

KOENIG SOLUTIONS • COURSE OUTLINE

DP-100 Exam Prep

Learn to set up Azure Machine Learning workspaces, train and optimise models, and deploy them into production.

4

MODULES

49

TOPICS

COURSE OVERVIEW

The Azure Data Scientist applies their knowledge of data science and machine learning to implement and run machine learning workloads on Azure; in particular, using Azure Machine Learning Service. This entails planning and creating a suitable working environment for data science workloads on Azure, running data experiments and training predictive models, managing and optimizing models, and deploying machine learning models into production.

WHO THIS COURSE IS FOR

Data Scientists

Machine Learning Engineers

Azure Solution Architects

Data Analysts

AI Developers

LEARNING OBJECTIVES

- Create and configure an Azure Machine Learning workspace
- Manage datastores, datasets and compute contexts
- Build training pipelines using Azure Machine Learning designer
- Run and monitor training scripts using the Azure Machine Learning SDK
- Generate and troubleshoot metrics from experiment runs
- Apply Automated ML and Hyperdrive to optimise models
- Interpret models using model explainers and manage data drift
- Deploy models as services and batch inferencing pipelines

4 Modules, 49 Topics

01

30-35%

Set up an Azure Machine Learning Workspace

CREATE AN AZURE MACHINE LEARNING WORKSPACE

- create an Azure Machine Learning workspace
- configure workspace settings
- manage a workspace by using Azure Machine Learning studio

MANAGE DATA OBJECTS IN AN AZURE MACHINE LEARNING WORKSPACE

- register and maintain datastores
- create and manage datasets

MANAGE EXPERIMENT COMPUTE CONTEXTS

- create a compute instance
- determine appropriate compute specifications for a training workload
- create compute targets for experiments and training

02

25-30%

Run Experiments and Train Models

CREATE MODELS BY USING AZURE MACHINE LEARNING DESIGNER

- create a training pipeline by using Azure Machine Learning designer
- ingest data in a designer pipeline
- use designer modules to define a pipeline data flow
- use custom code modules in designer

RUN TRAINING SCRIPTS IN AN AZURE MACHINE LEARNING WORKSPACE

- create and run an experiment by using the Azure Machine Learning SDK
- configure run settings for a script
- consume data from a dataset in an experiment by using the Azure Machine Learning SDK

GENERATE METRICS FROM AN EXPERIMENT RUN

- log metrics from an experiment run
- retrieve and view experiment outputs
- use logs to troubleshoot experiment run errors

AUTOMATE THE MODEL TRAINING PROCESS

- create a pipeline by using the SDK
- pass data between steps in a pipeline
- run a pipeline
- monitor pipeline runs

Optimize and Manage Models

USE AUTOMATED ML TO CREATE OPTIMAL MODELS

- use the Automated ML interface in Azure Machine Learning studio
- use Automated ML from the Azure Machine Learning SDK
- select pre-processing options
- determine algorithms to be searched
- define a primary metric
- get data for an Automated ML run
- retrieve the best model

USE HYPERDRIVE TO TUNE HYPERPARAMETERS

- select a sampling method
- define the search space
- define the primary metric
- define early termination options
- find the model that has optimal hyperparameter values

USE MODEL EXPLAINERS TO INTERPRET MODELS

- select a model interpreter
- generate feature importance data

MANAGE MODELS

- register a trained model
- monitor model usage
- monitor data drift

Deploy and Consume Models

CREATE PRODUCTION COMPUTE TARGETS

- consider security for deployed services
- evaluate compute options for deployment

DEPLOY A MODEL AS A SERVICE

- configure deployment settings
- consume a deployed service
- troubleshoot deployment container issues

CREATE A PIPELINE FOR BATCH INFERENCE

- publish a batch inferencing pipeline
- run a batch inferencing pipeline and obtain outputs

PUBLISH A DESIGNER PIPELINE AS A WEB SERVICE

- create a target compute resource
- configure an Inference pipeline
- consume a deployed endpoint

NOTES

- The bullets that appear below each of the skills measured are intended to illustrate how the exam assesses that skill. This list is not definitive or exhaustive.
- Most questions cover features that are General Availability (GA). The exam may contain questions on Preview features if those features are commonly used.
- This exam was updated on December 8, 2020. The document includes the current exam guide reflecting the changes made on that date (registering and maintaining datastores; consuming data from a dataset rather than a data store in an experiment via the SDK; removal of 'choose an estimator for a training experiment'; selecting pre-processing options rather than scaling functions and pre-processing options; monitoring model usage rather than model history).

WHAT IS INCLUDED

Everything You Need to Succeed

- Koenig course material for every module
- Delivery by certified instructors and subject matter experts
- Hands-on exercises throughout the course
- Course completion certificate from Koenig

ABOUT KOENIG

Who We Are

Established in 1993, Koenig Solutions is a globally recognised training provider with 35+ technology partnerships, 5,000+ training programmes and a network of certified instructors and subject matter experts, supporting individuals and enterprise teams worldwide.

FLEXIBLE DELIVERY

Learn the Way That Suits You

Classroom Training

Live Online Classes (ILO)

Fly-Me-A-Trainer (FMAT)

Flexi (Self-Paced)

1-on-1 Training

Ready to Get Started?

Talk to a learning advisor about schedules, group training and customised delivery.

Visit koenig-solutions.com

✉ info@koenig-solutions.com