

Microsoft DP-100 Checklist

Azure Data Scientist Associate

General Overview

- 180 minutes
- 60 questions (45 Q&A, 15 about Case Studies)

Not on the exam

- R
- Power BI
- Publishing Azure ML models

General Machine Learning

- Evaluation of Linear Regression
- Evaluation of Classification
- positive/negative skew
- Poisson regression
- Fisher's exact test
- Pearson
- PCA
- deep learning
- Neural Networks (RNN vs CNN vs ...)

Azure ML

- Modeling Building
- Model Evaluation
- SMOTE
- MICE
- feature extraction
- missing data questions
- AutoML

Azure Products

- Azure Machine Learning Service
- Blob storage - specifically how to get data in/out
- Azure Notebooks
- Azure Cognitive Services (high level)
- Kubernetes
- HDInsight
- Data Science Virtual Machine

Python

- Scikit-learn
- Azure Machine Learning SDK for Python
- hyperparameters