



Vendor: Microsoft

Exam Code: AI-900

Exam Name: Microsoft Azure AI Fundamentals

Version: DEMO

QUESTION 1

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 Features are 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 2

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 3

Hotspot Question

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

Answer Area

Predicting how many vehicles will travel across a bridge on a given day is an example of

	▼
classification.	
clustering.	
regression.	

Answer:

Answer Area

Predicting how many vehicles will travel across a bridge on a given day is an example of

	▼
classification.	
clustering.	
regression.	

Explanation:

Regression is a machine learning task that is used to predict the value of the label from a set of related features.

Reference:

<https://docs.microsoft.com/en-us/dotnet/machine-learning/resources/tasks>

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

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:

The six guiding principles are:

1. Fairness
2. Inclusiveness
3. Transparency
4. Privacy and Security
5. Reliability and Safety
6. Accountability

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

QUESTION 6

Your company wants to build a recycling machine for bottles. The recycling machine must automatically identify bottles of the correct shape and reject all other items.

Which type of AI workload should the company use?

- A. anomaly detection
- B. conversational AI
- C. computer vision
- D. natural language processing

Answer: C

Explanation:

Azure's Computer Vision service gives you access to advanced algorithms that process images and return information based on the visual features you're interested in. For example, Computer Vision can determine whether an image contains adult content, find specific brands or objects, or find human faces.

Reference:

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

QUESTION 7

Drag and Drop Question

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

Answer:

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

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: Accountability

The people who design and deploy AI systems must be accountable for how their systems operate. Organizations should draw upon industry standards to develop accountability norms. These norms can ensure that AI systems are not the final authority on any decision that impacts people's lives and that humans maintain meaningful control over otherwise highly autonomous AI systems.

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>

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

QUESTION 8

Hotspot Question

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.

		Actual	
		1	0
Predicted	1	11	5
	0	1033	13951

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.

Answer Area

There are [answer choice] correctly predicted positives.

There are [answer choice] false negatives.

Answer:

Answer Area

There are [answer choice] correctly predicted positives.

There are [answer choice] false negatives.

Explanation:

Box 1: 11

	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 8

You have insurance claim reports that are stored as text.
You need to extract key terms from the reports to generate summaries.
Which type of AI workload should you use?

- A. conversational AI
- B. anomaly detection
- C. natural language processing
- D. computer vision

Answer: C

Explanation:

Key phrase extraction is the concept of evaluating the text of a document, or documents, and then identifying the main talking points of the document(s). Key phrase extraction is a part of Text Analytics. The Text Analytics service is a part of the Azure Cognitive Services offerings that can perform advanced natural language processing over raw text.

<https://docs.microsoft.com/en-us/learn/modules/analyze-text-with-text-analytics-service/2-get-started-azure>

QUESTION 9

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

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

Answer: A

Explanation:

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.

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

QUESTION 10

You use drones to identify where weeds grow between rows of crops to send an Instruction for the removal of the weeds.
This is an example of which type of computer vision?

- A. scene segmentation
- B. optical character recognition (OCR)
- C. object detection

Answer: C

Explanation:

Object detection is similar to tagging, but the API returns the bounding box coordinates for each tag applied. 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.

Reference:

<https://docs.microsoft.com/en-us/ai-builder/object-detection-overview>

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/overview-ocr>

<https://docs.microsoft.com/en-us/azure/azure-video-analyzer/video-analyzer-for-media-docs/video-indexer-overview>

QUESTION 11

In which two scenarios can you use a speech synthesis solution? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. an automated voice that reads back a credit card number entered into a telephone by using a numeric keypad
- B. generating live captions for a news broadcast
- C. extracting key phrases from the audio recording of a meeting
- D. an AI character in a computer game that speaks audibly to a player

Answer: AD

Explanation:

Azure Text to Speech is a Speech service feature that converts text to lifelike speech.

Incorrect Answers:

C: Extracting key phrases is not speech synthesis.

Reference:

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

QUESTION 12

You need to provide content for a business chatbot that will help answer simple user queries.

What are three ways to create question and answer text by using QnA Maker? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Generate the questions and answers from an existing webpage.
- B. Use automated machine learning to train a model based on a file that contains the questions.
- C. Manually enter the questions and answers.
- D. Connect the bot to the Cortana channel and ask questions by using Cortana.
- E. Import chat content from a predefined data source.

Answer: ACE

Explanation:

Automatic extraction

Extract question-answer pairs from semi-structured content, including FAQ pages, support websites, excel files, SharePoint documents, product manuals and policies.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/concepts/content-types>

QUESTION 13

You have a frequently asked questions (FAQ) PDF file.

You need to create a conversational support system based on the FAQ.

Which service should you use?

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

Answer: A

Explanation:

QnA Maker is a cloud-based API service that lets you create a conversational question-and-answer layer over your existing data. Use it to build a knowledge base by extracting questions and answers from your semi-structured content, including FAQs, manuals, and documents.

Reference:

<https://azure.microsoft.com/en-us/services/cognitive-services/qna-maker/>

QUESTION 14

You need to develop a web-based AI solution for a customer support system. Users must be able to interact with a web app that will guide them to the best resource or answer.

Which service should you use?

- A. Custom Vision
- B. QnA Maker
- C. Translator Text
- D. Face

Answer: B

Explanation:

QnA Maker is a cloud-based API service that lets you create a conversational question-and-answer layer over your existing data. Use it to build a knowledge base by extracting questions and answers from your semi-structured content, including FAQs, manuals, and documents. Answer users' questions with the best answers from the QnAs in your knowledge base—automatically. Your knowledge base gets smarter, too, as it continually learns from user behavior.

Incorrect Answers:

A: Azure Custom Vision is a cognitive service that lets you build, deploy, and improve your own image classifiers. An image classifier is an AI service that applies labels (which represent classes) to images, according to their visual characteristics. Unlike the Computer Vision service,

Custom Vision allows you to specify the labels to apply.

D: Azure Cognitive Services Face Detection API: At a minimum, each detected face corresponds to a faceRectangle field in the response. This set of pixel coordinates for the left, top, width, and height mark the located face. Using these coordinates, you can get the location of the face and its size. In the API response, faces are listed in size order from largest to smallest.

Reference:

<https://azure.microsoft.com/en-us/services/cognitive-services/qna-maker/>

QUESTION 15

You are authoring a Language Understanding (LUIS) application to support a music festival. You want users to be able to ask questions about scheduled shows, such as: "Which act is playing on the main stage?"

The question "Which act is playing on the main stage?" is an example of which type of element?

- A. an intent
- B. an utterance
- C. a domain
- D. an entity

Answer: B

Explanation:

Utterances are input from the user that your app needs to interpret.

Reference:

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

QUESTION 16

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

Answer: B

Explanation:

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 17

You are building a tool that will process images from retail stores and identify the products of competitors.

The solution will use a custom model.

Which Azure Cognitive Services service should you use?

- A. Custom Vision

- B. Form Recognizer
- C. Face
- D. Computer Vision

Answer: A

Explanation:

Azure Custom Vision is an image recognition service that lets you build, deploy, and improve your own image identifier models. An image identifier applies labels (which represent classifications or objects) to images, according to their detected visual characteristics. Unlike the Computer Vision service, Custom Vision allows you to specify your own labels and train custom models to detect them.

QUESTION 18

Which two of these sources can you translate from one language into another? Each correct answer presents a complete solution. Choose the correct answers

- A. Text
- B. Speech
- C. Image
- D. Video
- E. Handwriting

Answer: AB

Explanation:

Text translation translates the text documents from one language into another language.

Speech translation translates spoken audio from one language into another language.

Converting handwriting to text is an example of Optical character recognition (OCR). OCR extracts handwritten text from an image. You can use text translation with the output from OCR. Extracting the text in an image is also an example of Optical character recognition (OCR). OCR can recognize individual shapes as letters, numerals, punctuation, and other elements of text. OCR extracts the recognizable text from an image. You can use text translation with the output from OCR.

Videos can be analyzed and the spoken audio transcribed into text with media services. You can use text translation with the output from media services.

QUESTION 19

You are processing photos of runners in a race.

You need to read the numbers on the runners' shirts to identity the runners in the photos.

Which type of computer vision should you use?

- A. facial recognition
- B. optical character recognition (OCR)
- C. image classification
- D. object detection

Answer: B

Explanation:

Optical character recognition (OCR) allows you to extract printed or handwritten text from images and documents.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/overview-ocr>

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14