

Microsoft.AI-900.vFeb-2024.by.Xakinato.110q

Number: AI-900  
Passing Score: 800.0  
Time Limit: 120.0  
File Version: 11.0

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



**Exam A**

**QUESTION 1**

DRAG DROP

Match the facial recognition tasks to the appropriate questions.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Tasks	Answer Area
grouping	Task Do two images of a face belong to the same person?
identification	Task Does this person look like other people?
similarity	Task Do all the faces belong together?
verification	Task Who is this person in this group of people?

Correct Answer:

Tasks	Answer Area
verification	Do two images of a face belong to the same person?
similarity	Does this person look like other people?
grouping	Do all the faces belong together?
identification	Who is this person in this group of people?

**Section:**

**Explanation:**

Box 1: verification

Face verification: Check the likelihood that two faces belong to the same person and receive a confidence score.

Box 2: similarity

Box 3: Grouping

Box 4: identification

Face detection: Detect one or more human faces along with attributes such as: age, emotion, pose, smile, and facial hair, including 27 landmarks for each face in the image.

Reference:

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

**QUESTION 2**

DRAG DROP

Match the types of computer vision workloads to the appropriate scenarios.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Workloads Types	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)	

Correct Answer:

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

Section:

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.

Reference:

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

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

**QUESTION 3**

DRAG DROP

Match the types of natural languages processing workloads to the appropriate scenarios.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Workloads Types	Answer Area
Entity recognition	Workload Type Extracts persons, locations, and organizations from the text
Key phrase extraction	Workload Type Evaluates text along a positive-negative scale
Language modeling	Workload Type Returns text translated to the specified target language
Sentiment analysis	
Natural language processing	
Translation	
Speech recognition and speech synthesis	

Correct Answer:

Workloads Types	Answer Area
Entity recognition	Entity recognition Extracts persons, locations, and organizations from the text
Key phrase extraction	Sentiment analysis Evaluates text along a positive-negative scale
Language modeling	Translation Returns text translated to the specified target language
Natural language processing	
Speech recognition and speech synthesis	

Section:

Explanation:

Box 1: Entity recognition

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.

Box 2: Sentiment analysis

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

Box 3: Translation

Using Microsoft's Translator text API

This versatile API from Microsoft can be used for the following:

Translate text from one language to another.

Transliterate text from one script to another.

Detecting language of the input text.

Find alternate translations to specific text.

Determine the sentence length.

Reference:

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

#### QUESTION 4

DRAG DROP

You plan to apply Text Analytics API features to a technical support ticketing system.

Match the Text Analytics API features to the appropriate natural language processing scenarios.

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

NOTE: Each correct selection is worth one point.

Select and Place:

API Features	Answer Area
Entity recognition	API Feature Understand how upset a customer is based on the text contained in the support ticket.
Key phrase extraction	API Feature Summarize important information from the support ticket.
Language detection	API Feature Extract key dates from the support ticket.
Sentiment analysis	

Correct Answer:

API Features	Answer Area
	Sentiment analysis Understand how upset a customer is based on the text contained in the support ticket.
	Key phrase extraction Summarize important information from the support ticket.
Language detection	Entity recognition Extract key dates from the support ticket.

Section:

Explanation:

Box1: Sentiment analysis

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

Box 2: Broad entity extraction

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.

Box 3: Entity Recognition

Named Entity Recognition: Identify and categorize entities in your text as people, places, organizations, date/time, quantities, percentages, currencies, and more. Well-known entities are also recognized and linked to more information on the web.

Reference:

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

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

#### QUESTION 5

You are building a tool that will process your company's product images 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

**Correct Answer: C**

**Section:**

**Explanation:**

Reference:

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



#### QUESTION 6

HOTSPOT

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

**Hot Area:**

### Answer Area

	▼
Accuracy	
Confidence	
Root Mean Square Error	
Sentiment	

\_\_\_\_\_ is the calculated probability of a correct image classification.

**Answer Area:**

## Answer Area

	▼
Accuracy	
Confidence	
Root Mean Square Error	
Sentiment	

is the calculated probability of a correct image classification.

**Section:**

**Explanation:**

Reference:

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

**QUESTION 7**

**HOTSPOT**

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

**Hot Area:**



## Answer Area

Ensuring an AI system does not provide a prediction when important fields contain unusual or missing values is

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

principle for responsible AI.

**Answer Area:**

### Answer Area

Ensuring an AI system does not provide a prediction when important fields contain unusual or missing values is  principle for responsible AI.

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

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai>

### QUESTION 8

HOTSPOT

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

Hot Area:

### Answer Area

Ensuring that the numeric variables in training data are on a similar scale is an example of

- 
- data ingestion.
- feature engineering.
- feature selection.
- model training.

Answer Area:





**Answer Area**

Ensuring that the numeric variables in training data are on a similar scale is an example of

	▼
data ingestion.	
feature engineering.	
feature selection.	
model training.	

**Section:**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-science-process/create-features>

**QUESTION 9**

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
Object detection can identify the location of a damaged product in an image.	<input type="radio"/>	<input type="radio"/>
Object detection can identify multiple instances of a damaged product in an image.	<input type="radio"/>	<input type="radio"/>
Object detection can identify multiple types of damaged products in an image.	<input type="radio"/>	<input type="radio"/>

**Answer Area:**

**Answer Area**

Statements	Yes	No
Object detection can identify the location of a damaged product in an image.	<input checked="" type="radio"/>	<input type="radio"/>
Object detection can identify multiple instances of a damaged product in an image.	<input type="radio"/>	<input checked="" type="radio"/>
Object detection can identify multiple types of damaged products in an image.	<input checked="" type="radio"/>	<input type="radio"/>

**Section:**

**Explanation:**

**QUESTION 10**

HOTSPOT

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

**Hot 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.

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.

Section:

Explanation:

#### QUESTION 11

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

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: C, D

Section:

#### QUESTION 12

HOTSPOT

Select the answer that correctly completes the sentence

Hot Area:

The logo for Vdumps.com, featuring a stylized orange 'V' followed by the word 'dumps' in a grey, lowercase, sans-serif font.

Answer Area

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

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

Answer Area:

Answer Area

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

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



Section:

Explanation:

**QUESTION 13**

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. QnA Maker
- B. Azure Bot Service
- C. Form Recognizer
- D. Anomaly Detector

**Correct Answer: A, B**

Section:

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>

**QUESTION 14**

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:

Statements	Yes	No
A restaurant can use a chatbot to empower customers to make reservations by using a website or an app.	<input type="checkbox"/>	<input type="checkbox"/>
A restaurant can use a chatbot to answer inquiries about business hours from a webpage.	<input type="checkbox"/>	<input type="checkbox"/>
A restaurant can use a chatbot to automate responses to customer reviews on an external website.	<input type="checkbox"/>	<input type="checkbox"/>



Answer Area:

Statements	Yes	No
A restaurant can use a chatbot to empower customers to make reservations by using a website or an app.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
A restaurant can use a chatbot to answer inquiries about business hours from a webpage.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
A restaurant can use a chatbot to automate responses to customer reviews on an external website.	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Section:

Explanation:

Reference:

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

**QUESTION 15**

You have a webchat bot that provides responses from a QnA Maker knowledge base. You need to ensure that the bot uses user feedback to improve the relevance of the responses over time. What should you use?

- A. key phrase extraction
- B. sentiment analysis
- C. business logic
- D. active learning

**Correct Answer: D**

**Section:**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/improve-knowledge-base>

**QUESTION 16**

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



**Correct Answer: B**

**Section:**

**Explanation:**

Reference:

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

**QUESTION 17**

**HOTSPOT**

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

**Hot Area:**

## Answer Area

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

Answer Area:



## Answer Area

Statements	Yes	No
A bot that responds to queries by internal users is an example of a conversational AI workload.	<input checked="" type="radio"/>	<input type="radio"/>
An application that displays images relating to an entered search term is an example of a conversational AI workload.	<input type="radio"/>	<input checked="" type="radio"/>
A web form used to submit a request to reset a password is an example of a conversational AI workload.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Reference:

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

#### QUESTION 18

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 19

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 20

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 21

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 22

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 23

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 24

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 25

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 26

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 27

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.

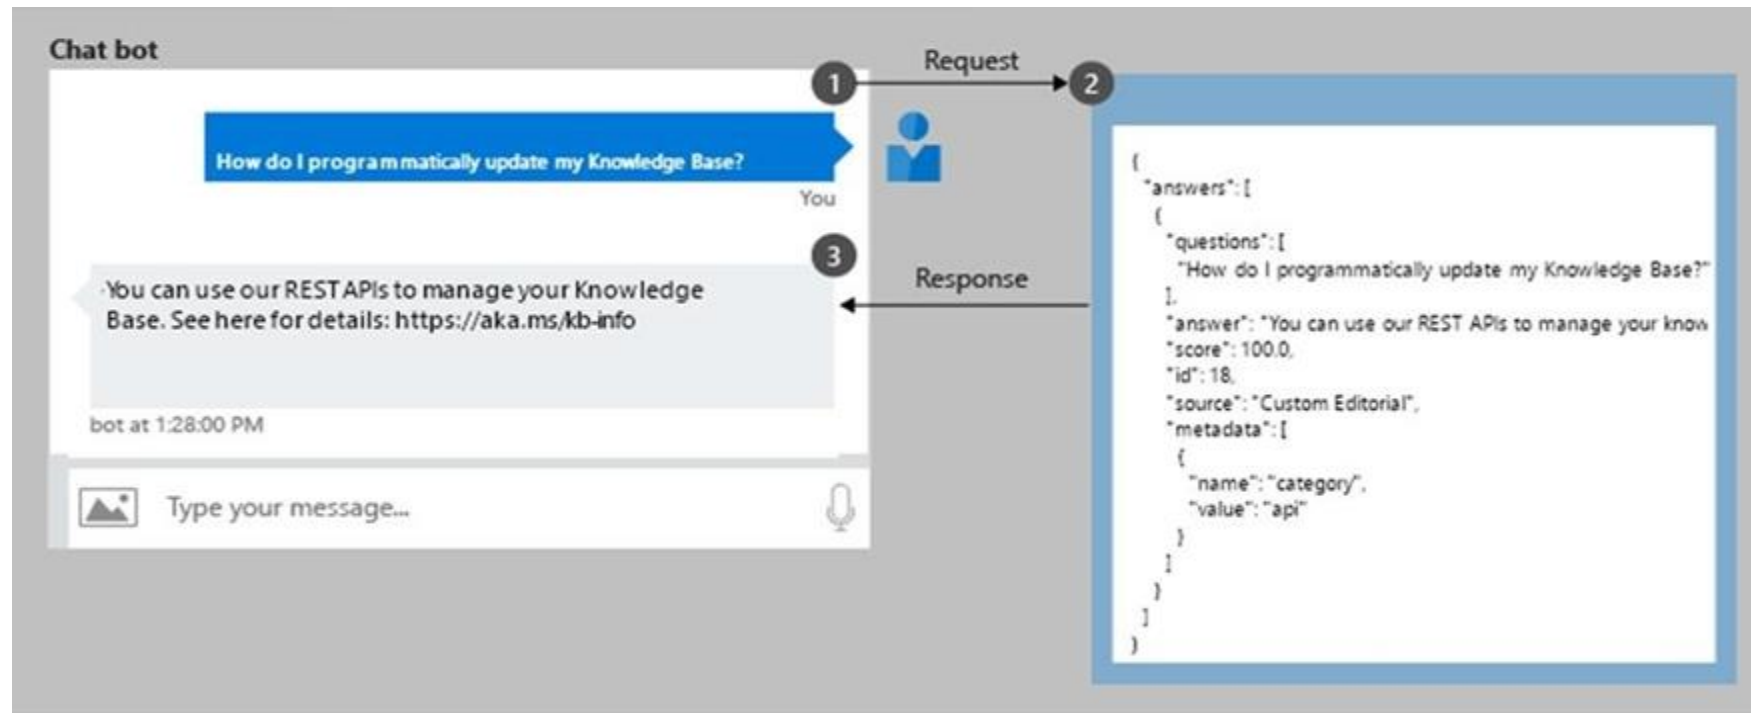
- A. a smart device in the home that responds to questions such as "What will the weather be like today?"
- B. a website that uses a knowledge base to interactively respond to users' questions
- C. assembly line machinery that autonomously inserts headlamps into cars
- D. monitoring the temperature of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

**Correct Answer: A, B**

**Section:**

#### QUESTION 28

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 29**

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 30**

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 31**

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 32**

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 33**

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
Forecasting housing prices based on historical data is an example of anomaly detection.	<input type="radio"/>	<input type="radio"/>
Identifying suspicious sign-ins by looking for deviations from usual patterns is an example of anomaly detection.	<input type="radio"/>	<input type="radio"/>
Predicting whether a patient will develop diabetes based on the patient's medical history is an example of anomaly detection.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Forecasting housing prices based on historical data is an example of anomaly detection.	<input type="radio"/>	<input checked="" type="radio"/>
Identifying suspicious sign-ins by looking for deviations from usual patterns is an example of anomaly detection.	<input checked="" type="radio"/>	<input type="radio"/>
Predicting whether a patient will develop diabetes based on the patient's medical history is an example of anomaly detection.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

QUESTION 34

HOTSPOT

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

Hot 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

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



**Section:**

**Explanation:**

**QUESTION 35**

**HOTSPOT**

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

**Hot Area:**

When developing an AI system for self-driving cars, the Microsoft for responsible AI should be applied to ensure consistent operation system during unexpected circumstances.

- inclusiveness
- accountability
- reliability and safety
- fairness

principle of the

**Answer Area:**

When developing an AI system for self-driving cars, the Microsoft for responsible AI should be applied to ensure consistent operation system during unexpected circumstances.

- inclusiveness
- accountability
- reliability and safety
- fairness

principle of the

**Section:**

**Explanation:**

Reliability and safety: To build trust, it's critical that AI systems operate reliably, safely, and consistently under normal circumstances and in unexpected conditions. These systems should be able to operate as they were originally designed, respond safely to unanticipated conditions, and resist harmful manipulation.

**Reference:**

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

**QUESTION 36**

**HOTSPOT**

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

**Hot 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 recognizer (OCR).
- semantic segmentation.



**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 recognizer (OCR).
- semantic segmentation.



**Section:**

**Explanation:**

Reference:

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



**QUESTION 37**

HOTSPOT

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

**Hot Area:**

**Answer Area**



- Feature engineering
- Feature selection
- Model evaluation
- Model training

\_\_\_\_\_ is used to generate additional features.



**Answer Area:**



Feature engineering

Feature selection

Model evaluation

Model training

is used to generate additional features.

**Section:**

**Explanation:**

Reference:

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

**QUESTION 38**

**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
	Organizing documents into groups based on similarities of the text contained in the documents is an example of clustering.	<input type="radio"/>	<input type="radio"/>
	Grouping similar patients based on symptoms and diagnostic test results is an example of clustering.	<input type="radio"/>	<input type="radio"/>
	Predicting whether a person will develop mild, moderate, or severe allergy symptoms based on pollen count is an example of clustering.	<input type="radio"/>	<input type="radio"/>

**Answer Area:**

### Answer Area

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

**Section:**

**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.

Reference:

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

**QUESTION 39**

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

**Hot Area:**

### Answer Area

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

Answer Area:



### Answer Area

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

Section:

Explanation:

Box 1: No

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

Box 2: Yes

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

Box 3: No

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

Reference:

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

**QUESTION 40**

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area	Statements	Yes	No
	You can use the Translator service to translate text between languages.	<input type="radio"/>	<input type="radio"/>
	You can use the Translator service to detect the language of a given text.	<input type="radio"/>	<input type="radio"/>
	You can use the Translator service to transcribe audible speech into text.	<input type="radio"/>	<input type="radio"/>

Answer Area:

### Answer Area

Statements	Yes	No
You can use the Translator service to translate text between languages.	<input checked="" type="radio"/>	<input type="radio"/>
You can use the Translator service to detect the language of a given text.	<input checked="" type="radio"/>	<input type="radio"/>
You can use the Translator service to transcribe audible speech into text.	<input type="radio"/>	<input checked="" type="radio"/>

#### Section:

#### Explanation:

The translator service provides multi-language support for text translation, transliteration, language detection, and dictionaries. Speech-to-Text, also known as automatic speech recognition (ASR), is a feature of Speech Services that provides transcription.

Reference:

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

<https://docs.microsoft.com/en-us/legal/cognitive-services/speech-service/speech-to-text/transparency-note>

#### QUESTION 41

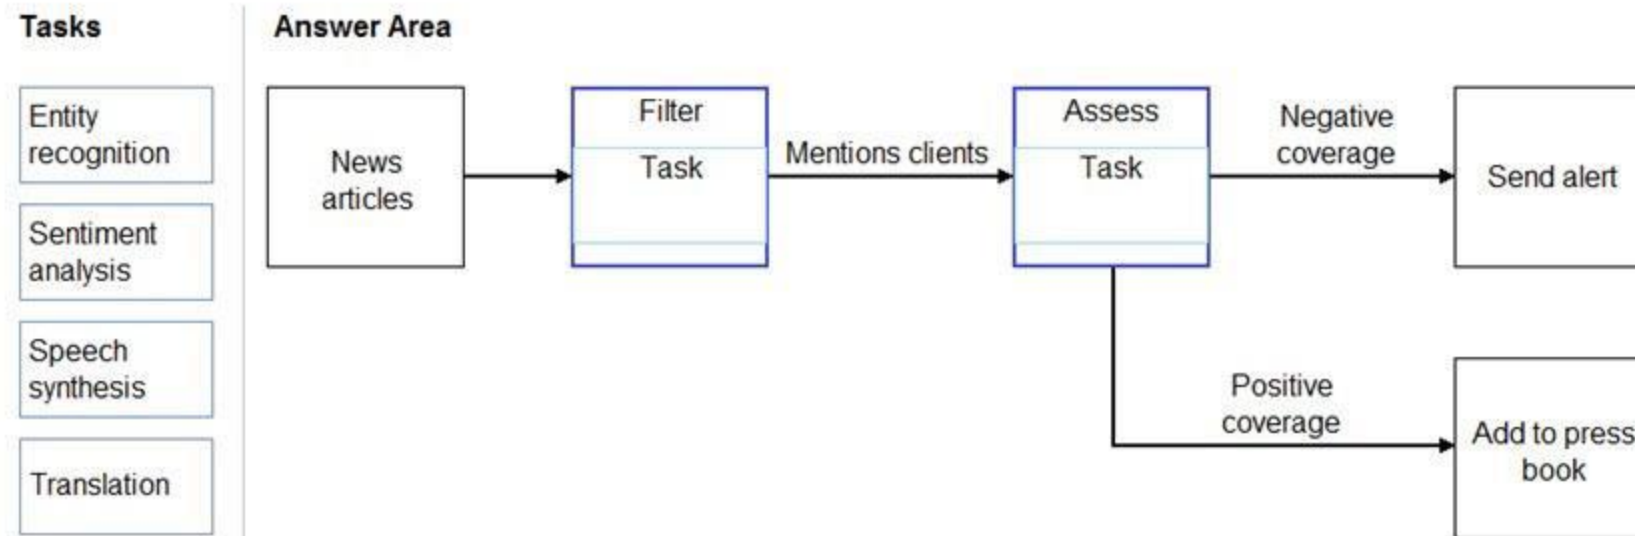
##### DRAG DROP

You need to scan the news for articles about your customers and alert employees when there is a negative article. Positive articles must be added to a press book.

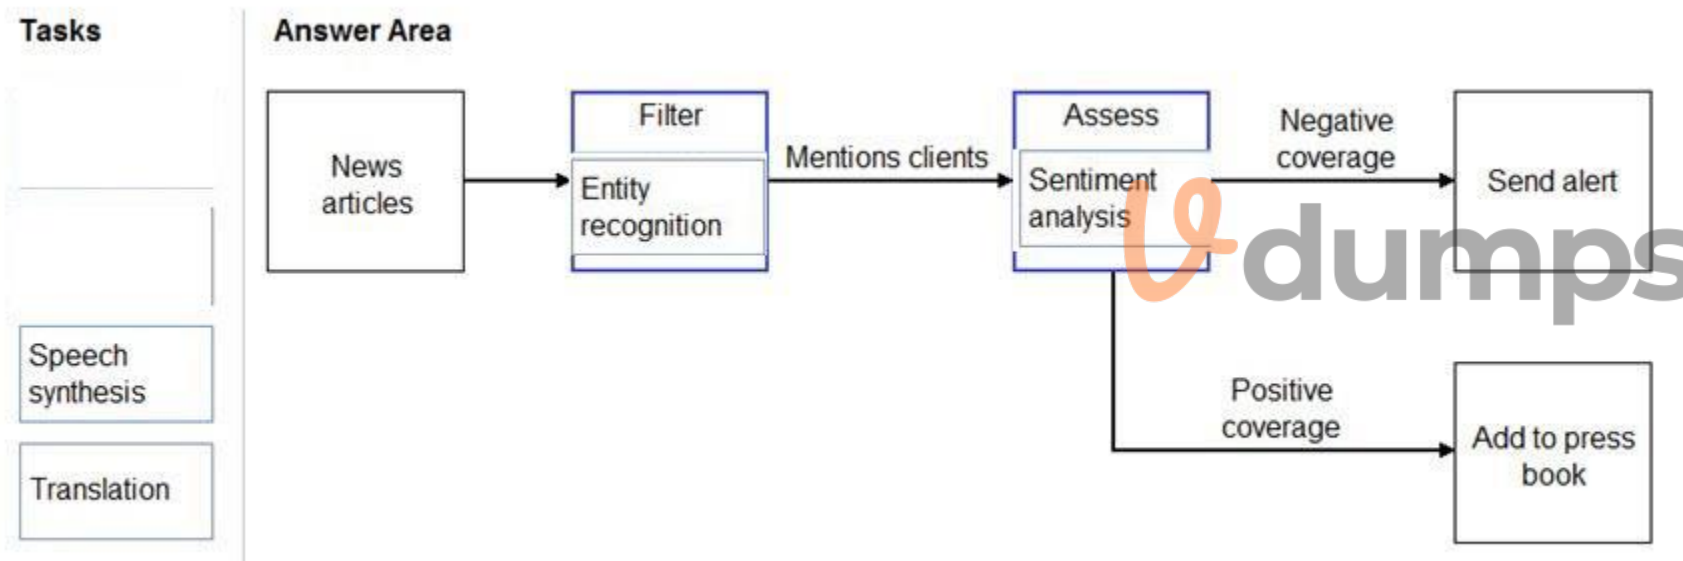
Which natural language processing tasks should you use to complete the process? To answer, drag the appropriate tasks to the correct locations. Each task may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

#### Select and Place:



Correct Answer:



**Section:**

**Explanation:**

Box 1: Entity recognition

the Named Entity Recognition module in Machine Learning Studio (classic), to identify the names of things, such as people, companies, or locations in a column of text.

Named entity recognition is an important area of research in machine learning and natural language processing (NLP), because it can be used to answer many real-world questions, such as:

Which companies were mentioned in a news article?

Does a tweet contain the name of a person? Does the tweet also provide his current location?

Were specified products mentioned in complaints or reviews?

Box 2: Sentiment Analysis

The Text Analytics API's Sentiment Analysis feature provides two ways for detecting positive and negative sentiment. If you send a Sentiment Analysis request, the API will return sentiment labels (such as "negative", "neutral" and "positive") and confidence scores at the sentence and document-level.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/named-entity-recognition>

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-to/text-analytics-how-to-sentiment-analysis>

**QUESTION 42**

You are building a knowledge base by using QnA Maker.

Which file format can you use to populate the knowledge base?

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

**Correct Answer: D**

**Section:**

**Explanation:**

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

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

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

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

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

Incorrect Answers:

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

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

Reference:

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

**QUESTION 43**

In which scenario should you use key phrase extraction?

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

**Correct Answer: C**

**Section:**

**QUESTION 44**

You have insurance claim reports that are stored as text.

You need to extract key terms from the reports to generate summaries.

Which type of AI workload should you use?

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

**Correct Answer: A**

**Section:**

**Explanation:**

Reference:



**QUESTION 45**

DRAG DROP

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.

Select and Place:

**Principles**

- Fairness
- Privacy and security
- Reliability and safety
- Transparency

**Answer Area**

Three empty dashed boxes for matching.

The system must not discriminate based on gender, race

Personal data must be visible only to approve

Automated decision-making processes must be recorded so that approved users can identify why a decision was made



Correct Answer:

**Principles**

- 
- 
- Reliability and safety
- 

**Answer Area**

Three boxes for matching, with 'Fairness', 'Privacy and security', and 'Transparency' placed in the first, second, and third boxes respectively.

The system must not discriminate based on gender, race

Personal data must be visible only to approve

Automated decision-making processes must be recorded so that approved users can identify why a decision was made

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai>

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



**QUESTION 46**

DRAG DROP

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.

Select and Place:

**Processes**

- data encryption
- model retraining
- model training
- data preparation
- model evaluation

**Answer Area**

Correct Answer:

**Processes**

- data encryption
- model retraining
- 
- 
- 

**Answer Area**

- data preparation
- model training
- model evaluation

**Section:**

**Explanation:**

Reference:

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

**QUESTION 47**

You are building an AI-based app.

You need to ensure that the app uses the principles for responsible AI.



Which two principles should you follow? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Implement an Agile software development methodology
- B. Implement a process of AI model validation as part of the software review process
- C. Establish a risk governance committee that includes members of the legal team, members of the risk management team, and a privacy officer
- D. Prevent the disclosure of the use of AI-based algorithms for automated decision making

**Correct Answer: B, C**

**Section:**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai>

<https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/3-implications-responsible-ai-practical>

#### QUESTION 48

HOTSPOT

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

**Hot Area:**

**Answer Area**

According to Microsoft's 

	▼
accountability	
fairness	
inclusiveness	
transparency	

 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,

accountability
fairness
inclusiveness
transparency

AI systems should **NOT** reflect biases from the data sets that are used to train the systems.

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai>

QUESTION 49

HOTSPOT

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

Hot Area:

Answer Area

Data values that influence the prediction of a model are called

dependant variables.
features.
identifiers.
labels.

Answer Area:

Data values that influence the prediction of a model are called

▼

- dependant variables.
- features.
- identifiers.
- labels.

**Section:**

**Explanation:**

Reference:

<https://www.baeldung.com/cs/feature-vs-label>

<https://machinelearningmastery.com/discover-feature-engineering-how-to-engineer-features-and-how-to-get-good-at-it/>

**QUESTION 50**

**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:**



**Statements**

**Yes**

**No**

Automated machine learning is the process of automating the time-consuming, iterative tasks of machine learning model development.

Automated machine learning can automatically infer the training data from the use case provided.

Automated machine learning works by running multiple training iterations that are scored and ranked by the metrics you specify.

Automated machine learning enables you to specify a dataset and will automatically understand which label to predict.

**Answer Area:**

Statements	Yes	No
Automated machine learning is the process of automating the time-consuming, iterative tasks of machine learning model development.	<input checked="" type="radio"/>	<input type="radio"/>
Automated machine learning can automatically infer the training data from the use case provided.	<input type="radio"/>	<input checked="" type="radio"/>
Automated machine learning works by running multiple training iterations that are scored and ranked by the metrics you specify.	<input checked="" type="radio"/>	<input type="radio"/>
Automated machine learning enables you to specify a dataset and will automatically understand which label to predict.	<input type="radio"/>	<input checked="" type="radio"/>

**Section:**

**Explanation:**

Box 1: Yes

Automated machine learning, also referred to as automated ML or AutoML, is the process of automating the time consuming, iterative tasks of machine learning model development. It allows data scientists, analysts, and developers to build ML models with high scale, efficiency, and productivity all while sustaining model quality.

Box 2: No

Box 3: Yes

During training, Azure Machine Learning creates a number of pipelines in parallel that try different algorithms and parameters for you. The service iterates through ML algorithms paired with feature selections, where each iteration produces a model with a training score. The higher the score, the better the model is considered to "fit" your data. It will stop once it hits the exit criteria defined in the experiment.

Box 4: No

Apply automated ML when you want Azure Machine Learning to train and tune a model for you using the target metric you specify. The label is the column you want to predict.

Reference:

<https://azure.microsoft.com/en-us/services/machine-learning/automatedml/#features>

**QUESTION 51**

DRAG DROP

Match the types of AI workloads to the appropriate scenarios.

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

NOTE: Each correct selection is worth one point.

**Select and Place:**

**Workload Types**

Anomaly detection

Computer vision

Machine Learning (Regression)

Natural language processing

**Answer Area**

Workload Type

Identify handwritten letters.

Workload Type

Predict the sentiment of a social media post.

Workload Type

Identify a fraudulent credit card payment.

Workload Type

Predict next month's toy sales.

**Correct Answer:**

Workload Types	Answer Area
	Computer vision Identify handwritten letters.
	Natural language processing Predict the sentiment of a social media post.
	Anomaly detection Identify a fraudulent credit card payment.
	Machine Learning (Regression) Predict next month's toy sales.

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/learn/paths/get-started-with-artificial-intelligence-on-azure/>

**QUESTION 52**

DRAG DROP

Match the types of machine learning to the appropriate scenarios.

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

NOTE: Each correct selection is worth one point.

Select and Place:

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

Correct Answer:

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

Section:

Explanation:

Box 1: Regression

In the most basic sense, regression refers to prediction of a numeric target.

Linear regression attempts to establish a linear relationship between one or more independent variables and a numeric outcome, or dependent variable.

You use this module to define a linear regression method, and then train a model using a labeled dataset. The trained model can then be used to make predictions.

Box 2: Classification

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

Box 3: Clustering

Clustering, in machine learning, is a method of grouping data points into similar clusters. It is also called segmentation.

Over the years, many clustering algorithms have been developed. Almost all clustering algorithms use the features of individual items to find similar items. For example, you might apply clustering to find similar people by demographics. You might use clustering with text analysis to group sentences with similar topics or sentiment.

Reference:

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

**QUESTION 53**

DRAG DROP

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.

Select and Place:

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	

Correct Answer:

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

**Section:****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.

Reference:

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

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

**QUESTION 54**

When you design an AI system to assess whether loans should be approved, the factors used to make the decision should be explainable.

This is an example of which Microsoft guiding principle for responsible AI?

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

**Correct Answer: A**

**Section:**

**Explanation:**

Achieving transparency helps the team to understand the data and algorithms used to train the model, what transformation logic was applied to the data, the final model generated, and its associated assets. This information offers insights about how the model was created, which allows it to be reproduced in a transparent way.

Incorrect Answers:

B: Inclusiveness mandates that AI should consider all human races and experiences, and inclusive design practices can help developers to understand and address potential barriers that could unintentionally exclude people. Where possible, speech-to-text, text-to-speech, and visual recognition technology should be used to empower people with hearing, visual, and other impairments. C: Fairness is a core ethical principle that all humans aim to understand and apply. This principle is even more important when AI systems are being developed. Key checks and balances need to make sure that the system's decisions don't discriminate or run a gender, race, sexual orientation, or religion bias toward a group or individual. D: A data holder is obligated to protect the data in an AI system, and privacy and security are an integral part of this system. Personal needs to be secured, and it should be accessed in a way that doesn't compromise an individual's privacy.

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai>

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/responsible-ai>

**QUESTION 55**

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**

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.

**Answer Area:****Answer Area****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.

**Section:****Explanation:**

Box 1: Yes

Achieving transparency helps the team to understand the data and algorithms used to train the model, what transformation logic was applied to the data, the final model generated, and its associated assets. This information offers insights about how the model was created, which allows it to be reproduced in a transparent way.

Box 2: No

A data holder is obligated to protect the data in an AI system, and privacy and security are an integral part of this system. Personal needs to be secured, and it should be accessed in a way that doesn't compromise an individual's privacy.

Box 3: No

Inclusiveness mandates that AI should consider all human races and experiences, and inclusive design practices can help developers to understand and address potential barriers that could unintentionally exclude people.

Where possible, speech-to-text, text-to-speech, and visual recognition technology should be used to empower people with hearing, visual, and other impairments.

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai>

**QUESTION 56**

HOTSPOT

Select the answer that correctly completes the sentence

Hot Area:

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

Answer Area:

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

Section:

Explanation:

**QUESTION 57**

DRAG DROP

You plan to deploy an Azure Machine Learning model by using the Machine Learning designer.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

**Actions**

- Evaluate the model against the original dataset.
- Ingest and prepare a dataset.
- Split the data randomly into training data and validation data.
- Train the model.
- Evaluate the model against the validation dataset.

**Answer Area**

- 
- 
- 
- 

**Correct Answer:**

**Actions**

- Evaluate the model against the original dataset.
- 
- 
- 
- 

**Answer Area**

- Ingest and prepare a dataset.
- Split the data randomly into training data and validation data.
- Train the model.
- Evaluate the model against the validation dataset.



**Section:**

**Explanation:**

**QUESTION 58**

**HOTSPOT**

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**

Regression models can be used to predict the sale price of auctioned items.

- Regression
- Classification
- Clustering
- Regression

**Answer Area:**

Answer Area

Regression models can be used to predict the sale price of auctioned items.

- Regression
- Classification
- Clustering
- Regression

Section:

Explanation:

QUESTION 59

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

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.

Answer Area:

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.

Section:

Explanation:

### QUESTION 60

For which two workloads can you use computer vision? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. creating photorealistic images by using three-dimensional models
- B. assigning the color pixels in an image to object names
- C. describing the contents of an image
- D. detecting inconsistencies and anomalies in a stream of data
- E. creating visual representations of numerical data

**Correct Answer: B, C**

**Section:**

### QUESTION 61

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 users geographical location. Which type of natural language processing workload should you use?

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

**Correct Answer: C**

**Section:**

### QUESTION 62

HOTSPOT

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**

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

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



**Answer Area:**

**Answer Area**

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

- 
- AUC metric
- 
- ROC curve
- threshold

**Section:**

**Explanation:**

**QUESTION 63**

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 64**

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 65**

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 66**

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 67**

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

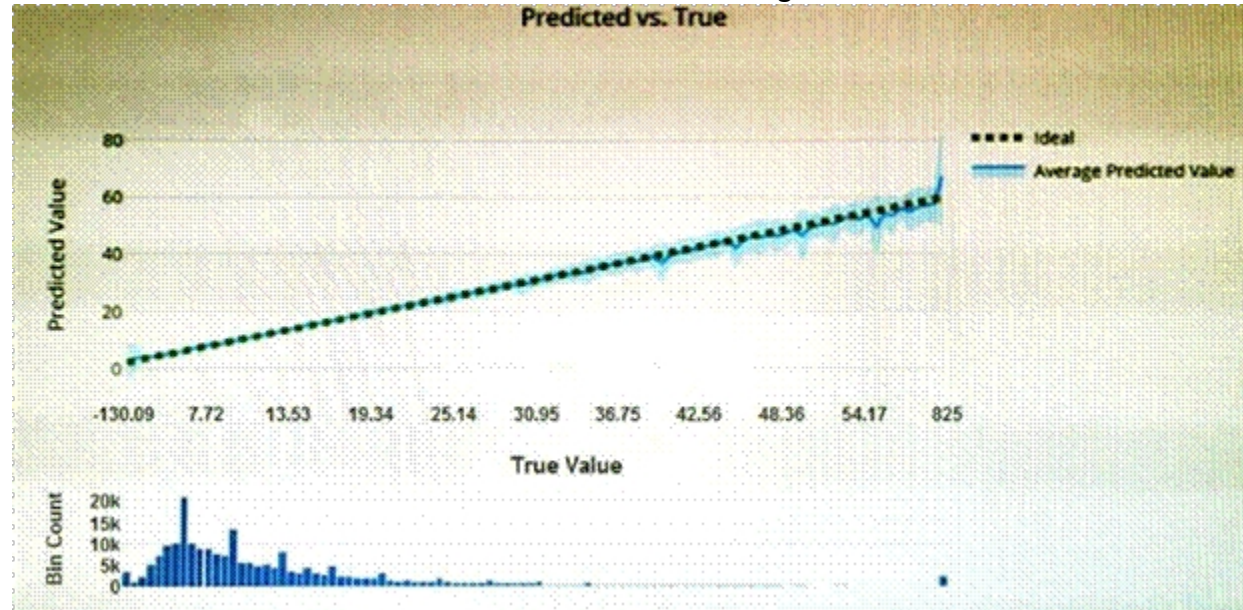
Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-configure-cross-validation-data-splits2>



### QUESTION 68

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 69

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 70



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 71

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 72

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 73**

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 74**

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 75**

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 76**

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 77

You have an Internet of Things (IoT) device that monitors engine temperature.

The device generates an alert if the engine temperature deviates from expected norms.

Which type of AI workload does the device represent?

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

**Correct Answer: C**

**Section:**

#### QUESTION 78

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

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

**Correct Answer: C**

**Section:**

#### QUESTION 79

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

A smart device in the home that responds to questions such as "When is my next appointment?" is an example of conversational AI.

An interactive webchat feature on a company website can be implemented by using Azure Bot Service.

Automatically generating captions for pre-recorded videos is an example of conversational AI.

Yes

No

Answer Area:

Answer Area

Statements

A smart device in the home that responds to questions such as "When is my next appointment?" is an example of conversational AI.

An interactive webchat feature on a company website can be implemented by using Azure Bot Service.

Automatically generating captions for pre-recorded videos is an example of conversational AI.

Yes

No

Section:

Explanation:

QUESTION 80

During the process of Machine Learning, when should you review evaluation metrics?

- A. After you clean the data.
- B. Before you train a model.
- C. Before you choose the type of model.
- D. After you test a model on the validation data.

Correct Answer: D

Section:

QUESTION 81

Which Azure Cognitive Services service can be used to identify documents that contain sensitive information?

- A. Custom Vision
- B. Conversational Language Understanding
- C. Form Recognizer

**Correct Answer: C**

**Section:**

**QUESTION 82**

HOTSPOT

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**

A historian can use  to digitize newspaper articles.

- facial analysis
- image classification
- object detection
- optical character recognition (OCR)

**Answer Area:**

**Answer Area**

A historian can use  to digitize newspaper articles.

- facial analysis
- image classification
- object detection
- optical character recognition (OCR)

**Section:**

**Explanation:**

**QUESTION 83**

HOTSPOT

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**

When building a regression model, labels must have a data type of .

- numeric.
- boolean.
- datetime.
- numeric.
- text.

**Answer Area:**

Answer Area

When building a regression model, labels must have a data type of

numeric.  
boolean.  
datetime.  
numeric.  
text.

Section:

Explanation:

**QUESTION 84**

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

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

**Correct Answer: B**

Section:

**QUESTION 85**

HOTSPOT

Select the answer that correctly completes the sentence.



Hot Area:

Answer Area

Detecting unusual temperature fluctuations for a large machine is an example of

an anomaly detection workload.  
a computer vision workload.  
a knowledge mining workload.  
a natural language processing (NLP) workload.  
an anomaly detection workload.

Answer Area:

Answer Area

Detecting unusual temperature fluctuations for a large machine is an example of

an anomaly detection workload.  
a computer vision workload.  
a knowledge mining workload.  
a natural language processing (NLP) workload.  
an anomaly detection workload.

Section:

Explanation:

**QUESTION 86**

HOTSPOT

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.

**Hot Area:**

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)

**Answer Area:**

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)

**Section:**

**Explanation:**

**QUESTION 87**

HOTSPOT

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

**Hot Area:**

Answer Area

Data values that used to make a prediction are called

features.  
 dependant variables.  
 features  
 identifiers.  
 labels.

**Answer Area:**

**Answer Area**

Data values that used to make a prediction are called

features.	▼
dependant variables.	
features	
identifiers.	
labels.	

**Section:**

**Explanation:**

**QUESTION 88**

HOTSPOT

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**

For 

▼
time constraints
feature engineering
MLflow models
model training

, you use a portion of a dataset to prepare a machine learning

model and retain the balance of the dataset to verify the results.

**Answer Area:**

**Answer Area**

For 

▼
time constraints
feature engineering
MLflow models
model training

, you use a portion of a dataset to prepare a machine learning

model and retain the balance of the dataset to verify the results.

**Section:**

**Explanation:**

**QUESTION 89**

You have a natural language processing (NLP) 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

**Correct Answer: C**

**Section:**

**QUESTION 90**

HOTSPOT

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**

The  service can be used to extract information from a driver's license to populate a database.

- Form Recognizer
- Computer Vision**
- Conversational Language Understanding
- Custom Vision
- Form Recognizer

**Answer Area:**

**Answer Area**

The  service can be used to extract information from a driver's license to populate a database.

- Form Recognizer
- Computer Vision**
- Conversational Language Understanding
- Custom Vision
- Form Recognizer

**Section:**

**Explanation:**

**QUESTION 91**

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 92

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 93

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

- classification.
- clustering.
- regression.

**Section:**

**Explanation:**

**QUESTION 94**

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

**Correct Answer: A**

**Section:**

**QUESTION 95**

HOTSPOT

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**

Creating a text transcript of a voice recording is an example of

- a computer vision workload.
- a knowledge mining workload.
- a natural language processing (NLP) workload.
- an anomaly detection workload.

Answer selections:

**Answer Area:**

**Answer Area**

Creating a text transcript of a voice recording is an example of

- a computer vision workload.
- a knowledge mining workload.
- a natural language processing (NLP) workload.
- an anomaly detection workload.

Answer selections:

**Section:**

**Explanation:**

**QUESTION 96**

You need to create a customer support solution to help customers access information. The solution must support email, phone, and live chat channels. Which type of AI solution should you use?

- A. natural language processing (NLP)
- B. computer vision
- C. machine learning
- D. chatbot

**Correct Answer: D**

**Section:**

**QUESTION 97**

You are building a chatbot that will use natural language processing (NLP) to perform the following actions based on the text input of a user:

- Accept customer orders.
- Retrieve support documents.
- Retrieve order status updates.

Which type of NLP should you use?

- A. sentiment analysis
- B. translation
- C. language modeling
- D. named entity recognition

**Correct Answer: D**

**Section:**



**QUESTION 98**

DRAG DROP

Match the Azure Cognitive Services service to the appropriate actions.

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

**Select and Place:**

Azure Cognitive Services	Answer Area
Language service	Service Convert spoken requests into text.
Speech	Service Identify the intent of a user's requests.
Translator	Service Apply intent to entities and utterances.

**Correct Answer:**

Azure Cognitive Services	Answer Area
Translator	Convert spoken requests into text.
Speech	Identify the intent of a user's requests.
Language service	Apply intent to entities and utterances.

**Section:**

Explanation:

QUESTION 99

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A historian can use  to digitize newspaper articles.

facial analysis
image classification
object detection
optical character recognition (OCR)

Answer Area:

Answer Area

A historian can use  to digitize newspaper articles.

facial analysis
image classification
object detection
optical character recognition (OCR)



Section:

Explanation:

QUESTION 100

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  type of machine learning.

classification
regression
clustering

Answer Area:

Answer Area

A banking system that predicts whether a loan will be repaid is an example of the \_\_\_\_\_ type of machine learning.

- classification
- regression
- clustering

Section:

Explanation:

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

QUESTION 101

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:

Answer Area

Statements	Yes	No
Labelling is the process of tagging training data with known values.	<input checked="" type="radio"/>	<input type="radio"/>
You should evaluate a model by using the same data used to train the model.	<input type="radio"/>	<input checked="" type="radio"/>
Accuracy is always the primary metric used to measure a model's performance.	<input type="radio"/>	<input checked="" type="radio"/>

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 102

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 103

HOTSPOT

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

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.
- conversational AI.
- 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.
- conversational AI.
- forecasting.

**Section:**

**Explanation:**

With Microsoft's Conversational AI tools developers can build, connect, deploy, and manage intelligent bots that naturally interact with their users on a website, app, Cortana, Microsoft Teams, Skype, Facebook Messenger, Slack, and more.

Reference:

<https://azure.microsoft.com/en-in/blog/microsoft-conversational-ai-tools-enable-developers-to-build-connect-and-manage-intelligent-bots>

**QUESTION 104**

**HOTSPOT**

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

**Hot Area:**





Answer Area



Statements

Yes

No

You can use QnA Maker to query an Azure SQL database.

You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.

The QnA Maker service can determine the intent of a user utterance.



Answer Area:

Answer Area



Statements

Yes

No

You can use QnA Maker to query an Azure SQL database.

You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.

The QnA Maker service can determine the intent of a user utterance.



Section:

Explanation:

Reference:

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

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/choose-natural-language-processing-service>

QUESTION 105

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

  
Statements

Yes      No

You can communicate with a bot by using email.  Yes  No

You can communicate with a bot by using Microsoft Teams.  Yes  No

You can communicate with a bot by using a webchat interface.  Yes  No

Answer Area:

Answer Area

  
Statements

Yes      No

You can communicate with a bot by using email.  Yes  No

You can communicate with a bot by using Microsoft Teams.  Yes  No

You can communicate with a bot by using a webchat interface.  Yes  No

Section:

Explanation:

Reference:  
<https://docs.microsoft.com/en-us/azure/bot-service/bot-service-manage-channels?view=azure-bot-service-4.0>

QUESTION 106

DRAG DROP

Match the types of AI workloads to the appropriate scenarios.

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

NOTE: Each correct selection is worth one point.

Select and Place:

**Workloads Types**

Anomaly detection

Computer vision

Conversational AI

Knowledge mining

Natural language processing

**Answer Area**

Workload Type An automated chat to answer questions about refunds and exchange

Workload Type Determining whether a photo contains a person

Workload Type Determining whether a review is positive or negative



**Correct Answer:**

**Workloads Types**

Anomaly detection

Computer vision

Conversational AI

Knowledge mining

Natural language processing

**Answer Area**

Conversational AI An automated chat to answer questions about refunds and exchange

Computer vision Determining whether a photo contains a person

Natural language processing Determining whether a review is positive or negative



**Section:**

**Explanation:**

Box 3: Natural language processing

Natural language processing (NLP) is used for tasks such as sentiment analysis, topic detection, language detection, key phrase extraction, and document categorization.

Reference:

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

**QUESTION 107**

DRAG DROP

Match the Microsoft guiding principles for responsible AI to the appropriate descriptions.

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

NOTE: Each correct selection is worth one point.

**Select and Place:**

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

Correct Answer:

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

Section:

Explanation:

Box 1: Reliability and safety

To build trust, it's critical that AI systems operate reliably, safely, and consistently under normal circumstances and in unexpected conditions. These systems should be able to operate as they were originally designed, respond safely to unanticipated conditions, and resist harmful manipulation.

Box 2: Accountability

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

Box 3: Privacy and security

As AI becomes more prevalent, protecting privacy and securing important personal and business information is becoming more critical and complex. With AI, privacy and data security issues require especially close attention because access to data is essential for AI systems to make accurate and informed predictions and decisions about people. AI systems must comply with privacy laws that require transparency about the collection, use, and storage of data and mandate that consumers have appropriate controls to choose how their data is used

Reference:

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

#### QUESTION 108

You have an Azure Machine Learning model that uses clinical data to predict whether a patient has a disease.

You clean and transform the clinical data.

You need to ensure that the accuracy of the model can be proven.

What should you do next?

- A. Train the model by using the clinical data.
- B. Split the clinical data into Two datasets.
- C. Train the model by using automated machine learning (automated ML).
- D. Validate the model by using the clinical data.

**Correct Answer: D**

**Section:**

**QUESTION 109**

Which Azure AI Language feature can be used to retrieve data, such as dates and people's names, from social media posts?

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

**Correct Answer: D**

**Section:**

**QUESTION 110**

HOTSPOT

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**



As part of the Microsoft responsible AI principles, customers must

obtain approval based on their intended usage	▼
commit to a minimum level of expenditure	
obtain approval based on their intended usage	
pay an upfront fee	
provide their credit card details	

before they can use Azure OpenAI.

**Answer Area:**

**Answer Area**

As part of the Microsoft responsible AI principles, customers must

obtain approval based on their intended usage	▼
commit to a minimum level of expenditure	
obtain approval based on their intended usage	
pay an upfront fee	
provide their credit card details	

before they can use Azure OpenAI.

**Section:**

**Explanation:**