



Free Questions for DP-100 by ebraindumps

Shared by Turner on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it as a result, these questions will not appear in the review screen.

You use Azure Machine Learning designer to load the following datasets into an experiment:

Dataset1

Age	Length	Width
3	22	13
7	11	96
18	32	85

Dataset2

Age	Length	Width
11	101	65
6	98	23
33	22	54
17	52	12

You need to create a dataset that has the same columns and header row as the input datasets and contains all rows from both input datasets.

Solution: Use the Apply Transformation module.

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

B

Question 2

Question Type: MultipleChoice

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You use Azure Machine Learning designer to load the following datasets into an experiment:

Dataset1

Age	Length	Width
3	22	13
7	11	96
18	32	85

Dataset2

Age	Length	Width
11	101	65
6	98	23
33	22	54
17	52	12

You need to create a dataset that has the same columns and header row as the input datasets and contains all rows from both input datasets.

Solution: Use the Join Data module.

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

B

Question 3

Question Type: MultipleChoice

You use Azure Machine Learning studio to analyze an mltable data asset containing a decimal column named column1. You need to verify that the column1 values are normally distributed.

Which statistic should you use?

Options:

A- Max

B- Type

C- Profile

D- Mean

Answer:

C

Question 4

Question Type: MultipleChoice

You create an Azure Machine Learning workspace. You use Azure Machine Learning designer to create a pipeline within the workspace. You need to submit a pipeline run from the designer.

What should you do first?

Options:

- A- Create a compute cluster.
- B- Create an attached compute resource.
- C- Select a model.
- D- Create an experiment.

Answer:

B

Question 5

Question Type: MultipleChoice

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it as a result, these questions will not appear in the review screen.

You train and register an Azure Machine Learning model.

You plan to deploy the model to an online end point.

You need to ensure that applications will be able to use the authentication method with a non-expiring artifact to access the model.

Solution:

Create a Kubernetes online endpoint and set the value of its auth-mode parameter to amyI Token. Deploy the model to the online endpoint.

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

B

Question 6

Question Type: MultipleChoice

You run Azure Machine Learning training experiments. The training scripts directory contains 100 files that includes a file named. amlignore. The directory also contains subdirectories named. /outputs and./logs.

There are 20 files in the training scripts directory that must be excluded from the snapshot to the compute targets. You create a file named. gift ignore in the root of the directory. You add the names of the 20 files to the. gift ignore file. These 20 files continue to be copied to the compute targets.

You need to exclude the 20 files. What should you do?

Options:

- A- Add the contents of the file named. amlignore to the file named. gift ignore.
- B- Move the file named. gift ignore to the. /logs directory.
- C- Copy the contents of the file named. gift ignore to the file named. amlignore.
- D- Move the file named. gift ignore to the. /outputs directory.

Answer:

C

Question 7

Question Type: MultipleChoice

You manage an Azure Machine Learning workspace.

You need to define an environment from a Docker image by using the Azure Machine Learning Python SDK v2.

Which parameter should you use?

Options:

A- conda_file

B- image

C- build

D- properties

Answer:

B

Question 8

Question Type: MultipleChoice

You manage an Azure Machine Learning workspace.

You must log multiple metrics by using MLflow.

You need to maximize logging performance.

What are two possible ways to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

Options:

- A- MLflowClient.log_batch
- B- mlflowlog_metrics
- C- mlflow.log_param
- D- mlflow.log. metric

Answer:

A, B

Question 9

Question Type: MultipleChoice

You have the following Azure subscriptions and Azure Machine Learning service workspaces:

Subscription	Workspace	Comment
385bdf5-4cef-4ad4-b977-3f86d92727c9	ml-default	This is the default subscription.
5a5891d1-557a-4234-9b83-2e90412b1068	ml-project	The information required to uniquely identify this workspace is stored in the file config.json in the same folder as the Python script.

You need to obtain a reference to the mi-protect workspace

Solution: Run the following Python code.

```
from azure.ai.ml import MLClient  
ws = MLClient.workspaces.get("ml-project")
```

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

A

Question 10

Question Type: MultipleChoice

You use Azure Machine Learning to train a model.

You must use Bayesian sampling to tune hyperparameters.

You need to select a learning_rate parameter distribution.

Which two distributions can you use? Each correct answer presents a complete solution.

NOTE Each correct selection is worth one point.

Options:

A- Normal

B- Uniform

C- Choice

D- LogUniform

E- QNormal

Answer:

B, C

Question 11

Question Type: MultipleChoice

You use the Azure Machine learning SDK v2 for Python and notebooks to train a model. You use Python code to create a compute target, an environment, and a training script. You need to prepare information to submit a training job.

Which class should you use?

Options:

A- MLClient

B- command

C- BuildContext

D- EndpointConnection

Answer:

B

Question 12

Question Type: MultipleChoice

You create an Azure Machine Learning workspace.

You must use the Python SDK v2 to implement an experiment from a Jupiter notebook in the workspace. The experiment must log string metrics.

You need to implement the method to log the string metrics.

Which method should you use?

Options:

A- mlflow.log-metric()

B- mlflow.log_artifact()

C- mlflow.log_dist()

D- mlflow.log_text()

Answer:

D

To Get Premium Files for DP-100 Visit

<https://www.p2pexams.com/products/dp-100>

For More Free Questions Visit

<https://www.p2pexams.com/microsoft/pdf/dp-100>

