



## Microsoft

### Exam Questions AI-900

Microsoft Azure AI Fundamentals (beta)

**NEW QUESTION 1**

- (Exam 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

**Answer: B**

**NEW QUESTION 2**

- (Exam 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.

**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>

**NEW QUESTION 3**

- (Exam Topic 1)

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.

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	
Privacy and security	Principle: Provide consumers with information and controls over the collection, use, and storage of their data.
Reliability and safety	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Reliability and safety

To build trust, it's critical that AI systems operate reliably, safely, and consistently under normal circumstances and in unexpected conditions. These systems should be able to operate as they were originally designed, respond safely to unanticipated conditions, and resist harmful manipulation.

Box 2: Fairness

Fairness: AI systems should treat everyone fairly and avoid affecting similarly situated groups of people in different ways. For example, when AI systems provide guidance on medical treatment, loan applications, or employment, they should make the same recommendations to everyone with similar symptoms, financial circumstances, or professional qualifications.

We believe that mitigating bias starts with people understanding the implications and limitations of AI predictions and recommendations. Ultimately, people should supplement AI decisions with sound human judgment and be held accountable for consequential decisions that affect others.

Box 3: Privacy and security

As AI becomes more prevalent, protecting privacy and securing important personal and business information is becoming more critical and complex. With AI, privacy and data security issues require especially close attention because access to data is essential for AI systems to make accurate and informed predictions and decisions about people. AI systems must comply with privacy laws that require transparency about the collection, use, and storage of data and mandate that consumers have appropriate controls to choose how their data is used

Reference:

<https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding-principles>

**NEW QUESTION 4**

- (Exam Topic 2)

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>

**NEW QUESTION 5**

- (Exam Topic 2)

You use Azure Machine Learning designer to publish an inference pipeline.

Which two parameters should you use to consume the pipeline? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. the model name
- B. the training endpoint
- C. the authentication key
- D. the REST endpoint

**Answer:** AD

**Explanation:**

A: The trained model is stored as a Dataset module in the module palette. You can find it under My Datasets. Azure Machine Learning designer lets you visually connect datasets and modules on an interactive canvas to create machine learning models.

D: You can consume a published pipeline in the Published pipelines page. Select a published pipeline and find the REST endpoint of it.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-run-batch-predictions-designer> <https://docs.microsoft.com/en-us/azure/machine-learning/concept-designer>

**NEW QUESTION 6**

- (Exam Topic 2)

To complete the sentence, select the appropriate option in the answer area.

**Answer Area**

A banking system that predicts whether a loan will be repaid is an example of the \_\_\_\_\_ type of machine learning.

- classification
  - regression
  - clustering

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

In the most basic sense, regression refers to prediction of a numeric target.

Example: Regression Model: A Boosted Decision Tree algorithm was used to create and train the model for predicting the repayment rate.

Reference:

<https://gallery.azure.ai/Experiment/Student-Loan-Repayment-Rate-Prediction>

**NEW QUESTION 10**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### AI-900 Practice Exam Features:

- \* AI-900 Questions and Answers Updated Frequently
- \* AI-900 Practice Questions Verified by Expert Senior Certified Staff
- \* AI-900 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* AI-900 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
[Order The AI-900 Practice Test Here](#)