



Free Questions for DP-100 by dumpssheet

Shared by Cervantes on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: OrderList

You manage an Azure Machine Learning workspace. You train a model named model1.

You must identify the features to modify for a differing model prediction result.

You need to configure the Responsible AI (RAI) dashboard for model1.

Which three actions should you perform in sequence? To answer move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Add the explanation component to the Responsible AI Insights dashboard.

Add the error analysis component to the Responsible AI Insights dashboard.

Add the causal component to the Responsible AI Insights dashboard.

Load and configure the Responsible AI Insights dashboard constructor component.

Add the counterfactuals component to the Responsible AI Insights dashboard.

Use the Gather Responsible AI Insights dashboard component to present the dashboard.



Answer Area

Answer:

Use the `continue_training_stopping_criteria`...

Question 2

Question Type: Hotspot

You manage an Azure Machine Learning workspace. You configure an automated machine learning regression training job by using the Azure Machine Learning Python SDK v2. You configure the regression job by using the following script:

```
regression_job.set_limits(  
    timeout_minutes = 60,  
    continue_training_stopping_criteria = 5)
```

Answer Area

Yes if the statement is true. Otherwise, select No.

Answer:

Statements

Yes

The job is terminated if the score is not improving in a specific number of iterations.

A maximum of five AutoML trials are run in parallel during the regression job.

One AutoML trial can run for 60 minutes before it is terminated.

The AutoML trial run can take up to 1 month before it terminates.

Question 3

Question Type: MultipleChoice

You create a binary classification model. You use the Fairlearn package to assess model fairness. You must eliminate the need to retrain the model. You need to implement the Fair learn package. Which algorithm should you use?

Options:

- A- fairlearn.reductions.ExponentiatedGradient
- B- fairlearn.reductions.GridSearch
- C- fairlearn.postprocessing.ThresholdOptimizer
- D- fairlearn.preprocessing.CorrelationRemover

Answer:

D

Question 4

Question Type: MultipleChoice

You have a dataset that contains records of patients tested for diabetes. The dataset includes the patient's age.

You plan to create an analysis that will report the mean age value from the differentially private data derived from the dataset.

You need to identify the epsilon value to use in the analysis that minimizes the risk of exposing the actual data.

Which epsilon value should you use?

Options:

A- -1.5

B- -0.5

C- 0.5

D- 1.5

Answer:

D

Question 5

Question Type: Hotspot

You manage an Azure Machine Learning workspace.

You must define the execution environments for your jobs and encapsulate the dependencies for your code.

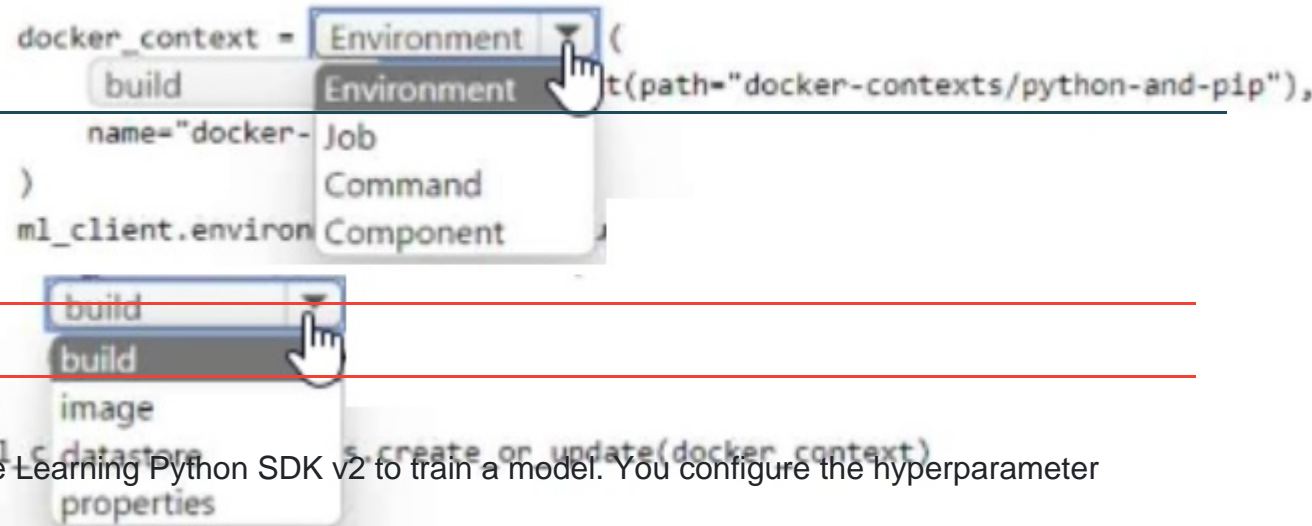
You need to configure the environment from a Docker build context.

How should you complete the code segment? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Answer:



Question 6

Question Type: Hotspot

You are using hyperparameter tuning in Azure Machine Learning Python SDK v2 to train a model. You configure the hyperparameter tuning experiment by running the following code:

```
from azure.ai.ml.sweep import Normal, Uniform

command_job_for_sweep = command_job(
    learning_rate=Normal(10, 3),
    keep_probability=Uniform(0.05, 0.1),
    batch_size=Choice(values=[16, 32, 64, 128]),
    number_of_hidden_layers=Choice(range(3,5))
)
```

s true. Otherwise, select No. NOTE: Each correct selection is worth one

Answer Area

Statements

Yes

Answer:

By defining sampling in this manner, every possible combination of the parameters will be tested.

Random values of the learning_rate parameter will be selected from a normal distribution with a mean of 10 and a standard deviation of 3.

The keep_probability parameter value will always be either **0.05** or **0.1**.

Random values for the number of hidden layers parameter will be selected from a normal distribution with a mean of 3 and a standard deviation of 5.

Question 7

Question Type: MultipleChoice

You build a data pipeline in an Azure Machine Learning workspace by using the Azure Machine Learning SDK for Python.

You need to run a Python script as a pipeline step.

Which two classes could you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

Options:

A- PythonScriptStep

B- AutoMLStep

C- CommandStep

D- StepRun

Answer:

A, C

Question 8

Question Type: MultipleChoice

You train and register an Azure Machine Learning model

You plan to deploy the model to an online endpoint

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 managed online endpoint with the default authentication settings. Deploy the model to the online endpoint.

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

B

Question 9

Question Type: MultipleChoice

You train and register an Azure Machine Learning model

You plan to deploy the model to an online endpoint

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 managed online endpoint and set the value of its auth.mode parameter to aml.token. Deploy the model to the online endpoint.

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

B

Question 10

Question Type: MultipleChoice

You have an Azure Machine Learning workspace. You build a deep learning model.

You need to publish a GPU-enabled model as a web service.

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

NOTE: Each correct selection is worth one point.

Options:

- A- Azure Kubernetes Service (AKS)
- B- Azure Container Instances (ACI)
- C- Local web service
- D- Azure Machine Learning compute clusters

Answer:

A, B

Question 11

Question Type: MultipleChoice

You are using Azure Machine Learning to monitor a trained and deployed model. You implement Event Grid to respond to Azure Machine Learning events.

Model performance has degraded due to model input data changes.

You need to trigger a remediation ML pipeline based on an Azure Machine Learning event.

Which event should you use?

Options:

- A- RunStatusChanged
- B- DatasetDriftDetected
- C- ModelDeployed
- D- RunCompleted

Answer:

B

Question 12

Question Type: MultipleChoice

You create an Azure Machine Learning workspace.

You must configure an event handler to send an email notification when data drift is detected in the workspace datasets. You must minimize development efforts.

You need to configure an Azure service to send the notification.

Which Azure service should you use?

Options:

A- Azure Function apps

B- Azure DevOps pipeline

C- Azure Automation runbook

D- Azure Logic Apps

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>

