



Implement a data science and machine learning solution for AI with Microsoft Fabric

Course DP-604T00: 1 day; Beginner; Instructor-Led

Introduction

Explore the data science process and learn how to train machine learning models to accomplish artificial intelligence in Microsoft Fabric.

Prerequisites

- You should be familiar with basic data concepts and terminology.

Course Outline

Module 1: Get started with data science in Microsoft Fabric

In Microsoft Fabric, data scientists can manage data, notebooks, experiments, and models while easily accessing data from across the organization and collaborating with their fellow data professionals.

In this module, you'll learn how to:

- Understand the data science process
- Train models with notebooks in Microsoft Fabric
- Track model training metrics with MLflow and experiments

Prerequisites

- Before starting this module, you should be familiar with the basic principles of machine learning.

Lessons

- Understand the data science process
- Explore and process data with Microsoft Fabric
- Train and score models with Microsoft Fabric
- Exercise - Explore data science in Microsoft Fabric

Module 2: Explore data for data science with notebooks in Microsoft Fabric

Microsoft Fabric notebooks serve as a comprehensive tool for data exploration, enabling users to uncover hidden patterns and relationships in their datasets.

In this module, you'll:

- Load data and perform initial data exploration.
- Gain knowledge about different types of data distributions.
- Understand the concept of missing data, and strategies to handle missing data effectively.
- Visualize data using various data visualization techniques and libraries.

Prerequisites

- Before starting this module, you should be familiar with the basic principles of machine learning.

Lessons

- Explore notebooks
- Load data for exploration
- Understand data distribution
- Check for missing data in notebooks
- Apply advanced data exploration techniques
- Visualize charts in notebooks
- Exercise: Use notebook for data exploration in Microsoft Fabric

Module 3: Preprocess data with Data Wrangler in Microsoft Fabric

Data Wrangler serves as a comprehensive tool for preprocessing data. It enables users to clean data, handle missing values, and transform features to build machine learning models.

In this module, you'll:

- Learn Data Wrangler features, and its role in the data science workflow.
- Perform different types of preprocessing operations in data science.
- Learn how to handle missing values, and imputation strategies.
- Use one-hot encoding and other techniques to convert categorical data into a format suitable for machine learning algorithms.

Prerequisites

- Before starting this module, you should be familiar with the basic principles of machine learning.

Lessons

- Understand Data Wrangler
- Perform data exploration
- Handle missing data
- Transform data with operators
- Exercise: Preprocess data with Data Wrangler in Microsoft Fabric

Module 4: Train and track machine learning models with MLflow in Microsoft Fabric

In Microsoft Fabric, data scientists can train models in notebooks, track their work in experiments, and manage their models with MLflow.

In this module, you'll learn how to:

- Train machine learning models with open-source frameworks
- Train models with notebooks in Microsoft Fabric
- Track model training metrics with MLflow and experiments in Microsoft Fabric

Prerequisites

- Before starting this module, you should be familiar with the data science process.

Lessons

- Understand how to train machine learning models
- Train and track models with MLflow and experiments
- Manage models in Microsoft Fabric
- Exercise - Train and track a model in Microsoft Fabric

Module 5: Generate batch predictions using a deployed model in Microsoft Fabric

Save and use your machine learning models in Microsoft Fabric to generate batch predictions and enrich your data.

In this module, you'll learn how to:

- Save a model in the Microsoft Fabric workspace
- Prepare a dataset for batch predictions
- Apply the model to dataset to generate new predictions
- Save the predictions to a Delta table

Prerequisites

- Before starting this module, you should be familiar with the data science process.

Lessons

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