Free Questions for DP-100 by certscare

Shared by Blankenship on 06-06-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You have a dataset that includes confidential dat

a. You use the dataset to train a model.

You must use a differential privacy parameter to keep the data of individuals safe and private.

You need to reduce the effect of user data on aggregated results.

What should you do?

Options:

A- Decrease the value of the epsilon parameter to reduce the amount of noise added to the data

B- Increase the value of the epsilon parameter to decrease privacy and increase accuracy

C- Decrease the value of the epsilon parameter to increase privacy and reduce accuracy

D- Set the value of the epsilon parameter to 1 to ensure maximum privacy

Answer:

C

Explanation:

Differential privacy tries to protect against the possibility that a user can produce an indefinite number of reports to eventually reveal sensitive data. A value known as epsilon measures how noisy, or private, a report is. Epsilon has an inverse relationship to noise or privacy. The lower the epsilon, the more noisy (and private) the data is.

https://docs.microsoft.com/en-us/azure/machine-learning/concept-differential-privacy

Question 2

Question Type: MultipleChoice

You create a binary classification model. The model is registered in an Azure Machine Learning workspace. You use the Azure Machine Learning Fairness SDK to assess the model fairness.

You develop a training script for the model on a local machine.

You need to load the model fairness metrics into Azure Machine Learning studio.

What should you do?

Options:

A- Implement the download_dashboard_by_upload_id function

B- Implement the creace_group_metric_sec function

C- Implement the upload_dashboard_dictionary function

D- Upload the training script

Answer:

C

Explanation:

import azureml.contrib.fairness package to perform the upload:

from azureml.contrib.fairness import upload_dashboard_dictionary, download_dashboard_by_upload_id

https://docs.microsoft.com/en-us/azure/machine-learning/how-to-machine-learning-fairness-aml

Question 3

Question Type: MultipleChoice

You use the Azure Machine Learning designer to create and run a training pipeline.

The pipeline must be run every night to inference predictions from a large volume of files. The folder where the files will be stored is defined as a dataset.

You need to publish the pipeline as a REST service that can be used for the nightly inferencing run.

What should you do?

Options:

A- Create a batch inference pipeline

B- Set the compute target for the pipeline to an inference cluster

C- Create a real-time inference pipeline

D- Clone the pipeline

Answer:

A

Explanation:

Azure Machine Learning Batch Inference targets large inference jobs that are not time-sensitive. Batch Inference provides cost-effective inference compute scaling, with unparalleled throughput for asynchronous applications. It is optimized for high-throughput, fire-and-forget inference over large collections of data.

You can submit a batch inference job by pipeline_run, or through REST calls with a published pipeline.

https://github.com/Azure/MachineLearningNotebooks/blob/master/how-to-use-azureml/machine-learning-pipelines/parallel-run/README.md

Question 4

Question Type: **MultipleChoice**

You train and register a machine learning model. You create a batch inference pipeline that uses the model to generate predictions from multiple data files.

You must publish the batch inference pipeline as a service that can be scheduled to run every night.

You need to select an appropriate compute target for the inference service.

Which compute target should you use?

Options:

A- Azure Machine Learning compute instance

B- Azure Machine Learning compute cluster

C- Azure Kubernetes Service (AKS)-based inference cluster

D- Azure Container Instance (ACI) compute target

Answer:

B

Explanation:

Azure Machine Learning compute clusters is used for Batch inference. Run batch scoring on serverless compute. Supports normal and low-priority VMs. No support for real-time inference.

https://docs.microsoft.com/en-us/azure/machine-learning/concept-compute-target

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 plan to use a Python script to run an Azure Machine Learning experiment. The script creates a reference to the experiment run context, loads data from a file, identifies the set of unique values for the label column, and completes the experiment run:

```
from azureml.core import Run
import pandas as pd

run = Run.get_context()
data = pd.read_csv('data.csv')
label_vals = data['label'].unique()
# Add code to record metrics here
run.complete()
```

The experiment must record the unique labels in the data as metrics for the run that can be reviewed later.

You must add code to the script to record the unique label values as run metrics at the point indicated by the comment.

Solution: Replace the comment with the following code:

run.log_list('Label Values', label_vals)

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

A

Explanation:

run.log_list log a list of values to the run with the given name using log_list.

Example: run.log_list('accuracies', [0.6, 0.7, 0.87])

Note:

Data= pd.read_csv('data.csv')

Data is read into a pandas.DataFrame, which is a two-dimensional, size-mutable, potentially heterogeneous tabular data.

label_vals =data['label'].unique

label_vals contains a list of unique label values.

https://www.element61.be/en/resource/azure-machine-learning-services-complete-toolbox-ai

https://docs.microsoft.com/en-us/python/api/azureml-core/azureml.core.run(class)

https://pandas.pydata.org/docs/reference/api/pandas.DataFrame.html

Question 6

Question Type: **MultipleChoice**

You have a Jupyter Notebook that contains Python code that is used to train a model.

You must create a Python script for the production deployment. The solution must minimize code maintenance.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

Options:

A- Refactor the Jupyter Notebook code into functions

B- Save each function to a separate Python file

C- Define a main() function in the Python script

D- Remove all comments and functions from the Python script

Answer:

A, C

Explanation:

https://www.guru99.com/learn-python-main-function-with-examples-understand-main.html

https://towardsdatascience.com/from-jupyter-notebook-to-deployment-a-straightforward-example-1838c203a437