QUESTION 1
A company employs a team of customer service agents to provide telephone and email support to customers.
The company develops a webchat bot to provide automated answers to common customer queries.
Which business benefit should the company expect as a result of creating the webchat bot solution?

A. increased sales
B. a reduced workload for the customer service agents
C. improved product reliability

Answer: B

QUESTION 2
For a machine learning progress, how should you split data for training and evaluation?

A. Use features for training and labels for evaluation.
B. Randomly split the data into rows for training and rows for evaluation.
C. Use labels for training and features for evaluation.
D. Randomly split the data into columns for training and columns for evaluation.

Answer: D
Explanation:
In Azure Machine Learning, the percentage split is the available technique to split the data. In this technique, random data of a given percentage will be split to train and test data.
Reference:
https://www.sqlshack.com/prediction-in-azure-machine-learning/

QUESTION 3
You build a machine learning model by using the automated machine learning user interface (UI).
You need to ensure that the model meets the Microsoft transparency principle for responsible AI.
What should you do?

A. Set Validation type to Auto.
B. Enable Explain best model.
C. Set Primary metric to accuracy.
D. Set Max concurrent iterations to 0.

Answer: B
Explanation:
Model Explain Ability.
Most businesses run on trust and being able to open the ML "black box" helps build transparency and trust. In heavily

regulated industries like healthcare and banking, it is critical to comply with regulations and best practices. One key aspect of this is understanding the relationship between input variables (features) and model output. Knowing both the magnitude and direction of the impact each feature (feature importance) has on the predicted value helps better understand and explain the model. With model explain ability, we enable you to understand feature importance as part of automated ML runs.
Reference:
https://azure.microsoft.com/en-us/blog/new-automated-machine-learning-capabilities-in-azure-machine-learning-service/

QUESTION 4
You are designing an AI system that empowers everyone, including people who have hearing, visual, and other impairments.
This is an example of which Microsoft guiding principle for responsible AI?

A. fairness
B. inclusiveness
C. reliability and safety
D. accountability

Answer: B
Explanation:
Inclusiveness: At Microsoft, we firmly believe everyone should benefit from intelligent technology, meaning it must incorporate and address a broad range of human needs and experiences. For the 1 billion people with disabilities around the world, AI technologies can be a game-changer.
Reference:
https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding-principles

QUESTION 5
You are building an AI system.
Which task should you include to ensure that the service meets the Microsoft transparency principle for responsible AI?

A. Ensure that all visuals have an associated text that can be read by a screen reader.
B. Enable autoscaling to ensure that a service scales based on demand.
C. Provide documentation to help developers debug code.
D. Ensure that a training dataset is representative of the population.

Answer: C
Explanation:
https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding-principles

QUESTION 6
Your company is exploring the use of voice recognition technologies in its smart home devices. The company wants to identify any barriers that might unintentionally leave out specific user groups.
This an example of which Microsoft guiding principle for responsible AI?

A. accountability
B. fairness
C. inclusiveness
D. privacy and security

Answer: C
Explanation:
https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding-principles

QUESTION 7
What are three Microsoft guiding principles for responsible AI? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. knowledgeability
B. decisiveness
C. inclusiveness
D. fairness
E. opinionatedness
F. reliability and safety

Answer: CDF
Explanation:
https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding-principles

QUESTION 8
You run a charity event that involves posting photos of people wearing sunglasses on Twitter.
You need to ensure that you only retweet photos that meet the following requirements:
- Include one or more faces.
- Contain at least one person wearing sunglasses.
What should you use to analyze the images?

A. the Verify operation in the Face service
B. the Detect operation in the Face service
C. the Describe Image operation in the Computer Vision service
D. the Analyze Image operation in the Computer Vision service

Answer: B
Explanation:
https://docs.microsoft.com/en-us/azure/cognitive-services/face/overview

QUESTION 9
Which metric can you use to evaluate a classification model?

A. true positive rate
B. mean absolute error (MAE)
C. coefficient of determination (R2)
D. root mean squared error (RMSE)

Answer: A
Explanation:
What does a good model look like?
An ROC curve that approaches the top left corner with 100% true positive rate and 0% false positive rate will be the best model. A random model would display as a flat line from the bottom left to the top right corner. Worse than random would dip below the y=x line.
Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/how-to-understand-automated-ml#classification

QUESTION 10
Which two components can you drag onto a canvas in Azure Machine Learning designer? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. dataset
B. compute
C. pipeline
D. module

AI-900 Exam Dumps AI-900 Exam Questions   AI-900 PDF Dumps   AI-900 VCE Dumps

https://www.braindump2go.com/ai-900.html

Answer: AD
Explanation:
You can drag-and-drop datasets and modules onto the canvas.
Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/concept-designer

QUESTION 11
You need to create a training dataset and validation dataset from an existing dataset.
Which module in the Azure Machine Learning designer should you use?

A. Select Columns in Dataset
B. Add Rows
C. Split Data
D. Join Data

Answer: C
Explanation:
A common way of evaluating a model is to divide the data into a training and test set by using Split Data, and then validate the model on the training data.
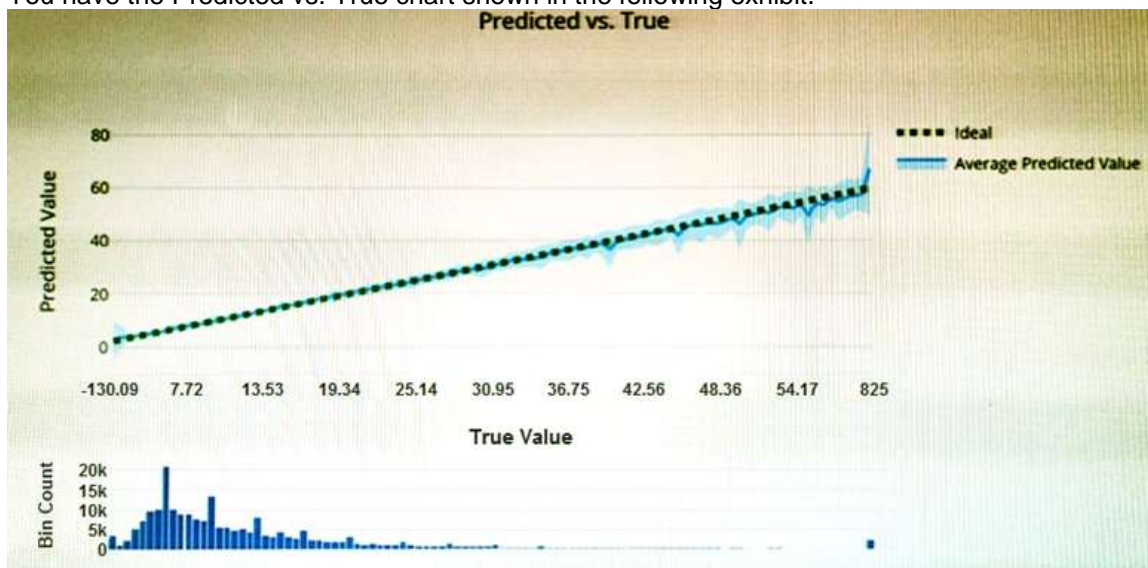Use the Split Data module to divide a dataset into two distinct sets.
The studio currently supports training/validation data splits Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/how-to-configure-cross-validation-data-splits2

QUESTION 12
You have the Predicted vs. True chart shown in the following exhibit.

Which type of model is the chart used to evaluate?

A. classification
B. regression
C. clustering

Answer: B
Explanation:
What is a Predicted vs. True chart?
Predicted vs. True shows the relationship between a predicted value and its correlating true value for a regression problem. This graph can be used to measure performance of a model as the closer to the y=x line the predicted values are, the better the accuracy of a predictive model.

AI-900 Exam Dumps   AI-900 Exam Questions   AI-900 PDF Dumps   AI-900 VCE Dumps

https://www.braindump2go.com/ai-900.html

Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/how-to-understand-automated-m

QUESTION 13
Which type of machine learning should you use to predict the number of gift cards that will be sold next month?

A. classification
B. regression
C. clustering

Answer: C
Explanation:
Clustering, in machine learning, is a method of grouping data points into similar clusters. It is also called segmentation.
Over the years, many clustering algorithms have been developed. Almost all clustering algorithms use the features of individual items to find similar items. For example, you might apply clustering to find similar people by demographics.
You might use clustering with text analysis to group sentences with similar topics or sentiment.
Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/machine-learning-initialize-model-clustering

QUESTION 14
You have a dataset that contains information about taxi journeys that occurred during a given period.
You need to train a model to predict the fare of a taxi journey.
What should you use as a feature?

A. the number of taxi journeys in the dataset
B. the trip distance of individual taxi journeys
C. the fare of individual taxi journeys
D. the trip ID of individual taxi journeys

Answer: B
Explanation:
The label is the column you want to predict. The identified Featuresare the inputs you give the model to predict the Label.
Example:
The provided data set contains the following columns:
vendor_id: The ID of the taxi vendor is a feature.
rate_code: The rate type of the taxi trip is a feature.
passenger_count: The number of passengers on the trip is a feature.
trip_time_in_secs: The amount of time the trip took. You want to predict the fare of the trip before the trip is completed. At that moment, you don't know how long the trip would take. Thus, the trip time is not a feature and you'll exclude this column from the model.
trip_distance: The distance of the trip is a feature.
payment_type: The payment method (cash or credit card) is a feature.
fare_amount: The total taxi fare paid is the label.
Reference:
https://docs.microsoft.com/en-us/dotnet/machine-learning/tutorials/predict-prices

QUESTION 15
You need to predict the sea level in meters for the next 10 years.
Which type of machine learning should you use?

A. classification
B. regression
C. clustering

Answer: B
Explanation:
In the most basic sense, regression refers to prediction of a numeric target.
Linear regression attempts to establish a linear relationship between one or more independent variables and a numeric outcome, or dependent variable.
You use this module to define a linear regression method, and then train a model using a labeled dataset.
The trained model can then be used to make predictions.
Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/linear-regression