



CompTIA

# CertyIQ

## Premium exam material

---

Get certification quickly with the CertyIQ Premium exam material.

Everything you need to prepare, learn & pass your certification exam easily. Lifetime free updates

First attempt guaranteed success.

<https://www.CertyIQ.com>

# About CertyIQ

We here at CertyIQ eventually got enough of the industry's greedy exam paid for. Our team of IT professionals comes with years of experience in the IT industry Prior to training CertyIQ we worked in test areas where we observed the horrors of the paywall exam preparation system.

The misuse of the preparation system has left our team disillusioned. And for that reason, we decided it was time to make a difference. We had to make In this way, CertyIQ was created to provide quality materials without stealing from everyday people who are trying to make a living.

## Doubt Support

We have developed a very scalable solution using which we are able to solve 400+ doubts every single day with an average rating of 4.8 out of 5.

<https://www.certyiq.com>

[Mail us on - certyiqofficial@gmail.com](mailto:certyiqofficial@gmail.com)



### Lifetime Free Updates

We provide lifetime free updates to our customers. To make life easier for our valued customers and fulfill their needs



### Free Exam PDF

You are sure to pass the exam completely free of charge



### Money Back Guarantee

We Provide 100% money back guarantee to our customer in case of any failure

John

October 19, 2022



Thanks you so much for your help. I scored 972 in my exam today. More than 90% were from your PDFs!

Dana

September 04, 2022



Thanks a lot for this updated AZ-900 Q&A. I just passed my exam and got 974, I followed both of your Az-900 videos and the 6 PDF, the PDFs are very much valid, all answers are correct. Could you please create a similar video/PDF for DP900, your content/PDF's is really awesome. The team did a really good job. Thank You 😊.

Ahamed Shibly

2 months ago



Customer support is really fast and helpful, I just finished my exam and this video along with the 6 PDF helped me pass! Definitely recommend getting the PDFs. Thank you!

October 22, 2022



Passed my exam today with 891 marks. Out of 52 questions, 51 were from certyiq PDFs including Contoso case study. Thank You certyiq team!

Henry Rome

2 months ago



These questions are real and 100 % valid. Thank you so much for your efforts, also your 4 PDFs are awesome, I passed the DP900 exam on 1 Sept. With 968 marks. Thanks a lot, buddy!

Esmaria

2 months ago



Simple easy to understand explanations. To anyone out there wanting to write AZ900, I highly recommend 6 PDF's. Thank you so much, appreciate all your hard work in having such great content. Passed my exam Today - 3 September with 942 score.

# Microsoft

(AI-900)

Microsoft Azure AI Fundamentals

Total: **246 Questions**

Link: <https://certyiq.com/papers/microsoft/ai-900>

### Question: 1

A company employs a team of customer service agents to provide telephone and email support to customers. The company develops a webchat bot to provide automated answers to common customer queries. Which business benefit should the company expect as a result of creating the webchat bot solution?

- A. increased sales
- B. a reduced workload for the customer service agents
- C. improved product reliability

**Answer: B**

#### Explanation:

Correct answer is B: a reduced workload for the customer service agents.

A webchat bot is designed to automate responses to common customer queries, which typically account for a significant portion of customer service interactions. By handling these straightforward queries, the bot reduces the number of cases that require manual intervention by human customer service agents. This results in a reduced workload for the customer service team, allowing them to focus on more complex or critical issues.

Why not the other options?

A. Increased sales:

While a webchat bot can potentially assist customers and provide information that might indirectly contribute to sales, its primary function is to handle queries, not directly drive sales. Sales increases would depend on other factors, such as product offerings, marketing, and customer satisfaction.

C. Improved product reliability:

Product reliability is typically enhanced through better design, testing, and manufacturing processes. A webchat bot does not directly impact the reliability of the products but can help resolve customer issues or guide them in product usage.

Thus, the main business benefit of introducing a webchat bot is streamlining customer service operations, leading to a reduced workload for customer service agents.

### Question: 2

For a machine learning progress, how should you split data for training and evaluation?

- A. Use features for training and labels for evaluation.
- B. Randomly split the data into rows for training and rows for evaluation.
- C. Use labels for training and features for evaluation.
- D. Randomly split the data into columns for training and columns for evaluation.

**Answer: B**

#### Explanation:

You split rows not columns:

The Split Data module is particularly useful when you need to separate data into training and testing sets. Use the Split Rows option if you want to divide the data into two parts. You can specify the percentage of data to



put in each split, but by default, the data is divided 50-50. You can also randomize the selection of rows in each group, and use stratified sampling.


Reference:

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

Question: 3

CertyIQ

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	

There are **[answer choice]** false negatives.

	▼
5	
11	
1,033	
13,951	

Answer:

## Answer Area

There are **[answer choice]** correctly predicted positives.

	▼
5	
11	
1,033	
13,951	

There are **[answer choice]** false negatives.

	▼
5	
11	
1,033	
13,951	

Explanation:

Box 1: 11 -

	Predicted	
	Positive	Negative
Actual True	<b>TP</b>	<b>FN</b>
Actual False	<b>FP</b>	<b>TN</b>

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: 4

CertyIQ

You build a machine learning model by using the automated machine learning user interface (UI). You need to ensure that the model meets the Microsoft transparency principle for responsible AI. What should you do?

- A. Set Validation type to Auto.
- B. Enable Explain best model.
- C. Set Primary metric to accuracy.
- D. Set Max concurrent iterations to 0.

**Answer: B**

**Explanation:**

Model Explain Ability.

Most businesses run on trust and being able to open the ML black box helps build transparency and trust. In heavily regulated industries like healthcare and banking, it is critical to comply with regulations and best practices. One key aspect of this is understanding the relationship between input variables (features) and model output. Knowing both the magnitude and direction of the impact each feature (feature importance) has on the predicted value helps better understand and explain the model. With model explain ability, we enable you to understand feature importance as part of automated ML runs.

Reference:

<https://azure.microsoft.com/en-us/blog/new-automated-machine-learning-capabilities-in-azure-machine-learning-service/>

**Question: 5**

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Statements	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:

**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"/>

**Explanation:**

A - Regression

B - Anomaly

C - Classification

Anomaly detection encompasses many important tasks in machine learning:

Identifying transactions that are potentially fraudulent.

Learning patterns that indicate that a network intrusion has occurred.

Finding abnormal clusters of patients.

Checking values entered into a system.

Reference:

**Question: 6**

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:**

**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

**Explanation:**

Reliability and safety:

AI systems need to be reliable and safe in order to be trusted. It is important for a system to perform as it was originally designed and for it to respond safely to new situations. Its inherent resilience should resist intended or unintended manipulation. Rigorous testing and validation should be established for operating conditions to ensure that the system responds safely to edge cases, and A/B testing and champion/challenger methods should be integrated into the evaluation process.

An AI system's performance can degrade over time, so a robust monitoring and model tracking process needs to be established to reactively and proactively measure the model's performance and retrain it, as necessary, to modernize it.

Reference:

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

## Question: 7

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

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

### 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.

### What is Natural Language Processing?

Natural language processing (NLP) is the area of AI that deals with creating software that understands written and spoken language.

NLP enables you to create software that can:

-Analyze text documents to extract key phrases and recognize entities (such as places, dates, or people). ie Text Analytics service

-Perform sentiment analysis to determine how positive or negative the language used in a document is. ie Text Analytics service

-Interpret spoken language, and synthesize speech responses. ie Speech service(speech to text and text to



speech)

-Automatically translate spoken or written phrases between languages. ie Text service(for text to text translation)/Speech service(for speech to text/speech translation)

Interpret commands and determine appropriate actions. ie Language Understanding(LUIS) service

Reference:

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

### Question: 8

CertyIQ

You are designing an AI system that empowers everyone, including people who have hearing, visual, and other impairments.

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

- A. fairness
- B. inclusiveness
- C. reliability and safety
- D. accountability

**Answer: B**

#### Explanation:

Inclusiveness: At Microsoft, we firmly believe everyone should benefit from intelligent technology, meaning it must incorporate and address a broad range of human needs and experiences. For the 1 billion people with disabilities around the world, AI technologies can be a game-changer.

Reference:

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

### Question: 9

CertyIQ

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

Accountability

Fairness

Inclusiveness

Privacy and security

Reliability and safety

#### Answer Area

Principle

Ensure that AI systems operate as they were originally designed, respond to unanticipated conditions, and resist harmful manipulation.

Principle

Implementing processes to ensure that decisions made by AI systems can be overridden by humans.

Principle

Provide consumers with information and controls over the collection, use, and storage of their data.

Answer:

Principles	Answer Area
Accountability	Reliability and safety
Fairness	
Inclusiveness	Accountability
Privacy and security	Privacy and security
Reliability and safety	

**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: 10**

CertyIQ

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:



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

**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: 11**

CertyIQ

You are building an AI system.

Which task should you include to ensure that the service meets the Microsoft transparency principle for responsible AI?

- A. Ensure that all visuals have an associated text that can be read by a screen reader.
- B. Enable autoscaling to ensure that a service scales based on demand.
- C. Provide documentation to help developers debug code.
- D. Ensure that a training dataset is representative of the population.

**Answer: C**

**Explanation:**

The Microsoft transparency principle for responsible AI emphasizes that AI systems should be understandable and explainable to users, developers, and other stakeholders. Transparency involves providing clear and comprehensive information about how an AI system works, its capabilities, and its limitations. This helps users trust the system and developers maintain and debug it effectively.

By providing documentation to help developers debug code, the AI system ensures transparency in its development and maintenance processes, aligning with Microsoft's responsible AI guidelines.

Why not the other options?

A. Ensure that all visuals have an associated text that can be read by a screen reader:

This aligns with accessibility principles, ensuring that the system is usable by people with disabilities. However, it does not directly address transparency.

B. Enable autoscaling to ensure that a service scales based on demand:

Autoscaling is related to performance and availability, not transparency. It ensures the system can handle varying loads but does not make the AI system more explainable or understandable.

D. Ensure that a training dataset is representative of the population:

This aligns with fairness and inclusivity principles to avoid bias but does not directly ensure transparency.

Thus, the task that best aligns with the transparency principle is providing documentation to help developers debug code.

Reference:

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

### Question: 12

CertyIQ

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.

Answer:

#### Workload Types

Anomaly detection

Computer vision

Machine Learning (Regression)

Natural language processing

#### 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.

Explanation:

Keywords:

Computer vision = identify (object) letters

NLP = sentiment

Anomaly Detection = fraud

Machine Learning (regression) = predict

Reference:

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

### Question: 13

CertyIQ

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

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

**Answer: C**

**Explanation:**

C - Inclusiveness. No one is left out (disabled, gender, ethnicity, LGBTQIA+ etc etc)

Reference:

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

#### Question: 14

**CertyIQ**

What are three Microsoft guiding principles for responsible AI? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. knowledgeability
- B. decisiveness
- C. inclusiveness
- D. fairness
- E. opinionatedness
- F. reliability and safety

**Answer: CDF**

**Explanation:**

The six guiding principles are:

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

Reference:

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

**Question: 15**

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:**

**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.

**Explanation:**

Object detection is correct. Semantic segmentation can seem tempting at first but that is more about classifying individual pixels based on their objects

Reference:

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

**Question: 16**

## 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:

### Answer Area

▼

Feature engineering
Feature selection
Model evaluation
Model training

is used to generate additional features.

**Explanation:**

**Answer is Feature Engineering.**

Explanation - feature engineering is applied first to generate additional features, and then feature selection is done to eliminate irrelevant, redundant, or highly correlated features.

Reference:

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

## Question: 17

CertyIQ

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

**Answer: B**

**Explanation:**

Face detect can be requested to detect also glasses attribute

Reference:

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

### Question: 18

CertyIQ

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

**Answer: A**

**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: 19

CertyIQ

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.	<input type="radio"/>	<input type="radio"/>
A triage bot that prioritizes insurance claims based on injuries is an example of the Microsoft reliability and safety principle for responsible AI.	<input type="radio"/>	<input type="radio"/>
An AI solution that is offered at different prices for different sales territories is an example of the Microsoft inclusiveness principle for responsible AI.	<input type="radio"/>	<input type="radio"/>

Answer:

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.	<input checked="" type="radio"/>	<input type="radio"/>
A triage bot that prioritizes insurance claims based on injuries is an example of the Microsoft reliability and safety principle for responsible AI.	<input type="radio"/>	<input checked="" type="radio"/>
An AI solution that is offered at different prices for different sales territories is an example of the Microsoft inclusiveness principle for responsible AI.	<input type="radio"/>	<input checked="" type="radio"/>

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>

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**

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

**Answer:**

**Principles**

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

**Answer Area**

- Fairness
- Privacy and security
- Transparency

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

**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>

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

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

**Question: 21**

CertyIQ

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

Answer:

### Processes

data encryption

model retraining



### Answer Area

data preparation

model training

model evaluation

**Explanation:**

Data Prep

Train Model

Evaluate Model

Reference:

### Question: 22

CertyIQ

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

**Answer: BC**

**Explanation:**

B ensures reliability and safety principle and C ensures privacy and security principle of AI.

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: 23

CertyIQ

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:**

## Answer Area

According to Microsoft's 

	▼
accountability	
fairness	
inclusiveness	
transparency	

 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.

**Explanation:**

Reference:

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

## Question: 24

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

## Answer Area

According to Microsoft's 

	▼
accountability	
fairness	
inclusiveness	
transparency	

 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.

**Answer:**

## Answer Area

According to Microsoft's 

	▼
accountability	
fairness	
inclusiveness	
transparency	

 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.

**Explanation:**

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.

Reference:

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

## Question: 25

CertyIQ

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

Knowledge mining

Natural language processing

### Answer Area

Workload type

An automated chatbot to answer questions about refunds and exchanges

Workload type

Determining whether a photo contains a person

Workload type

Determining whether a review is positive or negative

Answer:

### Workload Types

Anomaly detection

Computer vision

Knowledge mining

Natural language processing

### Answer Area

Knowledge mining

An automated chatbot to answer questions about refunds and exchanges

Computer vision

Determining whether a photo contains a person

Natural language processing

Determining whether a review is positive or negative

### Explanation:

Box 1: Knowledge mining -

You can use Azure Cognitive Search's knowledge mining results and populate your knowledge base of your chatbot.

Box 2: Computer vision -

Box 3: Natural language processing

Natural language processing (NLP) is used for tasks such as sentiment analysis.

Reference:

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

## Question: 26

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

Feature engineering

Feature selection

Model deployment

Model evaluation

Model training

### Answer Area

Task

Examining the values of a confusion matrix

Task

Splitting a date into month, day, and year fields

Task

Picking temperature and pressure to train a weather model

### Answer:

#### Learning Types

Feature engineering

Feature selection

Model deployment

Model evaluation

Model training

#### Answer Area

Model evaluation

Examining the values of a confusion matrix

Feature engineering

Splitting a date into month, day, and year fields

Feature selection

Picking temperature and pressure to train a weather model

### 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: 27

CertyIQ

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:

#### Answer Area

Data values that influence the prediction of a model are called

	▼
dependant variables.	
features.	
identifiers.	
labels.	

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: 28

CertyIQ

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

**Answer: B**

**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: 29**

**CertyIQ**

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

**Answer: B**

**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:

### Question: 30

CertyIQ

You have a dataset that contains information about taxi journeys that occurred during a given period. You need to train a model to predict the fare of a taxi journey. What should you use as a feature?

- A. the number of taxi journeys in the dataset
- B. the trip distance of individual taxi journeys
- C. the fare of individual taxi journeys
- D. the trip ID of individual taxi journeys

#### Answer: B

#### Explanation:

The label is the column you want to predict. The identified Features are the inputs you give the model to predict the Label.

Example:

The provided data set contains the following columns:

vendor\_id: The ID of the taxi vendor is a feature.

rate\_code: The rate type of the taxi trip is a feature.

passenger\_count: The number of passengers on the trip is a feature. trip\_time\_in\_secs: The amount of time the trip took. You want to predict the fare of the trip before the trip is completed. At that moment, you don't know how long the trip would take. Thus, the trip time is not a feature and you'll exclude this column from the model. trip\_distance: The distance of the trip is a feature. payment\_type: The payment method (cash or credit card) is a feature. fare\_amount: The total taxi fare paid is the label.

Reference:

<https://docs.microsoft.com/en-us/dotnet/machine-learning/tutorials/predict-prices>

### Question: 31

CertyIQ

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

#### Answer: B

#### 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: 32**

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

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

Answer:

**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"/>

**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: 33

CertyIQ

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:

### 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

Explanation:

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

**classification** as the prediction whether the loan will be repaid or not -- whether means class - categorical answer as yes or no - hence classification

### Question: 34

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:

### 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"/>

### 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: 35

CertyIQ

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

**Answer: A**

#### Explanation:

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

Reference:

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

### Question: 36

CertyIQ

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:**

#### 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

#### 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: 37

CertyIQ

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

**Answer: CD**

#### 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: 38

CertyIQ

HOTSPOT -

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

Hot Area:

#### Answer Area

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

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



Answer:

### Answer Area

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

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

### Explanation:

For Prod - AKS and for Dev/Test - use Azure Container Service

To perform real-time inferencing, you must deploy a pipeline as a real-time endpoint.

Real-time endpoints must be deployed to an Azure Kubernetes Service cluster.

Reference:

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

### Question: 39

CertyIQ

HOTSPOT -

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

Hot Area:

### Answer Area

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

- classification.
- clustering.
- regression.

Answer:

## Answer Area

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

	▼
classification.	
clustering.	
regression.	

### Explanation:

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

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

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

Incorrect Answers:

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

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

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

Reference:

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

## Question: 40

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

Statements	Yes	No
Azure Machine Learning designer provides a drag-and-drop visual canvas to build, test, and deploy machine learning models.	<input type="radio"/>	<input type="radio"/>
Azure Machine Learning designer enables you to save your progress as a pipeline draft.	<input type="radio"/>	<input type="radio"/>
Azure Machine Learning designer enables you to include custom JavaScript functions.	<input type="radio"/>	<input type="radio"/>

Answer:

### Answer Area

Statements	Yes	No
Azure Machine Learning designer provides a drag-and-drop visual canvas to build, test, and deploy machine learning models.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Machine Learning designer enables you to save your progress as a pipeline draft.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Machine Learning designer enables you to include custom JavaScript functions.	<input type="radio"/>	<input checked="" type="radio"/>

#### Explanation:

Box 1: Yes -

Azure Machine Learning designer lets you visually connect datasets and modules on an interactive canvas to create machine learning models.

Box 2: Yes -

With the designer you can connect the modules to create a pipeline draft.

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

Box 3: No -

Reference:

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

#### Question: 41

CertyIQ

HOTSPOT -

You have the following dataset.



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

You plan to use the dataset to train a model that will predict the house price categories of houses. What are Household Income and House Price Category? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

Household Income:

House Price Category:

Answer:

## Answer Area

Household Income:

House Price Category:

**Explanation:**

Feature = Input. Label = Output.

Reference:

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

HOTSPOT -

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

Hot Area:

### Answer Area

Azure Machine Learning designer lets you create machine learning models by

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

Answer:

### Answer Area

Azure Machine Learning designer lets you create machine learning models by

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

Explanation:

Reference:

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

### Question: 43

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

Statements	Yes	No
Automated machine learning provides you with the ability to include custom Python scripts in a training pipeline.	<input type="radio"/>	<input type="radio"/>
Automated machine learning implements machine learning solutions without the need for programming experience.	<input type="radio"/>	<input type="radio"/>
Automated machine learning provides you with the ability to visually connect datasets and modules on an interactive canvas.	<input type="radio"/>	<input type="radio"/>

Answer:

### Answer Area

Statements	Yes	No
Automated machine learning provides you with the ability to include custom Python scripts in a training pipeline.	<input type="radio"/>	<input checked="" type="radio"/>
Automated machine learning implements machine learning solutions without the need for programming experience.	<input checked="" type="radio"/>	<input type="radio"/>
Automated machine learning provides you with the ability to visually connect datasets and modules on an interactive canvas.	<input type="radio"/>	<input checked="" type="radio"/>

### Explanation:

No - Automated machine learning only requires you to choose between Python SDK and studio web experience.

Yes - Automated machine learning is a no code solution.

No - This is done in the Azure Machine Learning studio web experience.

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

Reference:

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

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

## Question: 44

CertyIQ

A medical research project uses a large anonymized dataset of brain scan images that are categorized into

predefined brain haemorrhage types.

You need to use machine learning to support early detection of the different brain haemorrhage types in the images before the images are reviewed by a person.

This is an example of which type of machine learning?

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

**Answer: C**

**Explanation:**

This is classification. There are multiple reasons it's classification

- The training data is already tagged as with the correct type of hemorrhage
- Classification can be done for more than two classes (which people seem to not realize based on the comments)
- You do clustering on a group of inputs. For example, the scans of 10 people. You can't cluster a single input. Clearly you get a new scan of a new patient and you want to know what that scan shows, you don't have a group of scans to cluster.
- Clustering gives NO labels. You just get groups and don't know what the label is, but in this question it's very clear they want to know the label that belongs to the new scan -> classification

Reference:

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

### Question: 45

CertyIQ

When training a model, why should you randomly split the rows into separate subsets?

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

**Answer: C**

**Explanation:**

First for training then evaluation.

He is referring to the famous train\_test\_split that everyone uses to split the dataset into train and test sets.

You lose statistical power by estimating on a subset  $n$  like 50 percent of original  $N$ . That's the price you pay for splitting data. For example a normally independently distributed (n.i.d)  $x$  has an estimator of arithmetic mean  $\bar{x}$  whose variance inversely related to  $N$ . Bigger sample is better for accuracy of a simple arithmetic mean.

### Question: 46

CertyIQ

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.

**Answer: AC**

**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: 47**

**CertyIQ**

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

**Answer: AC**

**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: 48**

**CertyIQ**

You are building a tool that will process images from retail stores and identify the products of competitors. The solution will use a custom model. Which Azure Cognitive Services service should you use?



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

**Answer: A**

**Explanation:**

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

Reference:

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

**Question: 49**

**CertyIQ**

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:**



## 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"/>

### 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: 50

CertyIQ

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 model.	<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:

## Answer Area

Statements	Yes	No
A validation set includes the set of input examples that will be used to train a model.	<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"/>

### 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: 51

CertyIQ

What are two metrics that you can use to evaluate a regression model? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. coefficient of determination (R<sup>2</sup>)
- B. F1 score
- C. root mean squared error (RMSE)
- D. area under curve (AUC)
- E. balanced accuracy

**Answer: AC**

**Explanation:**

**Small Tip:** If its Regression model then remembers 'R' and the corresponding answer will be R<sup>2</sup> and RMSE.

A: R-squared (R<sup>2</sup>), or Coefficient of determination represents the predictive power of the model as a value between -inf and 1.00. 1.00 means there is a perfect fit, and the fit can be arbitrarily poor so the scores can be negative.

C: RMS-loss or Root Mean Squared Error (RMSE) (also called Root Mean Square Deviation, RMSD), measures the difference between values predicted by a model and the values observed from the environment that is

being modeled.

Incorrect Answers:

B: F1 score also known as balanced F-score or F-measure is used to evaluate a classification model.

D: aucROC or area under the curve (AUC) is used to evaluate a classification model.

Reference:

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

### Question: 52

CertyIQ

HOTSPOT -

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

Hot Area:

#### Answer Area

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

	▼
classification.	
clustering.	
regression.	

Answer:

#### Answer Area

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

	▼
classification.	
clustering.	
regression.	

**Explanation:**

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

Reference:

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

### Question: 53

CertyIQ

DRAG DROP -

You need to use Azure Machine Learning designer to build a model that will predict automobile prices. Which type of modules should you use to complete the model? To answer, drag the appropriate modules to the correct locations. Each module may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

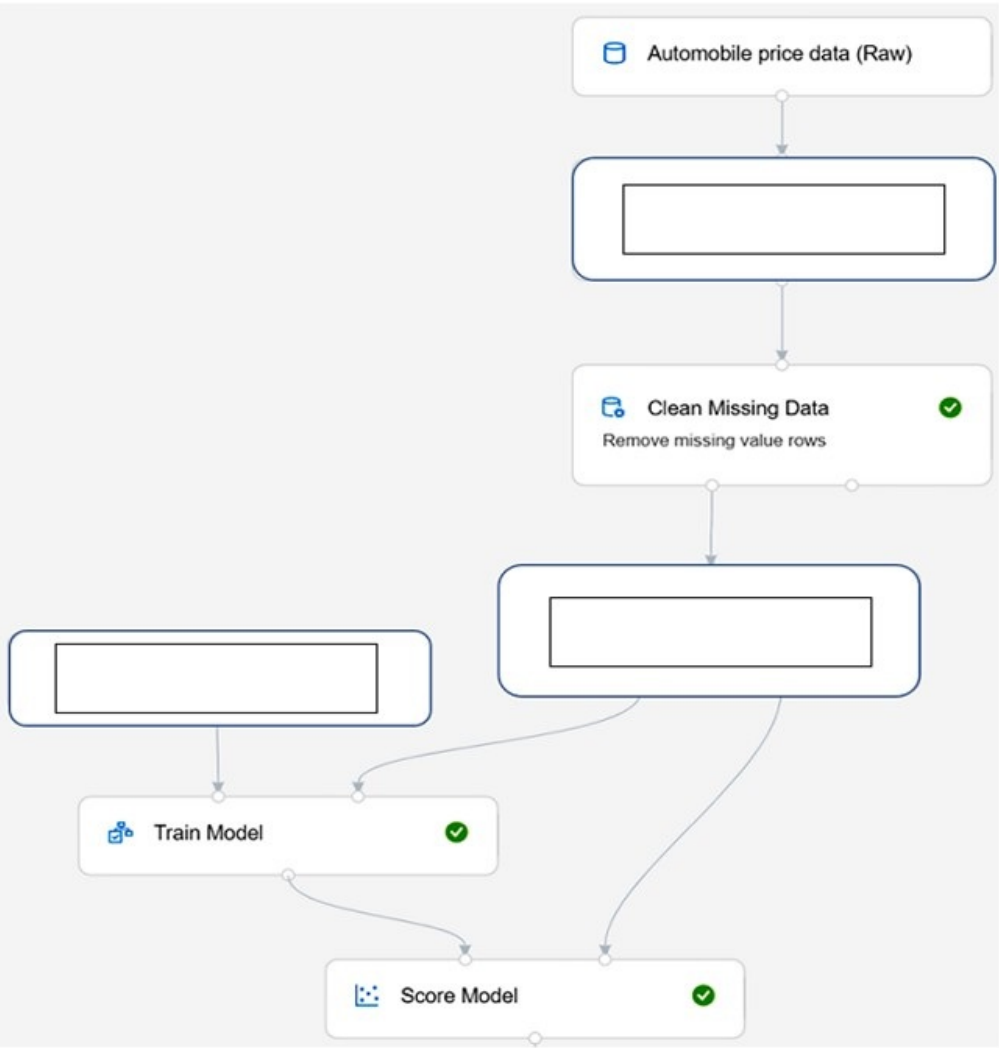
NOTE: Each correct selection is worth one point.

Select and Place:

Modules

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

Answer Area



Answer:

## Modules

Convert to CSV

K-Means Clustering

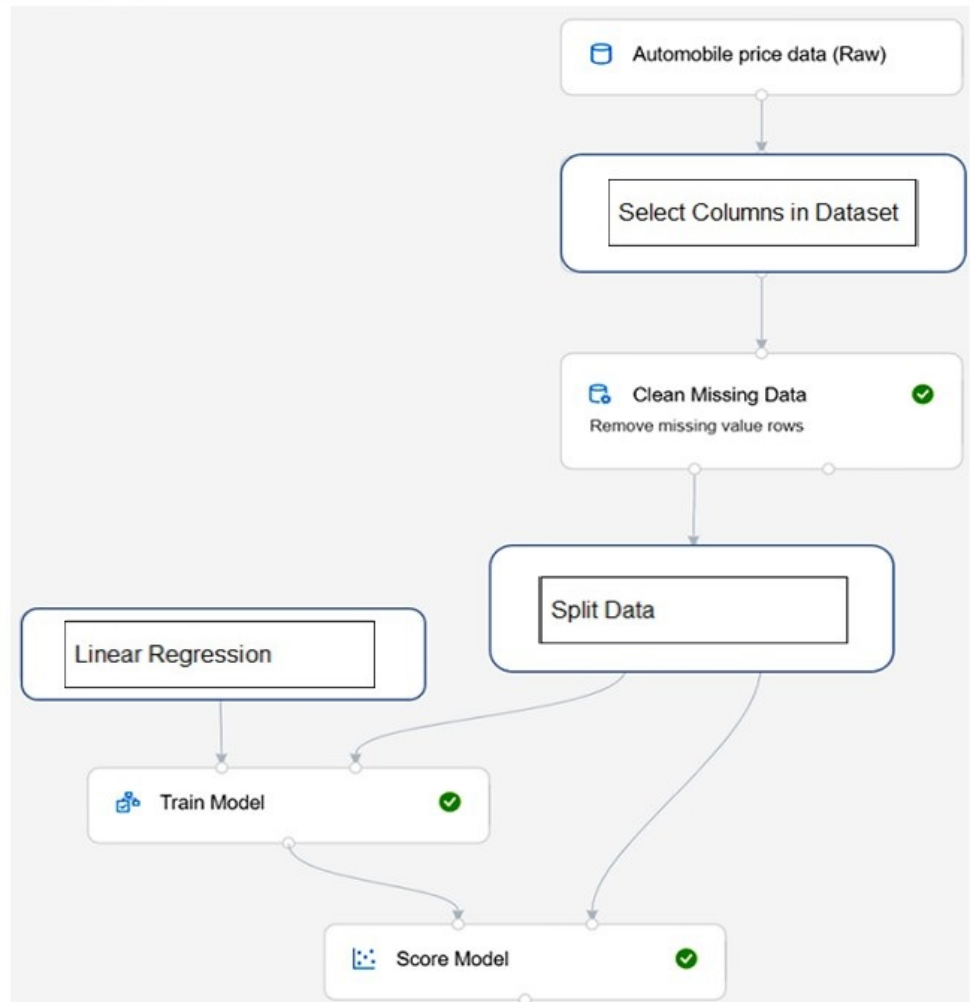
Linear Regression

Split Data

Select Columns in Dataset

Summarize Data

## Answer Area




### Explanation:

Box 1: Select Columns in Dataset

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

Example:

 Automobile price data (Raw)

 Select Columns in Dataset

Exclude normalized losses

 Clean Missing Data

Remove missing value rows

Box 2: Split data -

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

Box 3: Linear regression -

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

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/tutorial-designer-automobile-price-train-score>

#### Question: 54

CertyIQ

Which type of machine learning should you use to identify groups of people who have similar purchasing habits?

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

**Answer: C**

**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

Reference:

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

### Question: 55

CertyIQ

HOTSPOT -

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

Hot Area:

#### Answer Area

	▼
Classification	
Clustering	
Regression	

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

Answer:

#### Answer Area

	▼
Classification	
Clustering	
Regression	

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

**Explanation:**

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

Reference:

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

### Question: 56

CertyIQ

Which metric can you use to evaluate a classification model?

- A. true positive rate
- B. mean absolute error (MAE)
- C. coefficient of determination (R2)
- D. root mean squared error (RMSE)

**Answer: A**

**Explanation:**

What does a good model look like?

An ROC curve that approaches the top left corner with 100% true positive rate and 0% false positive rate will

be the best model. A random model would display as a flat line from the bottom left to the top right corner. Worse than random would dip below the  $y=x$  line.

MAE, RMSE and R2 are metrics for regression:

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

Reference:

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

### Question: 57

CertyIQ

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

**Answer: AD**

**Explanation:**

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

Azure Machine Learning designer lets you visually connect datasets and modules on an interactive canvas to create machine learning models.

Reference:

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

### Question: 58

CertyIQ

You need to create a training dataset and validation dataset from an existing dataset. Which module in the Azure Machine Learning designer should you use?

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

**Answer: C**

**Explanation:**

A common way of evaluating a model is to divide the data into a training and test set by using Split Data, and then validate the model on the training data.

Use the Split Data module to divide a dataset into two distinct sets.

The studio currently supports training/validation data splits

Reference:

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

## Question: 59

CertyIQ

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

Classification

Clustering

Regression

### Answer Area

Learning Type

Predict how many minutes late a flight will arrive based on the amount of snowfall at an airport.

Learning Type

Segment customers into different groups to support a marketing department.

Learning Type

Predict whether a student will complete a university course.

### Answer:

#### Learning Types

Classification

Clustering

Regression

#### Answer Area

Regression

Predict how many minutes late a flight will arrive based 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.

### 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: 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.

Box 3: Classification -

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

Reference:

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

**Question: 60**

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:****Answer Area**

	▼
Accuracy	
Confidence	
Root Mean Square Error	
Sentiment	

is the calculated probability of a correct image classification.

**Explanation:**

Confidence is the right answer.

..."The probability score of the object classification (which you can interpret as the confidence of the predicted class being correct)"...

Source: <https://docs.microsoft.com/en-us/learn/modules/detect-objects-images-custom-vision/1a-what-is-object-detection>

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/getting-started-build-a-classifier>**Question: 61**

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:

## 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.

Explanation:

Reference:

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

## Question: 62

CertyIQ

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:

## 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.	

Explanation:

Feature engineering is the correct answer.

"In Azure Machine Learning, data-scaling and normalization techniques are applied to make feature engineering easier. Collectively, these techniques and this feature engineering are called featurization in automated ML experiments."

Feature selection is only about selection.

Modifying features = Feature engineering

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-configure-auto-features>

Reference:

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

### Question: 63

CertyIQ

HOTSPOT -

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

Hot Area:

#### Answer Area

Assigning classes to images before training a classification model is an example of

	▼
evaluation.	
feature engineering	
hyperparameter tuning.	
labeling.	

Answer:

#### Answer Area

Assigning classes to images before training a classification model is an example of

	▼
evaluation.	
feature engineering	
hyperparameter tuning.	
labeling.	

Explanation:

Reference:

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

### Question: 64

CertyIQ

HOTSPOT -

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



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

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

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

Answer:

## Answer Area

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

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/component-reference/filter-based-feature-selection>

### Question: 65

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

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

Answer:

### Answer Area

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

Explanation:

Reference:

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

## Question: 66

CertyIQ

Which two actions are performed during the data ingestion and data preparation stage of an Azure Machine Learning process? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Calculate the accuracy of the model.
- B. Score test data by using the model.
- C. Combine multiple datasets.
- D. Use the model for real-time predictions.
- E. Remove records that have missing values.

Answer: CE

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/concept-data-ingestion> <https://docs.microsoft.com/en-us/azure/architecture/data-science-process/prepare-data>

### Question: 67

CertyIQ

You need to predict the animal population of an area.  
Which Azure Machine Learning type should you use?

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

**Answer: A**

**Explanation:**

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

Reference:

<https://docs.microsoft.com/en-us/learn/modules/create-regression-model-azure-machine-learning-designer/1-introduction>

### Question: 68

CertyIQ

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

NOTE: Each correct selection is worth one point.

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

**Answer: AB**

**Explanation:**

**Use Azure Machine Learning designer for customizing using Python and R code.**

"Build and train machine learning models with state-of-the art machine learning and deep learning algorithms, including those for computer vision, text analytics, recommendations, and anomaly detection. Drag and drop modules for no-code models or customize using Python and R code." as per the link given:

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

Reference:

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

### Question: 69

CertyIQ

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
For a regression model, labels must be numeric.	<input type="radio"/>	<input type="radio"/>
For a clustering model, labels must be used.	<input type="radio"/>	<input type="radio"/>
For a classification model, labels must be numeric.	<input type="radio"/>	<input type="radio"/>

Answer:

### Answer Area

Statements	Yes	No
For a regression model, labels must be numeric.	<input checked="" type="radio"/>	<input type="radio"/>
For a clustering model, labels must be used.	<input type="radio"/>	<input checked="" type="radio"/>
For a classification model, labels must be numeric.	<input type="radio"/>	<input checked="" type="radio"/>

#### Explanation:

Box 1: Yes -

For regression problems, the label column must contain numeric data that represents the response variable. Ideally the numeric data represents a continuous scale.

Box 2: No -

K-Means Clustering -

Because the K-means algorithm is an unsupervised learning method, a label column is optional.

If your data includes a label, you can use the label values to guide selection of the clusters and optimize the model.

If your data has no label, the algorithm creates clusters representing possible categories, based solely on the data.

Box 3: No -

For classification problems, the label column must contain either categorical values or discrete values. Some examples might be a yes/no rating, a disease classification code or name, or an income group. If you pick a noncategorical column, the component will return an error during training.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/component-reference/train-model> <https://docs.microsoft.com/en-us/azure/machine-learning/component-reference/k-means-clustering>

Your company wants to build a recycling machine for bottles. The recycling machine must automatically identify bottles of the correct shape and reject all other items.  
Which type of AI workload should the company use?

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

**Answer: C**

**Explanation:**

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

Reference:

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

**Question: 71**

**CertyIQ**

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

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

**Answer:**

## Answer Area

Statements	Yes	No
When creating an object detection model in the Custom Vision service, you must choose a classification type of either <b>Multilabel</b> or <b>Multiclass</b> .	<input type="radio"/>	<input checked="" type="radio"/>
You can create an object detection model in the Custom Vision service to find the location of content within an image.	<input checked="" type="radio"/>	<input type="radio"/>
When creating an object detection model in the Custom Vision service, you can select from a set of predefined domains.	<input checked="" type="radio"/>	<input type="radio"/>

### Explanation:

Reference:

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

## Question: 72

CertyIQ

In which two scenarios can you use the Form Recognizer service? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

### Answer: AD

### Explanation:

The Form Recognizer service in Microsoft Azure is an AI-powered document analysis tool that helps extract structured data and key-value pairs from documents such as invoices, receipts, forms, and business cards.

Correct Scenarios:

A. Extract the invoice number from an invoice:

Reason: Form Recognizer is designed to analyze documents like invoices and extract specific fields, such as invoice numbers, dates, or total amounts. This is one of its primary use cases.

D. Identify the retailer from a receipt:

Reason: Form Recognizer can extract key information from receipts, such as the retailer's name, transaction dates, or itemized totals. This aligns with its capabilities.

Incorrect Scenarios:

B. Translate a form from French to English:

Reason: Translating text is not a feature of Form Recognizer. Translation tasks fall under services like Azure



Translator or similar language translation tools.

C. Find an image of a product in a catalog:

Reason: Form Recognizer processes text and data from documents but does not work with image-based object detection. For identifying images, you would use services like Computer Vision or Custom Vision.

Key Tip:

Use Form Recognizer for scenarios that involve extracting structured data from documents like invoices, receipts, or forms. Remember: it's about data extraction from documents, not translation or image recognition.

Reference:

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

### Question: 73

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

## Answer Area

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

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

Answer:

## Answer Area

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

▼

forecasting.

computer vision.

conversational AI.

anomaly detection.

### Explanation:

#### Reference:

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

## Question: 74

CertyIQ

### HOTSPOT -

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

You are tagging new photos of the employees.

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

Statements	Yes	No
The Face service can be used to perform facial recognition for employees	<input type="radio"/>	<input type="radio"/>
The Face service will be more accurate if you provide more sample photos of each employee from different angles.	<input type="radio"/>	<input type="radio"/>
If an employee is wearing sunglasses, the Face service will always fail to recognize the employee.	<input type="radio"/>	<input type="radio"/>

### Answer:

## Answer Area

Statements	Yes	No
The Face service can be used to perform facial recognition for employees	<input checked="" type="radio"/>	<input type="radio"/>
The Face service will be more accurate if you provide more sample photos of each employee from different angles.	<input checked="" type="radio"/>	<input type="radio"/>
If an employee is wearing sunglasses, the Face service will always fail to recognize the employee.	<input type="radio"/>	<input checked="" type="radio"/>

### Explanation:

Reference:

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

### Question: 75

CertyIQ

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)

**Answer: D**

### 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: 76

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

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

Answer:

### Answer Area

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

#### Explanation:

Box 1: Yes -

Custom Vision functionality can be divided into two features. Image classification applies one or more labels to an image. Object detection is similar, but it also returns the coordinates in the image where the applied label(s) can be found.

Box 2: Yes -

The Custom Vision service uses a machine learning algorithm to analyze images. You, the developer, submit groups of images that feature and lack the characteristics in question. You label the images yourself at the time of submission. Then, the algorithm trains to this data and calculates its own accuracy by testing itself on those same images.

Box 3: No -

Custom Vision service can be used only on graphic files.

Reference:

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

## Question: 77

CertyIQ

You are processing photos of runners in a race.

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

Which type of computer vision should you use?

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

**Answer: B**

**Explanation:**

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

Reference:

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

**Question: 78**

CertyIQ

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:

**Machine Learning Types**

- Facial detection
- Facial recognition
- Image classification
- Object detection
- Optical character recognition (OCR)
- Semantic segmentation

**Answer Area**

- Machine Learning Type Separate images of polar bears and brown bears.
- Machine Learning Type Determine the location of a bear in a photo.
- Machine Learning Type Determine which pixels in an image are part of a bear.

**Answer:**

**Machine Learning Types**

- Facial detection
- Facial recognition
- Image classification
- Object detection
- Optical character recognition (OCR)
- Semantic segmentation

**Answer Area**

- Image classification Separate images of polar bears and brown bears.
- Object detection Determine the location of a bear in a photo.
- Semantic segmentation Determine which pixels in an image are part of a bear.

**Explanation:**

Box 1: Image classification -

Image classification is a supervised learning problem: define a set of target classes (objects to identify in images), and train a model to recognize them using labeled example photos.

### Box 2: Object detection -

Object detection is a computer vision problem. While closely related to image classification, object detection performs image classification at a more granular scale. Object detection both locates and categorizes entities within images.

### Box 3: Semantic Segmentation -

Semantic segmentation achieves fine-grained inference by making dense predictions inferring labels for every pixel, so that each pixel is labeled with the class of its enclosing object or region.

Reference:

<https://developers.google.com/machine-learning/practica/image-classification> <https://docs.microsoft.com/en-us/dotnet/machine-learning/tutorials/object-detection-model-builder> <https://nanonets.com/blog/how-to-do-semantic-segmentation-using-deep-learning/>

### Question: 79

CertyIQ

You use drones to identify where weeds grow between rows of crops to send an instruction for the removal of the weeds.

This is an example of which type of computer vision?

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

**Answer: A**

#### Explanation:

Object detection is similar to tagging, but the API returns the bounding box coordinates for each tag applied. For example, if an image contains a dog, cat and person, the Detect operation will list those objects together with their coordinates in the image.

Incorrect Answers:

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

C: Scene segmentation determines when a scene changes in video based on visual cues. A scene depicts a single event and it's composed by a series of consecutive shots, which are semantically related.

Reference:

<https://docs.microsoft.com/en-us/ai-builder/object-detection-overview> <https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/overview-ocr> <https://docs.microsoft.com/en-us/azure/azure-video-analyzer/video-analyzer-for-media-docs/video-indexer-overview>

### Question: 80

CertyIQ

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

grouping

identification

similarity

verification

## Answer Area

Task

Do two images of a face belong to the same person?

Task

Does this person look like other people?

Task

Do all the faces belong together?

Task

Who is this person in this group of people?

## Answer:

### Tasks

grouping

identification

similarity

verification

### 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?

## 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.

~Face verification:

The Verify API does an authentication against two detected faces or from one detected face to one person object. Practically, it evaluates whether two faces belong to the same person.

~Person identification:

The Identify API is used to identify a detected face against a database of people (facial recognition search). This feature might be useful for automatic image tagging in photo management software. You create the database in advance, and you can edit it over time.

## Tips to remember: -

verification = same person?

similarity = look like?

grouping = belong together?

identification = who is this person?

Reference:

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

## Question: 81

CertyIQ

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

Facial recognition

Image classification

Object detection

Optical character recognition (OCR)

### Answer Area

Workload Type

Identify celebrities in images.

Workload Type

Extract movie title names from movie poster images.

Workload Type

Locate vehicles in images.

Answer:

### Workloads Types

Facial recognition

Image classification

Object detection

Optical character recognition (OCR)

### Answer Area

Facial recognition

Identify celebrities in images.

Optical character recognition (OCR)

Extract movie title names from movie poster images.

Object detection

Locate vehicles in images.

### 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: Object 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: 82

CertyIQ

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

**Answer: B**

**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: 83

CertyIQ

HOTSPOT -

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

Hot Area:

**Answer Area**

You can use the  service to train an object detection model by using your own images.

- |                 |
|-----------------|
| Computer Vision |
| Custom Vision   |
| Form Recognizer |
| Video Indexer   |

Answer:

### Answer Area

You can use the  service to train an object detection model by using your own images.

Computer Vision

Custom Vision

Form Recognizer

Video Indexer

### Explanation:

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

Note: The Custom Vision service uses a machine learning algorithm to apply labels to images. You, the developer, must submit groups of images that feature and lack the characteristics in question. You label the images yourself at the time of submission. Then the algorithm trains to this data and calculates its own accuracy by testing itself on those same images. Once the algorithm is trained, you can test, retrain, and eventually use it to classify new images according to the needs of your app. You can also export the model itself for offline use.

Incorrect Answers:

Computer Vision:

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

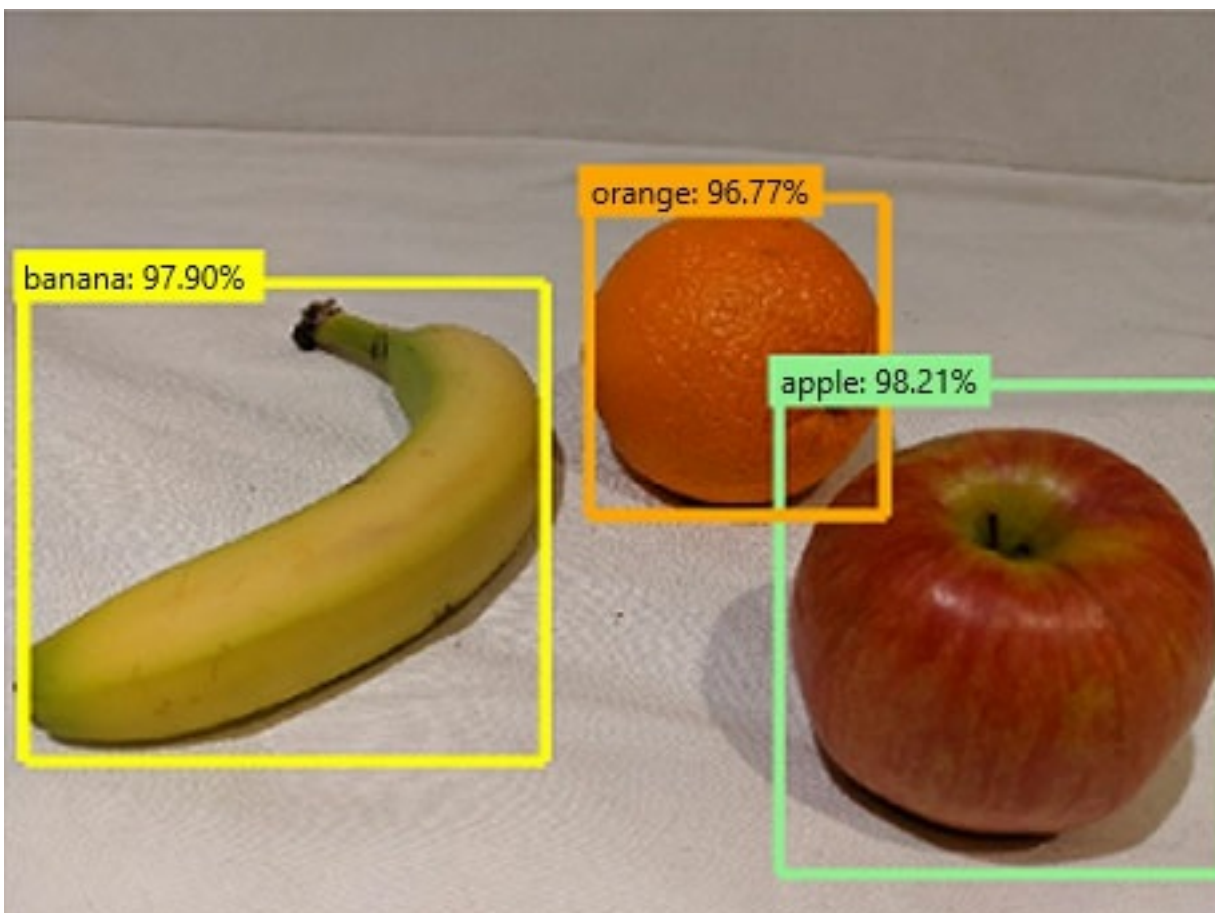
Reference:

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

### Question: 84

CertyIQ

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

**Answer: A**

**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: 85**

**CertyIQ**

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.

**Answer: BC**

**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: 86**

**CertyIQ**

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

**Answer: C**

**Explanation:**

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

Incorrect Answers:

- A: This is Regression.
- B: This is Clustering.
- D: This is Regression.

Reference:

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

**Question: 87**

**CertyIQ**

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.



**Answer: BC**

**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: 88**

**CertyIQ**

You need to build an image tagging solution for social media that tags images of your friends automatically. Which Azure Cognitive Services service should you use?

- A. Face
- B. Form Recognizer
- C. Text Analytics
- D. Computer Vision

**Answer: A**

**Explanation:**

More aligned towards face

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/face/overview> <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/howtodetectfacesinimage>

**Question: 89**

**CertyIQ**

In which two scenarios can you use the Form Recognizer service? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

**Answer: AC**

**Explanation:**

Correct Scenarios:

A. Identify the retailer from a receipt:

Why: Form Recognizer is designed to analyze receipts and extract key details such as the retailer's name,

address, transaction amounts, or dates. This is a common use case for the service.

C. Extract the invoice number from an invoice:

Why: Another key feature of Form Recognizer is its ability to process invoices and extract fields like invoice numbers, dates, amounts, or customer details. This aligns perfectly with its capabilities.

Incorrect Scenarios:

B. Translate from French to English:

Why Not: Translation is not a feature of Form Recognizer. Language translation tasks are handled by services like Azure Translator, not Form Recognizer.

D. Find images of products in a catalog:

Why Not: Form Recognizer processes text and structured data from documents, not images or product catalogs. Tasks involving images are better suited for services like Computer Vision or Custom Vision.

Key Tip:

Use Form Recognizer for extracting text-based data from documents like receipts, invoices, or forms. Focus on tasks related to data extraction and ignore tasks involving translation or image analysis.

Reference:

<https://docs.microsoft.com/en-us/azure/applied-ai-services/form-recognizer/overview?tabs=v2-1>

## Question: 90

CertyIQ

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

grouping

identification

similarity

verification

### Answer Area

Task

Do two images of a face belong to the same person?

Task

Does this person look like other people?

Task

Who is this person in this group of people?

Answer:

### Tasks

grouping

identification

similarity

verification

### Answer Area

verification

Do two images of a face belong to the same person?

similarity

Does this person look like other people?

identification

Who is this person in this group of people?

#### Explanation:

Box 1: verification -

Identity verification -

Modern enterprises and apps can use the Face identification and Face verification operations to verify that a user is who they claim to be.

Box 2: similarity -

The Find Similar operation does face matching between a target face and a set of candidate faces, finding a smaller set of faces that look similar to the target face.

This is useful for doing a face search by image.

The service supports two working modes, matchPerson and matchFace. The matchPerson mode returns similar faces after filtering for the same person by using the Verify API. The matchFace mode ignores the same-person filter. It returns a list of similar candidate faces that may or may not belong to the same person.

Box 3: identification -

Face identification can address "one-to-many" matching of one face in an image to a set of faces in a secure repository. Match candidates are returned based on how closely their face data matches the query face. This scenario is used in granting building or airport access to a certain group of people or verifying the user of a device.

Reference:

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

### Question: 91

CertyIQ

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

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

**Answer: D**

#### Explanation:

Describe images with human-readable language

Computer Vision can analyze an image and generate a human-readable phrase that describes its contents.

The algorithm returns several descriptions based on different visual features, and each description is given a confidence score. The final output is a list of descriptions ordered from highest to lowest confidence.

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

Reference:

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

### Question: 92

CertyIQ

Which service should you use to extract text, key/value pairs, and table data automatically from scanned documents?

- A. Custom Vision
- B. Face
- C. Form Recognizer
- D. Language

**Answer: C**

#### Explanation:

Form Recognizer applies advanced machine learning to accurately extract text, key-value pairs, tables, and structures from documents.

Reference:

<https://azure.microsoft.com/en-us/services/form-recognizer/>

### Question: 93

CertyIQ

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

#### Answer Area

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

extracts text from handwritten documents.

**Answer:**

#### Answer Area

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

extracts text from handwritten documents.

#### Explanation:

Handwriting OCR (optical character recognition) is the process of automatically extracting handwritten information from paper, scans and other low-quality digital documents.

Reference:

<https://vidado.ai/handwriting-ocr>

### Question: 94

CertyIQ

You are developing a solution that uses the Text Analytics service. You need to identify the main talking points in a collection of documents. Which type of natural language processing should you use?

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

**Answer: B**

**Explanation:**

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

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

Reference:

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

### Question: 95

CertyIQ

In which two scenarios can you use speech recognition? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

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

**Answer: BD**

**Explanation:**

B, D is correct as LUIS interpret meaning of text whereas text analytics is for sentiment or key Phrase extraction.

Reference:

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

### Question: 96

CertyIQ

HOTSPOT -

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

Hot Area:

## Answer Area

While presenting at a conference, your session is transcribed into subtitles for the audience. This is an example of

- sentiment analysis.
- speech recognition.
- speech synthesis.
- translation.

Answer:

## Answer Area

While presenting at a conference, your session is transcribed into subtitles for the audience. This is an example of

- sentiment analysis.
- speech recognition.
- speech synthesis.
- translation.

Explanation:

Reference:

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

## Question: 97

CertyIQ

You need to build an app that will read recipe instructions aloud to support users who have reduced vision. Which version service should you use?

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

Answer: C

Explanation:

Speech is actually Text-to-Speech

Speech Recognition is actually Speech-to-Text.



Reference:

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

**Question: 98**

**CertyIQ**

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 Speech service to transcribe a call to text.	<input type="radio"/>	<input type="radio"/>
You can use the Text Analytics service to extract key entities from a call transcript.	<input type="radio"/>	<input type="radio"/>
You can use the Speech service to translate the audio of a call to a different language.	<input type="radio"/>	<input type="radio"/>

**Answer:**

**Answer Area**

Statements	Yes	No
You can use the Speech service to transcribe a call to text.	<input checked="" type="radio"/>	<input type="radio"/>
You can use the Text Analytics service to extract key entities from a call transcript.	<input checked="" type="radio"/>	<input type="radio"/>
You can use the Speech service to translate the audio of a call to a different language.	<input checked="" type="radio"/>	<input type="radio"/>

**Explanation:**

Transcribe a call to text -Speech Service: Speech to Text Service

Extract call Transcription to find key entity - Text Analytic : Entity Recognition

Translate a call to different language : Speech Service : Speech Translations

Thus Y/ Y /Y

Reference:

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

<https://azure.microsoft.com/en-gb/services/cognitive-services/speech-services/>

### Question: 99

CertyIQ

Your website has a chatbot to assist customers.

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

Which type of AI workload should you use?

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

**Answer: D**

**Explanation:**

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

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

Reference:

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

### Question: 100

CertyIQ

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

**Answer: AB**

**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: 101

CertyIQ

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

NOTE: Each correct selection is worth one point.

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

**Answer: AD**

**Explanation:**

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

Incorrect Answers:

C: Extracting key phrases is not speech synthesis.

Reference:

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

**Question: 102**

**CertyIQ**

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:**

## 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"/>

### 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: 103

CertyIQ

### 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.

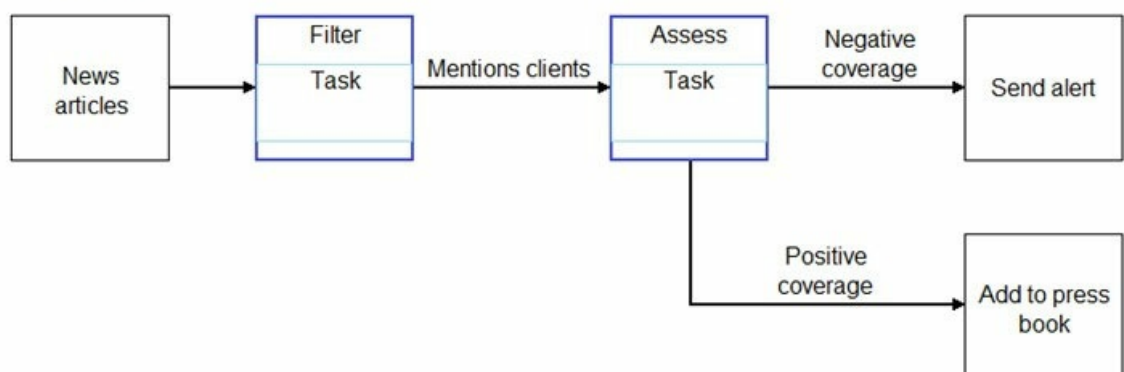
NOTE: Each correct selection is worth one point.

Select and Place:

### Tasks

- Entity recognition
- Sentiment analysis
- Speech synthesis
- Translation

### Answer Area



### Answer:

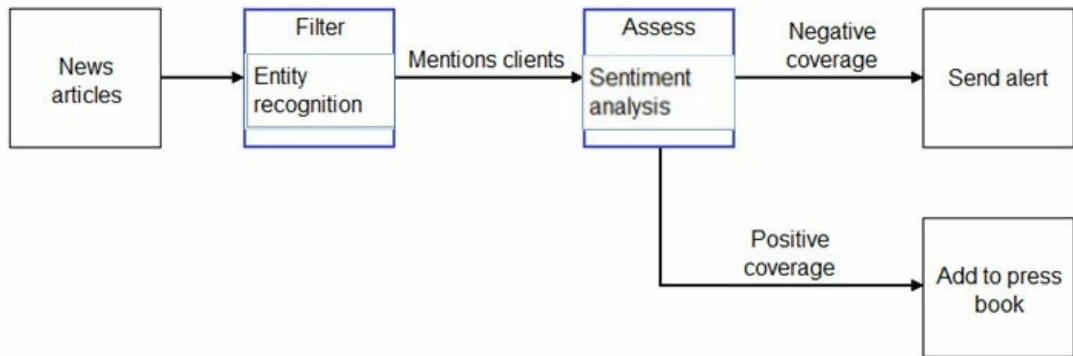
**Tasks**

Entity recognition

Sentiment analysis

Speech synthesis

Translation

**Answer Area****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-tos/text-analytics-how-to-sentiment-analysis>

**Question: 104****CertyIQ**

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

**Answer: D**

**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: 105

CertyIQ

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

**Answer: C**

**Explanation:**

C - "Use key phrase extraction to quickly identify the main concepts in text. For example, in the text "The food was delicious and the staff were wonderful.", key phrase extraction will return the main topics: "food" and "wonderful staff"."

### Question: 106

CertyIQ

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

**Answer: A**

**Explanation:**

NLP can extract key terms from the reports to generate summaries.

Reference:

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



**Question: 107**

HOTSPOT -

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

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:

**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.	

**Explanation:**

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: 108**

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

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

**Answer: D**

**Explanation:**

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

Reference:

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

**Question: 109**

**CertyIQ**

You are developing a chatbot solution in Azure.  
Which service should you use to determine a user's intent?

- A. Translator
- B. QnA Maker
- C. Speech
- D. Language Understanding (LUIS)

**Answer: D**

**Explanation:**

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

Design your LUIS model with categories of user intentions called intents. Each intent needs examples of user utterances. Each utterance can provide data that needs to be extracted with machine-learning entities.

Reference:

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

**Question: 110**

**CertyIQ**

You need to make the written press releases of your company available in a range of languages.  
Which service should you use?

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

**Answer: A**

**Explanation:**

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

Reference:

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

**Question: 111****CertyIQ**

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Statements	Yes	No
The Text Analytics service can identify in which language text is written.	<input type="radio"/>	<input type="radio"/>
The Text Analytics service can detect handwritten signatures in a document.	<input type="radio"/>	<input type="radio"/>
The Text Analytics service can identify companies and organizations mentioned in a document.	<input type="radio"/>	<input type="radio"/>

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

Statements	Yes	No
The Text Analytics service can identify in which language text is written.	<input checked="" type="radio"/>	<input type="radio"/>
The Text Analytics service can detect handwritten signatures in a document.	<input type="radio"/>	<input checked="" type="radio"/>
The Text Analytics service can identify companies and organizations mentioned in a document.	<input checked="" type="radio"/>	<input type="radio"/>

**Explanation:**

The Text Analytics API is a cloud-based service that provides advanced natural language processing over raw text, and includes four main functions: sentiment analysis, key phrase extraction, named entity recognition, and language detection.

Box 1: Yes -

You can detect which language the input text is written in and report a single language code for every document submitted on the request in a wide range of languages, variants, dialects, and some regional/cultural languages. The language code is paired with a score indicating the strength of the score.

Box 2: No -

Box 3: Yes -

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/cognitive-services/text-analytics/overview>

CertyIQ

### Question: 112

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

Entity recognition

Key phrase extraction

Language modeling

Sentiment analysis

Translation

Speech recognition and speech synthesis

#### Answer Area

Workload Type

Extracts persons, locations, and organizations from the text

Workload Type

Evaluates text along a positive-negative scale

Workload Type

Converts text to a different language

#### Answer:

##### Workloads Types

Entity recognition

Key phrase extraction

Language modeling

Sentiment analysis

Translation

Speech recognition and speech synthesis

##### Answer Area

Entity recognition

Extracts persons, locations, and organizations from the text

Sentiment analysis

Evaluates text along a positive-negative scale

Translation

Converts text to a different language

#### 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: 113

CertyIQ

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
Monitoring online service reviews for profanities is an example of natural language processing.	<input type="radio"/>	<input type="radio"/>
Identifying brand logos in an image is an example of natural languages processing.	<input type="radio"/>	<input type="radio"/>
Monitoring public news sites for negative mentions of a product is an example of natural language processing.	<input type="radio"/>	<input type="radio"/>

Answer:

### Answer Area

Statements	Yes	No
Monitoring online service reviews for profanities is an example of natural language processing.	<input checked="" type="radio"/>	<input type="radio"/>
Identifying brand logos in an image is an example of natural languages processing.	<input type="radio"/>	<input checked="" type="radio"/>
Monitoring public news sites for negative mentions of a product is an example of natural language processing.	<input checked="" type="radio"/>	<input type="radio"/>

### Explanation:

Box 1: Yes -

Content Moderator is part of Microsoft Cognitive Services allowing businesses to use machine assisted moderation of text, images, and videos that augment human review.

The text moderation capability now includes a new machine-learning based text classification feature which uses a trained model to identify possible abusive, derogatory or discriminatory language such as slang, abbreviated words, offensive, and intentionally misspelled words for review.



Box 2: No -

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

Box 3: Yes -

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

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

Reference:

<https://azure.microsoft.com/es-es/blog/machine-assisted-text-classification-on-content-moderator-public-preview/> <https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/natural-language-processing>

### Question: 114

CertyIQ

You are developing a natural language processing solution in Azure. The solution will analyze customer reviews and determine how positive or negative each review is.

This is an example of which type of natural language processing workload?

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

**Answer: B**

**Explanation:**

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

Reference:

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

### Question: 115


CertyIQ

You use natural language processing to process text from a Microsoft news story.

You receive the output shown in the following exhibit.



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



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

Which type of natural languages processing was performed?

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

**Answer: A**

**Explanation:**

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

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

Reference:

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

**Question: 116**

**CertyIQ**

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

Entity recognition

Key phrase extraction

Language detection

Sentiment analysis

### Answer Area

API Feature

Understand how upset a customer is based on the text contained in the support ticket.

API Feature

Summarize important information from the support ticket.

API Feature

Extract key dates from the support ticket.

### Answer:

#### API Features

Entity recognition

Key phrase extraction

Language detection

Sentiment analysis

#### 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.

Entity recognition

Extract key dates from the support ticket.

### 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: 117

CertyIQ

You are authoring a Language Understanding (LUIS) application to support a music festival.

You want users to be able to ask questions about scheduled shows, such as: 'Which act is playing on the main stage?'

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

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

D. an entity

**Answer: B**

**Explanation:**

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

Reference:

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

### Question: 118

CertyIQ

You build a QnA Maker bot by using a frequently asked questions (FAQ) page. You need to add professional greetings and other responses to make the bot more user friendly. What should you do?

- A. Increase the confidence threshold of responses
- B. Enable active learning
- C. Create multi-turn questions
- D. Add chit-chat

**Answer: D**

**Explanation:**

Chitchat is correct. You can choose between different personas.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/chit-chat-knowledge-base?tabs=v1>

### Question: 119

CertyIQ

You need to develop a chatbot for a website. The chatbot must answer users' questions based on the information in the following documents:

- A product troubleshooting guide in a Microsoft Word document
- A frequently asked questions (FAQ) list on a webpage

Which service should you use to process the documents?

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

**Answer: D**

**Explanation:**

QnA Maker is a service designed to build conversational question-and-answer experiences. It processes information from documents like FAQs, Word files, and webpages to create a knowledge base for chatbots.

Why D is Correct:

It extracts and organizes information from FAQs and documents.

Perfect for creating a chatbot that answers based on specific content.

Why Not the Others?

A. Azure Bot Service: Manages chatbot hosting but doesn't process documents.

B. Language Understanding: Focuses on interpreting user intent, not extracting Q&A content.

C. Text Analytics: Analyzes text for sentiment or key phrases but doesn't create Q&A systems.

Key Tip: Use QnA Maker for chatbots needing Q&A data from documents and webpages.

Reference:

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

### Question: 120

CertyIQ

You are building a Language Understanding model for an e-commerce business. You need to ensure that the model detects when utterances are outside the intended scope of the model. What should you do?

- A. Test the model by using new utterances
- B. Add utterances to the None intent
- C. Create a prebuilt task entity
- D. Create a new model

**Answer: B**

**Explanation:**

The None intent is filled with utterances that are outside of your domain.

In Language Understanding (LUIS) models, the None intent is used to handle utterances that are outside the model's intended scope. Adding irrelevant or out-of-scope utterances to the None intent ensures the model can correctly identify and disregard such inputs.

Why B is Correct:

Adding examples of out-of-scope utterances to the None intent helps the model distinguish between valid and irrelevant inputs.

Why Not the Others?

A. Test the model by using new utterances: Testing helps validate the model but doesn't directly teach it to handle out-of-scope utterances.

C. Create a prebuilt task entity: Entities are for extracting data, not detecting out-of-scope inputs.

D. Create a new model: A new model is unnecessary; handling out-of-scope utterances is a built-in feature using the None intent.

Key Tip: Use the None intent to classify and handle irrelevant or out-of-scope utterances in LUIS models.

Reference:

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

### Question: 121

CertyIQ

Which two scenarios are examples of a natural language processing workload? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

**Answer: BC**

#### Explanation:

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: 122

CertyIQ

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

**Answer: CD**

#### Explanation:

speech-to-text and language modeling. You need to use language modeling to determine the intent of the utterance and to perform an action based on that intent.

The AI solution for controlling smart devices using verbal commands involves the following NLP workloads:

C. Speech-to-text: Converts spoken commands into text so they can be processed by the AI system.

D. Language modeling: Helps the system understand and interpret the meaning of the transcribed text (verbal commands).

Why Not the Others?

- A. Text-to-speech: Converts text to spoken words but is not used for understanding verbal commands.
- B. Key phrase extraction: Identifies important phrases in text but isn't directly related to command processing.
- E. Translation: Converts text from one language to another, which isn't required here.

Key Tip: For verbal command processing, focus on Speech-to-text (conversion of spoken words) and Language modeling (understanding the command).

## Question: 123

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
The Language service can identify in which language text is written.	<input type="radio"/>	<input type="radio"/>
The Language service can detect handwritten signatures in a document.	<input type="radio"/>	<input type="radio"/>
The Language service can identify companies and organizations mentioned in a document.	<input type="radio"/>	<input type="radio"/>

### Answer:

#### Answer Area

Statements	Yes	No
The Language service can identify in which language text is written.	<input checked="" type="radio"/>	<input type="radio"/>
The Language service can detect handwritten signatures in a document.	<input type="radio"/>	<input checked="" type="radio"/>
The Language service can identify companies and organizations mentioned in a document.	<input checked="" type="radio"/>	<input type="radio"/>

### Explanation:

Box 1: Yes -

Azure Cognitive Service for Language provides features including:

\* Language detection: This pre-configured feature evaluates text, and determines the language it was written in. It returns a language identifier and a score that indicates the strength of the analysis.

Box 2: No -

Handwritten detection is part of OCR (Optical Character Recognition).

Box 3: Yes -

Azure Cognitive Service for Language provides features including:

\* Named Entity Recognition (NER): This pre-configured feature identifies entities in text across several pre-defined categories.

Note: Named entity recognition is a natural language processing technique that can automatically scan entire



articles and pull out some fundamental entities in a text and classify them into predefined categories. Entities may be,  
Organizations,  
Quantities,  
Monetary values,  
Percentages, and more.

People's names -

Company names -  
Geographic locations (Both physical and political)

Product names -

Dates and times -

Amounts of money -

Names of events -

Reference:

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

## Question: 124

CertyIQ

DRAG DROP -

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

Match the Azure Cognitive Services to the appropriate tasks.

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

NOTE: Each correct selection is worth one point.

Select and Place:

### Services

Speech

Language service

Translator Text

### Answer Area

Convert a user's speech to text

Identify a user's intent

Provide a spoken response to the user

Answer:

### Services

Speech

Language service

Translator Text

### Answer Area

Speech

Language service

Speech

Convert a user's speech to text

Identify a user's intent

Provide a spoken response to the user

**Explanation:**

Box 1: Speech -

The Speech service provides speech-to-text and text-to-speech capabilities with an Azure Speech resource. You can transcribe speech to text with high accuracy, produce natural-sounding text-to-speech voices, translate spoken audio, and use speaker recognition during conversations.

Box 2: Language service -

Build applications with conversational language understanding, a Cognitive Service for Language feature that understands natural language to interpret user goals and extracts key information from conversational phrases. Create multilingual, customizable intent classification and entity extraction models for your domain-specific keywords or phrases across 96 languages.

Box 3: Speech -

Incorrect:

Not Translator text: Text translation is a cloud-based REST API feature of the Translator service that uses neural machine translation technology to enable quick and accurate source-to-target text translation in real time across all supported languages.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/overview> <https://azure.microsoft.com/en-us/services/cognitive-services/conversational-language-understanding/> <https://docs.microsoft.com/en-us/azure/cognitive-services/translator/text-translation-overview>

**Question: 125****CertyIQ**

You need to make the written press releases of your company available in a range of languages. Which service should you use?

- A. Speech
- B. Language
- C. Translator
- D. Personalizer

**Answer: C****Explanation:**

Translator, an AI service for real-time document and text translation.

Translate text instantly or in batches across more than 100 languages, powered by the latest innovations in machine translation. Support a wide range of use cases, such as translation for call centers, multilingual conversational agents, or in-app communication.

Reference:

<https://azure.microsoft.com/en-us/services/cognitive-services/translator/4>

**Question: 126****CertyIQ**

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. anomaly detection
- B. natural language processing

- C. computer vision
- D. knowledge mining

**Answer: B**

**Explanation:**

Key phrase extraction is one of the features offered by Azure Cognitive Service for Language, a collection of machine learning and AI algorithms in the cloud for developing intelligent applications that involve written language. Use key phrase extraction to quickly identify the main concepts in text. For example, in the text "The food was delicious and the staff were wonderful.", key phrase extraction will return the main topics: "food" and "wonderful staff".

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/language-service/key-phrase-extraction/overview>

**Question: 127**

**CertyIQ**

You need to build an app that will read recipe instructions aloud to support users who have reduced vision. Which version service should you use?

- A. Language service
- B. Translator
- C. Speech
- D. Personalizer

**Answer: C**

**Explanation:**

Speech, a managed service offering industry-leading speech capabilities such as speech-to-text, text-to-speech, speech translation, and speaker recognition.

Reference:

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

**Question: 128**

**CertyIQ**

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

**Answer: D**

**Explanation:**

Active learning.

1.Bot gets the answer from the knowledge base with the GenerateAnswer API, using the top property to get a number of answers.

2.Bot determines explicit feedback: Using your own custom business logic, filter out low scores. In the bot or client-application, display list of possible answers to the user and get user's selected answer.

3.Bot sends selected answer back to QnA Maker with the Train API.

Reference:

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

### Question: 129

CertyIQ

You are developing a conversational AI solution that will communicate with users through multiple channels including email, Microsoft Teams, and webchat.

Which service should you use?

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

**Answer: B**

**Explanation:**

The Azure Bot Service is specifically designed to build conversational AI solutions that can integrate with multiple communication channels such as email, Microsoft Teams, and webchat. It provides tools to create, manage, and deploy chatbots seamlessly across these platforms.

Why B is Correct:

Azure Bot Service allows integration with various channels and handles the communication logic needed for conversational AI.

Why Not the Others?

A. Text Analytics: Extracts insights from text (e.g., sentiment, key phrases) but doesn't support building or deploying conversational bots.

C. Translator: Provides translation services but does not enable chatbot development.

D. Form Recognizer: Extracts data from documents and forms but is unrelated to conversational AI.

Key Tip: Use Azure Bot Service for developing and deploying chatbots across multiple communication channels.

Reference:

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

### Question: 130

CertyIQ

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:

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 type="radio"/>	<input checked="" type="radio"/>

Explanation:

Y/N/N

1. A bot that responds to queries by internal users is an example of a conversational AI workload.

- **Answer: Yes**

- Conversational AI includes chatbots or virtual assistants that interact with users via text or speech. A bot responding to internal queries falls under this category.

2. An application that displays images relating to an entered search term is an example of a conversational AI workload.

- **Answer: No**

- This is not conversational AI. It is a search or recommendation system, which is different from conversational workloads.

3. A web form used to submit a request to reset a password is an example of a conversational AI workload.

- **Answer: No**

- A web form is a standard user interface component and does not involve conversational AI.



Key Tip:

Conversational AI involves systems capable of natural language interaction, like chatbots or virtual assistants. Regular applications, forms, or search systems do not fall under this category.

Reference:

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

### Question: 131

CertyIQ

You need to provide content for a business chatbot that will help answer simple user queries. What are three ways to create question and answer text by using QnA Maker? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

**Answer: ACE**

**Explanation:**

Automatic extraction -

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



Reference:

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

### Question: 132

CertyIQ

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

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

Which service should you use?

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

**Answer: A**

**Explanation:**

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

Reference:

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

### Question: 133

CertyIQ

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

**Answer: BC**

**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: 134

CertyIQ

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

**Answer: AB**

**Explanation:**

A - smart device

B - website

C&D don't look like as chat-bots.

Conversational AI workloads involve systems that can understand, process, and respond to human input (usually in natural language, either through text or speech).

Why A and B are Correct:

A. A smart device in the home that responds to questions such as "What will the weather be like today?":

This is a clear example of conversational AI. The smart device listens to verbal queries and provides spoken responses based on information, which is a typical function of virtual assistants (e.g., Alexa, Google Assistant).

B. A website that uses a knowledge base to interactively respond to users' questions:

This is another example of conversational AI. It involves an interactive Q&A system (such as a chatbot) that uses a knowledge base to respond to user queries in natural language.

Why Not C and D:

C. Assembly line machinery that autonomously inserts headlamps into cars:

This is an example of automation, not conversational AI. It involves mechanical processes, not natural language interaction.

D. Monitoring the temperature of machinery to turn on a fan when the temperature reaches a specific threshold:

This is an example of environmental monitoring and automation. It doesn't involve conversational interaction or natural language processing, so it isn't considered conversational AI.

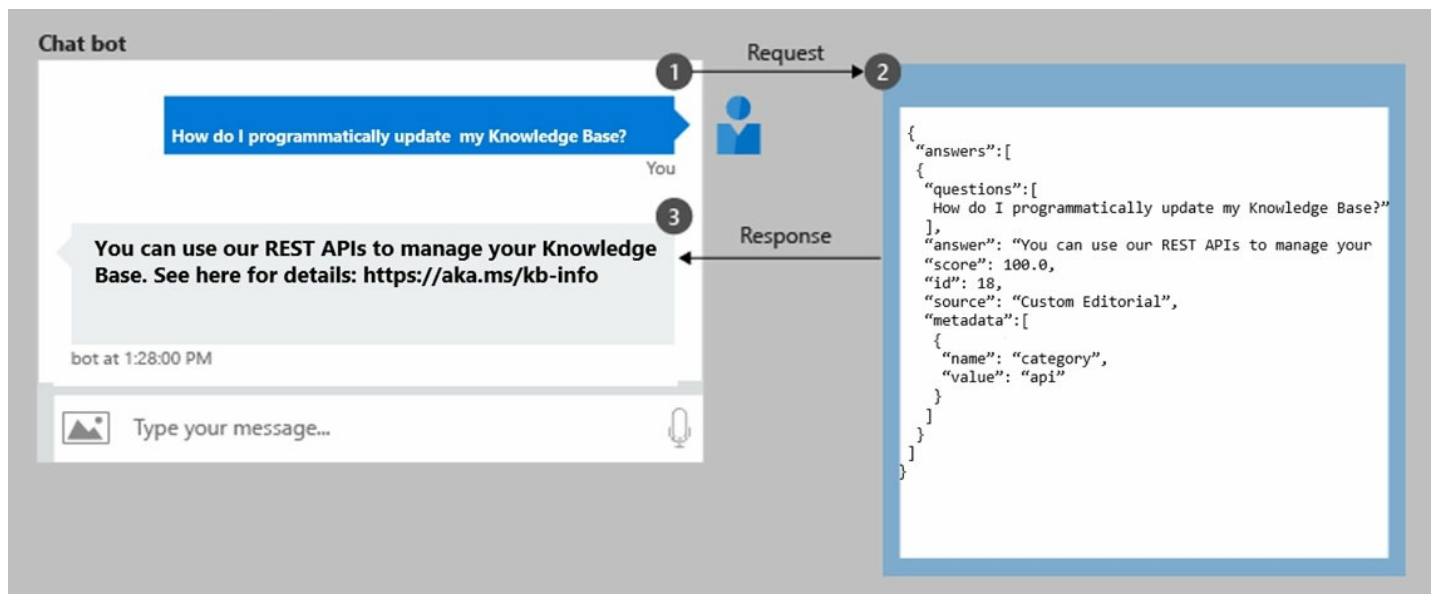
**Key Tip:**

Conversational AI systems allow users to interact with machines or software using natural language (e.g., chatbots, virtual assistants).

### Question: 135

**CertyIQ**

You have the process shown in the following exhibit.



Which type of 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

**Answer: B**

**Explanation:**

In the provided diagram, the process shows the interaction between a user and a bot, which responds to user queries using a knowledge base. The bot's response is based on a structured data format (JSON) that includes a specific answer with metadata.

Why B is Correct (Chatbot):

The bot responds to a user query ("How do I programmatically update my Knowledge Base?") and provides an answer.

The response comes from a knowledge base that the bot can access, which is typically part of a chatbot solution.

The interaction shown — where a user asks a question and gets a structured, informative response from a knowledge base — is characteristic of a chatbot system that handles user queries.

Why Not the Others?

A. Sentiment analysis solution:

Sentiment analysis involves understanding the sentiment or emotion behind text (e.g., positive, negative, neutral). This scenario focuses on information retrieval and not sentiment analysis.

C. Machine learning model:

The solution shown doesn't involve training or predictive models but is instead using predefined responses from a knowledge base, which is common in chatbots.

D. Computer vision application:

Computer vision is used for analyzing and processing visual data (images or videos). There is no indication of any visual data or image processing in this scenario.

**Key Tip:**

The diagram shows a chatbot interacting with a knowledge base to provide answers to queries, which is a classic conversational AI scenario.

**Question: 136****CertyIQ**

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

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

**Answer: B****Explanation:**

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

**Incorrect Answers:**

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

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

**Reference:**

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

**Question: 137****CertyIQ**

Which AI service should you use to create a bot from a frequently asked questions (FAQ) document?

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

**Answer: A****Explanation:**

QnA Maker is the AI service designed specifically to create bots that can answer questions based on a frequently asked questions (FAQ) document. It allows you to upload FAQs or other knowledge base documents and turns them into a question-answering bot.

Why A is Correct:

QnA Maker is ideal for building conversational AI systems like chatbots that respond to user queries using a pre-existing knowledge base (such as FAQs).

Why Not the Others?

B. Language Understanding (LUIS):

LUIS is used to build language models that understand user intent and extract entities, not for creating a Q&A system from a document.

C. Text Analytics:

Text Analytics is used for tasks like sentiment analysis, key phrase extraction, or entity recognition, but it doesn't specifically turn FAQs into a bot.

D. Speech:

Speech is used for converting speech to text or vice versa, but it doesn't build bots from FAQs or documents.

**Key Tip:**

Use QnA Maker when you need to create a bot that answers questions from documents, such as FAQs. It's specifically built for question-answering systems.

**Question: 138**

**CertyIQ**

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:**

## 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.	

### 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: 139

CertyIQ

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.

### Answer: D

### Explanation:

A webchat bot refers to a chatbot integrated into a website, allowing users to interact with it in real-time via a chat interface. The bot typically answers questions or performs actions based on the user's input.

Why D is Correct:

This scenario clearly describes a chatbot on a website that answers common questions about events and ticket purchases, which is a typical use of a webchat bot.

Why Not the Others?

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:

This is an example of sentiment analysis and not a chatbot, as it focuses on analyzing reviews rather than



engaging in conversation.

B. Translate into English questions entered by customers at a kiosk so that the appropriate person can call the customers back:

This involves translation and human routing, not a conversational interaction with a bot.

C. Accept questions through email, and then route the email messages to the correct person based on the content of the message:

This describes email routing and automation but not a chatbot interaction on a website.

**Key Tip:**

A webchat bot directly interacts with users through a chat interface on a website, making it ideal for answering questions in real-time.

**Question: 140**

**CertyIQ**

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.	<input type="radio"/>	<input type="radio"/>
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	<input type="radio"/>	<input type="radio"/>
The QnA Maker service can determine the intent of a user utterance.	<input type="radio"/>	<input type="radio"/>

**Answer:**

## Answer Area

Statements	Yes	No
You can use QnA Maker to query an Azure SQL database.	<input type="radio"/>	<input checked="" type="radio"/>
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	<input checked="" type="radio"/>	<input type="radio"/>
The QnA Maker service can determine the intent of a user utterance.	<input type="radio"/>	<input checked="" type="radio"/>

### Explanation:

Language Understanding (LUIS) and QnA Maker solve different issues. LUIS determines the intent of a user's text (known as an utterance), while QnA Maker determines the answer to a user's text (known as a query).

### 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: 141

CertyIQ

### 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 Cortana.	<input type="radio"/>	<input type="radio"/>
You can communicate with a bot by using Microsoft Teams.	<input type="radio"/>	<input type="radio"/>
You can communicate with a bot by using a webchat interface.	<input type="radio"/>	<input type="radio"/>

### Answer:

## Answer Area

Statements	Yes	No
You can communicate with a bot by using Cortana.	<input checked="" type="radio"/>	<input type="radio"/>
You can communicate with a bot by using Microsoft Teams.	<input checked="" type="radio"/>	<input type="radio"/>
You can communicate with a bot by using a webchat interface.	<input checked="" type="radio"/>	<input type="radio"/>

### Explanation:

You can connect a bot via:

Alexa, Office 365 email, Facebook, Kik, LINE, Teams, Skype for Business, Slack, Telegram, WeChat, Webex

Reference:

The Channels List for Bots: <https://learn.microsoft.com/en-us/azure/bot-service/bot-service-manage-channels?view=azure-bot-service-4.0>

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

### Question: 142

CertyIQ

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

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

Answer:

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

### Explanation:

Box 1: Yes

Box 2: Yes

Box 3: yes

1. A restaurant can use a chatbot to empower customers to make reservations by using a website or an app.

Answer: Yes

Chatbots can assist users in booking tables directly through websites or apps by interacting with them in real-time.

Tip: Use chatbots to simplify reservation processes, improving customer convenience.

2. A restaurant can use a chatbot to answer inquiries about business hours from a webpage.

Answer: Yes

Chatbots can provide instant answers to FAQs, including business hours, directly from a webpage.

Tip: Program chatbots to handle common customer queries to save time and improve accessibility.

3. A restaurant can use a chatbot to automate responses to customer reviews on an external website.

Answer: Yes

Some external platforms may allow chatbot integration or APIs to post automated responses to customer reviews. However, this depends on the platform's rules and permissions. For instance, Google and Facebook support limited automation via APIs, while others may not.

Key Tip: If the platform allows it, you can integrate chatbots to respond to reviews. Always ensure responses are professional and platform-compliant.

### Question: 143

CertyIQ

Which two scenarios are examples of a conversational AI workload? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

**Answer: BC**

**Explanation:**

B: A bot is an automated software program designed to perform a particular task. Think of it as a robot without a body.

C: Automated customer interaction is essential to a business of any size. In fact, 61% of consumers prefer to communicate via speech, and most of them prefer self-service. Because customer satisfaction is a priority for all businesses, self-service is a critical facet of any customer-facing communications strategy.

Incorrect Answers:

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

Reference:

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

**Question: 144**

**CertyIQ**

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Statements	Yes	No
Azure Bot Service and Azure Cognitive Services can be integrated.	<input type="radio"/>	<input type="radio"/>
Azure Bot Service engages with customers in a conversational manner.	<input type="radio"/>	<input type="radio"/>
Azure Bot Service can import frequently asked questions (FAQ) to question and answer sets.	<input type="radio"/>	<input type="radio"/>

**Answer:**



## Answer Area

Statements	Yes	No
Azure Bot Service and Azure Cognitive Services can be integrated.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Bot Service engages with customers in a conversational manner.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Bot Service can import frequently asked questions (FAQ) to question and answer sets.	<input type="radio"/>	<input checked="" type="radio"/>

### Explanation:

Box 1: Yes -

Azure bot service can be integrated with the powerful AI capabilities with Azure Cognitive Services.

Box 2: Yes -

Azure bot service engages with customers in a conversational manner.

Box 3: No -

The QnA Maker service creates knowledge base, not question and answers sets.

Note: You can use the QnA Maker service and a knowledge base to add question-and-answer support to your bot. When you create your knowledge base, you seed it with questions and answers.

Reference:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-tutorial-add-qna>

### Question: 145

CertyIQ

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 webchat bot can interact with users visiting a website	<input type="radio"/>	<input type="radio"/>
Automatically generating captions for pre-recorded videos is an example of conversational AI	<input type="radio"/>	<input type="radio"/>
A smart device in the home that responds to questions such as "What will the weather like today?" is an example of conversational AI	<input type="radio"/>	<input type="radio"/>



Answer:

## Answer Area

Statements	Yes	No
A webchat bot can interact with users visiting a website	<input checked="" type="radio"/>	<input type="radio"/>
Automatically generating captions for pre-recorded videos is an example of conversational AI	<input type="radio"/>	<input checked="" type="radio"/>
A smart device in the home that responds to questions such as “What will the weather like today?” is an example of conversational AI	<input checked="" type="radio"/>	<input type="radio"/>

### Explanation:

1. A webchat bot can interact with users visiting a website.

Answer: Yes

Webchat bots are designed to engage with website visitors, answer queries, and perform tasks like providing information or collecting user input.

Tip: Use webchat bots to improve user experience and provide quick customer support.

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

Answer: No

Generating captions involves speech-to-text processing, which is not conversational AI. Conversational AI focuses on interactive communication with users, like chatbots or virtual assistants.

Tip: Remember, conversational AI requires interaction or dialogue, not just text generation.

3. A smart device in the home that responds to questions such as “What will the weather be like today?” is an example of conversational AI.

Answer: Yes

Smart devices like Alexa or Google Assistant use conversational AI to interpret user queries and provide responses.

Tip: Smart assistants are classic examples of conversational AI, as they allow real-time interaction.

### Reference:

<https://docs.microsoft.com/en-us/azure/architecture/reference-architectures/ai/conversational-bot>

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

**Question: 146**

You have a knowledge base of frequently asked questions (FAQ).  
You create a bot that uses the knowledge base to respond to customer requests.  
You need to identify what the bot can perform without adding additional skills.  
What should you identify?

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

**Answer: C****Explanation:**

A. Register customer purchases

Incorrect: Registering purchases involves transactional actions and requires integrating with a system or database to log the purchases. This is beyond the scope of an FAQ-based bot.

B. Register customer complaints

Incorrect: Registering complaints may require creating records in a complaint management system, which is not achievable with just an FAQ knowledge base.

C. Answer questions from multiple users simultaneously

Correct: An FAQ-based bot is designed to provide answers to predefined questions and can handle multiple users simultaneously without requiring additional skills or integrations.

D. Provide customers with return materials authorization (RMA) numbers

Incorrect: Providing RMA numbers usually involves accessing backend systems to generate or retrieve these numbers, which requires additional skills and integrations.

Correct Answer: C

Key Tip:

An FAQ-based bot excels in answering predefined questions efficiently for multiple users but does not perform complex tasks like managing transactions or accessing external systems.

Reference:

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

**Question: 147**

HOTSPOT -

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

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

Answer:

### Answer Area

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

Explanation:

Box 1: Yes -

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

Note: Connect Cortana Channels -

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

Etc.

Box 2: Yes -

QnA Maker is an easy-to-use web-based service that makes it easy to power a question-answer application or chatbot from semi-structured content like FAQ documents and product manuals. With QnA Maker, developers can build, train, and publish question and answer bots in minutes.

Box 3: Yes -

Reference:

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

Question: 148

HOTSPOT -

CertyIQ

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
Chatbots can only be built by using custom code.	<input type="radio"/>	<input type="radio"/>
The Azure Bot Service provides services that can be used to host conversational bots.	<input type="radio"/>	<input type="radio"/>
Bots built by using the Azure Bot Service can communicate with Microsoft Teams users.	<input type="radio"/>	<input type="radio"/>

**Answer:**

**Answer Area**

Statements	Yes	No
Chatbots can only be built by using custom code.	<input type="radio"/>	<input checked="" type="radio"/>
The Azure Bot Service provides services that can be used to host conversational bots.	<input checked="" type="radio"/>	<input type="radio"/>
Bots built by using the Azure Bot Service can communicate with Microsoft Teams users.	<input checked="" type="radio"/>	<input type="radio"/>

**Explanation:**

Box 1: No -

Build conversational experiences with Power Virtual Agents and Azure Bot Service

Azure Bot Service provides an integrated development environment for bot building. Its integration with Power Virtual Agents, a fully hosted low-code platform, enables developers of all technical abilities build conversational AI bots"no code needed.

Box 2: Yes -

Box 3: Yes -

You can configure your bot to communicate with people via Microsoft Teams.

Reference:

<https://azure.microsoft.com/en-us/services/bot-services/#overview> <https://docs.microsoft.com/en-us/azure/bot-service/channel-connect-teams>

**Question: 149**

**CertyIQ**

HOTSPOT -

Select the answer that correctly completes the sentence.

Hot Area:

**Answer Area**

Computer vision capabilities can be deployed to

develop a text-based chatbot for a website.
identify anomalous customer behavior on an online store.
integrate a facial recognition feature into an app.
suggest automated responses to incoming email.

**Answer:**

### Answer Area

Computer vision capabilities can be deployed to

develop a text-based chatbot for a website.
identify anomalous customer behavior on an online store.
integrate a facial recognition feature into an app.
suggest automated responses to incoming email.

### 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.

\* Optical Character Recognition (OCR)

\* Spatial Analysis

\* Image Analysis

The Image Analysis service extracts many visual features from images, such as objects, faces, adult content, and auto-generated text descriptions. Follow the Image Analysis quickstart to get started.

Reference:

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

### Question: 150

CertyIQ

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

**Answer: B**

### Explanation:

B is the answer.

<https://learn.microsoft.com/en-us/azure/machine-learning/component-reference/split-data?view=azureml-api-2>

Use the Split Data component to divide a dataset into two distinct sets.

This component is useful when you need to separate data into training and testing sets. You can also customize the way that data is divided. Some options support randomization of data. Others are tailored for a certain data type or model type.

### Question: 151

CertyIQ

Which statement is an example of a Microsoft responsible AI principle?

- A. AI systems must use only publicly available data
- B. AI systems must be transparent and inclusive
- C. AI systems must keep personal details public

D. AI systems must protect the interests of the company

**Answer: B**

**Explanation:**

Transparency and inclusiveness are the Microsoft AI principle.

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

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.

### Question: 152

CertyIQ

DRAG DROP -

Match the types of natural language 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 match is worth one point.

Select and Place:

#### Workload types

Entity recognition

Key phrase extraction

Language modeling

Sentiment analysis

Speech recognition and speech synthesis

Translation

#### Answer Area

Extracts persons, locations, and organizations from the text.

Evaluates text along a positive-negative scale.

Converts text to a different language.

**Answer:**

#### Workload types

Entity recognition

Key phrase extraction

Language modeling

Sentiment analysis

Speech recognition and speech synthesis

Translation

#### Answer Area

Entity recognition

Extracts persons, locations, and organizations from the text.

Sentiment analysis

Evaluates text along a positive-negative scale.

Translation

Converts text to a different language.

**Explanation:**

1. Entity recognition
2. Sentiment analysis
3. Translation



<https://learn.microsoft.com/en-us/azure/cognitive-services/language-service/named-entity-recognition/overview>

Named Entity Recognition (NER) is one of the features offered by Azure Cognitive Service for Language, a collection of machine learning and AI algorithms in the cloud for developing intelligent applications that involve written language. The NER feature can identify and categorize entities in unstructured text. For example: people, places, organizations, and quantities.

<https://learn.microsoft.com/en-us/azure/cognitive-services/language-service/sentiment-opinion-mining/overview>

Sentiment analysis and opinion mining are features offered by Azure Cognitive Service for Language, a collection of machine learning and AI algorithms in the cloud for developing intelligent applications that involve written language. These features help you find out what people think of your brand or topic by mining text for clues about positive or negative sentiment, and can associate them with specific aspects of the text.

### Question: 153

CertyIQ

You need to reduce the load on telephone operators by implementing a chatbot to answer simple questions with predefined answers.

Which two AI services should you use to achieve the goal? Each correct answers presents part of the solution.

NOTE: Each correct selection is worth one point.

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

**Answer: BC**

**Explanation:**

B. Azure Bot Service

C. Language Service - QnA is replaced by Language Service

To implement a chatbot for answering simple questions with predefined answers and reduce the load on telephone operators, you should use:

B. Azure Bot Service - Azure Bot Service is a platform for creating and managing chatbots, making it a suitable choice for building a chatbot to handle simple questions.

C. Language Service - Language services, including natural language processing capabilities, are essential for understanding user questions and providing relevant responses. This can be integrated into your chatbot to improve its conversational abilities.

Azure Machine Learning (A) and Translator (D) are not typically used as primary components for building chatbots for this specific task. While Azure Machine Learning can be used for more complex machine learning scenarios, it's not necessary for simple question and answer chatbots. Azure Translator is primarily used for language translation tasks and doesn't directly address the goal of reducing the load on telephone operators with a chatbot.

### Question: 154

DRAG DROP

Match the principles of 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. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

#### Principles

Fairness

Inclusiveness

Privacy and security

Reliability and safety

#### Answer Area

AI systems must consistently operate as intended, even under unexpected conditions.

AI systems must protect and secure personal and businesses information.

#### Answer:

##### Principles

Fairness

Inclusiveness

Privacy and security

Reliability and safety

##### Answer Area

Reliability and safety AI systems must consistently operate as intended, even under unexpected conditions.

Privacy and security AI systems must protect and secure personal and businesses information.

#### Explanation:

1. Reliability and safety

2. Privacy and security

<https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai#reliability-and-safety>

AI systems need to be reliable and safe in order to be trusted. It's important for a system to perform as it was originally designed and for it to respond safely to new situations. Its inherent resilience should resist intended or unintended manipulation. Rigorous testing and validation should be established for operating conditions to ensure that the system responds safely to edge cases, and A/B testing and champion/challenger methods should be integrated into the evaluation process.

### Question: 155

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

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

**Answer: D**

**Explanation:**

D is the answer.

<https://learn.microsoft.com/en-us/training/modules/use-automated-machine-learning/5-machine-learning-steps>

You can think of the steps in a machine learning process as:

- Prepare data: Identify the features and label in a dataset. Pre-process, or clean and transform, the data as needed.
- Train model: Split the data into two groups, a training and a validation set. Train a machine learning model using the training data set. Test the machine learning model for performance using the validation data set.
- Evaluate performance: Compare how close the model's predictions are to the known labels.
- Deploy a predictive service: After you train a machine learning model, you can deploy the model as an application on a server or device so that others can use it.

**Question: 156**

**CertyIQ**

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

**Answer: B**

**Explanation:**

B. privacy and security

Using data obtained without permission breaches the Microsoft principle for responsible AI related to "privacy and security." This principle emphasizes the importance of respecting the privacy of individuals and securing their data. Data should be collected and used in a legal and ethical manner, and obtaining data without proper consent or permission can lead to privacy and security violations.

**Question: 157**

**CertyIQ**

HOTSPOT

-

Select the answer that correctly completes the sentence.

## Answer Area

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

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

Answer:

### 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

**Explanation:**

a reliability and safety

<https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai#reliability-and-safety>

AI systems need to be reliable and safe in order to be trusted. It's important for a system to perform as it was originally designed and for it to respond safely to new situations. Its inherent resilience should resist intended or unintended manipulation. Rigorous testing and validation should be established for operating conditions to ensure that the system responds safely to edge cases, and A/B testing and champion/challenger methods should be integrated into the evaluation process.

**Question: 158**

DRAG DROP

CertyIQ

Match the services to the appropriate descriptions.

To answer, drag the appropriate service from the column on the left to its description on the right. Each service may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Services	Answer Area
Azure Storage	
Azure Bot Service	<input type="text"/> Enables the use of natural language to query a knowledge base.
Language Service	<input type="text"/> Enables the real-time transcription of speech-to-text.
Speech	

Answer:

Services	Answer Area
Azure Storage	
Azure Bot Service	
Language Service	<input type="text"/> Enables the use of natural language to query a knowledge base.
Speech	<input type="text"/> Enables the real-time transcription of speech-to-text.

Explanation:

1. Language Service

2. Speech

<https://learn.microsoft.com/en-us/training/modules/build-faq-chatbot-qna-maker-azure-bot-service/2-get-started-knowledge-base>

The Language service includes a custom question answering feature that enables you to create a knowledge base of question and answer pairs that can be queried using natural language input.

<https://learn.microsoft.com/en-us/azure/cognitive-services/speech-service/overview>

The Speech service provides speech to text and text to speech capabilities with an Speech resource. You can transcribe speech to text with high accuracy, produce natural-sounding text to speech voices, translate spoken audio, and use speaker recognition during conversations.

Question: 159

CertyIQ

Which machine learning technique can be used for anomaly detection?

- A. A machine learning technique that classifies objects based on user supplied images.
- B. A machine learning technique that understands written and spoken language.
- C. A machine learning technique that classifies images based on their contents.
- D. A machine learning technique that analyzes data over time and identifies unusual changes.

**Answer: D**

**Explanation:**

D is the answer.

<https://learn.microsoft.com/en-us/azure/cognitive-services/anomaly-detector/overview>

Anomaly Detector is an AI service with a set of APIs, which enables you to monitor and detect anomalies in your time series data with little machine learning (ML) knowledge, either batch validation or real-time inference.

### Question: 160

CertyIQ

You have an AI-based loan approval system.

During testing, you discover that the system has a gender bias.

Which responsible AI principle does this violate?

- A. accountability
- B. reliability and safety
- C. transparency
- D. fairness

**Answer: D**

**Explanation:**

D is the answer.

<https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai?view=azureml-api-2#fairness-and-inclusiveness>

AI systems should treat everyone fairly and avoid affecting similarly situated groups of people in different ways. For example, when AI systems provide guidance on medical treatment, loan applications, or employment, they should make the same recommendations to everyone who has similar symptoms, financial circumstances, or professional qualifications.

### Question: 161

CertyIQ

You are developing a system to predict the prices of insurance for drivers in the United Kingdom.

You need to minimize bias in the system.

What should you do?

- A. Remove information about protected characteristics from the data before sampling.



- B. Take a training sample that is representative of the population in the United Kingdom.
- C. Create a training dataset that uses data from global insurers.
- D. Take a completely random training sample.

**Answer: B**

**Explanation:**

B. Take a training sample that is representative of the population in the United Kingdom.

To minimize bias in the system, it's important that your training data is representative of the population you're modeling. This helps ensure that the model's predictions are valid for the full range of drivers in the United Kingdom.

While option A (Remove information about protected characteristics from the data before sampling) could help in some cases to reduce direct discrimination, it might not be sufficient to minimize all types of biases, as some of these characteristics might be indirectly encoded in the remaining features.

Option C (Create a training dataset that uses data from global insurers) may introduce more bias since driving conditions, laws, and demographics vary greatly by country.

Option D (Take a completely random training sample) could still introduce bias if the original data pool is not representative of the population you're interested in.

## Question: 162

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.

**Answer Area**

Azure Machine Learning designer lets you create machine learning models by

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

**Answer:**

**Answer Area**

Azure Machine Learning designer lets you create machine learning models by

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

**Explanation:**

1. adding and connecting modules on a visual canvas

<https://learn.microsoft.com/en-us/azure/machine-learning/concept-designer?view=azureml-api-2>

Azure Machine Learning designer is a drag-and-drop UI interface to build pipeline in Azure Machine Learning.

**Question: 163**

You have a dataset.

You need to build an Azure Machine Learning classification model that will identify defective products.

What should you do first?

- A. Load the dataset.
- B. Create a clustering model.
- C. Split the data into training and testing datasets.
- D. Create a classification model.

**Answer: A**

**Explanation:**

First at all you need to load the dataset

When building a machine learning model, the process typically follows a sequence of steps. The first step is to load the dataset so you can perform subsequent tasks like preprocessing, splitting, and model training. Without loading the data, you cannot proceed with the remaining steps.

Why the other options are incorrect:

B. Create a clustering model:

Incorrect because clustering is not relevant here. The goal is classification, not clustering.

C. Split the data into training and testing datasets:

This step comes after loading the dataset. Splitting cannot occur before the data is loaded.

D. Create a classification model:

Creating the classification model is one of the later steps after loading and preparing the data.

**Key Tip:**

Always load and inspect your dataset as the first step in any machine learning workflow. This ensures the data is accessible and ready for preprocessing and analysis.

**Question: 164**

You use Azure Machine Learning designer to build a model pipeline.

What should you create before you can run the pipeline?

- A. a registered model
- B. a compute resource
- C. a Jupyter notebook

**Answer: B**

### Explanation:

In Azure Machine Learning Designer, a compute resource is required to execute a pipeline. The compute resource provides the necessary processing power to run the pipeline steps, such as training models, performing data processing, or evaluating results.

Why the other options are incorrect:

A. A registered model

Incorrect: A registered model is created after training and registering the model in Azure Machine Learning. It is not a prerequisite for running the pipeline.

C. A Jupyter notebook

Incorrect: Jupyter notebooks are used for coding and experimentation, but they are not necessary for running pipelines in the Azure Machine Learning designer.

Key Tip:

Always ensure you have a compute resource (like an Azure ML Compute Instance or Compute Cluster) configured and available before running pipelines in Azure Machine Learning Designer.

### Question: 165

CertyIQ

DRAG DROP

Match the tool to the Azure Machine Learning task.

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

NOTE: Each correct selection is worth one point.

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

### Answer:

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

**Explanation:**

1. Azure portal
2. ML designer
3. Automated ML

<https://learn.microsoft.com/en-us/azure/machine-learning/quickstart-create-resources?view=azureml-api-2>

<https://learn.microsoft.com/en-us/azure/machine-learning/concept-designer?view=azureml-api-2>

Azure Machine Learning designer is a drag-and-drop UI interface to build pipeline in Azure Machine Learning.

**Question: 166****CertyIQ**

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

**Answer: C****Explanation:**

A chatbot is the most suitable AI solution for providing a customer support solution across multiple channels such as email, phone, and live chat. Chatbots can integrate with different communication platforms and provide automated responses to customer queries using pre-defined logic or AI capabilities.

Why the other options are incorrect:

A. Machine learning:

Incorrect: While machine learning can be a component of a chatbot, it is not a direct solution for customer support. ML is a general technique for building predictive models and patterns in data.

B. Computer vision:

Incorrect: Computer vision is used for image and video analysis, such as object detection or image recognition, and is not relevant for a text- or voice-based customer support solution.

D. Natural Language Processing (NLP):

Incorrect: NLP is a technique used by chatbots to understand and process human language. However, it is not the complete solution but rather a core component of a chatbot.

**Key Tip:**

Use chatbots for multi-channel customer support solutions and enhance them with NLP for better understanding of customer queries.

## Question: 167

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.

### Workload Types

Anomaly detection

Computer vision

Machine Learning (Clustering)

Natural language processing

### Answer Area

Workload Type

Identify handwritten letters.

Workload Type

Predict the sentiment of a social media post.

Workload Type

Identify an unusual credit card payment.

Workload Type

Group animals based on multiple measurements.

Answer:

### Answer Area

Computer vision

Identify handwritten letters.

Natural language processing

Predict the sentiment of a social media post.

Anomaly detection

Identify an unusual credit card payment.

Machine Learning (Clustering)

Group animals based on multiple measurements.

Explanation:

1. Computer vision
2. Natural language processing
3. Anomaly detection
4. Machine learning (Clustering)

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

OCR or Optical Character Recognition is also referred to as text recognition or text extraction. Machine-learning based OCR techniques allow you to extract printed or handwritten text from images, such as posters, street signs and product labels, as well as from documents like articles, reports, forms, and invoices. The text is typically extracted as words, text lines, and paragraphs or text blocks, enabling access to digital version of the scanned text. This eliminates or significantly reduces the need for manual data entry.

**Question: 168****CertyIQ**

Predicting how many vehicles will travel across a bridge on a give day is an example of \_\_\_\_\_.

Select the answer that correctly completes the sentence.

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

**Answer: A****Explanation:**

Predicting how many vehicles will travel across a bridge involves forecasting a numerical value, which is the hallmark of a regression problem. Regression is used in machine learning to predict continuous outcomes, such as sales, temperatures, or traffic volume.

Why the other options are incorrect:

B. Translation:

Incorrect: Translation is a Natural Language Processing (NLP) task that converts text from one language to another. It is unrelated to predicting numerical values.

C. Classification:

Incorrect: Classification is used for predicting discrete categories (e.g., "defective" vs. "non-defective" or "yes" vs. "no"). It does not deal with continuous numerical outputs.

D. Clustering:

Incorrect: Clustering groups data into clusters based on similarities but does not predict numerical values. It is used for exploratory data analysis, not forecasting.

Key Tip:

When predicting a numerical value, always consider regression methods such as linear regression, decision trees, or neural networks designed for regression tasks.

**Question: 169****CertyIQ**

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

Select the answer that correctly completes the sentence.

- A. dataset
- B. labels
- C. variables

**Answer: C**



### Explanation:

Labels is indeed the output. Would prefer to say features as well, but features and variables refer to the same thing in the context of machine learning, which is the input data used to train the model. A dataset is a collection of data points, each of which contains one or more features. So would go for C.

In a machine learning model, the input data that is used to predict the output is referred to as variables (or features). These are the independent variables or predictors that the model uses to learn patterns and make predictions.

Why the other options are incorrect:

A. Dataset:

Incorrect: A dataset refers to the entire collection of data, including both input variables and target labels. It is not just the inputs but the whole dataset.

B. Labels:

Incorrect: Labels are the output or target values in supervised learning, representing the result the model is trying to predict. They are not the inputs.

Key Tip:

In machine learning, variables (or features) are the attributes or columns in your dataset that are used to make predictions. These are the inputs that influence the model's output.

### Question: 170

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

### Answer Area

Using Recency, Frequency, and Monetary (RFM) values to identify segments of a customer base is an example of

- clustering.
- regression.
- classification.
- regularization.

Answer:

## Answer Area

Using Recency, Frequency, and Monetary (RFM) values to identify segments of a customer base is an example of

clustering.

regression.

classification.

regularization.

### Explanation:

#### Clustering:

RFM analysis does not involve training a machine learning model to predict a specific outcome, it does involve clustering data based on specific criteria, which makes it more of a form of clustering rather than classification. So would go for clustering.

### Question: 171

CertyIQ

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.

#### Actions

Train the model.

Split the data randomly into training data and validation data.

Evaluate the model against the original dataset.

Evaluate the model against the validation dataset.

Ingest and prepare a dataset.

#### Answer area



### 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.



### Explanation:

1. Ingest and prepare dataset
2. Split data randomly into training and validation data
3. Train model.
4. Evaluate model against validation dataset

<https://learn.microsoft.com/en-us/training/modules/use-automated-machine-learning/5-machine-learning-steps>

## Question: 172

CertyIQ

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.

### Answer Area

Statements	Yes	No
Organizing documents into groups based on different usage statistics 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:

## Answer Area

Statements	Yes	No
Organizing documents into groups based on different usage statistics 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"/>

### Explanation:

YYN is the answer.

<https://learn.microsoft.com/en-us/training/modules/create-clustering-model-azure-machine-learning-designer/2-clustering-scenarios>

Clustering is a form of machine learning that is used to group similar items into clusters based on their features. For example, a researcher might take measurements of penguins, and group them based on similarities in their proportions.

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

Classification is a form of machine learning that is used to predict which category, or class, an item belongs to. This machine learning technique can be applied to binary and multi-class scenarios. For example, a health clinic might use the characteristics of a patient (such as age, weight, blood pressure, and so on) to predict whether the patient is at risk of diabetes. In this case, the characteristics of the patient are the features, and the label is a binary classification of either 0 or 1, representing non-diabetic or diabetic.

### Question: 173

CertyIQ


HOTSPOT

-

Select the answer that correctly completes the sentence.

## Answer Area

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




A dropdown menu with a downward arrow icon. The menu is open, showing four options: 'boolean.', 'datetime.', 'numeric.', and 'text.'.

Answer:

### Answer Area

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



A dropdown menu with a downward arrow icon. The menu is open, showing four options: 'boolean.', 'datetime.', 'numeric.', and 'text.'. The 'numeric.' option is highlighted with a red rectangular box.

Explanation:

**numeric**

<https://learn.microsoft.com/en-us/training/modules/create-regression-model-azure-machine-learning-designer/2-regression-scenarios>

Regression is a form of machine learning used to understand the relationships between variables to predict a desired outcome. Regression predicts a numeric label or outcome based on variables, or features. For example, an automobile sales company might use the characteristics of a car (such as engine size, number of seats, mileage, and so on) to predict its likely selling price. In this case, the characteristics of the car are the features, and the selling price is the label.

## Question: 174

CertyIQ

You need to create a clustering model and evaluate the model by using Azure Machine Learning designer.

What should you do?

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

**Answer: A**

**Explanation:**

Split the original dataset into a dataset for training and a dataset for testing. Use the testing dataset for evaluation.

In machine learning, it's essential to evaluate your model using data it hasn't seen during training to assess its performance accurately. This approach helps in understanding how the model generalizes to new, unseen data.

Correct Answer: A. Split the original dataset into a dataset for training and a dataset for testing. Use the testing dataset for evaluation.

**Explanation:**

Training Dataset: Used to train the model, allowing it to learn patterns and relationships within the data.

Testing Dataset: Used to evaluate the model's performance after training. This dataset should not overlap with the training data to ensure an unbiased assessment of the model's generalization capability.

**Why Not the Other Options?**

B. Use the original dataset for training and evaluation: Using the same dataset for both training and evaluation can lead to overfitting, where the model performs well on the training data but poorly on new, unseen data.

C. Split the original dataset into a dataset for features and a dataset for labels. Use the features dataset for evaluation: This approach doesn't separate the data into training and testing sets, which is crucial for unbiased evaluation.

D. Split the original dataset into a dataset for training and a dataset for testing. Use the training dataset for evaluation: Evaluating the model on the training dataset can result in overly optimistic performance metrics and doesn't provide insight into how the model will perform on new data.

**Key Tip:** Always use a separate testing dataset to evaluate your model's performance to ensure it generalizes well to new, unseen data.

**Question: 175**

**CertyIQ**

You have a dataset that contains the columns shown in the following table.

Name	Type
ColumnA	Integer
ColumnB	Numeric
ColumnC	Numeric
ColumnD	Numeric
ColumnE	Numeric

You have a machine learning model that predicts the value of ColumnE based on the other numeric columns.



Which type of model is this?

- A. analysis
- B. clustering
- C. regression

**Answer: C**

**Explanation:**

regression

<https://learn.microsoft.com/en-us/training/modules/create-regression-model-azure-machine-learning-designer/2-regression-scenarios>

Regression is a form of machine learning used to understand the relationships between variables to predict a desired outcome. Regression predicts a numeric label or outcome based on variables, or features. For example, an automobile sales company might use the characteristics of a car (such as engine size, number of seats, mileage, and so on) to predict its likely selling price. In this case, the characteristics of the car are the features, and the selling price is the label.

### Question: 176

CertyIQ

You need to track multiple versions of a model that was trained by using Azure Machine Learning.

What should you do?

- A. Explain the model.
- B. Register the model.
- C. Register the training data.
- D. Provision an inference cluster.

**Answer: B**

**Explanation:**

Register the model.

To effectively track multiple versions of a model trained using Azure Machine Learning, you should register the model. Model registration allows you to store and version your models within your Azure Machine Learning workspace, facilitating organized management and easy retrieval of different model versions.

Correct Answer: B. Register the model.

Explanation:

- A. Explain the model: This involves interpreting the model's decisions and is not related to version tracking.
- B. Register the model: Correct. Registering the model enables version control and management within Azure Machine Learning.
- C. Register the training data: While registering training data is important for reproducibility, it doesn't directly facilitate version tracking of the model itself.
- D. Provision an inference cluster: This step is related to deploying the model for inference and doesn't pertain

to version tracking.

**Key Tip:** Always register your models in Azure Machine Learning to maintain a clear version history, ensuring reproducibility and effective model management.

### Question: 177

CertyIQ

You need to identify groups of rows with similar numeric values in a dataset.

Which type of machine learning should you use?

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

**Answer: A**

**Explanation:**

clustering

Clustering is a type of unsupervised learning used to group rows of data with similar patterns or numeric values without needing predefined labels. It identifies natural groupings or patterns in the dataset.

Why the other options are incorrect:

B. Regression

Regression is used for predicting a continuous numeric value based on input features. For example, predicting house prices or temperature. It's not suitable for identifying groups of similar rows, as it focuses on finding a relationship between variables, not grouping.

C. Classification

Classification is a supervised learning method used to assign predefined categories or labels to data points. For example, classifying emails as "spam" or "not spam." It requires labeled data, so it's not applicable when grouping similar rows without prior labels.

### Question: 178

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

## Answer Area

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

clustering  
regression  
classification

Answer:

## Answer Area

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

clustering  
regression  
classification

Explanation:

Classification, because there would be an answer like yes or no and not a value like \$10.000

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

Classification is a form of machine learning that is used to predict which category, or class, an item belongs to. This machine learning technique can be applied to binary and multi-class scenarios. For example, a health clinic might use the characteristics of a patient (such as age, weight, blood pressure, and so on) to predict whether the patient is at risk of diabetes. In this case, the characteristics of the patient are the features, and the label is a binary classification of either 0 or 1, representing non-diabetic or diabetic.

Question: 179

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

### Answer Area

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

- Classification
- Clustering
- Regression

### Answer:

#### Answer Area

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

- Classification
- Clustering
- Regression

### Explanation:

<https://learn.microsoft.com/en-us/training/modules/create-regression-model-azure-machine-learning-designer/2-regression-scenarios>

Regression is a form of machine learning used to understand the relationships between variables to predict a desired outcome. Regression predicts a numeric label or outcome based on variables, or features. For example, an automobile sales company might use the characteristics of a car (such as engine size, number of seats, mileage, and so on) to predict its likely selling price. In this case, the characteristics of the car are the features, and the selling price is the label.

### Question: 180

CertyIQ

A historian can use \_\_\_\_\_ to digitize newspaper articles.

Select the answer that correctly completes the sentence.

- A. Object detection
- B. Facial recognition
- C. Image classification
- D. Optical character recognition (OCR)

### Answer: D

### Explanation:

Optical Character Recognition (OCR) is the technology used to digitize printed or handwritten text, such as newspaper articles. It converts images of text into machine-readable text, making it searchable and editable.

Why the other options are incorrect:

- A. Object detection

Object detection identifies and locates objects (like cars, faces, or animals) in an image but doesn't focus on extracting text.

- B. Facial recognition

Facial recognition is specifically designed to identify or verify people's faces, which is unrelated to digitizing

text.

### C. Image classification

Image classification categorizes entire images into predefined labels (e.g., "cat" or "dog") but doesn't extract or digitize text from images.

#### Key Tip:

Use OCR when you need to extract and digitize text from images, such as newspapers, books, or handwritten documents.

## Question: 181

CertyIQ

### 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.

#### 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:

##### 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 checked="" type="radio"/>	<input type="radio"/>
Object detection can identify multiple types of damaged products in an image.	<input checked="" type="radio"/>	<input type="radio"/>

#### Explanation:

YYY

Here are the statements with their corresponding answers:

Object detection can identify the location of a damaged product in an image.

**True:** Object detection can be used to locate and outline the position of a specific object or region in an image, including a damaged product.

Object detection can identify multiple instances of a damaged product in an image.

**True:** Object detection is capable of identifying and locating multiple instances of a specified object, such as multiple damaged products in an image.

Object detection can identify multiple types of damaged products in an image.

**True:** Object detection can identify and classify different types of objects, including multiple types of damaged products in an image, if it has been trained to do so.

## Question: 182

CertyIQ

You need to create a model that labels a collection of your personal digital photographs.

Which Azure Cognitive Services service should you use?

- A. Form Recognizer
- B. Custom Vision
- C. Language
- D. Computer Vision

**Answer: B**

**Explanation:**

B. Custom Vision

With Custom Vision, users can upload their own images and labels, and train a model to recognize specific objects or patterns

Custom Vision is the correct choice because it allows you to train a custom model tailored to label or classify your personal digital photographs. You can define your own categories (labels) and train the model using your specific dataset.

**Why the other options are incorrect:**

A. Form Recognizer

Form Recognizer extracts data from structured documents like forms and invoices. It's not used for labeling photographs.

C. Language

The Language service is designed for natural language processing tasks, such as text analysis or language understanding, not image labeling.

D. Computer Vision

Computer Vision provides general pre-built image analysis capabilities (e.g., detecting objects or describing images), but it doesn't let you create a custom model tailored to your specific labels.

**Key Tip:**

Use Custom Vision when you need to create a tailored image classification model for your unique dataset. Use Computer Vision for general-purpose image analysis..



**Question: 183**

HOTSPOT

-

Select the answer that correctly completes the sentence.

**Answer Area**

▼

Object detection

Image description

Image classification

Optical character recognition (OCR)

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

**Answer:**

**Answer Area**

▼

Object detection

Image description

Image classification

Optical character recognition (OCR)

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

**Explanation:**

Object detection

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

Object detection is similar to tagging, but the API returns the bounding box coordinates (in pixels) for each object found in the image. For example, if an image contains a dog, cat and person, the Detect operation will list those objects 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 object in an image.

**Question: 184**

HOTSPOT

-

Select the answer that correctly completes the sentence.

## Answer Area

Identifying whether a kiosk user is annoyed by monitoring a video feed from the kiosk is an example of

- face detection.
- facial analysis.
- facial recognition.
- optical character recognition (OCR).

Answer:

## Answer Area

Identifying whether a kiosk user is annoyed by monitoring a video feed from the kiosk is an example of

- face detection.
- facial analysis.
- facial recognition.
- optical character recognition (OCR).

Explanation:

Microsoft has retired facial recognition capabilities that can be used to try to infer emotional states and identity attributes which, if misused, can subject people to stereotyping, discrimination or unfair denial of services. These include capabilities that predict emotion, gender, age, smile, facial hair, hair and makeup. Read more about this decision [here](https://azure.microsoft.com/en-us/blog/responsible-ai-investments-and-safeguards-for-facial-recognition/).

<https://azure.microsoft.com/en-us/blog/responsible-ai-investments-and-safeguards-for-facial-recognition/>

## Question: 185

CertyIQ

DRAG DROP

-

Match the Azure Cognitive Services 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 selection is worth one point.

## Services

Custom Vision

Face

Form Recognizer

## Answer Area

Identify objects in an image.

Automatically import data from an invoice to a database.

Identify people in an image.

Answer:

### Answer Area

Custom Vision

Identify objects in an image.

Form Recognizer

Automatically import data from an invoice to a database.

Face

Identify people in an image.

Explanation:

1. Custom Vision
2. Form Recognized
3. Face

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

Azure Custom Vision is an image recognition service that lets you build, deploy, and improve your own image identifier models. An image identifier applies labels to images, according to their visual characteristics. Each label represents a classification or object. Unlike the Computer Vision service, Custom Vision allows you to specify your own labels and train custom models to detect them.

## Question: 186

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

## Answer Area

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

- analysis.
- detection.
- recognition.

Answer:

### Answer Area

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

- analysis.
- detection.
- recognition.

Explanation:

Answer : Analysis

An AI solution that helps photographers to take better pictures by providing feedback on exposure, noise, and occlusion is an example of facial "analysis" or "photo analysis."

In this context, the AI system is analyzing the photo quality and characteristics, such as exposure (brightness and contrast), noise (graininess or pixel-level disturbances), and occlusion (obstructions or unwanted objects). The AI's analysis can then provide feedback to the photographer to improve the composition and settings, resulting in better photographs. While this specific example focuses on photo quality aspects, facial analysis may also involve analyzing and detecting faces within images, recognizing facial expressions, emotions, or attributes, but that doesn't seem to be the focus in this particular case.

## Question: 187

CertyIQ

Your company manufactures widgets.

You have 1,000 digital photos of the widgets.

You need to identify the location of the widgets within the photos.

What should you use?

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

**Answer: B**

**Explanation:**

Custom Vision object detection

Custom Vision Object Detection is the best choice because it can locate and identify specific objects (like widgets) within images by drawing bounding boxes around them. This helps pinpoint their location in the photos.

**Why the other options are incorrect:**

A. Computer Vision Spatial Analysis

Spatial analysis is typically used for analyzing movement and presence in physical spaces (e.g., people counting in a room). It's not meant for identifying specific objects in images.

C. Computer Vision Image Analysis

Image analysis provides a general description of the image (e.g., "This image contains a car and a tree") but does not identify the exact location of objects.

D. Custom Vision Classification

Classification assigns labels to entire images (e.g., "This is an image of a widget") but does not locate objects or provide bounding boxes.

**Key Tip:**

Use Object Detection when you need to identify and locate objects in images. Use Classification for labeling entire images without location details.

### Question: 188

CertyIQ

You need to convert handwritten notes into digital text.

Which type of computer vision should you use?

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

**Answer: B**

**Explanation:**

optical character recognition (OCR)

Optical Character Recognition (OCR) is the correct choice because it is designed to convert handwritten or printed text into digital, machine-readable text. It's widely used for digitizing documents, notes, and forms.

**Why the other options are incorrect:**

A. Facial detection

Facial detection identifies and locates human faces in images but does not process or convert text.

### C. Image classification

Image classification assigns labels to entire images based on their content (e.g., "dog" or "cat") but does not recognize or extract text.

### D. Object detection

Object detection identifies and locates objects within images (e.g., a car or a chair) but is not suitable for text recognition.

#### Key Tip:

Use OCR for tasks involving text extraction from handwritten or printed sources, such as notes, books, or forms.

## Question: 189

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

#### Answer Area

Image classification

Image description

Object detection

Optical character recognition (OCR)

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

#### Answer:

#### Answer Area

Image classification

Image description

Object detection

Optical character recognition (OCR)

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

#### Explanation:

Object detection

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

Object detection is similar to tagging, but the API returns the bounding box coordinates (in pixels) for each



object found in the image. For example, if an image contains a dog, cat and person, the Detect operation will list those objects 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 object in an image.

### Question: 190

CertyIQ

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

**Answer: D**

**Explanation:**

optical character recognition (OCR)

Optical Character Recognition (OCR) is the correct choice because it can scan and extract text from receipts or invoices (e.g., amounts, dates, item descriptions) and convert it into digital, machine-readable text to store in the app.

**Why the other options are incorrect:**

A. Face detection

Face detection identifies and locates human faces in images but is unrelated to scanning expenses or extracting text.

B. Image classification

Image classification categorizes entire images (e.g., "receipt" or "not a receipt") but does not extract or digitize the text from the receipt.

C. Object detection

Object detection identifies and locates objects in an image (e.g., a car or a phone) but cannot recognize or extract text from receipts.

**Key Tip:**

Use OCR whenever you need to scan, extract, and digitize text from physical documents like receipts, invoices, or forms.

### Question: 191

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

Answer Area

You can use the  service to train an object detection model by using your own images.

Computer Vision

Custom Vision

Form Recognizer

Azure Video Analyzer for Media

Answer:

Answer Area

You can use the  service to train an object detection model by using your own images.

Computer Vision

Custom Vision

Form Recognizer

Azure Video Analyzer for Media

Explanation:

Right answer is Custom vision.

Question: 192

CertyIQ

HOTSPOT

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

Answer Area

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

analysis.

detection.

recognition.

Answer:

Answer Area

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

analysis.

detection.

recognition.

Explanation:

Answer : Analysis

An AI solution that helps photographers to take better pictures by providing feedback on exposure, noise, and

occlusion is an example of facial "analysis" or "photo analysis."

In this context, the AI system is analyzing the photo quality and characteristics, such as exposure (brightness and contrast), noise (graininess or pixel-level disturbances), and occlusion (obstructions or unwanted objects). The AI's analysis can then provide feedback to the photographer to improve the composition and settings, resulting in better photographs. While this specific example focuses on photo quality aspects, facial analysis may also involve analyzing and detecting faces within images, recognizing facial expressions, emotions, or attributes, but that doesn't seem to be the focus in this particular case.

## Question: 193

CertyIQ

DRAG DROP

Match the Azure Cognitive Services to the appropriate AI workloads.

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

NOTE: Each correct match is worth one point.

### Services

Custom Vision

Face

Form Recognized

### Answer Area

Identify objects in an image.

Automatically import data from an invoice to a database.

Identify people in an image.

### Answer:

#### Answer Area

Custom Vision

Identify objects in an image.

Form Recognized

Automatically import data from an invoice to a database.

Face

Identify people in an image.

### Explanation:

1. Custom Vision
2. Form Recognized

### 3. Face

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

Azure Custom Vision is an image recognition service that lets you build, deploy, and improve your own image identifier models. An image identifier applies labels to images, according to their visual characteristics. Each label represents a classification or object. Unlike the Computer Vision service, Custom Vision allows you to specify your own labels and train custom models to detect them.

#### Question: 194

CertyIQ

You need to implement a pre-built solution that will identify well-known brands in digital photographs.

Which Azure Cognitive Services service should you use?

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

**Answer: D**

**Explanation:**

Computer Vision

Computer Vision is the correct choice because it provides pre-built functionality to analyze images and identify well-known brands, logos, or objects in digital photographs. It includes capabilities like image tagging, brand detection, and general image analysis.

**Why the other options are incorrect:**

A. Custom Vision

Custom Vision is used for creating custom models to classify or detect objects specific to your needs. It requires training with your dataset and is not pre-built for brand recognition.

B. Form Recognizer

Form Recognizer is designed to extract data from structured documents like forms or receipts, but it does not analyze photographs or identify brands.

C. Face

The Face service is specifically for detecting and analyzing human faces, such as identifying facial features or verifying identities. It cannot recognize brands or logos.

**Key Tip:**

Use Computer Vision for pre-built solutions like brand/logo recognition, image tagging, and content description without the need for custom training.

**Question: 195**

Natural language processing can be used to \_\_\_\_\_.

Select the answer that correctly completes the sentence.

- A. Analyze video content
- B. Generate speech
- C. Classify email messages as work-related or personal.
- D. Classify images

**Answer: C**

**Explanation:**

Classify email messages as work-related or personal.

Natural Language Processing (NLP) is a field of AI that enables machines to understand, interpret, and process human language. It can classify email messages based on their content (e.g., work-related or personal), making C the correct choice.

**Why the other options are incorrect:**

A. Analyze video content

Analyzing video content falls under computer vision, not NLP, as it deals with visual data rather than text or language.

B. Generate speech

Generating speech is related to speech synthesis (text-to-speech), which is a different domain of AI focused on converting text to audio, not text analysis.

D. Classify images

Classifying images is a task for computer vision, which analyzes visual data, not language or text.

**Key Tip:**

Use NLP for tasks involving language processing, such as text classification, sentiment analysis, and language translation. For images or video, rely on computer vision techniques.

**Question: 196**

You plan to develop a bot that will enable users to query a knowledge base by using natural language processing.

Which two services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

**Answer: AB**

**Explanation:**

A. Language Service

B. Azure Bot Service

To develop a bot that enables users to query a knowledge base using natural language processing, you need the following services:

A. Language Service

The Language Service (formerly known as Text Analytics or QnA Maker) processes and understands natural language queries. It enables querying a knowledge base by interpreting user input in natural language.

B. Azure Bot Service

Azure Bot Service provides the framework and tools to build, deploy, and manage the bot itself, allowing users to interact with the bot via chat or other interfaces.

**Why the other options are incorrect:**

C. Form Recognizer

Form Recognizer is designed to extract information from structured documents like forms or invoices, not to process natural language queries or enable bot interactions.

D. Anomaly Detector

Anomaly Detector is used for identifying anomalies in time-series data, which is unrelated to bots or natural language processing.

**Key Tip:**

Use Language Service for natural language understanding and Azure Bot Service for building conversational bots. These two services work together to create intelligent, user-friendly bot solutions.

## Question: 197

CertyIQ

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.

**Answer Area**

**Statements**

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

**Yes**

☐

**No**

☐

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

☐

☐

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

☐

☐



## Answer:

### Answer Area

#### Statements

- The following service call will accept English text as an input and output Italian and French text.  
/translate?from=it&to=fr&to=en
- The following service call will accept English text as an input and output Italian and French text.  
/translate?from=en&to=fr&to=it
- The Translator service can be used to translate documents from English to French.

Yes

No

☐☒☒☐☒☐

## Explanation:

NYN is the answer.

<https://learn.microsoft.com/en-us/azure/cognitive-services/translator/document-translation/overview>

Document Translation is a cloud-based feature of the Azure Translator service and is part of the Azure Cognitive Service family of REST APIs. The Document Translation API can be used to translate multiple and complex documents across all supported languages and dialects, while preserving original document structure and data format.

## Question: 198

CertyIQ

An app that analyzes social media posts to identify their tone is an example of which type of natural language processing (NLP) workload?

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

## Answer: A

### Explanation:

sentiment analysis

Sentiment Analysis is the correct answer because it determines the tone or emotional intent of a text, such as whether a social media post expresses positivity, negativity, or neutrality. This is a common application of NLP in analyzing user opinions and attitudes.

### Why the other options are incorrect:

B. Speech recognition

Speech recognition converts spoken words into text. It does not analyze the tone or sentiment of written posts.

C. Key phrase extraction

Key phrase extraction identifies important phrases or keywords in a text (e.g., "product quality," "delivery time"). It doesn't focus on the tone or emotional intent.

D. Entity recognition

Entity recognition identifies specific entities like names, dates, or locations in a text (e.g., "Microsoft," "New York"). It doesn't assess the tone or sentiment.

**Key Tip:**

Use sentiment analysis to evaluate the emotional tone in text, such as social media posts, reviews, or customer feedback.

**Question: 199****CertyIQ**

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. named entity recognition
- C. translation
- D. language modeling

**Answer: B****Explanation:**

named entity recognition

Named Entity Recognition (NER) is an NLP technique used to identify and extract specific entities in text, such as names, dates, product names, or order IDs. In this case, NER would be appropriate for recognizing entities like:

Customer orders: Identify product names, quantities, or order details.

Support documents: Recognize document titles or topics.

Order status updates: Extract order IDs or customer names from user input.

Why the other options are incorrect:

**A. Sentiment Analysis:**

Sentiment analysis identifies the emotional tone (positive, negative, or neutral) in text, which is not relevant to the actions described.

**C. Translation:**

Translation converts text from one language to another. It does not help with understanding and acting on user input.

**D. Language Modeling:**

Language modeling predicts the next word or sequence of words in a sentence. While useful for chatbot generation, it is not the primary technique for extracting specific entities like order details.

**Key Tip:**

Use Named Entity Recognition when your chatbot needs to identify and act on specific elements (e.g., names, order IDs, or document titles) from user input.

**Question: 200**

CertyIQ

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 selection is worth one point.

**Azure Cognitive Services**

Language service

Speech

Translator

**Answer Area**

Convert spoken requests into text.

Identify the intent of a user's requests.

Apply intent to entities and utterances.

**Answer:**

**Azure Cognitive Services**

Language service

Speech

Translator

**Answer Area**

Speech

Convert spoken requests into text.

Language service

Identify the intent of a user's requests.

Language service

Apply intent to entities and utterances.

**Explanation:**

### 1. Speech → Convert spoken requests into text

- The Speech service is used for speech-to-text functionality, which converts spoken words into textual data.

### 2. Language service → Identify the intent of a user's requests

- The Language service (formerly known as LUIS - Language Understanding Intelligent Service) is designed to identify the intent behind user inputs (e.g., "Book a flight").

### 3. Language service → Apply intent to entities and utterances

- The Language service also applies the identified intent to specific entities and utterances to provide more context for the request.

## Question: 201

CertyIQ

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.

#### Answer Area

Statements	Yes	No
A webchat bot can interact with users visiting a website.	<input type="radio"/>	<input type="radio"/>
Automatically generating captions for pre-recorded videos is an example of natural language processing.	<input type="radio"/>	<input type="radio"/>
A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of natural language processing.	<input type="radio"/>	<input type="radio"/>

#### Answer:

Statements	Yes	No
A webchat bot can interact with users visiting a website.	<input checked="" type="radio"/>	<input type="radio"/>
Automatically generating captions for pre-recorded videos is an example of natural language processing.	<input checked="" type="radio"/>	<input type="radio"/>
A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of natural language processing.	<input checked="" type="radio"/>	<input type="radio"/>

#### Explanation:

1. Yes

2. Yes

3. Yes

1. A webchat bot can interact with users visiting a website:

Yes, a webchat bot can be implemented on a website to provide automated responses and interact with users in a conversational manner.

2. Speech and Conversational AI workloads come under NLP.

3. A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of natural language processing:

Yes, a smart device that can understand and respond to natural language queries or commands, such as asking about the weather, relies on natural language processing techniques to interpret and process the user's input and provide a relevant response. (Chat GPT)

### Question: 202

CertyIQ

You have a website that includes customer reviews.

You need to store the reviews in English and present the reviews to users in their respective language by recognizing each user's geographical location.

Which type of natural language processing workload should you use?

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

**Answer: D**

**Explanation:**

This should be translation. for this specific use case, the best option would be the Azure Translator Text API.

The Translator Text API is a cloud-based machine translation service that can translate text between multiple languages in real-time. It supports a wide range of languages and provides language detection capabilities that can automatically identify the language of the input text. It can also detect the user's location based on their IP address or browser settings and automatically translate the reviews to the user's preferred language.

### Question: 203

CertyIQ

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.

## Answer Area

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

Answer:

## Answer Area

Statements	Yes	No
Chatbots can support voice input.	<input checked="" type="radio"/>	<input type="radio"/>
A separate chatbot is required for each communication channel.	<input type="radio"/>	<input checked="" type="radio"/>
Chatbots manage conversation flows by using a combination of natural language and constrained option responses.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

<https://learn.microsoft.com/en-us/azure/bot-service/bot-service-overview?view=azure-bot-service-4.0#connect>

Bot Framework does most of the work necessary to send and receive messages from all of these different platforms — your bot application receives a unified, normalized stream of messages regardless of the number and type of channels it's connected to.

Question: 204

CertyIQ

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.

## Answer Area

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

Answer:

## Answer Area

Statements	Yes	No
A bot that responds to queries by internal users is an example of a natural language processing workload.	<input checked="" type="radio"/>	<input type="radio"/>
A mobile application that displays images relating to an entered search term is an example of a natural language processing 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 natural language processing workload.	<input type="radio"/>	<input checked="" type="radio"/>

### Explanation:

YNN. Second one is Azure cognitive search

1 ->A bot that responds to queries by internal users is an example of a natural language processing (NLP) workload.

Correct Answer: Yes

This involves understanding and generating human language, which is the core of NLP. Responding to queries requires processing and interpreting language.

2 ->A mobile application that displays images relating to an entered search term is an example of a natural language processing workload.



Correct Answer: No

While the app accepts a search term (text), it focuses on retrieving and displaying images, not processing or analyzing the language itself. This is more aligned with search engine or image processing technologies, not NLP.

3 -> A web form used to submit a request to reset a password is an example of a natural language processing workload.

Correct Answer: No

Submitting a password reset request typically involves filling out a form and does not involve any language processing or understanding. It's a simple data input/output task.

**Important Tip:**

When identifying NLP workloads, look for tasks involving understanding, interpreting, generating, or translating human language. Examples include chatbots, translation tools, sentiment analysis, and text summarization.

**Question: 205**

**CertyIQ**

You have a solution that analyzes social media posts to extract the mentions of city names and the city names discussed most frequently.

Which type of natural language processing (NLP) workload does the solution use?

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

**Answer: D**

**Explanation:**

Entity recognition is the correct answer because it identifies and extracts specific entities, such as city names, from a text. This workload allows you to detect predefined categories like locations, dates, or people mentioned in social media posts.

**Why the other options are incorrect:**

A. Speech recognition

Speech recognition converts spoken words into text but does not analyze or extract entities from the text.

B. Sentiment analysis

Sentiment analysis determines the emotional tone of text (e.g., positive, negative, or neutral) but does not extract specific entities like city names.

C. Key phrase extraction

Key phrase extraction identifies important phrases in the text but does not classify or specifically extract entities like city names.

**Key Tip:**

Use entity recognition for extracting specific items like names, places, dates, or other defined entities from text. This is essential when analyzing content for structured information.

**Question: 206**

CertyIQ

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.

**Answer Area**

Statements	Yes	No
You can use Language Service's question answering to query an Azure SQL database.	<input type="radio"/>	<input type="radio"/>
You should use Language Service's question answering when you want a knowledge base to provide the same answer to different users who submit similar questions.	<input type="radio"/>	<input type="radio"/>
Language Service's question answering can determine the intent of a user utterance.	<input type="radio"/>	<input type="radio"/>

**Answer:**

**Answer Area**

Statements	Yes	No
You can use Language Service's question answering to query an Azure SQL database.	<input type="radio"/>	<input checked="" type="radio"/>
You should use Language Service's question answering when you want a knowledge base to provide the same answer to different users who submit similar questions.	<input checked="" type="radio"/>	<input type="radio"/>
Language Service's question answering can determine the intent of a user utterance.	<input type="radio"/>	<input checked="" type="radio"/>

**Explanation:**

No. Language Service's question answering is primarily designed to extract answers from pre-defined knowledge bases or documents, not directly query databases.

Yes. Language Service's question answering can provide consistent answers to similar questions by leveraging the knowledge base and understanding the intent of the user's question.

No. Language Service's question answering is focused on extracting answers from text and documents, rather than determining the intent of user utterances. Intent recognition is typically handled by other NLP components or services

**Question: 207**

You are developing a solution that uses the Language service.

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

Which type of natural language processing should you use?

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

**Answer: D**

**Explanation:**

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

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

Reference:

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

**Question: 208**

DRAG DROP

-

You are designing a system that will generate insurance quotes automatically.

Match the Microsoft responsible AI principles to the appropriate requirements.

To answer, drag the appropriate principle 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.

NOTE: Each correct match is worth one point.

**Principles****Answer Area**

Accountability

Fairness

Inclusiveness

Privacy and security

Reliability and safety

Transparency

A customer's personal information must be visible only to staff who are involved in the decision-making process:

The decision-making process must be recorded so that staff can identify the reasoning behind a particular quote:

The system must be accessible to customers who use screen readers or other assistive technology:

Answer:

### Principles

### Answer Area

Accountability

Fairness

Inclusiveness

Privacy and security

Reliability and safety

Transparency

A customer's personal information must be visible only to staff who are involved in the decision-making process:

The decision-making process must be recorded so that staff can identify the reasoning behind a particular quote:

The system must be accessible to customers who use screen readers or other assistive technology:

Privacy and security

Transparency

Inclusiveness

### Explanation:

Privacy and security: A customer's personal information must be visible only to staff who are involved in the decision-making process to ensure the privacy and security of sensitive data. Transparency: The decision-making process must be recorded so that staff can identify the reasoning behind a particular quote, promoting transparency and accountability. Inclusiveness: The system must be accessible to customers who use screen readers or other assistive technology, ensuring inclusiveness and providing equal access to all users.

### Question: 209

CertyIQ

Which type of natural language processing (NLP) entity is used to identify a phone number?

- A.regular expression
- B.machine-learned
- C.list
- D.Pattern.any

Answer: A

### Explanation:

A is the answer.

<https://learn.microsoft.com/en-us/azure/cognitive-services/luis/reference-entity-regular-expression?tabs=V2A> regular expression entity extracts an entity based on a regular expression pattern you provide.

IMHO Regular expression is the correct answer. Regular expressions are commonly used to define patterns for identifying phone numbers. By defining a regular expression pattern that matches the format of a phone number, you can effectively extract phone numbers from text using NLP techniques.

### Question: 210

CertyIQ

HOTSPOT

-

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

## Answer Area

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

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

Answer:

## 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

Explanation:

**Correct Answer: Object Detection**

**Why?** Object detection involves identifying the presence of objects (like vehicles) in an image and specifying their locations, often by drawing bounding boxes. This is exactly what the question describes.

**Why Other Options Are Incorrect:**

**Image Classification:** This assigns a label to the entire image (e.g., "vehicle"), without locating it in the image. It doesn't involve bounding boxes.

**Optical Character Recognition (OCR):** OCR extracts text from images, not objects like vehicles.

**Facial Detection:** This specifically identifies human faces in an image, not vehicles or other objects.

**Important Tip:**

**Key Difference:** Object detection provides both identification and localization (bounding boxes), whereas image classification only identifies the presence of an object in the entire image.

## Question: 211

CertyIQ

Your company is exploring the use of voice recognition technologies in its smart home devices. The company wants to identify any barriers that might unintentionally leave out specific user groups.

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

- A.accountability
- B.fairness

- C.privacy and security
- D.inclusiveness

**Answer: D**

**Explanation:**

inclusiveness is a correct answer.

The scenario involves ensuring that voice recognition technologies do not unintentionally exclude specific user groups. This aligns with the principle of Inclusiveness.

Correct Answer: D. Inclusiveness

Why? Inclusiveness focuses on designing AI systems that are accessible and usable by people of all backgrounds, abilities, and demographics. Ensuring no user group is unintentionally left out directly reflects this principle.

**Why Other Options Are Incorrect:**

A. Accountability:

Accountability refers to ensuring that AI systems have clear responsibilities and oversight mechanisms. While important, this scenario is more about inclusive design than assigning responsibility.

B. Fairness:

Fairness is about mitigating bias and ensuring equitable outcomes for all users. While related, the scenario is broader in scope, addressing inclusivity rather than focusing solely on fairness.

C. Privacy and Security:

This principle pertains to safeguarding user data and ensuring the confidentiality of personal information. It is not the primary concern in this context.

**Important Tip:**

Inclusiveness ensures that AI systems are designed to accommodate the widest range of users, making them accessible and equitable for all. This principle is critical in diverse and global applications like smart home technologies.

**Question: 212**

**CertyIQ**

HOTSPOT

-

You have a large dataset that contains motor vehicle sales data.

You need to train an automated machine learning (automated ML) model to predict vehicle sale values based on the type of vehicle.

Which task should you select? To answer, select the appropriate task in the answer area.

NOTE: Each correct selection is worth one point.



## Answer Area

### Select task and settings

Select the machine learning task type for the experiment. To fine tune the experiment, choose additional configuration or featurization settings.



#### Classification

To predict one of several categories in the target column. yes/no, blue, red, green.



#### Regression

To predict continuous numeric values.



#### Time series forecasting

To predict values based on time.



#### Natural Language Processing (preview)

Predict based on text-only data types using multi-class or multi-label classification.



#### Computer Vision (preview)

Multi-class or multi-label image classification, object detection, and instance segmentation.

 [View additional configuration settings](#)  [View featurization settings](#)

Answer:

## Answer Area

### Select task and settings

Select the machine learning task type for the experiment. To fine tune the experiment, choose additional configuration or featurization settings.



#### Classification

To predict one of several categories in the target column. yes/no, blue, red, green.



#### Regression

To predict continuous numeric values.



#### Time series forecasting

To predict values based on time.



#### Natural Language Processing (preview)

Predict based on text-only data types using multi-class or multi-label classification.



#### Computer Vision (preview)

Multi-class or multi-label image classification, object detection, and instance segmentation.

 [View additional configuration settings](#)  [View featurization settings](#)



### Explanation:

The scenario involves predicting **vehicle sale values**, which are **continuous numeric values**. This aligns with the **Regression** task.

### Correct Task Selection:

#### Regression:

**Why?** Regression is used to predict continuous numeric values, such as prices, sales, or other measurements. Since vehicle sale values are numeric, this is the appropriate choice.

### Why Other Options Are Incorrect:

#### Classification:

Used for predicting categories or labels (e.g., "Car Type: SUV, Sedan, Truck"), not continuous values.

#### Time Series Forecasting:

Used for predicting values over time (e.g., future sales trends), but the scenario does not specify a time-based prediction.

#### Natural Language Processing (NLP):

Used for text-based tasks (e.g., sentiment analysis, text classification), irrelevant to predicting numeric values.

#### Computer Vision:

Used for image-based tasks like object detection or image classification, unrelated to the scenario.

### Important Tip:

Always match the task type to the nature of the prediction target:

**Continuous numeric values → Regression**

**Categorical values → Classification**

**Time-based trends → Time Series Forecasting**

### Question: 213

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

### Answer Area

When evaluating the performance of a model, the

AUC metric  
confusion matrix  
ROC curve  
threshold

displays the predicted and actual positives and negatives by using a grid of 0 and 1 values.

Answer:

## Answer Area

When evaluating the performance of a model, the

▼

AUC metric

confusion matrix

ROC curve

threshold

displays the predicted and actual positives and negatives by using a grid of 0 and 1 values.

### Explanation:

confusion matrix

<https://learn.microsoft.com/en-us/azure/machine-learning/how-to-understand-automated-ml?view=azureml-api-2#confusion-matrix-for-a-good-model>

## Question: 214

CertyIQ

You need to convert receipts into transactions in a spreadsheet. The spreadsheet must include the date of the transaction, the merchant, the total spent, and any taxes paid.

Which Azure AI service should you use?

- A. Custom Vision
- B. Form Recognizer
- C. Face
- D. Language

### Answer: B

#### Explanation:

The task involves extracting structured data (such as date, merchant, total spent, and taxes) from receipts and converting it into a spreadsheet. This aligns with the capabilities of Azure Form Recognizer.

Correct Answer: B. Form Recognizer

Why? Azure Form Recognizer is specifically designed to extract structured data from forms, receipts, and invoices. It can identify key-value pairs and tables in scanned or digital documents, making it ideal for converting receipt data into a spreadsheet.

#### Why Other Options Are Incorrect:

A. Custom Vision:

Used for image classification and object detection tasks, not for extracting structured text or data from documents.

C. Face:

Focused on facial recognition tasks such as identifying or verifying individuals, which is unrelated to this scenario.

D. Language:

Used for text-based natural language processing tasks like sentiment analysis, translation, or entity recognition. While it deals with language, it doesn't specialize in extracting structured data from forms.

**Important Tip:**

Use Azure Form Recognizer whenever you need to extract and analyze data from structured or semi-structured documents like receipts, invoices, or business forms.

**Question: 215****CertyIQ**

HOTSPOT

-

Select the answer that correctly completes the sentence.

**Answer Area**

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

	▼
classification	
clustering	
regression	

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

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

	▼
classification	
clustering	
regression	

**Explanation:**

Predicting how many vehicles will travel across a bridge on a given day is an example of regression because it involves predicting a numeric value (the number of vehicles).

**Why the other options are incorrect:**

Classification

Classification is used to predict categories or labels (e.g., "low traffic" vs. "high traffic"), not numeric values.

Clustering

Clustering groups data points based on similarity but does not predict values. It's unsupervised learning used for finding patterns or groups in data.

**Key Tip:**

Use regression when predicting continuous numeric values (e.g., temperature, sales, or traffic volume).

**Question: 216**

CertyIQ

HOTSPOT

-

Select the answer that correctly completes the sentence.

**Answer Area**

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

▼

features  
functions  
labels  
instances

Answer:

**Answer Area**

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

▼

features  
functions  
labels  
instances

**Explanation:**

Correct answer is Features.

The inputs of a machine learning model are called features.

Reference:

<https://learn.microsoft.com/en-us/dotnet/machine-learning/how-does-ml-dotnet-work#basic>

**Question: 217**

CertyIQ

You have a security system that analyzes images from CCTV to provide authorized staff entry into restricted area.

Which type of computer vision does the system use?

- A. optical character recognition (OCR)
- B. semantic segmentation
- C. facial detection and facial recognition
- D. image analysis

**Answer: C**

**Explanation:**

facial detection and facial recognition.

The security system analyzes images to determine whether the staff is authorized to enter a restricted area. This process involves identifying (facial detection) and verifying (facial recognition) individuals, which aligns with the option:

Correct Answer: C. Facial detection and facial recognition

Why?

Facial detection identifies the presence of a face in the image.

Facial recognition verifies or identifies the individual by comparing the detected face with stored authorized records.

This combination is specifically designed for systems like security authentication and access control.

**Why Other Options Are Incorrect:**

A. Optical Character Recognition (OCR):

OCR is used to extract text from images, such as reading license plates or documents. It is not relevant to identifying or verifying faces.

B. Semantic Segmentation:

Semantic segmentation is used for pixel-level categorization in images, such as segmenting objects in an image. It doesn't involve face detection or recognition.

D. Image Analysis:

Image analysis is a broader category that involves understanding image content (e.g., colors, objects, etc.). While it may include face detection, it does not specifically handle face verification or recognition.

**Important Tip:**

Facial detection and recognition are widely used in access control systems, ensuring security by verifying identities in real-time. Always associate these terms with applications like surveillance, authentication, and biometric verification.

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

**Answer: AE**

**Explanation:**

A. assigning the color pixels in an image to object names

E. describing the contents of an image

Computer vision involves processing and analyzing images to extract meaningful information. Let's evaluate each option to determine the correct answers.

Correct Answers:

A. Assigning the color pixels in an image to object names

Why? This is an example of semantic segmentation, a computer vision task where each pixel in an image is classified as part of a specific object or category.

E. Describing the contents of an image

Why? This is an example of image analysis, where computer vision systems generate descriptions or captions for the contents of an image, such as identifying objects, scenes, or activities.

**Why Other Options Are Incorrect:**

B. Detecting inconsistencies and anomalies in a stream of data

This is not related to computer vision. It falls under data analytics or anomaly detection in time-series data.

C. Creating visual representations of numerical data

This is related to data visualization, not computer vision. Tasks like creating graphs or charts involve visualization tools, not image analysis.

D. Creating photorealistic images by using three-dimensional models

This relates to 3D rendering or computer graphics, not computer vision. Computer vision analyzes images, while 3D rendering generates images.

**Important Tip:**

Computer vision focuses on understanding and interpreting visual data from images or videos. Common tasks include object detection, image classification, semantic segmentation, and generating image descriptions.

## Question: 219

You have an app that identifies the coordinates of a product in an image of a supermarket shelf.

Which service does the app use?

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

**Answer: B**

### Explanation:

Custom Vision object detection is a correct answer.

The app identifies the coordinates of a product in an image of a supermarket shelf. This means it detects the product and provides its location in the image, typically as bounding boxes. This task aligns with object detection.

Correct Answer: B. Custom Vision object detection

Why?

Object detection identifies objects in an image and provides their coordinates (bounding boxes).

Custom Vision object detection is specifically designed to detect and locate objects based on custom-trained models.

### Why Other Options Are Incorrect:

A. Custom Vision classification:

Classification assigns a label to an entire image (e.g., "contains product X") but does not identify object locations or coordinates.

C. Computer Vision Read:

The "Read" feature in Computer Vision is used for extracting text from images, not detecting objects or providing coordinates.

D. Computer Vision optical character recognition (OCR):

OCR extracts text from images (e.g., price tags or product labels), but it does not detect objects or their locations.

### Important Tip:

For tasks involving detecting and localizing objects in an image, use object detection. If the task is about identifying objects without location or handling text extraction, other services like classification or OCR are more appropriate.



## Question: 220

HOTSPOT

-

Select the answer that correctly completes the sentence.

### Answer Area


A traffic monitoring system that collects vehicle registration numbers from CCTV footage is an example of  in the Computer Vision service.

image classification  
object detection  
spatial Analysis  
text extraction

Answer:

### Answer Area


A traffic monitoring system that collects vehicle registration numbers from CCTV footage is an example of  in the Computer Vision service.

image classification  
object detection  
spatial Analysis  
text extraction

Explanation:

the answer is text extraction Since the goal is to extract the vehicle registration numbers from the CCTV footage, text extraction using OCR algorithms would be the most appropriate choice. OCR algorithms can recognize and extract text information from images, making it possible to retrieve the registration numbers from the captured video frames. This enables automated monitoring and analysis of the traffic data by extracting the relevant textual information.

A traffic monitoring system that collects vehicle registration numbers from CCTV footage is an example of text extraction because the system identifies and extracts the alphanumeric text (license plate numbers) from the images.

**Why the other options are incorrect:**

Image classification

Image classification categorizes entire images into predefined classes (e.g., "car" or "truck") but does not extract text.

Object detection

Object detection identifies and locates objects (e.g., cars) in an image but does not handle text extraction from those objects.

Spatial Analysis

Spatial analysis is used to analyze movement or presence in physical spaces (e.g., people counting) and is not designed for extracting text.

**Key Tip:**

Use text extraction (OCR) when dealing with tasks involving reading and digitizing text from images or videos, such as license plates or documents.

**Question: 221**

**CertyIQ**

You need to build an image tagging solution for social media that tags images of your friends automatically.

Which Azure Cognitive Services service should you use?

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

**Answer: A**

**Explanation:**

Face is correct.

A. Face is the correct answer because the Face service is specifically designed to detect and recognize human faces in images. It can match faces to a database of your friends and automatically tag their names in the images.

**Why the other options are incorrect:**

**B. Form Recognizer**

Form Recognizer is used for extracting information from structured documents like forms or receipts. It's not relevant for tagging images of people.

**C. Language**

The Language service is designed for natural language processing tasks, such as analyzing or processing text. It doesn't handle image or face recognition.

**D. Computer Vision**

While Computer Vision provides general image analysis (e.g., object detection, scene description), it doesn't specialize in identifying and tagging specific faces. For facial recognition, the Face service is required.

**Key Tip:**

Use the Face service when you need to detect, recognize, or analyze human faces, such as tagging friends in


**Question: 222**

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.


**Answer Area**

A historian can use  to digitize newspaper articles

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

**Answer:**

**Answer Area**

A historian can use  to digitize newspaper articles

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

**Explanation:**

Optical Character recognition(OCR).

A historian can use optical character recognition (OCR) to digitize newspaper articles because OCR is a technology designed to extract text from scanned documents or images. It converts printed or handwritten text into machine-readable digital text.

**Why the other options are incorrect:**

Facial analysis

Facial analysis is used to detect and analyze human faces (e.g., identifying age or emotions). It does not extract text from documents.

Image classification

Image classification categorizes entire images based on their content (e.g., "newspaper" or "book"). It does not extract text.

## Object detection

Object detection identifies objects (e.g., cars or buildings) within an image, but it is not used for extracting or processing text.

### Key Tip:

Use OCR for digitizing printed or handwritten text from newspapers, documents, or historical archives.

## Question: 223

CertyIQ

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.

## Answer Area

Locate the birds:

Automated captioning  
Image classification  
Object detection  
Optical character recognition (OCR)

Identify the species of the birds:

Automated captioning  
Image classification  
Object detection  
Optical character recognition (OCR)

Answer:

## Answer Area

Locate the birds:

Automated captioning  
Image classification  
**Object detection**  
Optical character recognition (OCR)

Identify the species of the birds:

Automated captioning  
**Image classification**  
Object detection  
Optical character recognition (OCR)

### Explanation:

Locate the Birds: Object detection.

Identify the Species of the birds: Image classification.

### Explanation for the Question:

1. Locate the birds:

Correct Answer: Object detection

Object detection identifies and locates objects within an image, such as birds, by drawing bounding boxes around them. This allows the system to find where the birds are in the image.

2. Identify the species of the birds:

Correct Answer: Image classification

Image classification categorizes an entire image into predefined classes (e.g., different bird species) based on its content. It does not locate the birds but determines what kind of bird species they are.

### Why the other options are incorrect:

Automated captioning:

Automated captioning generates descriptive text for an image, but it doesn't locate or classify specific objects like birds.

Optical character recognition (OCR):

OCR extracts text from images but is not applicable for locating or identifying birds.

**Key Tip:**

Use object detection for locating objects in images.

Use image classification for identifying or categorizing objects within an image.

**Question: 224****CertyIQ**

You have a solution that reads manuscripts in different languages and categorizes the manuscripts based on topic.

Which types of natural language processing (NLP) workloads does the solution use?

- A. speech recognition and entity recognition
- B. speech recognition and language modeling
- C. translation and key phrase extraction
- D. translation and sentiment analysis

**Answer: C****Explanation:**

translation and key phrase extraction.

speech recognition would be required only if we were dealing with audio input.

The correct answer is C. Translation and key phrase extraction because:

Translation is needed to read manuscripts in different languages and convert them into a common language for further processing.

Key phrase extraction is used to identify important terms or topics within the text, enabling the categorization of the manuscripts by topic.

**Why the other options are incorrect:**

- A. Speech recognition and entity recognition

Speech recognition converts spoken language into text, which is irrelevant since manuscripts are already written.

Entity recognition focuses on extracting specific entities like names or locations, not categorizing topics.

- B. Speech recognition and language modeling

Speech recognition is not relevant as the task involves manuscripts (written text).

Language modeling predicts the likelihood of text sequences but does not categorize text by topic.

- D. Translation and sentiment analysis

Translation is correct, but sentiment analysis measures the emotional tone (positive/negative) of the text, which is unrelated to categorizing topics.

**Key Tip:**

For tasks involving text in multiple languages and categorization by topic, focus on translation for language standardization and key phrase extraction to identify the main themes.

### Question: 225

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.

## 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:

## 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.

**Explanation:**

The interactive answering of questions entered by a user as part of an application is an example of natural language processing (NLP). NLP is used to understand, process, and respond to human language in a meaningful way, which is the core functionality of interactive question answering.

**Why the other options are incorrect:**

Anomaly detection:

This is used to identify unusual patterns or deviations in data, such as fraud detection, not for answering



questions.

Computer vision:

This focuses on analyzing and interpreting visual data (images and videos), not processing textual queries.

Forecasting:

Forecasting predicts future outcomes based on historical data, such as weather predictions or sales forecasts, but it does not handle language interaction.

**Key Tip:**

Use natural language processing (NLP) for any task involving understanding, interpreting, or responding to human language. Examples include chatbots, virtual assistants, and text analysis.

**Question: 226**

**CertyIQ**

You have 100 instructional videos that do NOT contain any audio. Each instructional video has a script.

You need to generate a narration audio file for each video based on the script.

Which type of workload should you use?

- A. language modeling
- B. speech recognition
- C. speech synthesis
- D. translation

**Answer: C**

**Explanation:**

You're converting text to speech - so Speech service's text-to-speech option <https://learn.microsoft.com/en-us/azure/cognitive-services/speech-service/text-to-speech>

The correct answer is C. Speech synthesis because speech synthesis converts text (the script) into spoken audio, generating narration for the instructional videos.

**Why the other options are incorrect:**

A. Language modeling:

Language modeling predicts the next word in a sequence or processes language for other NLP tasks, but it does not generate audio.

B. Speech recognition:

Speech recognition converts spoken words into text, which is the opposite of what's needed here (text to speech).

D. Translation:

Translation converts text from one language to another. It does not create audio files from scripts.

**Key Tip:**

Use speech synthesis (also known as text-to-speech) to generate audio narrations from written text, especially for applications like instructional videos, audiobooks, or accessibility features.

**Question: 227****CertyIQ**

HOTSPOT

-

Select the answer that correctly completes the sentence.

**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:****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

**Explanation:**

The correct answer is classify email messages as work-related or personal because natural language processing (NLP) can analyze and understand the content of emails to categorize them based on their context.

**Why the other options are incorrect:**

Predict the number of future car rentals:

This task involves forecasting or regression analysis, not NLP.

Predict which website visitors will make a transaction:

This is a use case for predictive analytics, possibly involving machine learning models like classification, but not NLP unless the prediction is based on textual input.

Stop a process in a factory when extremely high temperatures are registered:

This task involves anomaly detection or monitoring, unrelated to NLP.

**Key Tip:**

Use natural language processing (NLP) for tasks that involve analyzing, understanding, or classifying text or speech, such as email classification, sentiment analysis, or chatbot responses.

**Question: 228****CertyIQ**

Which AI service can you use to extract intent from a user input such as “Call me back later”?

- A. Azure Cognitive Search
- B. Translator
- C. Language
- D. Speech

**Answer: C****Explanation:**

C. Language

You can use the Azure Cognitive Services Language service to extract intent from a user input like "Call me back later." This service provides natural language processing (NLP) capabilities, including intent recognition, and can be used to understand the meaning and intent behind user text input. It's commonly used for tasks like chatbots, virtual assistants, and other applications that involve processing and understanding user-generated text. The other options, Azure Cognitive Search (A), Translator (B), and Speech (D), are not specifically designed for intent extraction from user inputs.

**Question: 229****CertyIQ**

You are building a Language Understanding model for an e-commerce business.

You need to ensure that the model detects when utterances are outside the intended scope of the model.

What should you do?

- A. Export the model
- B. Add utterances to the None intent
- C. Create a prebuilt task entity
- D. Create a new model

**Answer: B****Explanation:**

B is the answer.

<https://learn.microsoft.com/en-us/azure/cognitive-services/language-service/conversational-language-understanding/concepts/none-intent> Every project in conversational language understanding includes a default None intent. The None intent is a required intent and can't be deleted or renamed. The intent is meant to categorize any utterances that do not belong to any of your other custom intents.

HOTSPOT

-

Select the answer that correctly completes the sentence.

## Answer Area

▼

Azure Health Bot  
Microsoft Bot Framework  
Power Virtual Agents

can be used to build no-code apps that use built-in natural language processing models

Answer:

## Answer Area

▼

Azure Health Bot  
Microsoft Bot Framework  
**Power Virtual Agents**

can be used to build no-code apps that use built-in natural language processing models

Explanation:

The correct answer is Power Virtual Agents, as it is a no-code platform that allows users to create intelligent chatbots using built-in natural language processing (NLP) models without requiring any coding expertise.

**Why the other options are incorrect:**

Azure Health Bot:

This is a specialized bot framework designed for healthcare scenarios, but it is not a general-purpose, no-code platform like Power Virtual Agents.

Microsoft Bot Framework:

While the Microsoft Bot Framework provides extensive tools and SDKs for building bots, it requires programming and is not a no-code solution.

**Key Tip:**

Use Power Virtual Agents for building chatbots quickly and easily with minimal technical expertise, especially for scenarios where no coding is desired.

**Question: 231**

HOTSPOT

-

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

NOTE: Each correct selection is worth one point.

**Answer Area**

Statements	Yes	No
A smart device in the home that responds to questions such as "When is my next appointment?" is an example of conversational AI.	<input type="radio"/>	<input type="radio"/>
An interactive webchat feature on a company website can be implemented by using Azure Bot Service.	<input type="radio"/>	<input type="radio"/>
Automatically generating captions for pre-recorded videos is an example of conversation AI.	<input type="radio"/>	<input type="radio"/>

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

Statements	Yes	No
A smart device in the home that responds to questions such as "When is my next appointment?" is an example of conversational AI.	<input checked="" type="radio"/>	<input type="radio"/>
An interactive webchat feature on a company website can be implemented by using Azure Bot Service.	<input checked="" type="radio"/>	<input type="radio"/>
Automatically generating captions for pre-recorded videos is an example of conversation AI.	<input type="radio"/>	<input checked="" type="radio"/>

**Explanation:**

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

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

Automatically generating captions for pre-recorded videos is an example of conversation AI. --> No.

**Question: 232**

What is an example of the Microsoft responsible AI principle of transparency?

- A. ensuring that opportunities are allocated equally to all applicants
- B. helping users understand the decisions made by an AI system
- C. ensuring that developers are accountable for the solutions they create
- D. ensuring that the privileged data of users is stored in a secure manner

**Answer: B**

**Explanation:**

The correct answer is B. helping users understand the decisions made by an AI system.

Transparency is a key principle in Microsoft's Responsible AI guidelines. It emphasizes that AI systems should provide clear and understandable information about how they function, including the decisions they make. This helps users trust the AI and ensures they can effectively use it.

**Why the other options are incorrect:**

A. ensuring that opportunities are allocated equally to all applicants:

This is an example of the Fairness principle, not Transparency. Fairness ensures AI systems do not create or reinforce biases.

C. ensuring that developers are accountable for the solutions they create:

This aligns with the principle of Accountability, where developers and organizations must take responsibility for their AI solutions.

D. ensuring that the privileged data of users is stored in a secure manner:

This relates to the Privacy and Security principle, which focuses on safeguarding user data.

**Key Tip:**

Transparency ensures users can understand and trust AI systems by providing clear explanations of how decisions are made.

**Question: 233**

**CertyIQ**

You need to provide customers with the ability to query the status of orders by using phones, social media, or digital assistants.

What should you use?

- A. an Azure Machine Learning model
- B. the Translator service
- C. a Form Recognizer model
- D. Azure Bot Service

**Answer: D**

**Explanation:**

D Azure Bot Service provides an integrated environment that is purpose-built for bot development, enabling



you to build, connect, test, deploy, and manage intelligent bots, all from one place. You can build bots that can interact naturally with users using a range of channels, including social media platforms and digital assistants, which fits your requirement.

#### Question: 234

CertyIQ

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

Which service should you use?

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

**Answer: A**

**Explanation:**

Bots created using Azure Bot Service can be integrated with voice-based solutions such as Cortana, Alexa, or Google Assistant. Azure Bot Service provides the necessary channels and interfaces to connect bots to these services, allowing them to process and respond to voice commands. This integration enables users to interact with the bots through natural language voice commands, extending the bot's capabilities to various voice-enabled platforms.

#### Question: 235

CertyIQ

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.

### Answer Area

Statements	Yes	No
An interactive webchat feature on a company website can be implemented by using Azure Bot Service.	<input type="radio"/>	<input type="radio"/>
Automatically generating captions for pre-recorded videos is an example of conversational AI	<input type="radio"/>	<input type="radio"/>
A smart device in the home that responds to questions such as <b>“When is my next appointment?”</b> is an example of conversational AI	<input type="radio"/>	<input type="radio"/>

**Answer:**



## Answer Area

### Statements

Yes

No

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

☐☒

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

☒☐

#### Explanation:

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

Answer: Yes

Azure Bot Service is designed to create conversational AI experiences, such as chatbots, which can be integrated into websites for customer interaction.

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

Answer: No

Automatically generating captions falls under **speech-to-text** processing, which is not conversational AI. Conversational AI involves interacting with users through natural language.

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

Answer: Yes

This is an example of conversational AI, as it involves interaction between the user and the device using natural language to provide relevant responses.

### Question: 236

CertyIQ

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

Answer: C

#### Explanation:

Azure Form Recognizer is designed to analyze and extract information from documents, including structured and unstructured data. It can identify and process sensitive information by analyzing the content of

documents and categorizing data, making it an ideal service for scenarios where identifying sensitive information is required.

**Other options:**

Custom Vision is used for image classification and object detection, not document analysis.

Conversational Language Understanding is used to understand and process conversational text but is not meant for document analysis or sensitive information detection.

**Question: 237**

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.

**Answer Area**

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

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

**Answer:**

**Answer Area**

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

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

**Explanation:**

Detecting unusual temperature fluctuations for a large machine is an example of **an anomaly detection workload**.

Anomaly detection refers to identifying data points or patterns that do not conform to expected behavior. This is commonly applied to monitor systems like machines, where detecting deviations (e.g., temperature changes) can indicate potential issues or failures.

The other workloads listed (computer vision, knowledge mining, and NLP) are unrelated to monitoring numerical or sensor data for anomalies.

**Question: 238**

CertyIQ

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

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

**Answer: B**

**Explanation:**

A smart device that responds to the question “What is the stock price of Contoso. Ltd.?” is an example of an AI workload related to natural language processing1. Natural language processing (NLP) is a branch of artificial intelligence that focuses on the interaction between computers and humans through natural language. It involves the ability of a computer program to understand, interpret, and generate human language in a valuable way. Knowledge mining is a process that involves extracting useful information from unstructured data sources such as text documents, images, and videos1. Computer vision refers to the field of computer science that focuses on enabling computers to gain high-level understanding from digital images or videos1. Anomaly detection is a technique used to identify patterns or data points that deviate significantly from the normal behavior or expected values1.

**Question: 239**

CertyIQ

DRAG DROP

Match the machine learning models to the appropriate descriptions.

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

NOTE: Each correct match is worth one point.

**Models**

Classification

Clustering

Regression

**Answer Area**

A supervised machine learning model used to predict numeric values.

A supervised machine learning model used to predict categories.

An unsupervised machine learning model used to group similar entities based on features.

**Answer:**

## Answer Area

Regression

A supervised machine learning model used to predict numeric values.

Classification

A supervised machine learning model used to predict categories.

Clustering

An unsupervised machine learning model used to group similar entities based on features.

### Explanation:

The correct matches for the machine learning models and their descriptions are:

Regression: A supervised machine learning model used to predict numeric values.

Classification: A supervised machine learning model used to predict categories.

Clustering: An unsupervised machine learning model used to group similar entities based on features.

### Explanation:

Regression is used in supervised learning tasks where the target variable is continuous (e.g., predicting house prices or temperatures).

Classification is used in supervised learning tasks where the target variable is categorical (e.g., identifying spam vs. non-spam emails).

Clustering is an unsupervised learning technique used to group data into clusters based on similarity (e.g., customer segmentation).

## Question: 240

CertyIQ

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

The solution must be trained on images provided by your company.

Which Azure AI service should you use?

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

### Answer: B

### Explanation:

B - Custom Vision is the correct answer, because you have to use your own images as input to train the model.

Custom Vision is the appropriate Azure AI service for this scenario because it allows you to train a custom

model on images provided by your company. This service is specifically designed for image classification and object detection tasks where custom datasets are used, such as identifying specific products in retail stores.

### Why not the others?

A. Form Recognizer: This is used for extracting structured data from documents like forms and receipts, not for identifying products in images.

C. Face: This is used for facial recognition tasks, such as detecting and identifying human faces, not for identifying products.

D. Computer Vision: While it provides general image analysis capabilities, it is not customizable like Custom Vision and wouldn't allow you to train a model specifically for identifying competitor products.

Custom Vision is best suited for scenarios requiring tailored object detection and classification models based on specific training data.

### Question: 241

CertyIQ

HOTSPOT

Select the answer that correctly completes the sentence.

### 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:

### 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.	

Explanation:

Predicting how many hours of overtime a delivery person will work based on the number of orders received is an example of a regression model. Regression is used for predicting numerical values based on input data, which in this case is the number of hours of overtime based on the number of orders.

**Question: 242****CertyIQ**

Predicting agricultural yields based on weather conditions and soil quality measurements is an example of which type of machine learning model?

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

**Answer: B****Explanation:**

Regression models are used to predict continuous numeric values, such as agricultural yields, based on input features like weather conditions and soil quality measurements. These models are ideal for problems where the output is a numeric quantity.

**Why not the others?**

- A. Classification: Classification is used to predict discrete categories (e.g., "healthy" or "unhealthy" crop), not numeric values like yield.
- C. Clustering: Clustering is an unsupervised learning approach used to group data points based on their features, not for predicting numeric outcomes.

Thus, regression is the appropriate choice for predicting agricultural yields.

**Question: 243****CertyIQ**

You need to identify street names based on street signs in photographs.

Which type of computer vision should you use?

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

**Answer: B****Explanation:**

B-OCR, because you are extracting text from images by identifying the characters.

OCR (Optical Character Recognition) is specifically designed to extract and recognize text from images, such as street names from street signs in photographs.

It processes the visual data to identify characters, letters, and words, making it the ideal choice for identifying street names.

### Why not the others?

A. Object detection: Object detection is used to identify and locate objects (e.g., cars, traffic lights) in an image but does not process text.

C. Image classification: Image classification assigns a label to the entire image (e.g., "street sign") but cannot extract specific text like street names.

D. Facial recognition: Facial recognition is used to identify or verify individuals' faces, which is unrelated to recognizing text.

Thus, OCR is the appropriate solution for identifying street names on street signs.

### Question: 244

CertyIQ

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.

Workload Types	Answer Area	
Image classification		Generate captions for images.
Object detection		Extract movie title names from movie poster images.
Optical character recognition (OCR)		Locate vehicles in images.

### Answer:

Workload Types	Answer Area	
Image classification	Image classification	Generate captions for images.
Object detection	Optical character recognition (OCR)	Extract movie title names from movie poster images.
Optical character recognition (OCR)	Object detection	Locate vehicles in images.

### Explanation:

Object Detection

(OCR)

Object detection



1. **Generate captions for images → Image classification:**

While this could be related to automated captioning, image classification is often used as part of creating captions by identifying what the image contains.

2. **Extract movie title names from movie poster images → Optical Character Recognition (OCR):**

OCR is specifically designed to extract text from images, such as the titles on movie posters.

3. **Locate vehicles in images → Object detection:**

Object detection identifies and locates specific objects (e.g., vehicles) within images.

**Question: 245**

CertyIQ

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

**Answer: B**

**Explanation:**

Computer Vision.

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

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

**Question: 246**

CertyIQ

You are developing a chatbot solution in Azure.

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

- A. Translator
- B. Language
- C. Azure Cognitive Search
- D. Speech

**Answer: B**

**Explanation:**

Language (formerly part of Language Understanding Intelligent Service or LUIS) is a service designed to analyze natural language text and determine a user's intent. This is essential for chatbots because understanding what the user wants to achieve is the first step in providing an appropriate response.

**Why not the other options?****A. Translator:**

Translator is used for language translation, not intent recognition.

**C. Azure Cognitive Search:**

This service is used to perform searches on structured and unstructured data but does not analyze user intent in conversations.

**D. Speech:**

Speech is used for speech-to-text and text-to-speech processing, but it does not determine intent from the user's input.

Using Language ensures your chatbot can recognize user intents and respond accordingly.

# Thank you

Thank you for being so interested in the premium exam material.  
I'm glad to hear that you found it informative and helpful.

If you have any feedback or thoughts on the bumps, I would love to hear them.  
Your insights can help me improve our writing and better understand our readers.

## Best of Luck

You have worked hard to get to this point, and you are well-prepared for the exam  
Keep your head up, stay positive, and go show that exam what you're made of!

Feedback

More Papers



Future is Secured  
100% Pass Guarantee



24/7 Customer Support  
Mail us - [certyiqofficial@gmail.com](mailto:certyiqofficial@gmail.com)



Free Updates  
Lifetime Free Updates!

Total: **246 Questions**

Link: <https://certyiq.com/papers/microsoft/ai-900>