

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

Overview

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.

Duration

4 Days

Who Should Attend

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.

Prerequisites

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 ScikitLearn, PyTorch, and TensorFlow.
- Working with containers

To gain these prerequisite skills, take the following free online training before attending the course:

- Explore Microsoft cloud concepts.
- Create machine learning models.
- Administer containers in Azure

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

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

Course Topics

<p>Module 1</p>	<p>Getting Started with Azure Machine Learning</p> <p>In this module, you will learn how to provision an Azure Machine Learning workspace and use it to manage machine learning assets such as data, compute, model training code, logged metrics, and trained models. You will learn how to use the web-based Azure Machine Learning studio interface as well as the Azure Machine Learning SDK and developer tools like Visual Studio Code and Jupyter Notebooks to work with the assets in your workspace.</p> <ul style="list-style-type: none"> • Introduction to Azure Machine Learning • Working with Azure Machine Learning <p>Lab: Create an Azure Machine Learning Workspace</p> <p>After completing this module, you will be able to</p> <ul style="list-style-type: none"> • Provision an Azure Machine Learning workspace • Use tools and code to work with Azure Machine Learning
<p>Module 2</p>	<p>No-Code Machine Learning</p> <p>This module introduces the Automated Machine Learning and Designer visual tools, which you can use to train, evaluate, and deploy machine learning models without writing any code.</p> <ul style="list-style-type: none"> • Automated Machine Learning • Azure Machine Learning Designer <p>Lab: Use Automated Machine Learning</p> <p>Lab: Use Azure Machine Learning Designer</p> <p>After completing this module, you will be able to</p> <ul style="list-style-type: none"> • Use automated machine learning to train a machine learning model • Use Azure Machine Learning designer to train a model • Storage, Azure Disk Storage, and Azure File Storage.

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

<p>Module 3</p>	<p>Running Experiments and Training Models</p> <p>In this module, you will get started with experiments that encapsulate data processing and model training code, and use them to train machine learning models.</p> <ul style="list-style-type: none"> • Introduction to Experiments • Training and Registering Models <p>Lab: Run Experiments Lab: Train Models</p> <p>After completing this module, you will be able to</p> <ul style="list-style-type: none"> • Run code-based experiments in an Azure Machine Learning workspace • Train and register machine learning models
<p>Module 4</p>	<p>Working with Data</p> <p>Data is a fundamental element in any machine learning workload, so in this module, you will learn how to create and manage datastores and datasets in an Azure Machine Learning workspace, and how to use them in model training experiments.</p> <ul style="list-style-type: none"> • Working with Datastores • Working with Datasets <p>Lab: Work with Data</p> <p>After completing this module, you will be able to</p> <ul style="list-style-type: none"> • Create and use datastores • Create and use datasets
<p>Module 5</p>	<p>Working with Compute</p> <p>One of the key benefits of the cloud is the ability to leverage compute resources on demand and use them to scale machine learning processes to an extent that would be infeasible on your own hardware. In this module, you'll learn how to manage experiment environments that ensure consistent runtime consistency for experiments, and how to create and use compute targets for experiment runs.</p> <ul style="list-style-type: none"> • Working with Environments • Working with Compute Targets <p>Lab: Work with Compute</p> <p>After completing this module, you will be able to</p> <ul style="list-style-type: none"> • Create and use environments • Create and use compute targets

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

Module 6

Orchestrating Operations with Pipelines

Now that you understand the basics of running workloads as experiments that leverage data assets and compute resources, it's time to learn how to orchestrate these workloads as pipelines of connected steps. Pipelines are key to implementing an effective Machine Learning Operationalization (ML Ops) solution in Azure, so you'll explore how to define and run them in this module.

- Introduction to Pipelines
- Publishing and Running Pipelines

Lab: Create a Pipeline

After completing this module, you will be able to

- Create pipelines to automate machine learning work
- Flows
- Publish and run pipeline services

Module 7

Deploying and Consuming Models

Models are designed to help decision making through predictions, so they're only useful when deployed and available for an application to consume. In this module learn how to deploy models for real-time inferencing, and for batch inferencing.

- Real-time Inferencing
- Batch Inferencing
- Continuous Integration and Delivery

Lab: Create a Real-time Inferencing Service

Lab: Create a Batch Inferencing Service

After completing this module, you will be able to

- Publish a model as a real-time inference service
- Publish a model as a batch inference service
- Describe techniques to implement continuous integration and delivery

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

<p>Module 8</p>	<p>Training Optimal Models</p> <p>By this stage of the course, you've learned the end-to-end process for training, deploying, and consuming machine learning models; but how do you ensure your model produces the best predictive outputs for your data? In this module, you'll explore how you can use hyperparameter tuning and automated machine learning to take advantage of cloud-scale compute and find the best model for your data.</p> <ul style="list-style-type: none"> • Hyperparameter Tuning • Automated Machine Learning <p>Lab: Tune Hyperparameters Lab: Use Automated Machine Learning from the SDK</p> <p>After completing this module, you will be able to</p> <ul style="list-style-type: none"> • Optimize hyperparameters for model training • Use automated machine learning to find the optimal model for your data
<p>Module 9</p>	<p>Responsible Machine Learning</p> <p>Data scientists have a duty to ensure they analyze data and train machine learning models responsibly, respecting individual privacy, mitigating bias, and ensuring transparency. This module explores some considerations and techniques for applying responsible machine learning principles.</p> <ul style="list-style-type: none"> • Differential Privacy • Model Interpretability • Fairness <p>Lab: Explore Differential provacy Lab: Interpret Models Lab: Detect and Mitigate Unfairness</p> <p>After completing this module, you will be able to</p> <ul style="list-style-type: none"> • Apply differential provacy to data analysis • Use explainers to interpret machine learning models • Evaluate models for fairness

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

Module 10

Responsible Machine Learning

After a model has been deployed, it's important to understand how the model is being used in production, and to detect any degradation in its effectiveness due to data drift. This module describes techniques for monitoring models and their data.

- Monitoring Models with Application Insights
- Monitoring Data Drift

Lab: Monitor a Model with Application Insights

Lab: Monitor Data Drift

After completing this module, you will be able to

- Use Application Insights to monitor a published model
- Monitor data drift