

Microsoft.AI-900.vSep-2024.by.Ando.98q

Number: AI-900
Passing Score: 800
Time Limit: 120
File Version: 17.0

Exam Code: AI-900
Exam Name: Microsoft Azure AI Fundamentals



Exam A

QUESTION 1

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)

Correct Answer: A

Section:

Explanation:

Translator is a cloud-based machine translation service you can use to translate text in near real-time through a simple REST API call. The service uses modern neural machine translation technology and offers statistical machine translation technology. Custom Translator is an extension of Translator, which allows you to build neural translation systems.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/translator/>

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

Section:

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>

QUESTION 3

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Counting the number of animals in an area based on a video feed is an example of

	▼
forecasting.	
computer vision.	
conversational AI.	
anomaly detection.	

Answer Area:

Answer Area

Counting the number of animals in an area based on a video feed is an example of

	▼
forecasting.	
computer vision.	
conversational AI.	
anomaly detection.	

Section:

Explanation:

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

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/intro-to-spatial-analysis-public-preview>

QUESTION 4

You have a knowledge base of frequently asked questions (FAQ).

You create a bot that uses the knowledge base to respond to customer requests.

You need to identify what the bot can perform without adding additional skills.
What should you identify?

- A. Register customer purchases.
- B. Register customer complaints.
- C. Answer questions from multiple users simultaneously.
- D. Provide customers with return materials authorization (RMA) numbers.

Correct Answer: C

Section:

Explanation:

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

QUESTION 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.

Hot Area:

Answer Area

Statements

The following service call will accept English text as an input and output Italian and French text.
/translate?from=it&to=fr&ss=1

The following service call will accept English text as an input and output Italian and French text.
/translate?from=en&to=fr&ss=15

The Translator service can be used to translate documents from English to French.

Yes No

Vdumps

Answer Area:

Answer Area

Statements

The following service call will accept English text as an input and output Italian and French text.
/translate?from=it&to=fr&ss=1

The following service call will accept English text as an input and output Italian and French text.
/translate?from=en&to=fr&ss=15

The Translator service can be used to translate documents from English to French.

Yes No

Vdumps

Section:

Explanation:

QUESTION 6

HOTSPOT

You have an Azure Machine Learning model that predicts product quality. The model has a training dataset that contains 50,000 records. A sample of the data is shown in the following table.

Date	Time	Mass (kg)	Temperature (C)	Quality Test
26/02/2021	15:31:07	2.108	62.5	Pass
26/02/2021	15:31:39	2.099	62.4	Pass
26/02/2021	02:32:21	2.098	66.4	Fail

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
Mass (kg) is a feature.	<input type="radio"/>	<input type="radio"/>
Quality Test is a label.	<input type="radio"/>	<input type="radio"/>
Temperature (C) is a label.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Mass (kg) is a feature.	<input checked="" type="radio"/>	<input type="radio"/>
Quality Test is a label.	<input checked="" type="radio"/>	<input type="radio"/>
Temperature (C) is a label.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:



QUESTION 7

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

Section:

Explanation:

Named Entity Recognition (NER) is the ability to identify different entities in text and categorize them into pre-defined classes or types such as: person, location, event, product, and organization.

In this question, the square brackets indicate the entities such as DateTime, PersonType, Skill.

Reference:

<https://docs.microsoft.com/en-in/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-entity-linking?tabs=version-3-preview>

QUESTION 8

You are developing a solution that uses the Text Analytics service.

You need to identify the main talking points in a collection of documents.

Which type of natural language processing should you use?

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

Correct Answer: B

Section:

Explanation:

Broad entity extraction: Identify important concepts in text, including key

Key phrase extraction/ Broad entity extraction: Identify important concepts in text, including key phrases and named entities such as people, places, and organizations.

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

QUESTION 9

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

Correct Answer: B, D

Section:

Explanation:

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

QUESTION 10

Your company is exploring the use of voice recognition technologies in its smart home devices. The company wants to identify any barriers that might unintentionally leave out specific user groups. This an example of which Microsoft guiding principle for responsible AI?

- A. accountability
- B. fairness
- C. inclusiveness
- D. privacy and security

Correct Answer: C

Section:

Explanation:

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

QUESTION 11

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

Correct Answer: C, D, F

Section:

Explanation:

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

QUESTION 12

You run a charity event that involves posting photos of people wearing sunglasses on Twitter.

You need to ensure that you only retweet photos that meet the following requirements:

Include one or more faces.

Contain at least one person wearing sunglasses.

What should you use to analyze the images?

- A. the Verify operation in the Face service
- B. the Detect operation in the Face service
- C. the Describe Image operation in the Computer Vision service
- D. the Analyze Image operation in the Computer Vision service

Correct Answer: B

Section:

Explanation:

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

QUESTION 13

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)

Correct Answer: A

Section:

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>

QUESTION 14

Which two components can you drag onto a canvas in Azure Machine Learning designer? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. dataset
- B. compute
- C. pipeline
- D. module

Correct Answer: A, D

Section:

Explanation:

You can drag-and-drop datasets and modules onto the canvas.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/concept-designer>

QUESTION 15

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

Correct Answer: C

Section:

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>

QUESTION 16

You have the Predicted vs. True chart shown in the following exhibit.





Which type of model is the chart used to evaluate?

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

Correct Answer: B

Section:

Explanation:

What is a Predicted vs. True chart?

Predicted vs. True shows the relationship between a predicted value and its correlating true value for a regression problem. This graph can be used to measure performance of a model as the closer to the $y=x$ line the predicted values are, the better the accuracy of a predictive model.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-understand-automated-m>

QUESTION 17

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

Section:

Explanation:

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 18

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

Correct Answer: B

Section:

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 19

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: C

Section:

QUESTION 20

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. Language Understanding
- D. Custom Vision

Correct Answer: A

Section:

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/>

QUESTION 21

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

Which two parameters should you use to access the web service? 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

Correct Answer: C, D

Section:

Explanation:

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

To consume the pipeline, you need:

The REST endpoint for your service

The Primary Key for your service

Reference:

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

QUESTION 22

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

Correct Answer: C

Section:

Explanation:

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

QUESTION 23

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

Correct Answer: C

Section:

QUESTION 24

You are evaluating whether to use a basic workspace or an enterprise workspace in Azure Machine Learning. What are two tasks that require an enterprise workspace? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Use a graphical user interface (GUI) to run automated machine learning experiments.
- B. Create a compute instance to use as a workstation.
- C. Use a graphical user interface (GUI) to define and run machine learning experiments from Azure Machine Learning designer.



D. Create a dataset from a comma-separated value (CSV) file.

Correct Answer: A, C

Section:

Explanation:

Note: Enterprise workspaces are no longer available as of September 2020. The basic workspace now has all the functionality of the enterprise workspace.

Reference: <https://www.azure.cn/en-us/pricing/details/machine-learning/> <https://docs.microsoft.com/en-us/azure/machine-learning/concept-workspace>

QUESTION 25

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: A, C

Section:

Explanation:

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.

QUESTION 26

You need to develop a mobile app for employees to scan and store their expenses while travelling.

Which type of computer vision should you use?

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

Correct Answer: D

Section:

Explanation:

Azure's Computer Vision API includes Optical Character Recognition (OCR) capabilities that extract printed or handwritten text from images. You can extract text from images, such as photos of license plates or containers with serial numbers, as well as from documents - invoices, bills, financial reports, articles, and more.

Reference:

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

QUESTION 27



You need to determine the location of cars in an image so that you can estimate the distance between the cars. Which type of computer vision should you use?

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

Correct Answer: B

Section:

Explanation:


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

QUESTION 28

You send an image to a Computer Vision API and receive back the annotated image shown in the exhibit.



Which type of computer vision was used?

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

Correct Answer: A

Section:**Explanation:**

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

QUESTION 29

What are two tasks that can be performed by using the Computer Vision service? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Train a custom image classification model.
- B. Detect faces in an image.
- C. Recognize handwritten text.
- D. Translate the text in an image between languages.

Correct Answer: B, C

Section:**Explanation:**

B: Azure's Computer Vision service provides developers with 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.

C: Computer Vision includes Optical Character Recognition (OCR) capabilities. You can use the new Read API to extract printed and handwritten text from images and documents.

Reference:

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

**QUESTION 30**

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: D

Section:**QUESTION 31**

What are two tasks that can be performed by using computer vision? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Predict stock prices.
- B. Detect brands in an image.
- C. Detect the color scheme in an image
- D. Translate text between languages.
- E. Extract key phrases.

Correct Answer: B, C

Section:

Explanation:

B: Identify commercial brands in images or videos from a database of thousands of global logos. You can use this feature, for example, to discover which brands are most popular on social media or most prevalent in media product placement.

C: Analyze color usage within an image. Computer Vision can determine whether an image is black & white or color and, for color images, identify the dominant and accent colors.

Reference:

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

QUESTION 32

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

Correct Answer: C

Section:

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 33

In which two scenarios can you use the Form Recognizer service? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Extract the invoice number from an invoice.
- B. Translate a form from French to English.
- C. Find image of product in a catalog.
- D. Identify the retailer from a receipt.

Correct Answer: A, D

Section:

Explanation:

Reference:

<https://azure.microsoft.com/en-gb/services/cognitive-services/form-recognizer/#features>

QUESTION 34

Your website has a chatbot to assist customers.

You need to detect when a customer is upset based on what the customer types in the chatbot.

Which type of AI workload should you use?

- A. anomaly detection
- B. semantic segmentation
- C. regression
- D. natural language processing

Correct Answer: D

Section:

Explanation:

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.

Reference:

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

QUESTION 35

Which AI service can you use to interpret the meaning of a user input such as "Call me back later?"

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

Correct Answer: D

Section:

Explanation:

Language Understanding (LUIS) is a cloud-based AI service, that applies custom machine-learning intelligence to a user's conversational, natural language text to predict overall meaning, and pull out relevant, detailed information.

Reference:

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

QUESTION 36

You are developing a Chatbot solution in Azure.

Which service should you use to determine a user's intent?

- A. Translator Text
- B. Azure Cognitive Search
- C. Speech
- D. Language Understanding (LUIS)

Correct Answer: B

Section:

Explanation:

QUESTION 37

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)

Correct Answer: C

Section:

Explanation:



Reference:

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

QUESTION 38

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 telephone answering service that has a pre-recorder message
- B. a chatbot that provides users with the ability to find answers on a website by themselves
- C. telephone voice menus to reduce the load on human resources
- D. a service that creates frequently asked questions (FAQ) documents by crawling public websites

Correct Answer: B, C

Section:

Explanation:

B: A bot is an automated software program designed to perform a particular task. Think of it as a robot without a body. C: Automated customer interaction is essential to a business of any size. In fact, 61% of consumers prefer to communicate via speech, and most of them prefer self-service. Because customer satisfaction is a priority for all businesses, selfservice is a critical facet of any customer-facing communications strategy.

Incorrect Answers:

D: Early bots were comparatively simple, handling repetitive and voluminous tasks with relatively straightforward algorithmic logic. An example would be web crawlers used by search engines to automatically explore and catalog web content.

Reference: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/big-data/ai-overview> <https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/interactive-voice-response-bot>

QUESTION 39

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-chat content from a predefined data source.

Correct Answer: A, C, E

Section:

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 40

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)

Correct Answer: A

Section:

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 41

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

- A. Text Analytics
- B. QnA Maker
- C. Azure Bot Service
- D. Translator Text

Correct Answer: B, C

Section:

Explanation:

Bots are a popular way to provide support through multiple communication channels. You can use the QnA Maker service and Azure Bot Service to create a bot that answers user questions.

Reference:

<https://docs.microsoft.com/en-us/learn/modules/build-faq-chatbot-qna-maker-azure-bot-service/>

QUESTION 42

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.

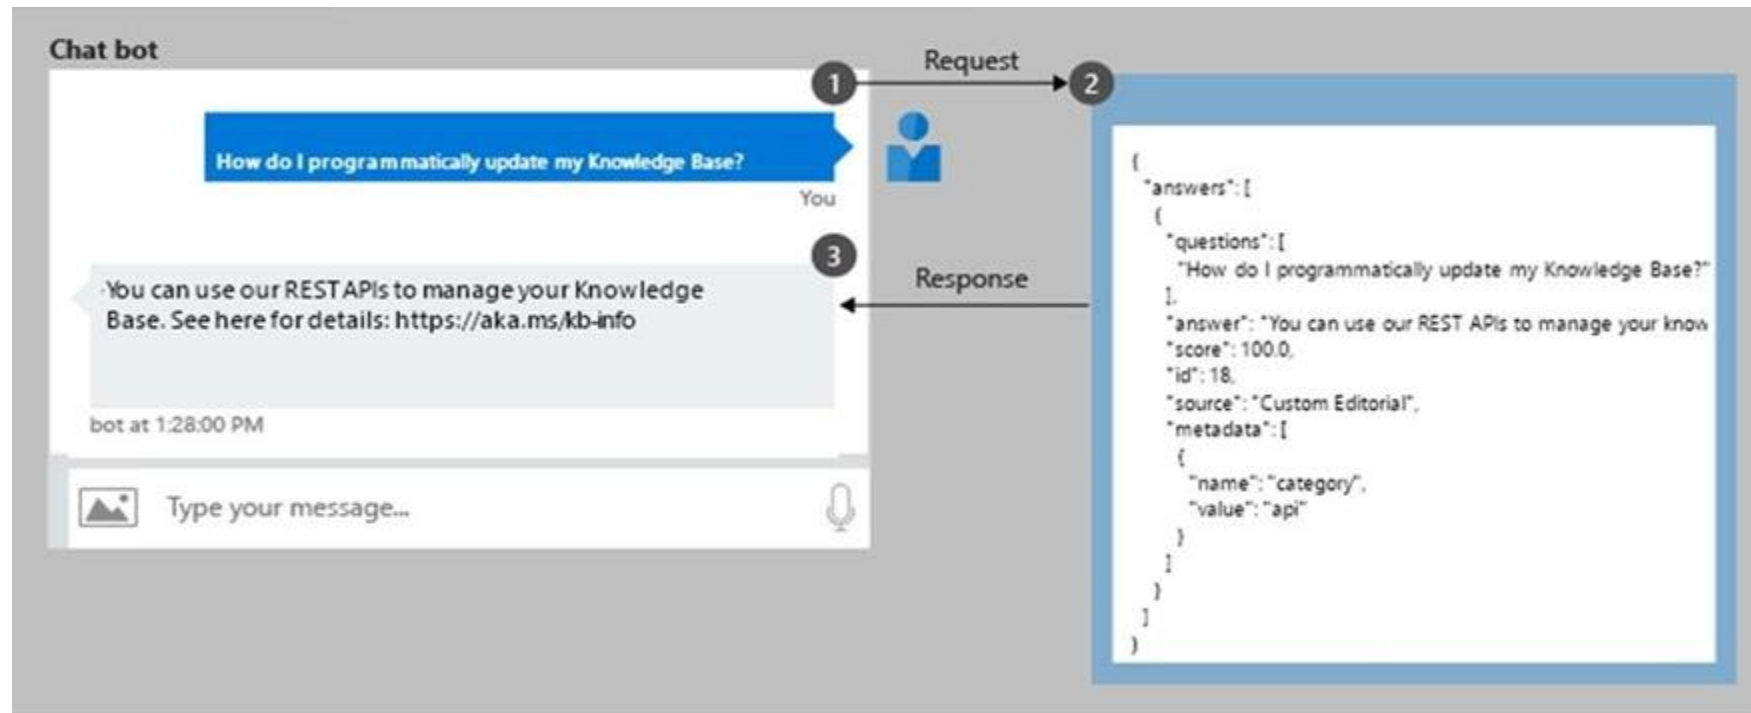
- A. a smart device in the home that responds to questions such as "What will the weather be like today?"
- B. a website that uses a knowledge base to interactively respond to users' questions
- C. assembly line machinery that autonomously inserts headlamps into cars
- D. monitoring the temperature of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

Correct Answer: A, B

Section:

QUESTION 43

You have the process shown in the following exhibit.



Which type AI solution is shown in the diagram?

- A. a sentiment analysis solution
- B. a chatbot
- C. a machine learning model
- D. a computer vision application

Correct Answer: B

Section:

QUESTION 44

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

Correct Answer: B

Section:

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 45

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

Section:

QUESTION 46

Which scenario is an example of a webchat bot?

- A. Determine whether reviews entered on a website for a concert are positive or negative, and then add a thumbs up or thumbs down emoji to the reviews.
- B. Translate into English questions entered by customers at a kiosk so that the appropriate person can call the customers back.
- C. Accept questions through email, and then route the email messages to the correct person based on the content of the message.
- D. From a website interface, answer common questions about scheduled events and ticket purchases for a music festival.

Correct Answer: D


Section:

QUESTION 47

HOTSPOT

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.



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.

Hot Area:

Answer Area

There are [answer choice] correctly predicted positives.

- 5
- 11
- 1,033
- 13,951

CEplus.com

There are [answer choice] false negatives.

- 5
- 11
- 1,033
- 13,951

CEplus.com

Answer Area:

Answer Area

There are [answer choice] correctly predicted positives.

- 5
- 11
- 1,033
- 13,951

CEplus.com

There are [answer choice] false negatives.

- 5
- 11
- 1,033
- 13,951

CEplus.com



Section:

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 48

HOTSPOT

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

Hot Area:

Answer Area

A banking system that predicts whether a loan will be repaid is an example of the

- classification
- regression
- clustering

Answer Area:

Answer Area

A banking system that predicts whether a loan will be repaid is an example of the

- classification
- regression
- clustering

Section:

Explanation:

Two-class classification provides the answer to simple two-choice questions such as Yes/No or True/False.

Vdumps

QUESTION 49

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
Labelling is the process of tagging training data with known values.	<input type="radio"/>	<input type="radio"/>
You should evaluate a model by using the same data used to train the model.	<input type="radio"/>	<input type="radio"/>
Accuracy is always the primary metric used to measure a model's performance.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements

Yes

No

Labelling is the process of tagging training data with known values.

You should evaluate a model by using the same data used to train the model.

Accuracy is always the primary metric used to measure a model's performance.

Section:

Explanation:

Box 1: Yes

In machine learning, if you have labeled data, that means your data is marked up, or annotated, to show the target, which is the answer you want your machine learning model to predict. In general, data labeling can refer to tasks that include data tagging, annotation, classification, moderation, transcription, or processing.

Box 2: No

Box 3: No

Accuracy is simply the proportion of correctly classified instances. It is usually the first metric you look at when evaluating a classifier. However, when the test data is unbalanced (where most of the instances belong to one of the classes), or you are more interested in the performance on either one of the classes, accuracy doesn't really capture the effectiveness of a classifier.

Reference:

<https://www.cloudfactory.com/data-labeling-guide>

<https://docs.microsoft.com/en-us/azure/machine-learning/studio/evaluate-model-performance>



QUESTION 50

HOTSPOT

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

Hot Area:

Answer Area

The ability to extract subtotals and totals from a receipt is a capability of the _____ service.

- Custom Vision
- Form Recognizer
- Ink Recognizer
- Text Analytics

Answer Area:

Answer Area

The ability to extract subtotals and totals from a receipt is a capability of the _____ service.

- Custom Vision
- Form Recognizer
- Ink Recognizer
- Text Analytics

Section:

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/>

QUESTION 51

HOTSPOT

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

Hot Area:

Answer Area

From Azure Machine Learning designer, to deploy a real-time inference pipeline as a service for others to consume, you must deploy the model to _____

- a local web service.
- Azure Container Instances.
- Azure Kubernetes Service (AKS).
- Azure Machine Learning compute.

Answer Area:

Answer Area

From Azure Machine Learning designer, to deploy a real-time inference pipeline as a service for others to consume, you must deploy the model to _____

- a local web service.
- Azure Container Instances.
- Azure Kubernetes Service (AKS).
- Azure Machine Learning compute.

Section:

Explanation:

To perform real-time inferencing, you must deploy a pipeline as a real-time endpoint. Real-time endpoints must be deployed to an Azure Kubernetes Service cluster.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/concept-designer#deploy>

QUESTION 52

HOTSPOT

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

Hot Area:

Answer Area

Predicting how many hours of overtime a delivery person will work based on the number of order received is an example of

classification.
clustering.
regression.

Answer Area:

Answer Area

Predicting how many hours of overtime a delivery person will work based on the number of order received is an example of

classification.
clustering.
regression.



Section:

Explanation:

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.

Incorrect Answers:

Classification is a machine learning method that uses data to determine the category, type, or class of an item or row of data. 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/algorithm-module-reference/linear-regression>

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

QUESTION 53

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

Azure Machine Learning designer provides a drag-and-drop visual canvas to build, test, and deploy machine learning models.

Azure Machine Learning designer enables you to save your progress as a pipeline draft.

Azure Machine Learning designer enables you to include custom JavaScript functions.



Answer Area:

Answer Area

Statements

Yes

No

Azure Machine Learning designer provides a drag-and-drop visual canvas to build, test, and deploy machine learning models.

Azure Machine Learning designer enables you to save your progress as a pipeline draft.

Azure Machine Learning designer enables you to include custom JavaScript functions.



Section:

Explanation:

Box 1: Yes

Azure Machine Learning designer lets you visually connect datasets and modules on an interactive canvas to create machine learning models. Box 2: Yes
With the designer you can connect the modules to create a pipeline draft.

As you edit a pipeline in the designer, your progress is saved as a pipeline draft.

Box 3: No

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/concept-designer>

QUESTION 54

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

You can fine-tune some Azure OpenAI models by using your own data.

Yes

No

Pretrained generative AI models are a component of Azure OpenAI.

To build a solution that complies with Microsoft responsible AI principles, you must build and train your own model.

Answer Area:

Answer Area

Statements

You can fine-tune some Azure OpenAI models by using your own data.

Yes

No

Pretrained generative AI models are a component of Azure OpenAI.

To build a solution that complies with Microsoft responsible AI principles, you must build and train your own model.

Section:

Explanation:

QUESTION 55

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



Answer Area

You can modify the

Temperature
Frequency penalty
Max response
Stop sequence
Temperature

parameter to produce more deterministic responses from a chat solution that uses the Azure OpenAI GPT-3.5 model.

Answer Area:

Answer Area

You can modify the

Temperature
Frequency penalty
Max response
Stop sequence
Temperature

parameter to produce more deterministic responses from a chat solution that uses the Azure OpenAI GPT-3.5 model.

Section:

Explanation:

QUESTION 56

DRAG DROP

Match the Azure AI 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:

Azure Cognitive Services

- ☰ Azure AI Language
- ☰ Azure AI Speech
- ☰ Azure AI Translator

Answer Area

- Convert spoken requests into text.
- Identify the intent of a user's requests.
- Apply intent to entities and utterances.

Correct Answer:

Azure Cognitive Services

- ☰ Azure AI Language
- ☰ Azure AI Speech
- ☰ Azure AI Translator

Answer Area

- ☰ Azure AI Speech Convert spoken requests into text.
- ☰ Azure AI Language Identify the intent of a user's requests.
- ☰ Azure AI Language Apply intent to entities and utterances.

Section:

Explanation:

QUESTION 57

You need to identify groups of rows with similar numeric values in a dataset. Which type of machine learning should you use?

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

Correct Answer: A

Section:

QUESTION 58

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
The Azure AI Custom Vision service can be used to detect objects in an image.	<input type="radio"/>	<input type="radio"/>
The Azure AI Custom Vision service requires that you provide your own data to train the model.	<input type="radio"/>	<input type="radio"/>
The Azure AI Custom Vision service can be used to analyze video files.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
The Azure AI Custom Vision service can be used to detect objects in an image.	<input checked="" type="radio"/>	<input type="radio"/>
The Azure AI Custom Vision service requires that you provide your own data to train the model.	<input checked="" type="radio"/>	<input type="radio"/>
The Azure AI Custom Vision service can be used to analyze video files.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

QUESTION 59

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:

House Price Category:

Answer Area:



Answer Area

Household Income:

House Price Category:

Section:

Explanation:

Reference:

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

QUESTION 60

HOTSPOT

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

Hot Area:

Answer Area

Azure Machine Learning designer lets you create machine learning models by

- adding and connecting modules on a visual canvas.
- automatically performing common data preparation tasks.
- automatically selecting an algorithm to build the most accurate model.
- using a code-first notebook experience.



Answer Area:

Answer Area

Azure Machine Learning designer lets you create machine learning models by

- adding and connecting modules on a visual canvas.
- automatically performing common data preparation tasks.
- automatically selecting an algorithm to build the most accurate model.
- using a code-first notebook experience.

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/concept-designer>

QUESTION 61

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

Automated machine learning provides you with the ability to include custom Python scripts in a training pipeline.

Automated machine learning implements machine learning solutions without the need for programming experience.

Automated machine learning provides you with the ability to visually connect datasets and modules on an interactive canvas.



Answer Area:

Answer Area



Statements

Yes

No

Automated machine learning provides you with the ability to include custom Python scripts in a training pipeline.

Automated machine learning implements machine learning solutions without the need for programming experience.

Automated machine learning provides you with the ability to visually connect datasets and modules on an interactive canvas.



dumps

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-designer-python>

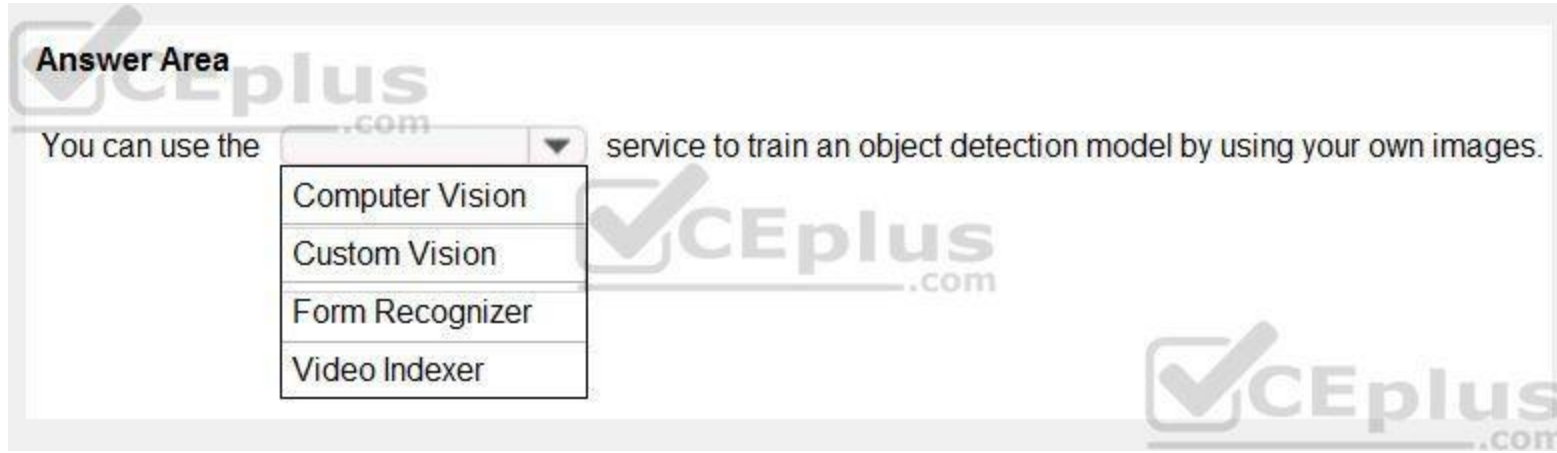
<https://docs.microsoft.com/en-us/azure/machine-learning/concept-automated-ml>

QUESTION 62

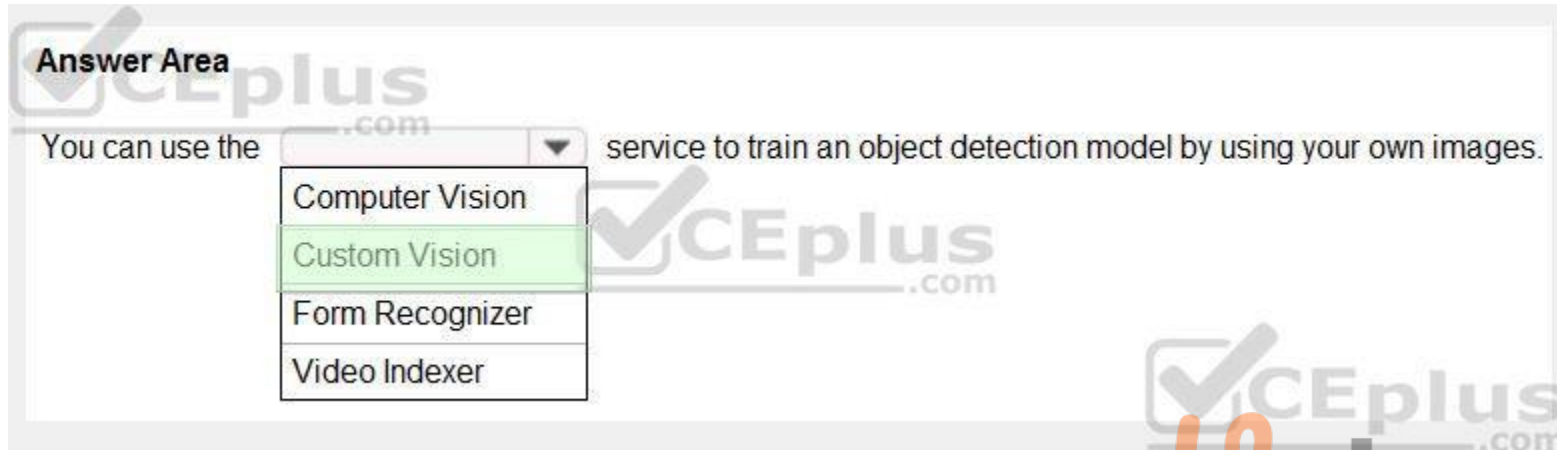
HOTSPOT

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

Hot Area:



Answer Area:



Section:

Explanation:

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.

Note: The Custom Vision service uses a machine learning algorithm to apply labels to images. You, the developer, must 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. Once the algorithm is trained, you can test, retrain, and eventually use it to classify new images according to the needs of your app. You can also export the model itself for offline use.

Incorrect Answers:

Computer Vision:

Azure's Computer Vision service provides developers with 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/custom-vision-service/home>

QUESTION 63

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

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.



Answer Area:

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.



Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/get-started-build-detector>

QUESTION 64

HOTSPOT

You have a database that contains a list of employees and their photos.

You are tagging new photos of the employees.

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:

Statements

Yes No

The Face service can be used to group all the employees who have similar facial characteristics.

The Face service will be more accurate if you provide more sample photos of each employee from different angles.

If an employee is wearing sunglasses, the Face service will always fail to recognize the employee.

Answer Area:

Statements

Yes No

The Face service can be used to group all the employees who have similar facial characteristics.

The Face service will be more accurate if you provide more sample photos of each employee from different angles.

If an employee is wearing sunglasses, the Face service will always fail to recognize the employee.

Section:

Explanation:

Reference:

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

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

QUESTION 65

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 train a regression model by using unlabeled data.	<input type="radio"/>	<input type="radio"/>
The classification technique is used to predict sequential numerical data over time.	<input type="radio"/>	<input type="radio"/>
Grouping items by their common characteristics is an example of clustering.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
You train a regression model by using unlabeled data.	<input type="radio"/>	<input checked="" type="radio"/>
The classification technique is used to predict sequential numerical data over time.	<input type="radio"/>	<input checked="" type="radio"/>
Grouping items by their common characteristics is an example of clustering.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/learn/modules/create-regression-model-azure-machine-learning-designer/5-create-training-pipeline> <https://docs.microsoft.com/en-us/learn/modules/create-classification-model-azure-machine-learningdesigner/introduction> <https://docs.microsoft.com/en-us/learn/modules/create-clustering-model-azure-machine-learningdesigner/1-introduction>

QUESTION 66

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: C, E

Section:

Explanation:

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 67

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In a machine learning model, the data that is used as inputs are called

features.
functions.
labels.
instances.

Answer Area:

Answer Area

In a machine learning model, the data that is used as inputs are called

features.
functions.
labels.
instances.

Section:

Explanation:

QUESTION 68

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

classification.
clustering.
regression.

Answer Area:

Answer Area

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

classification.
clustering.
regression.

Section:
Explanation:

QUESTION 69

DRAG DROP

Match the tool to the Azure Machine Learning task.

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

Select and Place:

Tools	Answer Area
Automated machine learning (automated ML)	Tool Create a Machine Learning workspace
The Azure portal	Tool Use a drag-and-drop interface used to train and deploy models
Machine Learning designer	Tool Use a wizard to select configurations for a machine learning run

Correct Answer:

Tools	Answer Area
	The Azure portal Create a Machine Learning workspace
	Machine Learning designer Use a drag-and-drop interface used to train and deploy models
	Automated machine learning (automated ML) Use a wizard to select configurations for a machine learning run

Section:
Explanation:

QUESTION 70

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Natural language processing can be used to

- classify email messages as work-related or personal.
- predict the number of future car rentals.
- predict which website visitors will make a transaction.
- stop a process in a factory when extremely high temperatures are registered.

Answer Area:

Answer Area

Natural language processing can be used to

- classify email messages as work-related or personal.
- predict the number of future car rentals.
- predict which website visitors will make a transaction.
- stop a process in a factory when extremely high temperatures are registered.

Section:

Explanation:

QUESTION 71

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

The interactive answering of questions entered by a user as part of an application is an example of

- anomaly detection.
- computer vision.
- natural language processing.
- forecasting.



Answer Area:

Answer Area

The interactive answering of questions entered by a user as part of an application is an example of

- anomaly detection.
- computer vision.
- natural language processing.
- forecasting.

Section:

Explanation:

QUESTION 72

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

- A. Recognize text.

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

Correct Answer: B

Section:

QUESTION 73

You plan to build a conversational AI solution that can be surfaced in Microsoft Teams, Microsoft Cortana, and Amazon Alexa

a. Which service should you use?

- A. Azure Bot Service
- B. Azure Cognitive Search
- C. Language service
- D. Speech

Correct Answer: A

Section:

QUESTION 74

Which two languages can you use to write custom code for Azure Machine Learning designer? Each correct answer presents a complete solution.

NOTE; Each correct selection is worth one point.

- A. C#
- B. Scala
- C. Python
- D. R

Correct Answer: C, D

Section:

QUESTION 75

You need to predict the animal population of an area.

Which Azure Machine Learning type should you use?

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

Correct Answer: C

Section:

QUESTION 76

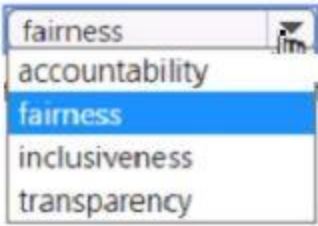
HOTSPOT

Select the answer that correctly completes the sentence.

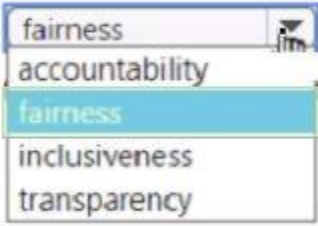
Hot Area:



Answer Area

According to Microsoft's  principle of responsible AI,
AI systems should **NOT** reflect biases from the data sets that are used to train the systems.

Answer Area:
Answer Area

According to Microsoft's  principle of responsible AI,
AI systems should **NOT** reflect biases from the data sets that are used to train the systems.

Section:
Explanation:

QUESTION 77

Which two scenarios are examples of a natural language processing workload? Each correct answer presents a complete solution.
NOTE; Each correct selection is worth one point.

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

Correct Answer: B, D

Section:

QUESTION 78

You need to create a model that labels a collection of your personal digital photographs.
Which Azure AI service should you use?

- A. Azure AI Language
- B. Azure AI Computer Vision
- C. Azure AI Document Intelligence
- D. Azure AI Custom Vision

Correct Answer: B

Section:

QUESTION 79

You need to build an app that will identify celebrities in images.
Which service should you use?

- A. Azure OpenAI Service
- B. Azure Machine Learning
- C. conversational language understanding (CLU)
- D. Azure AI Vision

Correct Answer: D

Section:

QUESTION 80

You have an app that identifies the coordinates of a product in an image of a supermarket shelf.
Which service does the app use?

- A. Azure AI Custom Vision object detection
- B. Azure AI Computer Vision Read
- C. Azure AI Computer Vision optical character recognition (OCR)
- D. Azure AI Custom Vision classification

Correct Answer: A

Section:

QUESTION 81

You need to convert receipts into transactions in a spreadsheet. The spreadsheet must include the date of the transaction, the merchant the total spent and any taxes paid.
Which Azure AI service should you use?

- A. Face
- B. Azure AI Language
- C. Azure AI Document Intelligence
- D. Azure AI Custom Vision

Correct Answer: C

Section:

QUESTION 82

DRAG DROP

You plan to use Azure Cognitive Services to develop a voice controlled personal assistant app.

Match the Azure Cognitive Services to the appropriate tasks.

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 selection is worth one point.

Select and Place:



Services

Azure AI Speech

Azure AI Language service

Azure AI Translator Text

Answer Area

Convert a user's speech to text.

Identify a user's intent.

Provide a spoken response to the user.

Correct Answer:

Services

Answer Area

Azure AI Speech

Convert a user's speech to text.

Azure AI Language service

Identify a user's intent.

Azure AI Translator Text

Provide a spoken response to the user.

Section:

Explanation:

QUESTION 83

DRAG DROP

Match the tasks to the appropriate machine learning models.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Models

- Classification
- Clustering
- Regression

Answer Area

- -
 -
- Assign categories to passengers based on demographic data.
- Predict the amount of consumed fuel based on flight distance.
- Predict whether a passenger will miss their flight based on demographic data.

Correct Answer:

Models

-
-
-

Answer Area

- Classification
 - Regression
 - Clustering
- Assign categories to passengers based on demographic data.
- Predict the amount of consumed fuel based on flight distance.
- Predict whether a passenger will miss their flight based on demographic data.



Section:

Explanation:

QUESTION 84

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

is an example of speech recognition.

- A voice-activated security key system
- Creating an audio commentary for a video recording
- Creating captions for a video recording
- Identifying key phrases in a video transcript

Answer Area:

Answer Area

A voice-activated security key system ▼
A voice-activated security key system
Creating an audio commentary for a video recording
Creating captions for a video recording
Identifying key phrases in a video transcript

is an example of speech recognition.

Section:

Explanation:

QUESTION 85

A smart device that responds to the question. 'What is the stock price of Contoso, Ltd.?' is an example of which AI workload?

- A. computer vision
- B. anomaly detection
- C. knowledge mining
- D. natural language processing

Correct Answer: C

Section:

QUESTION 86

What is an advantage of using a custom model in Form Recognizer?

- A. Only a custom model can be deployed on-premises.
- B. A custom model can be trained to recognize a variety of form types.
- C. A custom model is less expensive than a prebuilt model.
- D. A custom model always provides higher accuracy.

Correct Answer: B

Section:

QUESTION 87

What are three stages in a transformer model? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. object detection
- B. embedding calculation
- C. tokenization
- D. next token prediction
- E. anonymization

Correct Answer: B, C, D

Section:



QUESTION 88

Which parameter should you configure to produce more verbose responses from a chat solution that uses the Azure OpenAI GPT-3.5 model?

- A. Presence penalty
- B. Temperature
- C. Stop sequence
- D. Max response

Correct Answer: B
Section:

QUESTION 89

Which three actions improve the quality of responses returned by a generative AI solution that uses GPT-3.5? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. Add grounding data to prompts.
- B. Provide additional examples to prompts.
- C. Modify tokenization.
- D. Add training data to prompts.
- E. Modify system messages.

Correct Answer: B, D, E
Section:



QUESTION 90

DRAG DROP

Match the Azure OpenAI large language model (LLM) process to the appropriate task.

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

NOTE: Each correct match is worth one point.

Select and Place:

Processes

Classifying

Generating

Summarizing

Translating

Answer Area

Detect the genre of a work of fiction.

Create a list of bullet points based on text input.

Create advertising slogans from a product description.

Correct Answer:

Processes

Four empty text input boxes stacked vertically, with the bottom-most box containing the text "Translating".

Answer Area

- Classifying Detect the genre of a work of fiction.
- Summarizing Create a list of bullet points based on text input.
- Generating Create advertising slogans from a product description.

Section:

Explanation:

QUESTION 91

HOTSPOT

Select the answer that correctly completes the sentence.



Hot Area:

Answer Area

You can use the

Dropdown menu with options: Azure AI Custom Vision, Azure AI Computer Vision (highlighted), Azure AI Custom Vision, Azure AI Document Intelligence, Azure Video Analyzer for Media.

service to train an object detection model by using your own images.

Answer Area:

Answer Area

You can use the

Dropdown menu with options: Azure AI Custom Vision, Azure AI Computer Vision (highlighted), Azure AI Custom Vision, Azure AI Document Intelligence, Azure Video Analyzer for Media.

service to train an object detection model by using your own images.

Section:

Explanation:

QUESTION 92

DRAG DROP

Match the computer vision service to the appropriate AI workload.

To answer, drag the appropriate service from the column on the left to its workload 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:

Services

Azure AI Custom Vision

Azure AI Document Intelligence

Azure AI Vision

Answer Area

Extract information from scanned forms and invoices.

Analyze images and video, and extract descriptions, tags, objects, and text.

Train custom image classification and object detection models by using your own images.

Correct Answer:

Services

Azure AI Custom Vision

Azure AI Document Intelligence

Azure AI Vision

Answer Area

Extract information from scanned forms and invoices.

Analyze images and video, and extract descriptions, tags, objects, and text.

Train custom image classification and object detection models by using your own images.

Section:

Explanation:

QUESTION 93

What can be used to complete a paragraph based on a sentence provided by a user?

- A. Azure AI Language
- B. Azure OpenAI
- C. Azure Machine Learning
- D. Azure AI Vision

Correct Answer: B

Section:

QUESTION 94

DRAG DROP

Match the AI solution to the appropriate task.

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

NOTE: Each correct match is worth one point.

Select and Place:

Solutions

- Computer vision
- Data mining
- Generative AI
- Text analytics

Answer Area

- Generate a caption from a given image.
- Generate an image from a given caption.
- Generate a 200-word summary from a 2,000-word article.

Correct Answer:

Solutions

- Data mining

Answer Area

- Computer vision generate a caption from a given image.
- Generative AI generate an image from a given caption.
- Text analytics generate a 200-word summary from a 2,000-word article.

Section:

Explanation:

QUESTION 95

Which OpenAI model does GitHub Copilot use to make suggestions for client-side JavaScript?

- A. GPT-4
- B. Codex
- C. DALL-E
- D. GPT-3

Correct Answer: B

Section:

QUESTION 96

Which format should you use to send requests to a REST API endpoint for Azure OpenAI?

- A. CSV
- B. JSON
- C. XML
- D. YAML

Correct Answer: B

Section:

QUESTION 97

You are building an AI-based loan approval app.

You need to ensure that the app documents why a loan is approved or rejected and makes the report available to the applicant.

This is an example of which Microsoft responsible AI principle?

- A. fairness
- B. inclusiveness
- C. transparency
- D. accountability

Correct Answer: C

Section:

QUESTION 98

Which Azure AI Document Intelligence prebuilt model should you use to extract parties and jurisdictions from a legal document?

- A. contract
- B. layout
- C. general document
- D. read

Correct Answer: C

Section:

