



DP-100: Designing and Implementing a Data Science Solution on Azure

Microsoft - Data & AI

- **Localidade:** Porto
 - **Data:** 17 Jun 2024
 - **Preço:** 1590 € (Os valores apresentados não incluem IVA. Oferta de IVA a particulares e estudantes.)
 - **Horário:** Manhã das 9h00 às 12h30
 - **Nível:** Intermédio
 - **Duração:** 28h
-

Sobre o curso

Learn how to operate machine learning solutions at cloud scale using Azure Machine Learning. This course teaches you to leverage your existing knowledge of Python and machine learning to manage data ingestion and preparation, model training and deployment, and machine learning solution monitoring in Microsoft Azure.

Destinatários

This course is designed for data scientists with existing knowledge of Python and machine learning frameworks like Scikit-Learn, PyTorch, and Tensorflow, who want to build and operate machine learning solutions in the cloud.

Objetivos

- Create and publish models without writing code
 - Prepare for Exam DP-100: Designing and Implementing a Data Science Solution on Azure, Exam AI-900: Microsoft Azure AI Fundamentals
 - Create and manage enterprise-ready ML solutions
 - Explore, prepare, and model data; and integrate with Azure Machine Learning
-

Pré-requisitos

Successful Azure Data Scientists start this role with a fundamental knowledge of cloud computing concepts, and experience in general data science and machine learning tools and techniques. Specifically:

- Creating cloud resources in Microsoft Azure.
- Using Python to explore and visualize data.
- Training and validating machine learning models using common frameworks like Scikit-Learn, PyTorch, and TensorFlow.
- Working with container

If you are completely new to data science and machine learning, please complete [Microsoft Azure AI Fundamentals](#) first.

Programa

- Design a data ingestion strategy for machine learning projects
- Design a machine learning model training solution
- Design a model deployment solution
- Explore Azure Machine Learning workspace resources and assets
- Explore developer tools for workspace interaction
- Make data available in Azure Machine Learning
- Work with compute targets in Azure Machine Learning
- Work with environments in Azure Machine Learning
- Find the best classification model with Automated Machine Learning
- Track model training in Jupyter notebooks with MLflow
- Run a training script as a command job in Azure Machine Learning
- Track model training with MLflow in jobs
- Run pipelines in Azure Machine Learning
- Perform hyperparameter tuning with Azure Machine Learning
- Deploy a model to a managed online endpoint
- Deploy a model to a batch endpoint

Design a data ingestion strategy for machine learning projects

Learn how to design a data ingestion solution for training data used in machine learning projects.

- Identify your data source and format
- Choose how to serve data to machine learning workflows

- Design a data ingestion solution

Design a machine learning model training solution

Learn how to design a model training solution for machine learning projects.

- Identify machine learning tasks
- Choose a service to train a model
- Choose between compute options

Design a model deployment solution

Learn how to design a model deployment solution and how the requirements of the deployed model can affect the way you train a model.

- Understand how a model will be consumed.
- Decide whether to deploy your model to a real-time or batch endpoint.

Explore Azure Machine Learning workspace resources and assets

As a data scientist, you can use Azure Machine Learning to train and manage your machine learning models. Learn what Azure Machine Learning is, and get familiar with all its resources and assets.

- Create an Azure Machine Learning workspace.
- Identify resources and assets.
- Train models in the workspace.

Explore developer tools for workspace interaction

Learn how you can interact with the Azure Machine Learning workspace. You can use the Azure Machine Learning studio, the Python SDK (v2), or the Azure CLI (v2).

- The Azure Machine Learning studio.
- The Python Software Development Kit (SDK).
- The Azure Command Line Interface (CLI).

Make data available in Azure Machine Learning

Learn about how to connect to data from the Azure Machine Learning workspace. You'll be introduced to datastores and data assets.

- Work with Uniform Resource Identifiers (URIs).
- Create and use datastores.
- Create and use data assets.

Work with compute targets in Azure Machine Learning

Learn how to work with compute targets in Azure Machine Learning. Compute targets allow you to run your machine learning workloads. Explore how and when you can use a compute instance or compute cluster.

- Choose the appropriate compute target.
- Create and use a compute instance.
- Create and use a compute cluster.

Work with environments in Azure Machine Learning

Learn how to use environments in Azure Machine Learning to run scripts on any compute target.

- Understand environments in Azure Machine Learning.
- Explore and use curated environments.
- Create and use custom environments.

Find the best classification model with Automated Machine Learning

Learn how to find the best classification model with automated machine learning (AutoML). You'll use the Python SDK (v2) to configure and run an AutoML job.

- Prepare your data to use AutoML for classification.
- Configure and run an AutoML experiment.
- Evaluate and compare models.

Track model training in Jupyter notebooks with MLflow

Learn how to use MLflow for model tracking when experimenting in notebooks.

- Configure to use MLflow in notebooks
- Use MLflow for model tracking in notebooks

Run a training script as a command job in Azure Machine Learning

Learn how to convert your code to a script and run it as a command job in Azure Machine Learning.

- Convert a notebook to a script.
- Test scripts in a terminal.
- Run a script as a command job.
- Use parameters in a command job.

Track model training with MLflow in jobs

Learn how to track model training with MLflow in jobs when running scripts.

- Use MLflow when you run a script as a job.
- Review metrics, parameters, artifacts, and models from a run.

Run pipelines in Azure Machine Learning

Learn how to create and use components to build pipeline in Azure Machine Learning. Run and schedule Azure Machine Learning pipelines to automate machine learning workflows.

- Create components.
- Build an Azure Machine Learning pipeline.
- Run an Azure Machine Learning pipeline.

Perform hyperparameter tuning with Azure Machine Learning

Learn how to perform hyperparameter tuning with a sweep job in Azure Machine Learning.

- Define a hyperparameter search space.
- Configure hyperparameter sampling.
- Select an early-termination policy.
- Run a sweep job.

Deploy a model to a managed online endpoint

Learn how to deploy models to a managed online endpoint for real-time inferencing.

- Use managed online endpoints.
- Deploy your MLflow model to a managed online endpoint.
- Deploy a custom model to a managed online endpoint.
- Test online endpoints.

Deploy a model to a batch endpoint

Learn how to deploy models to a batch endpoint. When you invoke a batch endpoint, you'll trigger a batch scoring job.

- Create a batch endpoint.

- Deploy your MLflow model to a batch endpoint.
- Deploy a custom model to a batch endpoint.
- Invoke batch endpoints.