



# AI-900<sup>Q&As</sup>

Microsoft Azure AI Fundamentals

## Pass Microsoft AI-900 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/ai-900.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





### QUESTION 1

#### HOTSPOT

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

Hot Area:

#### Answer Area

<input type="text"/>	▼	is used to generate additional features.
Feature engineering		
Feature selection		
Model evaluation		
Model training		

Correct Answer:

#### Answer Area

<input type="text"/>	▼	is used to generate additional features.
Feature engineering		
Feature selection		
Model evaluation		
Model training		

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/team-data-science-process/create-features>

---

### QUESTION 2

You are developing a natural language processing solution in Azure. The solution will analyze customer reviews and determine how positive or negative each review is. This is an example of which type of natural language processing workload?

- A. language detection
- B. sentiment analysis
- C. key phrase extraction



D. entity recognition

Correct Answer: B

Sentiment Analysis is the process of determining whether a piece of writing is positive, negative or neutral.

Reference: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/natural-language-processing>

---

### QUESTION 3

#### 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
You can use QnA Maker to query an Azure SQL database.	<input type="radio"/>	<input type="radio"/>
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	<input type="radio"/>	<input type="radio"/>
The QnA Maker service can determine the intent of a user utterance.	<input type="radio"/>	<input type="radio"/>

Correct Answer:



### Answer Area

Statements	Yes	No
You can use QnA Maker to query an Azure SQL database.	<input type="radio"/>	<input checked="" type="radio"/>
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	<input checked="" type="radio"/>	<input type="radio"/>
The QnA Maker service can determine the intent of a user utterance.	<input type="radio"/>	<input checked="" type="radio"/>

Reference: <https://docs.microsoft.com/en-gb/azure/cognitive-services/qnamaker/concepts/data-sources-and-content>  
<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/choose-natural-language-processing-service>

### QUESTION 4

#### 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
A validation set includes the set of input examples that will be used to train a model.	<input type="radio"/>	<input type="radio"/>
A validation set can be used to determine how well a model predicts labels.	<input type="radio"/>	<input type="radio"/>
A validation set can be used to verify that all the training data was used to train the model.	<input type="radio"/>	<input type="radio"/>

Correct Answer:



## Answer Area

Statements	Yes	No
A validation set includes the set of input examples that will be used to train a mode.	<input type="radio"/>	<input checked="" type="radio"/>
A validation set can be used to determine how well a model predicts labels.	<input checked="" type="radio"/>	<input type="radio"/>
A validation set can be used to verify that all the training data was used to train the model.	<input type="radio"/>	<input checked="" type="radio"/>

Box 1: No

The validation dataset is different from the test dataset that is held back from the training of the model.

Box 2: Yes

A validation dataset is a sample of data that is used to give an estimate of model skill while tuning model's hyperparameters.

Box 3: No

The Test Dataset, not the validation set, used for this. The Test Dataset is a sample of data used to provide an unbiased evaluation of a final model fit on the training dataset.

Reference:

<https://machinelearningmastery.com/difference-test-validation-datasets/>

---

## QUESTION 5

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

Correct Answer: B

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>

---

### QUESTION 6

Which two actions are performed during the data ingestion and data preparation stage of an Azure Machine Learning process? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Calculate the accuracy of the model.
- B. Score test data by using the model.
- C. Combine multiple datasets.
- D. Use the model for real-time predictions.
- E. Remove records that have missing values.

Correct Answer: CE

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/concept-data-ingestion>  
<https://docs.microsoft.com/en-us/azure/architecture/data-science-process/prepare-data>

---

### QUESTION 7

DRAG DROP

Match the Azure Cognitive Services service to the appropriate actions. To answer, drag the appropriate service from the column on the left to its action on the right. Each service may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Select and Place:



language service	Convert spoken requests into text. Service
Speech	Identify the intent of a user's requests Service
Transtator	Apply intent to entities and utterances Service

Correct Answer:

language service	Convert spoken requests into text. Speech
Speech	Identify the intent of a user's requests language service
Transtator	Apply intent to entities and utterances language service

Box 1: Speech

Custom Speech: Code-free automated machine learning for speech recognition

Speech to text is a Speech service feature that accurately transcribes spoken audio to text.

Make spoken audio actionable

Quickly and accurately transcribe audio to text in more than 100 languages and variants. Customize models to enhance accuracy for domain-specific terminology. Get more value from spoken audio by enabling search or analytics on

transcribed text or facilitating action—all in your preferred programming language.

Box 2: Language service



Add intents to your LUIS app to identify groups of questions or commands that have the same intention.

Note: Language understanding (LU) is a very centric component to enable conversational services such as bots, IoT experiences, analytics, and others. In a spoken dialog system, LU converts from the words in a sentence into a machine-

readable meaning representation, typically indicating the intent of the sentence and any present entities. For example, consider a physical fitness domain, with a dialog system embedded in a wearable device like a watch. This dialog system

could recognize intents like StartActivity and StopActivity, and could recognize entities like ActivityType. In the user input "begin a jog", the goal of LU is to identify the intent as StartActivity, and identify the entity ActivityType= "jog".

Box 3: Language service

Intent compared to entity

The intent represents the action the application should take for the user, based on the entire utterance. An utterance can have only one top-scoring intent, but it can have many entities.

Create an intent when the user's intention would trigger an action in your client application, like a call to the checkweather() function from the table above. Then create entities to represent parameters required to execute the action.

Reference: <https://azure.microsoft.com/en-us/services/cognitive-services/speech-to-text>

<https://azure.microsoft.com/en-us/blog/luis-ai-automated-machine-learning-for-custom-language-understanding/>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/concepts/intents>

---

## QUESTION 8

### DRAG DROP

Match the types of computer vision 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**

- Facial recognition
- Image classification
- Object detection
- Optical character recognition (OCR)

**Answer Area**

- Workload Type Identify celebrities in images.
- Workload Type Extract movie title names from movie poster images.
- Workload Type Locate vehicles in images.

Correct Answer:

**Workloads Types**

- Image classification

**Answer Area**

- Facial recognition Identify celebrities in images.
- Optical character recognition (OCR) Extract movie title names from movie poster images.
- Object detection Locate vehicles in images.

Box 1: Facial recognition Face detection that perceives faces and attributes in an image; person identification that matches an individual in your private repository of up to 1 million people; perceived emotion recognition that detects a range of facial expressions like happiness, contempt, neutrality, and fear; and recognition and grouping of similar faces in images.

Box 2: OCR

Box 3: Object detection Object detection is similar to tagging, but the API returns the bounding box coordinates (in pixels) for each object found. For example, if an image contains a dog, cat and person, the Detect operation will list those objects together with their coordinates in the image. You can use this functionality to process the relationships between the objects in an image. It also lets you determine whether there are multiple instances of the same tag in an image.

The Detect API applies tags based on the objects or living things identified in the image. There is currently no formal relationship between the tagging taxonomy and the object detection taxonomy. At a conceptual level, the Detect API only finds objects and living things, while the Tag API can also include contextual terms like "indoor", which can't be localized with bounding boxes.

Reference: <https://azure.microsoft.com/en-us/services/cognitive-services/face/>

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-object-detection>

**QUESTION 9**

You have an AI solution that provides users with the ability to control smart devices by using verbal commands.



Which two types of natural language processing (NLP) workloads does the solution use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. text-to-speech
- B. key phrase extraction
- C. speech-to-text
- D. language modeling
- E. translation

Correct Answer: BC

Key phrase extraction is one of the features offered by Azure Cognitive Service for Language, a collection of machine learning and AI algorithms in the cloud for developing intelligent applications that involve written language. Use key phrase extraction to quickly identify the main concepts in text. For example, in the text "The food was delicious and the staff were wonderful.", key phrase extraction will return the main topics: "food" and "wonderful staff".

Reference: <https://docs.microsoft.com/en-us/azure/cognitive-services/language-service/key-phrase-extraction/overview>

---

#### QUESTION 10

Which AI service should you use to create a bot from a frequently asked questions (FAQ) document?

- A. QnA Maker
- B. Language Understanding (LUIS)
- C. Text Analytics
- D. Speech

Correct Answer: A

---

#### QUESTION 11

DRAG DROP

Match the types of machine learning to the appropriate scenarios.

To answer, drag the appropriate machine learning type from the column on the left to its scenario on the right. Each machine learning type may be used once, more than once, or not at all.

NOTE: Each correct selection is worth one point.

Select and Place:



### Learning Types

- Classification
- Clustering
- Regression

### Answer Area

- Learning Type: Predict how many minutes late a flight will arrive basen on the amount of snowfall at an airport.
- Learning Type: Segment customers into different groups to support a marketing department.
- Learning Type: Predict whether a student will complete a university course.

Correct Answer:

### Learning Types

- Regression
- Clustering
- Classification

### Answer Area

- Regression: Predict how many minutes late a flight will arrive basen on the amount of snowfall at an airport.
- Clustering: Segment customers into different groups to support a marketing department.
- Classification: Predict whether a student will complete a university course.

Box 1: Regression

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.

Box 2: Classification

Classification is a machine learning method that uses data to determine the category, type, or class of an item or row of data.

Box 3: Clustering 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/linear-regression>

## QUESTION 12

You are building a Language Understanding model for an e-commerce business.

You need to ensure that the model detects when utterances are outside the intended scope of the model.



What should you do?

- A. Test the model by using new utterances
- B. Add utterances to the None intent
- C. Create a prebuilt task entity
- D. Create a new model

Correct Answer: B

The None intent is filled with utterances that are outside of your domain.

Reference: <https://docs.microsoft.com/en-us/azure/cognitive-services/LUIS/luis-concept-intent>

### QUESTION 13

You need to predict the income range of a given customer by using the following dataset.

First Name	Last Name	Age	Education Level	Income Range
Orlando	Gee	45	University	25,000-50,000
Keith	Harris	36	High school	25,000-50,000
Donna	Carreras	52	University	50,000-75,000
Janet	Gates	21	University	75,000-100,000
Lucy	Harrington	68	High school	50,000-75,000

Which two fields should you use as features? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Education Level
- B. Last Name
- C. Age
- D. Income Range
- E. First Name

Correct Answer: AC

First Name, Last Name, Age and Education Level are features. Income range is a label (what you want to predict). First Name and Last Name are irrelevant in that they have no bearing on income. Age and Education level are the features you should use.