

DP-604T00: Implement a data science and machine learning solution for AI with Microsoft Fabric

Course Description

This course is designed to build practical skills in implementing data science and machine learning solutions using Microsoft Fabric. The course explores the complete end-to-end data science process, from understanding and exploring data to preparing and transforming datasets for analysis. Students will learn how to train, evaluate, and track machine learning models, as well as how to deploy those models and generate predictions using Microsoft Fabric tools and capabilities.

Course Duration

1 day

Prerequisites

You should be familiar with basic data concepts and terminology.

Course Outline

1 - Introduction to end-to-end analytics using Microsoft Fabric

- Explore end-to-end analytics with Microsoft Fabric
- Explore data teams and Microsoft Fabric
- Enable and use Microsoft Fabric
- Module assessment

2 - Get started with data science in Microsoft Fabric

- Understand the data science process
- Explore and process data with Microsoft Fabric
- Train and score models with Microsoft Fabric
- Module assessment

3 - Explore data for data science with notebooks in Microsoft Fabric

- Explore notebooks
- Load data for exploration
- Understand data distribution
- Check for missing data in notebooks
- Apply advanced data exploration techniques
- Visualize charts in notebooks
- Module assessment

4 - Preprocess data with Data Wrangler in Microsoft Fabric

- Understand Data Wrangler
- Perform data exploration
- Handle missing data
- Transform data with operators
- Module assessment

5 - Train and track machine learning models with MLflow in Microsoft Fabric

- Understand how to train machine learning models
- Train and track models with MLflow and experiments

- Manage models in Microsoft Fabric
- Module assessment

6 - Generate batch predictions using a deployed model in Microsoft Fabric

- Customize the model's behavior for batch scoring
- Prepare data before generating predictions
- Generate and save predictions to a Delta table
- Module assessment

Who Should Attend

This course is intended for data professionals and practitioners who regularly work with machine learning models and are responsible for building, evaluating, and deploying data science solutions. Students should already be familiar with the data science process, Python, and common open-source machine learning frameworks such as scikit-learn.