vector search with Databricks Vector Search
- Improving response relevance with reranking strategies

Module 4 – Agentic AI Frameworks

- Exploring frameworks such as LangChain, LlamaIndex and Haystack
- Designing systems capable of solving complex tasks
- Multi-step reasoning and tool calling patterns

Module 5 – Fine-tuning LLMs

- Preparing training data and optimising an Azure OpenAI model
- Adapting LLMs to specific business needs
- Cost-performance trade-offs and parameter-efficient fine-tuning

Module 6 – Evaluating LLM Performance

- Comparing LLM evaluation with traditional ML
- Using standard metrics and approaches like LLM-as-a-judge
- Continuous evaluation in production

Module 7 – Responsible AI and Security

- Identifying risks and applying mitigation strategies
- Implementing security tools to protect AI systems
- Compliance, content safety and Unity Catalog governance

Module 8 – Deploying Generative AI in Production

- Understanding deployment models and the transition from MLOps to LLMOps
- Managing the model lifecycle with MLflow and Unity Catalog
- Mosaic AI Model Serving and Foundation Model APIs

Module 9 – Practical Workshop

- Creating an LLM-based application
- Implementing a RAG pipeline end-to-end
- Deploying a generative model in production

TEACHING METHODS

Format and Delivery

The training is delivered remotely via an interactive virtual classroom. It can also be delivered on-site, with content customised to match the needs of your professional project. The theory/practice split is approximately 40%/60%.

MFE-IT Ultra-Personalised Format

Each session accommodates between 1 and 3 participants, ensuring highly individualised support. A preliminary interview allows us to tailor the content to each participant's profile. Inter-company sessions are guaranteed from just 1 registrant (except in cases of force majeure).

Skills Assessment

Throughout the training, the trainer assesses participant progress through multiple-choice questions, role-playing exercises and hands-on work. At the end, a certificate of achievement is issued to each participant.

Post-Training Support

For one month following the training, each participant can contact MFE-IT trainers with questions about implementing acquired knowledge. A response is provided by email or telephone within 48 working hours.

Accessibility

MFE-IT is committed to welcoming people with disabilities. Contact: contact@mfe-it.com.

PRACTICAL INFORMATION

Trainer Resources

- Structured demonstrations aligned with the detailed programme
- Exercise briefs and solutions throughout the training
- A ready-to-use technical environment for practical workshops
- Trainer validation of acquired knowledge at the end of each workshop
- Digital reference documents

Certification and Validation

At the end of the training, a certificate is sent by email specifying the objectives, nature, duration and assessment results. A completion certificate can also be provided on request.

Benefits for Participants

- Train from your workplace or home, with no travel required
- Benefit from an expert trainer-consultant on the subject
- Enjoy an ultra-personalised format (1 to 3 participants)
- Continue training even in the event of unforeseen circumstances

Benefits for the Organisation

- Optimise the training budget by reducing travel and accommodation costs
- Offer quality training to all employees, regardless of location
- Reduce absence time linked to travel
- Support team upskilling in all contexts