

# Designing And implementing A Data Science solution on Azure

Duration	Delivery Method	Level
1 day	Online / Instructor Led	Beginner

## Introduction:

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.

## Pre-requisites:

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.

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

## Course Content

### Module 1: Getting started with Azure machine learning

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

### Module 2: Visual Tools for machine Learning

- Automated Machine Learning
- Azure Machine Learning Designer

### **Module 3: Running Experiments and training models**

- Introduction to Experiments
- Training and Registering 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.

- Working with Datastores
- Working with Datasets

### **Module 5: Working with Compute**

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.

- Working with Environments
- Working with 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 Operationalisation (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

### **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

## **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.

- Hyperparameter Tuning
- Automated Machine Learning

## **Module 9: Responsible Machine learning**

Data scientists have a duty to ensure they analyse 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.

- Differential Privacy
- Model Interpretability
- 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.

- Monitoring Models with Application Insights
- Monitoring Data Drift