Question 1

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 have an Azure Machine Learning workspace. You connect to a terminal session from the Notebooks page in Azure Machine Learning studio.

You plan to add a new Jupyter kernel that will be accessible from the same terminal session.

You need to perform the task that must be completed before you can add the new kernel.

Solution: Delete the Python 3.6 - AzureML kernel.

Does the solution meet the goal?

Options:

A- Yes

B- No

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 have an Azure Machine Learning workspace. You connect to a terminal session from the Notebooks page in Azure Machine Learning studio.

You plan to add a new Jupyter kernel that will be accessible from the same terminal session.

You need to perform the task that must be completed before you can add the new kernel.

Solution: Create an environment.

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

A

Question 3

Question Type: MultipleChoice

You have an Azure Machine Learning workspace.

You plan to run a job to tram a model as an MLflow model output.

You need to specify the output mode of the MLflow model.

Which three modes can you specify? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

Options:

A- rw_mount

B- ro mount

C- upload

D- download

E- direct

Answer:

B, C, E

Question 4

Question Type: MultipleChoice

You have an Azure Machine Learning workspace.

You plan to use the workspace to set up automated machine learning training for an image classification model.

You need to choose the primary metric to optimize the model training.

Which primary metric should you choose?

Options:

A- r2_score

B- mean_absolute_error

C- accuracy

D- root_mean_squared_log_error

Answer:

C

Question 5

You have an Azure Machine Learning workspace named WS1.

You plan to use the Responsible AI dashboard to assess MLflow models that you will register in WS1.

You need to identify the library you should use to register the MLflow models.

Which library should you use?

Question 6

You have an Azure Machine Learning workspace.

You plan to tune a model hyperparameter when you train the model.

You need to define a search space that returns a normally distributed value.

Which parameter should you use?

Question 7

Question Type: **Hotspot**

You use Azure Machine Learning to train a machine learning model.

You use the following training script in Python to perform logging:

```
import mlflow
mlflow.log_metric("accuracy", float(val_accuracy))
```
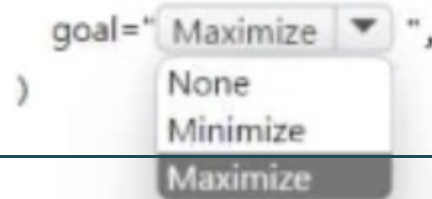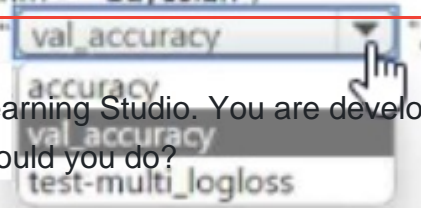
You need to provide the primary metric and goal you want hyperparameter tuning to optimize.

NOTE: Each correct selection is worth one point.

Answer Area

Answer:

```
from azure.ai.ml.sweep import Uniform, Choice
command_job_for_sweep = command_job(
    learning_rate=Uniform(min_value=0.05, max_value=0.1),
    batch_size=Choice(values=[16, 32, 64, 128]),
)
sweep_job = command_job_for_sweep.sweep(
    compute="cpu-cluster",
    sampling_algorithm = "bayesian",
    primary_metric="[ val_accuracy    ▼ ]",
```

| accuracy |
| val_accuracy |
| test-multi_logloss |

```
    goal="[ Maximize   ▼ ]",
)
```

| None |
| Minimize |
| Maximize |

Question 8

Question Type: MultipleChoice

You create a workspace to include a compute instance by using Azure Machine Learning Studio. You are developing a Python SDK v2 notebook in the workspace. You need to use Intellisense in the notebook. What should you do?

Options:

A- Start the compute instance.

B- Run a %pip magic function on the compute instance.

C- Run a !pip magic function on the compute instance.

D- Stop the compute instance.

Answer:

B

Question 9

Question Type: OrderList

You manage an Azure Machine Learning workspace named workspace1 with a compute instance named compute1. You connect to compute! by using a terminal window from wofkspace1. You create a file named "requirements.txt" containing Python dependencies to include Jupyler.

You need to add a new Jupyter kernel to compute1.

Which four commands should you use? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Commands		Answer Area
jupyter run		
conda create -n "python_env"		
conda activate "python_env"	>	
conda install -r "requirements.txt"	<	
ipython kernel install --user --name="python_env"		

Answer:

conda create -n "python_env"
conda activate "python_env"
conda install -r "requirements.txt"
ipython kernel install --user --name="python_env"

Question 10

Question Type: MultipleChoice

You train and publish a machine teaming model.

You need to run a pipeline that retrains the model based on a trigger from an external system.

What should you configure?

Options:

A- Azure Data Catalog

B- Azure Batch

C- Azure logic App

Answer:

C

Question 11

Question Type: **MultipleChoice**

You have an Azure Machine Learning (ML) model deployed to an online endpoint.

You need to review container logs from the endpoint by using Azure Ml Python SDK v2. The logs must include the console log from the inference server with print/log statements from the models scoring script.

What should you do first?

Options:

A- Create an instance of the the MLClient class.

B- Create an instance of the OnlineDeploymentOperations class.

C- Connect by using SSH to the inference server.

D- Connect by using Docker tools to the inference server.

Answer:

A

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

20% DISCOUNT