

70-774.exam.20q

Number: 70-774
Passing Score: 0
Time Limit: 120 min

70-774

Perform Cloud Data Science with Azure Machine Learning

Exam A

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are working on an Azure Machine Learning experiment.

You have the dataset configured as shown in the following table.

Model	Mean absolute error (MAE)
Boosted decision tree	.2
Relative absolute error (RAE)	.43

You need to ensure that you can compare the performance of the models and add annotations to the results.

Solution: You consolidate the output of the Score Model modules by using the Add Rows module, and then use the Execute R Script module.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://msdn.microsoft.com/en-us/library/azure/dn905915.aspx>

QUESTION 2

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are working on an Azure Machine Learning experiment.

You have the dataset configured as shown in the following table.

Model	Mean absolute error (MAE)
Boosted decision tree	.2
Relative absolute error (RAE)	.43

You need to ensure that you can compare the performance of the models and add annotations to the results.

Solution: You connect the Score Model modules from each trained model as inputs for the Evaluate Model module, and then save the results as a dataset.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://msdn.microsoft.com/en-us/library/azure/dn905915.aspx>

QUESTION 3

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are working on an Azure Machine Learning experiment.

You have the dataset configured as shown in the following table.

Model	Mean absolute error (MAE)
Boosted decision tree	.2
Relative absolute error (RAE)	.43

You need to ensure that you can compare the performance of the models and add annotations to the results.

Solution: You connect the Score Model modules from each trained model as inputs for the Evaluate Model module, and use the Execute R Script module.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://msdn.microsoft.com/en-us/library/azure/dn905915.aspx>

QUESTION 4

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are working on an Azure Machine Learning experiment.

You have the dataset configured as shown in the following table.

Model	Mean absolute error (MAE)
Boosted decision tree	.2
Relative absolute error (RAE)	.43

You need to ensure that you can compare the performance of the models and add annotations to the results.

Solution: You save the output of the Score Model modules as a combined set, and then use the Project Columns module to select the MAE.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://msdn.microsoft.com/en-us/library/azure/dn905915.aspx>

QUESTION 5

You have data about the following:

- Users
- Movies
- User ratings of the movies

You need to predict whether a user will like a particular movie.

Which Matchbox recommender should you use?

- A. Item Recommendation
- B. Related Items
- C. Rating Prediction
- D. Related Users

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://msdn.microsoft.com/en-us/library/azure/dn905970.aspx#RatingPredictionOptions>

QUESTION 6

You have the following three training datasets for a restaurant:

- User features
- Item features

- Ratings of items by users

You must recommend restaurant to a particular user based only on the users features.

You need to use a Matchbox Recommender to make recommendations.

How many input parameters should you specify?

- A. 1
- B. 2
- C. 3
- D. 4

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://msdn.microsoft.com/en-us/library/azure/dn905987.aspx>

QUESTION 7

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

A travel agency named Margie's Travel sells airline tickets to customers in the United States.

Margie's Travel wants you to provide insights and predictions on flight delays. The agency is considering implementing a system that will communicate to its customers as the flight departure nears about possible delays due to weather conditions. The flight data contains the following attributes:

- DepartureDate: The departure date aggregated at a per hour granularity
- Carrier: The code assigned by the IATA and commonly used to identify a carrier
- OriginAirportID: An identification number assigned by the USDOT to identify a unique airport (the flight's origin)
- DestAirportID: An identification number assigned by the USDOT to identify a unique airport (the flight's destination)
- DepDel: The departure delay in minutes
- DepDel30: A Boolean value indicating whether the departure was delayed by 30 minutes or more (a value of 1 indicates that the departure was delayed by 30 minutes or more)

The weather data contains the following attributes: AirportID, ReadingDate (YYYY/MM/DD HH), SkyConditionVisibility, WeatherType, WindSpeed, StationPressure, PressureChange, and HourlyPrecip.

You need to use historical data about on-time flight performance and the weather data to predict whether the departure of a scheduled flight will be delayed by more than 30 minutes.

Which method should you use?

- A. clustering
- B. linear regression
- C. classification
- D. anomaly detection

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://gallery.cortanaintelligence.com/Experiment/Binary-Classification-Flight-delay-prediction-3>

QUESTION 8

DRAG DROP

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

A travel agency named Margie's Travel sells airline tickets to customers in the United States.

Margie's Travel wants you to provide insights and predictions on flight delays. The agency is considering implementing a system that will communicate to its customers as the flight departure nears about possible delays due to weather conditions. The flight data contains the following attributes:

- **DepartureDate:** The departure date aggregated at a per hour granularity
- **Carrier:** The code assigned by the IATA and commonly used to identify a carrier
- **OriginAirportID:** An identification number assigned by the USDOT to identify a unique airport (the flight's origin)
- **DestAirportID:** An identification number assigned by the USDOT to identify a unique airport (the flight's destination)
- **DepDel:** The departure delay in minutes
- **DepDel30:** A Boolean value indicating whether the departure was delayed by 30 minutes or more (a value of 1 indicates that the departure was delayed by 30 minutes or more)

The weather data contains the following attributes: **AirportID**, **ReadingDate** (YYYY/MM/DD HH), **SkyCondition**, **Visibility**, **WeatherType**, **WindSpeed**, **StationPressure**, **PressureChange**, and **HourlyPrecip**.

You need to remove the bias and to identify the columns in the input dataset that have the greatest predictive power.

Which module should you use for each requirement? To answer, drag the appropriate modules to the correct requirements. Each module may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Modules

Cross-validate Model
Evaluate Model
Filter and Sample
Filter Based Feature Selection Module
Parameter Sweep
Tune Model Hyperparameters

Answer Area

Remove bias:	Module
Identify the columns that have the greatest predictive power:	Module

Correct Answer: