
DP-100T01-AC Designing and Implementing a Data Science Solution on Azure

Overview

Course Duration: 3 Days

About This Course

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.

Audience Profile

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.

Course Outline

Module 1: Getting Started with Azure Machine Learning 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. **Lesson**

- Introduction to Azure Machine Learning
- Working with Azure Machine Learning

Lab: Create an Azure Machine Learning Workspace After completing this module, you will be able to

- Provision an Azure Machine Learning workspace
- Use tools and code to work with Azure Machine Learning

Module 2: Visual Tools for Machine Learning 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. **Lesson**

- Automated Machine Learning
- Azure Machine Learning Designer

Lab: Use Automated Machine Learning Lab: Use Azure Machine Learning Designer After completing this module, you will be able to

- Use automated machine learning to train a machine learning model
- Use Azure Machine Learning designer to train a model

Module 3: Running Experiments and Training Models 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. **Lesson**

- Introduction to Experiments
- Training and Registering Models

Lab: Run Experiments Lab: Train Models After completing this module, you will be able to

- Run code-based experiments in an Azure Machine Learning workspace
- Train and register machine learning models

Module 4: Working with Data 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. **Lesson**

- Working with Datastores
- Working with Datasets

Lab: Work with Data After completing this module, you will be able to

- Create and use datastores
- Create and use datasets

Module 5: Working with Computer 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. **Lesson**

- Working with Environments
- Working with Compute Targets

Lab: Work with Computer After completing this module, you will be able to

- Create and use environments
- Create and use compute targets

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. **Lesson**

- 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 workflows
- 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. **Lesson**

- 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

Module 8: Training Optimal Models 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. **Lesson**

- Hyperparameter Tuning
- Automated Machine Learning

Lab: Tune Hyperparameters **Lab: Use Automated Machine Learning from the SDK** After completing this module, you will be able to

- Optimize hyperparameters for model training
- Use automated machine learning to find the optimal model for your data

Module 9: Responsible Machine Learning 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. **Lesson**

- Differential Privacy
- Model Interpretability
- Fairness

Lab: Explore Differential privacy **Lab: Interpret Models** **Lab: Detect and Mitigate Unfairness** After completing this module, you will be able to

- Apply differential privacy to data analysis
- Use explainers to interpret machine learning models
- Evaluate models for fairness

Module 10: Monitoring Models 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. **Lesson**

- 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

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 Scikit-Learn, PyTorch, and TensorFlow.
- Working with containers