

Exam DP-100 Preview & DP-100 Latest Braindumps Ebook



2026 Latest PracticeMaterial DP-100 PDF Dumps and DP-100 Exam Engine Free Share: <https://drive.google.com/open?id=1F3c-QfM-9oi8JUcc84I4tS6CEIUjacF2>

Our DP-100 practice materials not only apply to students, but also apply to office workers; not only apply to veterans in the workplace, but also apply to newly recruited newcomers. Our DP-100 study materials use a very simple and understandable language, to ensure that all people can learn and understand. Our DP-100 real test also allows you to avoid the boring of textbook reading, but let you master all the important knowledge in the process of doing exercises. And the high pass rate of our DP-100 exam questions is more than 98%. Why not have a try on our DP-100 study guide?

Microsoft DP-100 Exam Syllabus Topics:

Section	Weight	Objectives

Topic 1: Explore data and run experiments	20-25%	- Run experiments • 1. Configure experiment runs • 2. Track runs with MLflow • 3. Use automated machine learning • 4. Define parameters and configurations - Explore and visualize data • 1. Identify features and relationships • 2. Detect anomalies and outliers • 3. Profile and validate data - Implement pipelines • 1. Pass data between steps • 2. Build reusable components • 3. Schedule and monitor pipelines • 4. Create and publish pipelines
Topic 2: Optimize language models for AI applications	25-30%	- Evaluate and improve models • 1. Optimize for accuracy and safety • 2. Test and evaluate responses • 3. Apply responsible generative AI - Implement generative AI solutions • 1. Build prompt flows • 2. Use Azure AI Foundry • 3. Apply prompt engineering - Optimize with Retrieval Augmented Generation • 1. Create vector stores and indexes • 2. Configure Azure AI Search • 3. Prepare and process data
Topic 3: Design and prepare a machine learning solution	20-25%	- Manage data assets • 1. Select storage services • 2. Create and maintain data assets • 3. Register and manage datastores - Manage compute resources • 1. Attach and monitor compute • 2. Create and configure compute targets • 3. Select environments - Design a machine learning solution • 1. Determine dataset structure and format • 2. Define compute specifications for workloads • 3. Plan model deployment requirements • 4. Select development approach - Manage Azure Machine Learning workspace • 1. Work with registries • 2. Use developer tools and CLI • 3. Create and configure workspace • 4. Set up Git integration

Topic 4: Train and deploy models	25-30%	- Deploy models • 1. Deploy to batch endpoints • 2. Deploy to online endpoints • 3. Configure compute and scaling • 4. Secure endpoints and manage access - Manage models • 1. Register and version models • 2. Interpret models and explain predictions • 3. Package and validate models - Train models • 1. Use HyperDrive for hyperparameter tuning • 2. Apply responsible AI principles • 3. Configure jobs and environments • 4. Run training scripts - Monitor and maintain models • 1. Update and retrain models • 2. Monitor performance and data drift • 3. Implement MLOps practices
----------------------------------	--------	---

>> Exam DP-100 Preview <<

DP-100 Latest Braindumps Ebook - DP-100 Latest Test Preparation

A variety of PracticeMaterial' Microsoft dumps are very helpful for the preparation to get assistance in this regard. It is designed exactly according to the exams curriculum. The use of test preparation exam questions helps them to practice thoroughly. Rely on material of the Free DP-100 Braindumps online (easily available) sample tests, and resource material available on our website. These free web sources are significant for DP-100 certification syllabus. Our website provides the sufficient material regarding DP-100 exam preparation.

Microsoft Designing and Implementing a Data Science Solution on Azure Sample Questions (Q236-Q241):

NEW QUESTION # 236

You have an Azure Machine Learning workspace. You are running an experiment on your local computer.

You need to use MLflow Tracking to store metrics and artifacts from your local experiment runs in the workspace.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.


The screenshot shows a Microsoft exam interface with a list of actions on the left and an answer area on the right. The actions are:

- Import MLflow and Workspace classes.
- Load the workspace.
- Retrieve the tracking URI and set the experiment name.
- Start a training run and activate the MLflow logging API.

Navigation arrows (back, forward, up, down) are visible between the actions and the answer area.

Answer:

Explanation:



- 1 - Import MLflow and Workspace classes.
- 2 - Load the workspace.
- 3 - Retrieve the tracking URI and set the experiment name.
- 4 - Start a training run and activate the MLflow logging API.

NEW QUESTION # 237

You run an experiment that uses an AutoMLConfig class to define an automated machine learning task with a maximum of ten model training iterations. The task will attempt to find the best performing model based on a metric named accuracy. You submit the experiment with the following code:

```
from azureml.core.experiment import Experiment
automl_experiment = Experiment(ws, 'automl_experiment')
automl_run = automl_experiment.submit(automl_config, show_output=True)
```

You need to create Python code that returns the best model that is generated by the automated machine learning task. Which code segment should you use?

- A. `best_model = automl_run.get_details()`
- B. `best_model = automl_run.get_file_names()[1]`
- C. `best_model = automl_run.get_metrics()`
- D. `best_model = automl_run.get_output()[1]`

Answer: D

Explanation:

The `get_output` method returns the best run and the fitted model.

Reference:

<https://notebooks.azure.com/azureml/projects/azureml-getting-started/html/how-to-use-azureml/automated-machine-learning/classification/auto-ml-classification.ipynb>

NEW QUESTION # 238

You are a data scientist building a deep convolutional neural network (CNN) for image classification.

The CNN model you build shows signs of overfitting.

You need to reduce overfitting and converge the model to an optimal fit.

Which two actions should you perform? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Add L1/L2 regularization.
- B. Add an additional dense layer with 512 input units.
- C. Reduce the amount of training data.
- D. Use training data augmentation.
- E. Add an additional dense layer with 64 input units.

Answer: A,C

Explanation:

B: Weight regularization provides an approach to reduce the overfitting of a deep learning neural network model on the training data and improve the performance of the model on new data, such as the holdout test set.

Keras provides a weight regularization API that allows you to add a penalty for weight size to the loss function.

Three different regularizer instances are provided; they are:

* L1: Sum of the absolute weights.

* L2: Sum of the squared weights.

* L1L2: Sum of the absolute and the squared weights.

D: Because a fully connected layer occupies most of the parameters, it is prone to overfitting. One method to reduce overfitting is dropout. At each training stage, individual nodes are either "dropped out" of the net with probability 1-p or kept with probability p, so that a reduced network is left; incoming and outgoing edges to a dropped-out node are also removed.

By avoiding training all nodes on all training data, dropout decreases overfitting.

References:


<https://machinelearningmastery.com/how-to-reduce-overfitting-in-deep-learning-with-weight-regularization/>

https://en.wikipedia.org/wiki/Convolutional_neural_network

NEW QUESTION # 239


You need to visually identify whether outliers exist in the Age column and quantify the outliers before the outliers are removed.

Which three Azure Machine Learning Studio modules should you use in sequence? To answer, move the appropriate modules from the list of modules to the answer area and arrange them in the correct order.



Answer:

Explanation:



1 - Create Scatterplot

2 - Summarize Data

3 - Clip Values

Reference:

<https://blogs.msdn.microsoft.com/azuredev/2017/05/27/data-cleansing-tools-in-azure-machine-learning/>

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/clip-values>

NEW QUESTION # 240

You have an Azure Machine Learning workspace named Workspace1. Workspace1 has a registered Mlflow model named model1 with PyFunc flavor. You plan to deploy model1 to an online endpoint named endpoint1 without egress connectivity by using Azure Machine Learning Python SDK v1. You have the following code:

```
blue_deployment = ManagedOnlineDeployment(  
    name="blue",  
    endpoint_name=endpoint1,  
    model=model1,  
    instance_type="Standard_F4s_v2",  
    instance_count=1  
)
```

You need to add a parameter to the ManagedOnlineDeployment object to ensure the model deploys successfully. Solution: Add the with_package parameter.

Does the solution meet the goal?

- Answer: A**

• • • • •

DP-100 Latest Braindumps Ebook: <https://www.practicematerial.com/DP-100-exam-materials.html>

- BTW, DOWNLOAD part of PracticeMaterial DP-100 dumps from Cloud Storage: <https://drive.google.com/open?id=1F3c-QfM-9oi8JUcc844tS6CEIUjacF2>