DBCLOUD
EXAMS

QUESTIONS & ANSWERS

verified



Technology Provider: Microsoft
Exam: Microsoft Azure AI Fundamentals
Number of questions in the database: 307
Exam Version: NOV.2026

Topic 1 - Single Topic

Question #1*Topic 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

Correct Answer: *B*

Respuesta por mayor votación de la comunidad

B (100%)

Question #2*Topic 1*

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.

Correct Answer: *B*
The Split Data module is particularly useful when you need to separate data into training and testing sets. Use the Split Rows option if you want to divide the data into two parts. You can specify the percentage of data to put in each split, but by default, the data is divided 50-50. You can also randomize the selection of rows in each group, and use stratified sampling.
Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/split-data

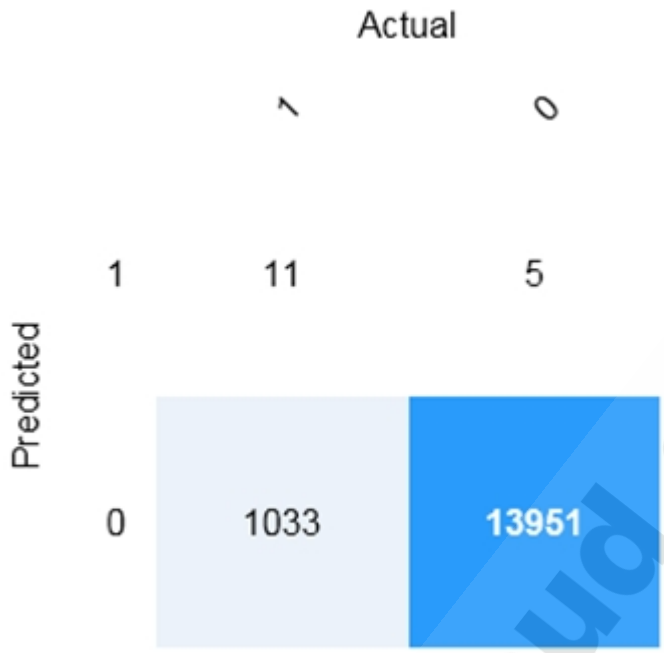
Respuesta por mayor votación de la comunidad

B (100%)

https://dbcloud.academy/

Question #3*Topic 1*

HOTSPOT -
You are developing a model to predict events by using classification.
You have a confusion matrix for the model scored on test data as shown in the following exhibit.



Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

https://dbcloud.academy/

Hot Area:

Answer Area

There are **[answer choice]** correctly predicted positives.

- 5
- 11
- 1,033
- 13,951

There are **[answer choice]** false negatives.

- 5
- 11
- 1,033
- 13,951

Correct

Answer Area

There are **[answer choice]** correctly predicted positives.

- 5
- **11**
- 1,033
- 13,951

There are **[answer choice]** false negatives.

- 5
- 11
- **1,033**
- 13,951

Answer:

Box 1: 11 -

https://dbcloud.academy/

	Predicted	
	Positive	Negative
Actual True	TP	FN
Actual False	FP	TN

TP = True Positive.
The class labels in the training set can take on only two possible values, which we usually refer to as positive or negative. The positive and negative instances that a classifier predicts correctly are called true positives (TP) and true negatives (TN), respectively. Similarly, the incorrectly classified instances are called false positives (FP) and false negatives (FN).

Box 2: 1,033 -

FN = False Negative -
Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/studio/evaluate-model-performance

Question #4*Topic 1*

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.

Correct Answer: *B*
Model Explain Ability.
Most businesses run on trust and being able to open the ML ɹ€black boxɹ€ 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/

https://dbcloud.academy/

Respuesta por mayor votación de la comunidad

B (100%)

Question #5*Topic 1*

HOTSPOT -
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Statements	Yes	No
Forecasting housing prices based on historical data is an example of anomaly detection.	○	○
Identifying suspicious sign-ins by looking for deviations from usual patterns is an example of anomaly detection.	○	○
Predicting whether a patient will develop diabetes based on the patient's medical history is an example of anomaly detection.	○	○

Correct Answer:

Answer Area

Statements	Yes	No
Forecasting housing prices based on historical data is an example of anomaly detection.	○	○
Identifying suspicious sign-ins by looking for deviations from usual patterns is an example of anomaly detection.	○	○
Predicting whether a patient will develop diabetes based on the patient's medical history is an example of anomaly detection.	○	○

Anomaly detection encompasses many important tasks in machine learning:
Identifying transactions that are potentially fraudulent.
Learning patterns that indicate that a network intrusion has occurred.
Finding abnormal clusters of patients.
Checking values entered into a system.
Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/anomaly-detection

https://dbcloud.academy/

Question #6*Topic 1*

HOTSPOT -
To complete the sentence, select the appropriate option in the answer area.
Hot Area:

Answer Area

The handling of unusual or missing values provided to an AI system is a consideration for the Microsoft [ ] principle for responsible AI.

- inclusiveness
- privacy and security
- reliability and safety
- transparency

Correct Answer:

Answer Area

The handling of unusual or missing values provided to an AI system is a consideration for the Microsoft [reliability and safety] principle for responsible AI.

Reliability and safety:
AI systems need to be reliable and safe in order to be trusted. It is important for a system to perform as it was originally designed and for it to respond safely to new situations. Its inherent resilience should resist intended or unintended manipulation. Rigorous testing and validation should be established for operating conditions to ensure that the system responds safely to edge cases, and A/B testing and champion/challenger methods should be integrated into the evaluation process.
An AI system's performance can degrade over time, so a robust monitoring and model tracking process needs to be established to reactively and proactively measure the model's performance and retrain it, as necessary, to modernize it.
Reference:
https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai

Question #7*Topic 1*

https://dbcloud.academy/

DRAG DROP -
Match the types of AI workloads to the appropriate scenarios.
To answer, drag the appropriate workload type from the column on the left to its scenario on the right. Each workload type may be used once, more than once, or not at all.
NOTE: Each correct selection is worth one point.
Select and Place:

Workloads Types	Answer Area	
Anomaly detection	Workload Type	An automated chat to answer questions about refunds and exchange
Computer vision	Workload Type	Determining whether a photo contains a person
Conversational AI	Workload Type	Determining whether a review is positive or negative
Knowledge mining		
Natural language processing		

Correct Answer:

Workloads Types	Answer Area	
Anomaly detection	Conversational AI	An automated chat to answer questions about refunds and exchange
Computer vision	Computer vision	Determining whether a photo contains a person
Conversational AI	Natural language processing	Determining whether a review is positive or negative
Knowledge mining		
Natural language processing		

Box 3: Natural language processing
Natural language processing (NLP) is used for tasks such as sentiment analysis, topic detection, language detection, key phrase extraction, and document categorization.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/natural-language-processing

Question #8*Topic 1*

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

https://dbcloud.academy/

- C. reliability and safety
- D. accountability

Correct Answer: *B*
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

Respuesta por mayor votación de la comunidad

B (100%)

Question #9*Topic 1*

DRAG DROP -
Match the Microsoft guiding principles for responsible AI to the appropriate descriptions.
To answer, drag the appropriate principle from the column on the left to its description on the right. Each principle may be used once, more than once, or not at all.
NOTE: Each correct selection is worth one point.
Select and Place:

Principles	Answer Area	
Accountability	Principle	Ensure that AI systems operate as they were originally designed, respond to unanticipated conditions, and resist harmful manipulation.
Fairness	Principle	Implementing processes to ensure that decisions made by AI systems can be overridden by humans.
Inclusiveness	Principle	Provide consumers with information and controls over the collection, use, and storage of their data.
Privacy and security		
Reliability and safety		

Correct Answer:

Principles	Answer Area	
Accountability	Reliability and safety	Ensure that AI systems operate as they were originally designed, respond to unanticipated conditions, and resist harmful manipulation.
Fairness	Accountability	Implementing processes to ensure that decisions made by AI systems can be overridden by humans.
Inclusiveness	Privacy and security	Provide consumers with information and controls over the collection, use, and storage of their data.
Privacy and security		
Reliability and safety		

https://dbcloud.academy/