Microsoft

Exam Questions AI-900

Microsoft Azure AI Fundamentals (beta)

NEW QUESTION 1
HOTSPOT - (Topic 5)
You have an app that identifies birds in images. The app performs the following tasks:
* Identifies the location of the birds in the image
* Identifies the species of the birds in the image
Which type of computer vision does each task use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Locate the birds:	Object detection ▼
	Automated captioning
	Image classification
	Object detection
	Optical character recognition (OCR)

Identify the species of the birds:	Image classification ▼
	Automated captioning
	Image classification
	Object detection
	Optical character recognition (OCR)

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Locate the birds:	Object detection ▼
	Automated captioning
	Image classification
	Object detection
	Optical character recognition (OCR)

Identify the species of the birds:	Image classification ▼
	Automated captioning
	Image classification
	Object detection
	Optical character recognition (OCR)

NEW QUESTION 2
- (Topic 5)
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

NEW QUESTION 3
FILL IN THE BLANK - (Topic 5)
To complete the sentence, select the appropriate option in the answer area. Computer vision capabilities can be Deployed to………………..

A. Mastered
B. Not Mastered

Answer: A

Explanation:

| Computer vision capabilities can be deployed to | integrate a facial recognition feature into an app. ▼ |

NEW QUESTION 4
HOTSPOT - (Topic 5)
Select the .

Answer Area

You can use the	Custom Vision ▼	service to train an object detection model by using your own images.
	Computer Vision	
	Custom Vision	
	Form Recognizer	
	Azure Video Analyzer for Media	

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

You can use the [Custom Vision ▼] service to train an object detection model by using your own images.

Computer Vision
Custom Vision
Form Recognizer
Azure Video Analyzer for Media
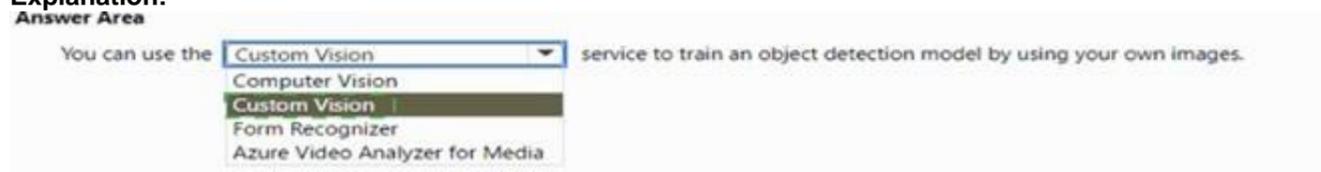
NEW QUESTION 5
- (Topic 5)
You are developing a conversational AI solution that will communicate with users through multiple channels including email, Microsoft Teams, and webchat.
Which service should you use?

A. Text Analytics
B. Azure Bot Service
C. Translator
D. Form Recognizer

Answer: B

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-service-overview-introduction?view=azure-bot-service-4.0

NEW QUESTION 6
- (Topic 5)
You have a website that includes customer reviews.
You need to store the reviews in English and present the reviews to users in their respective language by recognizing each user's geographical location.
Which type of natural language processing workload should you use?

A. translation
B. language modeling
C. key phrase extraction
D. speech recognition

Answer: C

NEW QUESTION 7
HOTSPOT - (Topic 5)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
A bot that responds to queries by internal users is an example of a conversational AI workload.	○	○
An application that displays images relating to an entered search term is an example of a conversational AI workload.	○	○
A web form used to submit a request to reset a password is an example of a conversational AI workload.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
A bot that responds to queries by internal users is an example of a conversational AI workload.	O	O
An application that displays images relating to an entered search term is an example of a conversational AI workload.	O	O
A web form used to submit a request to reset a password is an example of a conversational AI workload.	O	O

NEW QUESTION 8
HOTSPOT - (Topic 5)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
A restaurant can use a chatbot to empower customers to make reservations by using a website or an app.	O	O
A restaurant can use a chatbot to answer inquiries about business hours from a webpage.	O	O
A restaurant can use a chatbot to automate responses to customer reviews on an external website.	O	O

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
A restaurant can use a chatbot to empower customers to make reservations by using a website or an app.	O	O
A restaurant can use a chatbot to answer inquiries about business hours from a webpage.	O	O
A restaurant can use a chatbot to automate responses to customer reviews on an external website.	O	O

NEW QUESTION 9
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence

Answer Area

Ensuring an AI system does not provide a prediction when important fields contain unusual or missing values is [a privacy and security ▼] principle for responsible AI.

an inclusiveness
a privacy and security
a reliability and safety
a transparency

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Ensuring an AI system does not provide a prediction when important fields contain unusual or missing values is [a privacy and security ▼] principle for responsible AI.

- an inclusiveness
- **a privacy and security**
- a reliability and safety
- a transparency

NEW QUESTION 10
HOTSPOT - (Topic 5)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Providing an explanation of the outcome of a credit loan application is an example of the Microsoft transparency principle for responsible AI.	○	○
A triage bot that prioritizes insurance claims based on injuries is an example of the Microsoft reliability and safety principle for responsible AI.	○	○
An AI solution that is offered at different prices for different sales territories is an example of the Microsoft inclusiveness principle for responsible AI.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
Providing an explanation of the outcome of a credit loan application is an example of the Microsoft transparency principle for responsible AI.	○	○
A triage bot that prioritizes insurance claims based on injuries is an example of the Microsoft reliability and safety principle for responsible AI.	○	○
An AI solution that is offered at different prices for different sales territories is an example of the Microsoft inclusiveness principle for responsible AI.	○	○

NEW QUESTION 10
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence.
Answer Area

When evaluating the performance of a model, the [confusion matrix ▼] displays the predicted and actual positives and negatives by using a grid of 0 and 1 values.

- AUC metric
- **confusion matrix**
- ROC curve
- threshold

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

When evaluating the performance of a model, the [confusion matrix ▼] displays the predicted and actual positives and negatives by using a grid of 0 and 1 values.

- AUC metric
- **confusion matrix**
- ROC curve
- threshold

NEW QUESTION 11
DRAG DROP - (Topic 5)

Match the services to the appropriate descriptions.
To answer, drag the appropriate service from the column on the left to its description on the right. Each service may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point

Services	Answer Area	
Azure Storage		Enables the use of natural language to query a knowledge base.
Language Understanding (LUIS)		Enables the real-time transcription of speech-to-text.
QnA Maker		
Speech		

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Services	Answer Area	
Azure Storage	QnA Maker	Enables the use of natural language to query a knowledge base.
Language Understanding (LUIS)	Language Understanding (LUIS)	Enables the real-time transcription of speech-to-text.
QnA Maker		
Speech		

NEW QUESTION 15
HOTSPOT - (Topic 5)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Chatbots can support voice input.	○	○
A separate chatbot is required for each communication channel.	○	○
Chatbots manage conversation flows by using a combination of natural language and constrained option responses.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
Chatbots can support voice input.	○	○
A separate chatbot is required for each communication channel.	○	○
Chatbots manage conversation flows by using a combination of natural language and constrained option responses.	○	○

NEW QUESTION 17
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence.

Answer Area

The [Form Recognizer ▼] service can be used to extract information from a driver's license to populate a database.

Computer Vision
Conversational Language Understanding
Custom Vision
Form Recognizer

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

The [Form Recognizer ▼] service can be used to extract information from a driver's license to populate a database.

Computer Vision
Conversational Language Understanding
Custom Vision
Form Recognizer
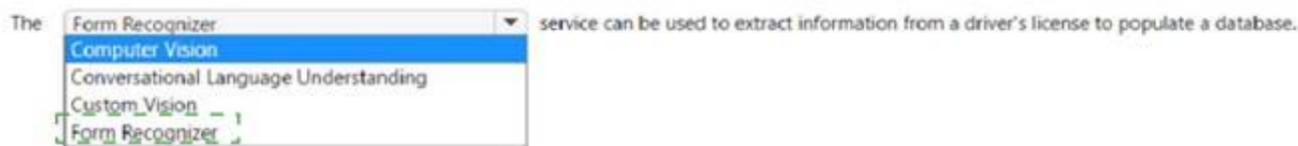
NEW QUESTION 18
- (Topic 5)
Which AI service should you use to create a bot from a frequently asked questions (FAQ) document?

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

Answer: A

NEW QUESTION 21
- (Topic 5)
Your company manufactures widgets.
You have 1.000 digital photos of the widgets.
You need to identify the location of the widgets within the photos. What should you use?

A. Computer Vision Spatial Analysis
B. Custom Vision object detection
C. Custom Vision classification
D. Computer Vision Image Analysis

Answer: B

NEW QUESTION 26
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence

Answer Area

[Object detection ▼] is used to identify multiple types of items in one image.

Image classification
Image description
Object detection
Optical character recognition (OCR)

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

[Object detection ▼] is used to identify multiple types of items in one image.

Image classification
Image description
Object detection
Optical character recognition (OCR)

NEW QUESTION 28
- (Topic 5)
Which two scenarios are examples of a conversational AI workload? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. a smart device in the home that responds to questions such as "What will the weather be like today?"
B. a website that uses a knowledge base to interactively respond to users' questions
C. assembly line machinery that autonomously inserts headlamps into cars
D. monitoring the temperature of machinery to turn on a fan when the temperature reaches a specificThreshold

Answer: AB

NEW QUESTION 31
- (Topic 5)
You have an Azure Machine Learning pipeline that contains a Split Data module. The Split Data module outputs to a Train Model module and a Score Model module. What is the function of the Split Data module?

A. selecting columns that must be included in the model
B. creating training and validation datasets
C. diverting records that have missing data
D. scaling numeric variables so that they are within a consistent numeric range

Answer: A

NEW QUESTION 35
HOTSPOT - (Topic 5)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE; Each correct selection is worth one point.

Answer Area

Statements	Yes	No
A restaurant can use a chatbot to answer queries through Cortana.	○	○
A restaurant can use a chatbot to answer inquiries about business hours from a webpage.	○	○
A restaurant can use a chatbot to automate responses to customer reviews on an external website.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
A restaurant can use a chatbot to answer queries through Cortana.	○	○
A restaurant can use a chatbot to answer inquiries about business hours from a webpage.	○	○
A restaurant can use a chatbot to automate responses to customer reviews on an external website.	○	○

NEW QUESTION 36
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence.

Answer Area

Using Recency, Frequency, and Monetary (RFM) values to identify segments of a customer base is an example of [classification. ▼]
clustering.
regression.
classification.
regularization.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Using Recency, Frequency, and Monetary (RFM) values to identify segments of a customer base is an example of [classification. ▼]
clustering.
regression.
classification.
regularization.

NEW QUESTION 38
- (Topic 5)
You have an AI-based loan approval system.
During testing, you discover that the system has a gender bias. Which responsible AI principle does this violate?

A. accountability
B. transparency

C. fairness
D. reliability and safety

Answer: C

NEW QUESTION 42
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence.

Answer Area

| Regression ▼ | models can be used to predict the sale price of auctioned items. |
| Classification |
| Clustering |
| **Regression** |

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

| Regression ▼ | models can be used to predict the sale price of auctioned items. |
| Classification |
| Clustering |
| **Regression** |

NEW QUESTION 44
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence.

Answer Area

In a machine learning model, the data that is used as inputs are called | labels. ▼ |
| features. |
| functions. |
| **labels.** |
| instances. |

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

In a machine learning model, the data that is used as inputs are called | labels. ▼ |
| features. |
| functions. |
| **labels.** |
| instances. |

NEW QUESTION 48
- (Topic 5)
You have a natural language processing (NIP) model that was created by using data obtained without permission.
Which Microsoft principle for responsible AI does this breach?

A. privacy and security
B. inclusiveness
C. transparency
D. reliability and safety

Answer: C

NEW QUESTION 52
- (Topic 5)
You need to track multiple versions of a model that was trained by using Azure Machine Learning. What should you do?

A. Provision an inference duster.

B. Explain the model.
C. Register the model.
D. Register the training data.

Answer: C

NEW QUESTION 57
HOTSPOT - (Topic 5)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE; Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Chatbots can only be built by using custom code.	○	○
The Azure Bot Service provides services that can be used to host conversational bots.	○	○
Bots built by using the Azure Bot Service can communicate with Microsoft Teams users.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Chatbots can only be built by using custom code.	○	[○]
The Azure Bot Service provides services that can be used to host conversational bots.	[○]	○
Bots built by using the Azure Bot Service can communicate with Microsoft Teams users.	[○]	○

NEW QUESTION 62
HOTSPOT - (Topic 5)
For each of the following statements, select Yes If the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

| Statements	Yes	No
Object detection can identify the location of a damaged product in an image.	○	○
Object detection can identify multiple instances of a damaged product in an image.	○	○
Object detection can identify multiple types of damaged products in an image.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Object detection can identify the location of a damaged product in an image.	[○]	○
Object detection can identify multiple instances of a damaged product in an image.	○	[○]
Object detection can identify multiple types of damaged products in an image.	[○]	○

NEW QUESTION 67
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence.

Answer Area

A historian can use [optical character recognition (OCR) ▼] to digitize newspaper articles.
- facial analysis
- image classification
- object detection
- optical character recognition (OCR)

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

A historian can use [optical character recognition (OCR) ▼] to digitize newspaper articles.
- facial analysis
- image classification
- object detection
- optical character recognition (OCR)
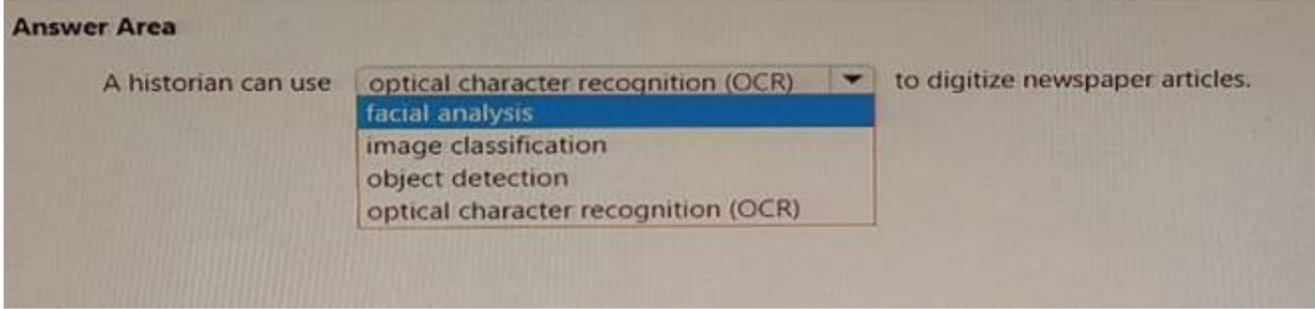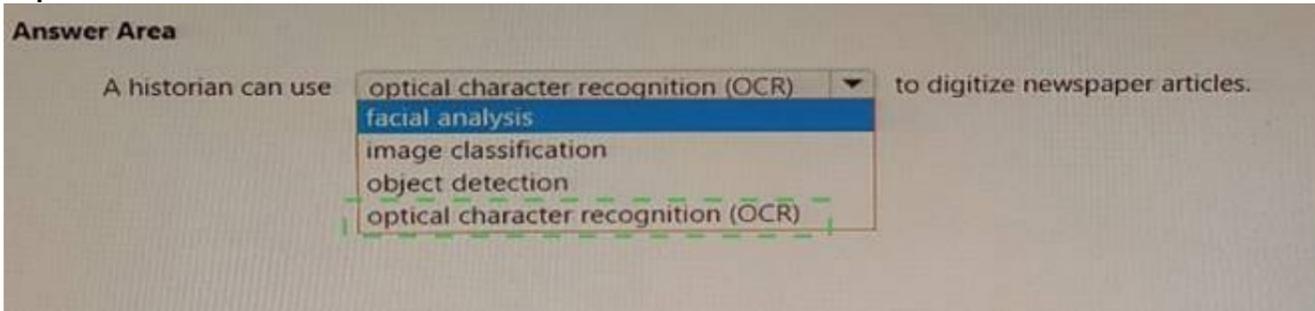
NEW QUESTION 70
HOTSPOT - (Topic 5)
Select the answer that correctly completes the sentence.

Answer Area

[Optical character recognition (OCR) ▼] extracts text from handwritten documents.
- Object detection
- Facial recognition
- Image classification
- Optical character recognition (OCR)

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

[Optical character recognition (OCR) ▼] extracts text from handwritten documents.
- Object detection
- Facial recognition
- Image classification
- Optical character recognition (OCR)

NEW QUESTION 74
HOTSPOT - (Topic 5)
To complete the sentence, select the appropriate option in the answer area.

Answer Area

Returning a bounding box that indicates the location of a vehicle in an image is an example of
- image classification.
- object detection.
- optical character recognition (OCR).
- facial detection.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Returning a bounding box that indicates the location of a vehicle in an image is an example of
- image classification.
- object detection.
- optical character recognition (OCR).
- facial detection.

NEW QUESTION 78
- (Topic 5)
You need to develop a chatbot for a website. The chatbot must answer users questions based on the information m the following documents
• A product troubleshooting guide m a Microsoft Word document
• A frequently asked questions (FAQ) list on a webpage Which service should you use to process the documents?

A. Language Undemanding
B. Text Analytics
C. Azure Bot Service
D. QnA Maker

Answer: D

NEW QUESTION 83
HOTSPOT - (Topic 5)
To complete the sentence, select the appropriate option in the answer area.

Answer Area

An AI solution that helps photographers take better portrait photographs by providing feedback on exposure, noise, and occlusion is an example of facial | recognition. ▼ |
analysis.
detection.
recognition.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

An AI solution that helps photographers take better portrait photographs by providing feedback on exposure, noise, and occlusion is an example of facial | recognition. ▼ |
analysis.
detection.
recognition.

NEW QUESTION 84
HOTSPOT - (Topic 5)
For each of the following statements. select Yes if the statement is true. Otherwise, select No. NOTE; Each correct selection is worth one point

Statements	Yes	No
The Custom Vision service can be used to detect objects in an image.		
The Custom Vision service requires that you provide your own data to train the model.		
The Custom Vision service can be used to analyze video files.		

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
The Custom Vision service can be used to detect objects in an image.	☐	
The Custom Vision service requires that you provide your own data to train the model.	☐	
The Custom Vision service can be used to analyze video files.		☐

NEW QUESTION 85
HOTSPOT - (Topic 5)
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.

Statements	Yes	No
You can communicate with a bot by using email.	○	○
You can communicate with a bot by using Microsoft Teams.	○	○
You can communicate with a bot by using a webchat interface.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
You can communicate with a bot by using email.	[○]	○
You can communicate with a bot by using Microsoft Teams.	[○]	○
You can communicate with a bot by using a webchat interface.	[○]	○

NEW QUESTION 89
HOTSPOT - (Topic 5)
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
A bot that responds to queries by internal users is an example of a natural language processing workload.	○	○
A mobile application that displays images relating to an entered search term is an example of a natural language processing workload.	○	○
A web form used to submit a request to reset a password is an example of a natural language processing workload.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

Statements	Yes	No
A bot that responds to queries by internal users is an example of a natural language processing workload.	[○]	○
A mobile application that displays images relating to an entered search term is an example of a natural language processing workload.	[○]	○
A web form used to submit a request to reset a password is an example of a natural language processing workload.	○	[○]

NEW QUESTION 92
HOTSPOT - (Topic 5)
To complete the sentence, select the appropriate option in the answer area.

Answer Area

Data values that used to make a prediction are called

features ▼
dependant variables.
features.
identifiers.
labels.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Data values that used to make a prediction are called [features ▼]
dependant variables.
features.
identifiers.
labels.
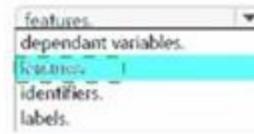
NEW QUESTION 95
- (Topic 4)
In which scenario should you use key phrase extraction?

A. translating a set of documents from English to German
B. generating captions for a video based on the audio track
C. identifying whether reviews of a restaurant are positive or negative
D. identifying which documents provide information about the same topics

Answer: D

NEW QUESTION 98
HOTSPOT - (Topic 4)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
You can use the Speech service to transcribe a call to text.	○	○
You can use the Text Analytics service to extract key entities from a call transcript.	○	○
You can use the Speech service to translate the audio of a call to a different language.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
You can use the Speech service to transcribe a call to text.	[○]	○
You can use the Text Analytics service to extract key entities from a call transcript.	[○]	○
You can use the Speech service to translate the audio of a call to a different language.	[○]	○

NEW QUESTION 103
- (Topic 5)
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 chit-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

NEW QUESTION 108
- (Topic 4)
In which two scenarios can you use speech recognition? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. an in-car system that reads text messages aloud
B. providing closed captions for recorded or live videos
C. creating an automated public address system for a train station
D. creating a transcript of a telephone call or meeting

Answer: BD

Explanation:
Reference:
https://azure.microsoft.com/en-gb/services/cognitive-services/speech-to-text/#features

NEW QUESTION 109
- (Topic 4)
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

Answer: B

Explanation:
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

NEW QUESTION 114
HOTSPOT - (Topic 4)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Monitoring online service reviews for profanities is an example of natural language processing.	○	○
Identifying brand logos in an image is an example of natural languages processing.	○	○
Monitoring public news sites for negative mentions of a product is an example of natural language processing.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: Yes
Content Moderator is part of Microsoft Cognitive Services allowing businesses to use machine assisted moderation of text, images, and videos that augment human review.
The text moderation capability now includes a new machine-learning based text classification feature which uses a trained model to identify possible abusive, derogatory or discriminatory language such as slang, abbreviated words, offensive, and intentionally misspelled words for review.
Box 2: No
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.
Box 3: Yes
Natural language processing (NLP) is used for tasks such as sentiment analysis, topic detection, language detection, key phrase extraction, and document categorization.
Sentiment Analysis is the process of determining whether a piece of writing is positive, negative or neutral.

NEW QUESTION 119
- (Topic 4)
You need to make the press releases of your company available in a range of languages. Which service should you use?

A. Translator Text
B. Text Analytics
C. Speech
D. Language Understanding (LUIS)

Answer: A

Explanation:
Press release is a written communication. Speech wouldn't make sense. Plus, the Speech service doesn't translate languages, it "translates" audio into text, and vice versa.
https://docs.microsoft.com/en-us/learn/modules/translate-text-with-translation-service/2- get-started-azure

NEW QUESTION 120
- (Topic 4)
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.
Reference:
https://azure.microsoft.com/en-in/services/cognitive-services/text-to-speech/

NEW QUESTION 123
- (Topic 4)
You plan to develop a bot that will enable users to query a knowledge base by using natural language processing.
Which two services should you include in the solution? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Language Service
B. Azure Bot Service
C. Form Recognizer
D. Anomaly Detector

Answer: AD

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-service-overview-introduction?view=azure-bot-service-4.0
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/choose-natural-language-processing-service

NEW QUESTION 126
- (Topic 4)
You need to build an app that will read recipe instructions aloud to support users who have reduced vision.
Which version service should you use?

A. Text Analytics
B. Translator Text
C. Speech
D. Language Understanding (LUIS)

Answer: C

Explanation:
Reference:
https://azure.microsoft.com/en-us/services/cognitive-services/text-to-speech/#features

NEW QUESTION 131
HOTSPOT - (Topic 3)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
When creating an object detection model in the Custom Vision service, you must choose a classification type of either **Multilabel** or **Multiclass**.	○	○
You can create an object detection model in the Custom Vision service to find the location of content within an image.	○	○
When creating an object detection model in the Custom Vision service, you can select from a set of predefined domains.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

Statements	Yes	No
When creating an object detection model in the Custom Vision service, you must choose a classification type of either **Multilabel** or **Multiclass**.	○	⊡
You can create an object detection model in the Custom Vision service to find the location of content within an image.	⊡	○
When creating an object detection model in the Custom Vision service, you can select from a set of predefined domains.	⊡	○

NEW QUESTION 133
- (Topic 2)
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:
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/overview

NEW QUESTION 138
- (Topic 2)
You need to create a training dataset and validation dataset from an existing dataset. Which module in the Azure Machine Learning designer should you use?

A. Select Columns in Dataset
B. Add Rows
C. Split Data
D. Join Data

Answer: C

Explanation:
A common way of evaluating a model is to divide the data into a training and test set by

using Split Data, and then validate the model on the training data. Use the Split Data module to divide a dataset into two distinct sets. The studio currently supports training/validation data splits
Reference:
https://docs.microsoft.com/en-us/azure/machine-learning/how-to-configure-cross- validation-data-splits2

NEW QUESTION 139
HOTSPOT - (Topic 2)
To complete the sentence, select the appropriate option in the answer area.

Answer Area

Data values that influence the prediction of a model are called [▼]

dependant variables.
features.
identifiers.
labels.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Data values that influence the prediction of a model are called [▼]

dependant variables.
features.
identifiers.
labels.

NEW QUESTION 144
DRAG DROP - (Topic 3)
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.

Workloads Types	Answer Area	
Facial recognition	Workload Type	Identify celebrities in images.
Image classification	Workload Type	Extract movie title names from movie poster images.
Object detection	Workload Type	Locate vehicles in images.
Optical character recognition (OCR)		

A. Mastered
B. Not Mastered

Answer: A

Explanation:
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: Objection 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.

NEW QUESTION 146
HOTSPOT - (Topic 3)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
The Custom Vision service can be used to detect objects in an image.	○	○
The Custom Vision service requires that you provide your own data to train the model.	○	○
The Custom Vision service can be used to analyze video files.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Custom Vision functionality can be divided into two features. Image classification applies one or more labels to an image. Object detection is similar, but it also returns the coordinates in the image where the applied label(s) can be found.
Box 2: Yes
The Custom Vision service uses a machine learning algorithm to analyze images. You, the developer, submit groups of images that feature and lack the characteristics in question. You label the images yourself at the time of submission. Then, the algorithm trains to this data and calculates its own accuracy by testing itself on those same images.
Box 3: No
Custom Vision service can be used only on graphic files.

NEW QUESTION 150
DRAG DROP - (Topic 2)
Match the machine learning tasks to the appropriate scenarios.
To answer, drag the appropriate task from the column on the left to its scenario on the right. Each task may be used once, more than once, or not at all.
NOTE: Each correct selection is worth one point.

Learning Types	Answer Area	
Feature engineering	Task	Examining the values of a confusion matrix
Feature selection	Task	Splitting a date into month, day, and year fields
Model deployment	Task	Picking temperature and pressure to train a weather model
Model evaluation		
Model training		

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Model evaluation
The Model evaluation module outputs a confusion matrix showing the number of true positives, false negatives, false positives, and true negatives, as well as ROC, Precision/Recall, and Lift curves.
Box 2: Feature engineering
Feature engineering is the process of using domain knowledge of the data to create features that help ML algorithms learn better. In Azure Machine Learning, scaling and normalization techniques are applied to facilitate feature engineering. Collectively, these techniques and feature engineering are referred to as featurization.
Note: Often, features are created from raw data through a process of feature engineering. For example, a time stamp in itself might not be useful for modeling until the information is transformed into units of days, months, or categories that are relevant to the problem, such as holiday versus working day.
Box 3: Feature selection
In machine learning and statistics, feature selection is the process of selecting a subset of relevant, useful features to use in building an analytical model. Feature selection helps narrow the field of data to the most valuable inputs. Narrowing the field of data helps reduce noise and improve training performance.

NEW QUESTION 155
- (Topic 2)

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

NEW QUESTION 159
HOTSPOT - (Topic 2)
To complete the sentence, select the appropriate option in the answer area.

Predicting how many vehicles will travel across a bridge on a given day is an example of [▼]
classification.
clustering.
regression.
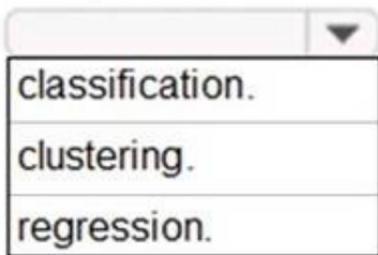
A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 162
- (Topic 2)
A medical research project uses a large anonymized dataset of brain scan images that are categorized into predefined brain haemorrhage types.
You need to use machine learning to support early detection of the different brain haemorrhage types in the images before the images are reviewed by a person.
This is an example of which type of machine learning?

A. clustering
B. regression
C. classification

Answer: C

Explanation:
Reference:
https://docs.microsoft.com/en-us/learn/modules/create-classification-model-azure-machine-learning-designer/introduction

NEW QUESTION 166
HOTSPOT - (Topic 2)
To complete the sentence, select the appropriate option in the answer area.

[▼] models can be used to predict the sale price of auctioned items.
Classification
Clustering
Regression

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 169
- (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: CD

Explanation:
https://docs.microsoft.com/en-in/learn/modules/create-regression-model-azure-machine-learning-designer/deploy-service

NEW QUESTION 172
- (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 174
- (Topic 2)
Which service should you use to extract text, key/value pairs, and table data automatically from scanned documents?

A. Form Recognizer
B. Text Analytics
C. Ink Recognizer
D. Custom Vision

Answer: A

Explanation:
Accelerate your business processes by automating information extraction. Form Recognizer applies advanced machine learning to accurately extract text, key/value pairs, and tables from documents. With just a few samples, Form Recognizer tailors its understanding to your documents, both on-premises and in the cloud. Turn forms into usable data at a fraction of the time and cost, so you can focus more time acting on the information rather than compiling it.
Reference:
https://azure.microsoft.com/en-us/services/cognitive-services/form-recognizer/

NEW QUESTION 177
HOTSPOT - (Topic 2)
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.

Statements	Yes	No
Organizing documents into groups based on similarities of the text contained in the documents is an example of clustering.	○	○
Grouping similar patients based on symptoms and diagnostic test results is an example of clustering.	○	○
Predicting whether a person will develop mild, moderate, or severe allergy symptoms based on pollen count is an example of clustering.	○	○

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Clustering is a machine learning task that is used to group instances of data into clusters that contain similar characteristics. Clustering can also be used to identify relationships in a dataset
Regression is a machine learning task that is used to predict the value of the label from a set of related features.

NEW QUESTION 182
- (Topic 2)
When training a model, why should you randomly split the rows into separate subsets?

A. to train the model twice to attain better accuracy
B. to train multiple models simultaneously to attain better performance
C. to test the model by using data that was not used to train the model

Answer: C

Explanation:
 The goal is to produce a trained (fitted) model that generalizes well to new, unknown data. The fitted model is evaluated using "new" examples from the held-out datasets (validation and test datasets) to estimate the model's accuracy in classifying new data.
https://en.wikipedia.org/wiki/Training,_validation,_and_test_sets#:~:text=Training%20datas et,-A%20training%20dataset&text=The%20goal%20is%20to%20produce,accuracy%20in%20classifying%20new%20data.

NEW QUESTION 185
DRAG DROP - (Topic 1)
You plan to deploy an Azure Machine Learning model as a service that will be used by client applications.
Which three processes should you perform in sequence before you deploy the model? To answer, move the appropriate processes from the list of processes to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 186
- (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.

Answer: B

Explanation:
https://docs.microsoft.com/en-us/azure/machine-learning/algorithm-module-reference/split-data

NEW QUESTION 189
HOTSPOT - (Topic 1)
To complete the sentence, select the appropriate option in the answer area.

Answer Area

Returning a bounding box that indicates the location of a vehicle in an image is an example of [dropdown]

- image classification.
- object detection.
- optical character recognizer (OCR).
- semantic segmentation.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Returning a bounding box that indicates the location of a vehicle in an image is an example of [dropdown]

- image classification.
- object detection.
- optical character recognizer (OCR).
- semantic segmentation.

NEW QUESTION 191
HOTSPOT - (Topic 1)
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

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.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
Box 2: Yes
Box 3: Yes
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.

NEW QUESTION 194
- (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
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

NEW QUESTION 195
DRAG DROP - (Topic 1)
Match the principles of responsible AI to appropriate requirements.
To answer, drag the appropriate principles from the column on the left to its requirement on the right. Each principle 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.

Principles	Answer Area	
Fairness		The system must not discriminate based on gender, race
Privacy and security		Personal data must be visible only to approve
Reliability and safety		Automated decision-making processes must be recorded so that approved users can identify why a decision was made
Transparency		

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Principles	Answer Area	
Fairness	Fairness	The system must not discriminate based on gender, race
Privacy and security	Privacy and security	Personal data must be visible only to approve
Reliability and safety	Transparency	Automated decision-making processes must be recorded so that approved users can identify why a decision was made
Transparency		

NEW QUESTION 198
- (Topic 1)
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:
Reference:
https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding- principles

NEW QUESTION 201
HOTSPOT - (Topic 1)
To complete the sentence, select the appropriate option in the answer 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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Reliability & Safety https://en.wikipedia.org/wiki/Tay_(bot)
"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. It's also important to be able to verify that these systems are behaving as intended under actual operating conditions. How they behave and the variety of conditions they can handle reliably and safely largely reflects the range of situations and circumstances that developers anticipate during design and testing. We believe that rigorous testing is essential during system development and deployment to ensure AI systems can respond safely in unanticipated situations and edge cases, don't have unexpected performance failures, and don't evolve in ways that are inconsistent with original expectations"

NEW QUESTION 202
HOTSPOT - (Topic 1)
To complete the sentence, select the appropriate option in the answer area.

Answer Area

| Feature engineering |
| Feature selection |
| Model evaluation |
| Model training |

is used to generate additional features.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

| Feature engineering |
| Feature selection |
| Model evaluation |
| Model training |

is used to generate additional features.

NEW QUESTION 204
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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 FirstTry

* AI-900 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The AI-900 Practice Test Here