

100% Money Back
Guarantee

Vendor:Microsoft

Exam Code:AI-900

Exam Name:Microsoft Azure AI Fundamentals

Version:Demo

QUESTION 1

What is a use case for classification?

- A. predicting how many cups of coffee a person will drink based on how many hours the person slept the previous night.
- B. analyzing the contents of images and grouping images that have similar colors
- C. predicting whether someone uses a bicycle to travel to work based on the distance from home to work
- D. predicting how many minutes it will take someone to run a race based on past race times

Correct Answer: C

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

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/algorithm-module-reference/linear-regression>

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/machine-learning-initialize-model-clustering>

QUESTION 2

You need to reduce the load on telephone operators by implementing a chatbot to answer simple questions with predefined answers. Which two AI services should you use to achieve the goal? Each correct answers presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Azure Machine Learning
- B. Azure Bot Service
- C. Language Service
- D. Translator

Correct Answer: AB

QUESTION 3

DRAG DROP

You need to use Azure Machine Learning designer to build a model that will predict automobile prices.

Which type of modules should you use to complete the model? To answer, drag the appropriate modules to the correct locations. 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.

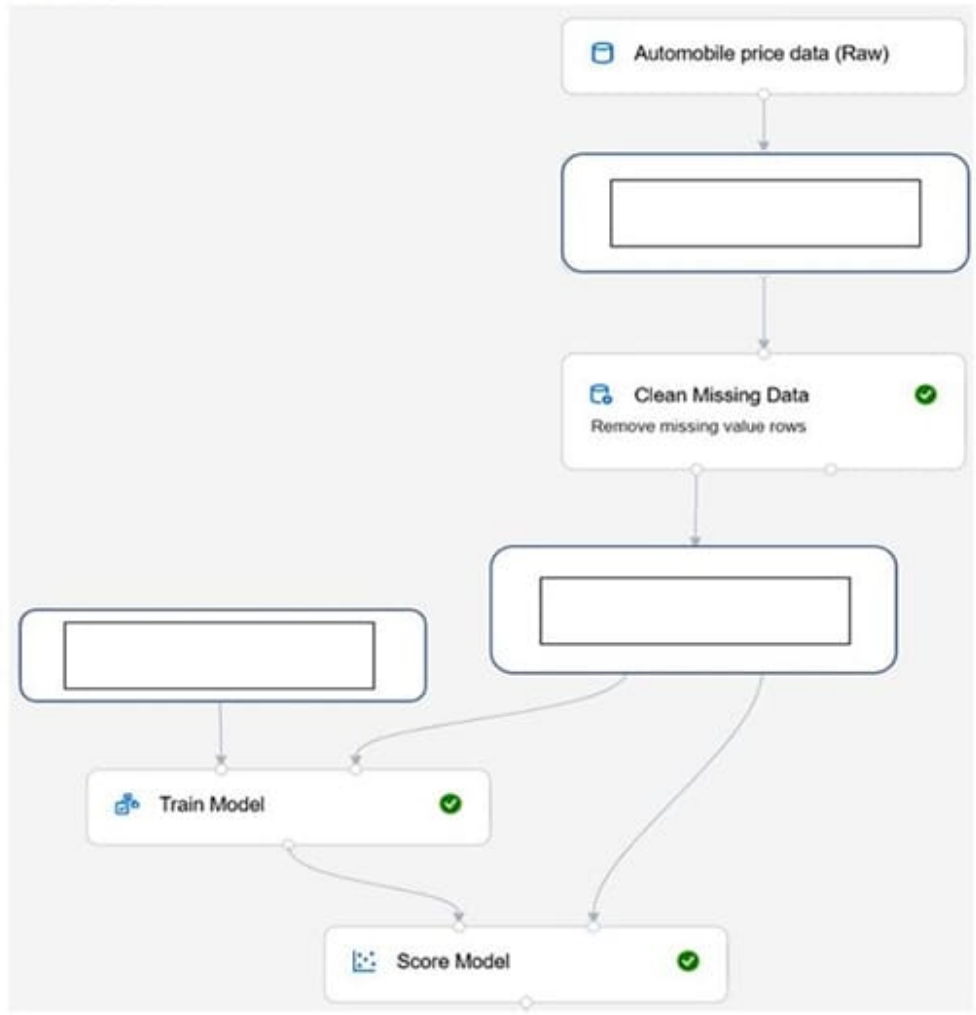
NOTE: Each correct selection is worth one point.

Select and Place:

Modules

- Convert to CSV
- K-Means Clustering
- Linear Regression
- Split Data
- Select Columns in Dataset
- Summarize Data

Answer Area



Correct Answer:

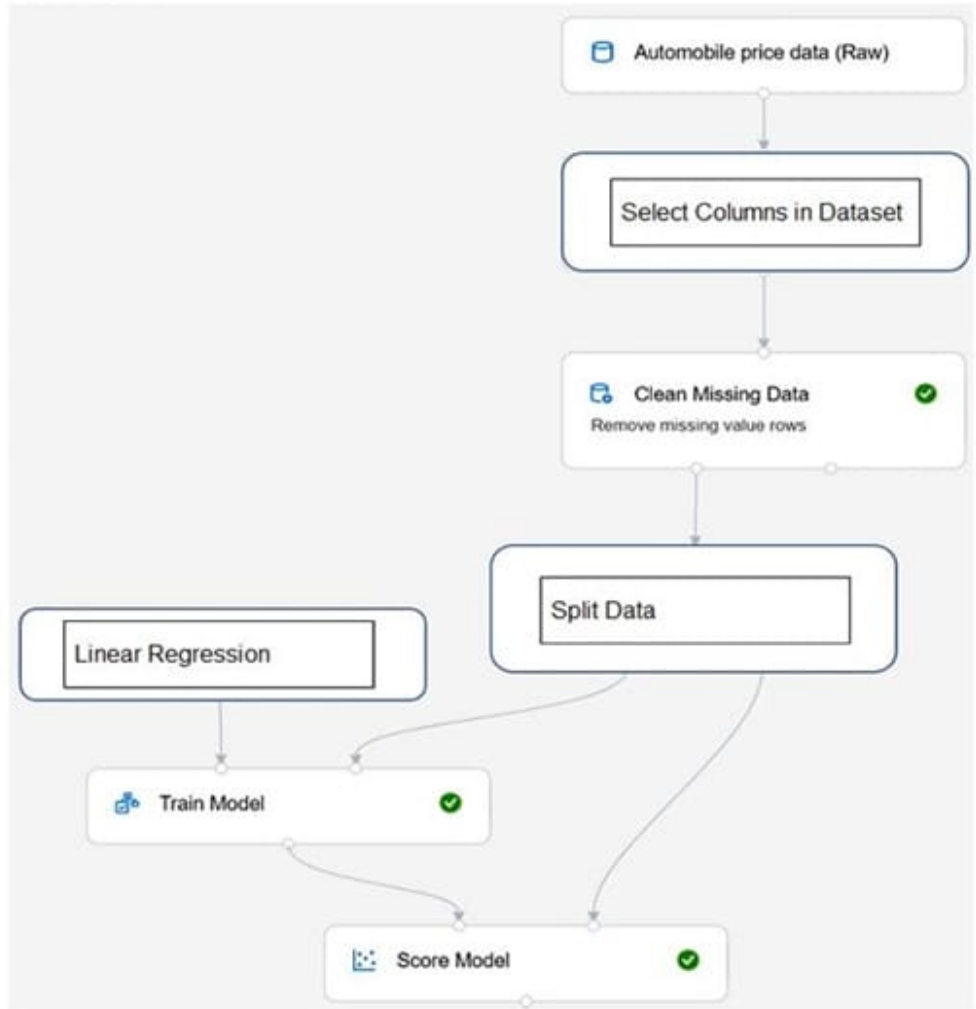
Modules

Convert to CSV

K-Means Clustering

Summarize Data

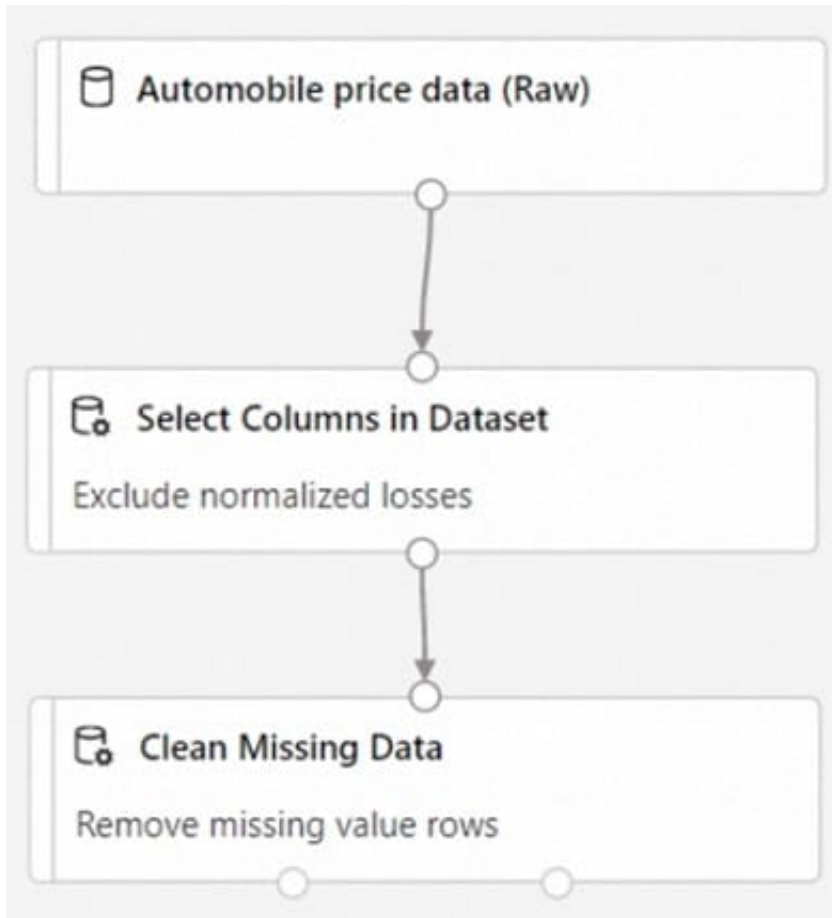
Answer Area



Box 1: Select Columns in Dataset

For Columns to be cleaned, choose the columns that contain the missing values you want to change. You can choose multiple columns, but you must use the same replacement method in all selected columns.

Example:



Box 2: Split data

Splitting data is a common task in machine learning. You will split your data into two separate datasets. One dataset will train the model and the other will test how well the model performed.

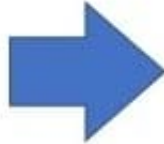
Box 3: Linear regression

Because you want to predict price, which is a number, you can use a regression algorithm. For this example, you use a linear regression model.

QUESTION 4

You use natural language processing to process text from a Microsoft news story. You receive the output shown in the following exhibit.

For weeks now, students and teachers have been settling into the uncharted routine of distance learning. Today I want to thank all of the educators who are connecting classrooms and classmates together in the sudden shift to remote learning. This change requires everyone working together and is unlike anything we've seen in the modern history of education. We've seen countries, school districts and universities move rapidly into remote learning environments with Microsoft Teams being used in 175 countries by 183,000 institutions.



now [DateTime]
students [PersonType]
teachers [PersonType]
distance learning [Skill]
Today [DateTime-Date]
educators [PersonType]
classrooms [Location]
classmates [PersonType]
remote learning [Skill]
history [Skill]
education [Skill]
remote learning [Skill]
Microsoft [Organization]
175 [Quantity-Number]
183,000 [Quantity-Number]

Which type of natural languages processing was performed?

- A. entity recognition
- B. key phrase extraction
- C. sentiment analysis
- D. translation

Correct Answer: A

You can provide the Text Analytics service with unstructured text and it will return a list of entities in the text that it recognizes. You can provide the Text Analytics service with unstructured text and it will return a list of entities in the text that it recognizes. The service can also provide links to more information about that entity on the web. An entity is essentially an item of a particular type or a category; and in some cases, subtype, such as those as shown in the following table. <https://docs.microsoft.com/en-us/learn/modules/analyze-text-with-text-analytics-service/2-get-started-azure>

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/overview>

QUESTION 5

HOTSPOT

You have the following dataset.

Household Income	Postal Code	House Price Category
20,000	55555	Low
23,000	20541	Middle
80,000	87960	High

You plan to use the dataset to train a model that will predict the house price categories of houses.

What are Household Income and House Price Category? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Household Income: ▼

- A feature
- A label

House Price Category: ▼

- A feature
- A label

Correct Answer:

Answer Area

Household Income: ▼

- A feature
- A label

House Price Category: ▼

- A feature
- A label

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/studio/interpret-model-results>

QUESTION 6

You need to create a clustering model and evaluate the model by using Azure Machine Learning designer. What should you do?

- A. Split the original dataset into a dataset for training and a dataset for testing. Use the testing dataset for evaluation.
- B. Use the original dataset for training and evaluation.
- C. Split the original dataset into a dataset for features and a dataset for labels. Use the features dataset for evaluation.
- D. Split the original dataset into a dataset for training and a dataset for testing. Use the training dataset for evaluation.

Correct Answer: A

Understand steps for clustering

You can think of the steps to train and evaluate a clustering machine learning model as:

1.

Prepare data: Identify the features and label in a dataset. Pre-process, or clean and transform, the data as needed.

2.

Train model: Split the data into two groups, a training and a validation set. Train a machine learning model using the training data set. Test the machine learning model for performance using the validation data set.

3.

Evaluate performance.

4.

Deploy a predictive service: After you train a machine learning model, you need to convert the training pipeline into a real-time inference pipeline. Then you can deploy the model as an application on a server or device so that others can use it.

Reference: <https://learn.microsoft.com/en-us/azure/machine-learning/component-reference/train-clustering-model>

QUESTION 7

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

Correct Answer: B

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 8

You are building a knowledge base by using QnA Maker. Which file format can you use to populate the knowledge base?

- A. PPTX
- B. XML
- C. ZIP
- D. PDF

Correct Answer: D

D: Content types of documents you can add to a knowledge base:

Content types include many standard structured documents such as PDF, DOC, and TXT.

Note: The tool supports the following file formats for ingestion:

.tsv: QnA contained in the format Question(tab)Answer.

.txt, .docx, .pdf: QnA contained as regular FAQ content--that is, a sequence of questions and answers.

Incorrect Answers:

A: PPTX is the default presentation file format for new PowerPoint presentations.

B: It is not possible to ingest xml file directly.

Reference: <https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/concepts/data-sources-and-content>

QUESTION 9

Which Computer Vision feature can you use to generate automatic captions for digital photographs?

- A. Recognize text.
- B. Identify the areas of interest.
- C. Detect objects.
- D. Describe the images.

Correct Answer: D

Describe images with human-readable language

Computer Vision can analyze an image and generate a human-readable phrase that describes its contents. The algorithm returns several descriptions based on different visual features, and each description is given a confidence score. The

final output is a list of descriptions ordered from highest to lowest confidence.

The image description feature is part of the Analyze Image API.

Reference: <https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-describing-images>

QUESTION 10

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 restaurant can use a chatbot to answer queries through Cortana.	<input type="checkbox"/>	<input type="checkbox"/>
A restaurant can use a chatbot to answer inquiries about business hours from a webpage.	<input type="checkbox"/>	<input type="checkbox"/>
A restaurant can use a chatbot to automate responses to customer reviews on an external website.	<input type="checkbox"/>	<input type="checkbox"/>

Correct Answer:

Answer Area

Statements	Yes	No
A restaurant can use a chatbot to answer queries through Cortana.	<input checked="" type="radio"/>	<input type="radio"/>
A restaurant can use a chatbot to answer inquiries about business hours from a webpage.	<input checked="" type="radio"/>	<input type="radio"/>
A restaurant can use a chatbot to automate responses to customer reviews on an external website.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: Yes

You can create and build a cortana bot using microsoft bot framework.

Note: Connect Cortana Channels

Login to Azure portal > Select the "All Resources" > Select Channels > Select Cortana icon. Let us start to configure the "Cortana "Channel and follow the below steps, at the end of this article you will be able to deploy the Bot into the

Cortana.

Etc.

Box 2: Yes

Language Service is an easy-to-use web-based service that makes it easy to power a question-answer application or chatbot from semi-structured content like FAQ documents and product manuals. With Language Service, developers can

build, train, and publish question and answer bots in minutes.

Box 3: Yes

Reference: <https://www.c-sharpcorner.com/article/create-and-build-a-cortana-bot-using-microsoft-bot-framework/>

QUESTION 11

You have a bot that identifies the brand names of products in images of supermarket shelves. Which service does the bot use?

- A. AI enrichment for Azure Search capabilities
- B. Computer Vision Image Analysis capabilities
- C. Custom Vision Image Classification capabilities
- D. Language Understanding capabilities

Correct Answer: C

Explanation:

Brand detection is a specialized mode of object detection that uses a database of thousands of global logos to identify commercial brands in images or video. You can use this feature, for example, to discover which brands are most popular

on social media or most prevalent in media product placement.

The Azure AI Vision service detects whether there are brand logos in a given image; if there are, it returns the brand name, a confidence score, and the coordinates of a bounding box around the logo.

The built-in logo database covers popular brands in consumer electronics, clothing, and more. If you find that the brand you're looking for is not detected by the Azure AI Vision service, you could also try creating and training your own logo detector using the Custom Vision service.

Reference:

<https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/concept-brand-detection>

QUESTION 12

Predicting how many vehicles will travel across a bridge on a give day is an example of _____. Select the answer that correctly completes the sentence.

- A. regression
- B. translation
- C. classification
- D. clustering

Correct Answer: A

Regression is a supervised machine learning technique used to predict numeric values.

Reference: <https://docs.microsoft.com/en-us/learn/modules/create-regression-model-azure-machine-learning-designer/>