MICROSOFT AI-900 Exam Questions

Total Questions: 300+
Demo Questions: 35
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit:
MICROSOFT AI-900 Exam Questions by Cert Empire

Question: 1

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

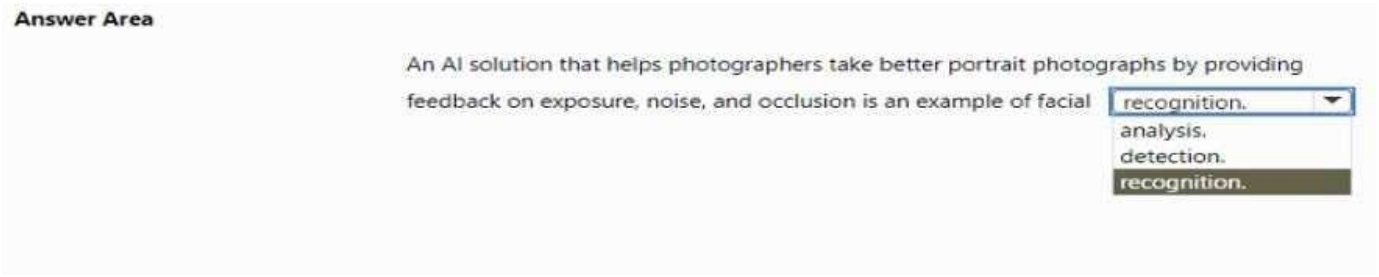

Answer:

analysis

Explanation:

The AI solution is evaluating specific attributes and quality metrics of the face (exposure, noise, occlusion), not simply locating it or identifying the person.

• Facial detection is the process of locating a face in an image, typically by providing bounding box coordinates.

• Facial recognition is the process of identifying or verifying a specific individual.

• Facial analysis (or face attribute analysis) is the process of extracting characteristics from a detected face. The scenario's requirements-evaluating exposure, noise, and occlusion-are explicit examples of face attribute analysis.

References:

Microsoft Azure AI Vision Documentation (Official Vendor Documentation)
Reference: "Face detection and attributes." Azure AI Vision documentation.
Citation: In the "Face attributes" section, the service description explicitly lists the attributes that can be returned from analysis. This includes "Occlusion... Exposure... Noise." This confirms that analyzing these specific photographic qualities is a function of face attribute analysis, which is distinct from mere detection or recognition.
ACM Computing Surveys (Peer-Reviewed Academic Publication)
Reference: Han, H., Wang, Z., Zhang, Z., & Shan, S. (2020). "A Survey on Face Attributes Analysis." ACM Computing Surveys, 53(2), Article 40, pp. 1-36.
Citation (DOI): https://doi.org/10.1145/3371909
Justification: This survey paper defines "Face Attributes Analysis" (FAA) as the task of "predicting a set of attributes (e.g., gender, age, expression, or... 'low quality' attributes) given a facial

image." The task described in the question (evaluating noise, exposure) falls directly under this academic definition.

IEEE Transactions on Pattern Analysis and Machine Intelligence (Peer-Reviewed Academic Publication)

Reference: Kumar, N., Belhumeur, P. N., & Nayar, S. K. (2009). "FaceTracer: A Search Engine for Large Collections of Images with Faces." IEEE Transactions on Pattern Analysis and Machine Intelligence, 31(5), 827-840.

Citation (DOI): https://doi.org/10.1109/TPAMI.2009.48

Justification: This foundational paper discusses descriptors for faces, distinguishing between "detection" (finding the face) and "analysis" of attributes, such as pose, illumination (exposure), and occlusion, which are used to describe the state of the face in the image rather than the identity of the person.

Question: 2

HOTSPOT Select the .

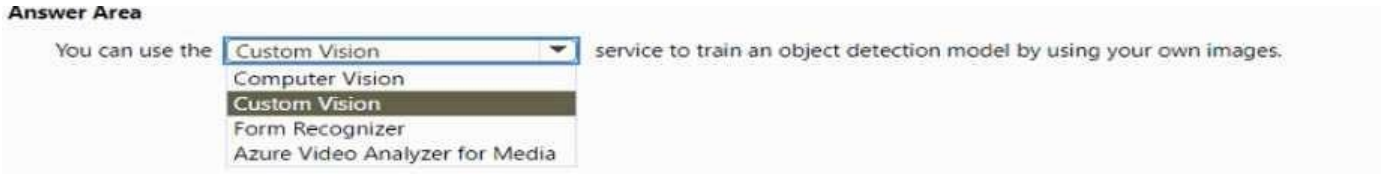

Answer:

Custom Vision

Explanation:

The Custom Vision service is the correct choice as it is designed to build, deploy, and improve custom image recognition models, including object detection, using your own images for training. The standard Computer Vision service provides pre-trained models for general-purpose object detection but does not allow training a custom model with your own images. Form Recognizer is used for document analysis (extracting text and tables), and Azure Video Analyzer for Media (now Video Indexer) is for analyzing video content, not training static image models.

References:

Microsoft. (n.d.). What is Custom Vision? Azure AI services documentation. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/ai-services/custom-vision-service/overview
Section: "What it does"
Quote: "Custom Vision is an image recognition service that lets you build, deploy, and improve your own image identifiers. ... You can train a project for one of two content types: ... Object detection models return the bounding box coordinates of where the detected object(s) in the image can be found."
Microsoft. (n.d.). Compare the Computer Vision and Custom Vision services. Azure AI services documentation. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/concept-object-detection-custom-vision
Section: "Compare object detection services"
Quote: "The Computer Vision service provides pre-trained models for general image analysis... The Custom Vision service... allows you to build your own models by providing your own labeled images. Use it if your application needs to... detect custom objects."
Microsoft. (n.d.). What is Azure Form Recognizer? Azure AI services documentation. Retrieved October 24, 2025, from
https://learn.microsoft.com/en-us/azure/ai-services/form-recognizer/overview
Section: "Introduction"
Quote: "Azure Form Recognizer is an applied AI service that... enables you to build automated data processing software... to extract text, key/value pairs, selection marks, tables, and structure

from your documents."
Microsoft. (n.d.). What is Azure Video Indexer? Azure AI services documentation. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/azure-video-indexer/video-indexer-overview
Section: "What is Azure Video Indexer?"
Quote: "Azure Video Indexer (AVI) is an artificial intelligence (AI) service... that offers you to extract deep insights... from your video and audio files..."

Question: 3

Your company manufactures widgets. You have 1.000 digital photos of the widgets. You need to identify the location of the widgets within the photos. What should you use?

A. Computer Vision Spatial Analysis

B. Custom Vision object detection

C. Custom Vision classification

D. Computer Vision Image Analysis

Answer:

B

Explanation:

The core requirement is to identify the location of specific objects ("widgets") within images. This task is known as object detection. The Azure AI Custom Vision service allows you to train a custom model for this purpose. An object detection project in Custom Vision not only identifies the presence of objects but also returns the coordinates of a bounding box that indicates their precise location within the image. This directly addresses the user's need.

Why Incorrect Options are Wrong:

A. Computer Vision Spatial Analysis: This service is designed to analyze real-time video streams to understand the movement and presence of people in physical environments, not to locate static objects in photos.

C. Custom Vision classification: Image classification predicts a label for an entire image (e.g., "this image contains a widget"). It does not provide the location or bounding box coordinates of the object within the image.

D. Computer Vision Image Analysis: The general Image Analysis service detects common, generic objects. Since "widgets" are specific to the company, a custom model trained with Custom Vision is the appropriate tool for accurate detection.

References:

1. Microsoft Learn, "What is Custom Vision?". Under the "What it does" section, it distinguishes between project types: "Object detection... the API returns the coordinates of the object(s) found in the image." This confirms that object detection provides location information.
Source: Microsoft. (n.d.). What is Custom Vision?. Azure AI services documentation. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/custom-vision-service/overview

2. Microsoft Learn, "AI-900: Explore object detection". This module in the official AI-900 learning path states, "Object detection models are trained to classify individual objects within an image,

and to identify their location with a bounding box." This directly contrasts with image classification, which "classifies the entire image."
Source: Microsoft. (n.d.). Explore object detection. AI-900: Microsoft Azure AI Fundamentals learning path. Retrieved from https://learn.microsoft.com/en-us/training/modules/explore-object-detection/
3. Microsoft Learn, "What is Spatial Analysis?". The documentation clearly states, "The Azure AI Vision Spatial Analysis service ingests video from cameras, detects the presence and movements of people in the video..." This shows its purpose is unrelated to the question's scenario.
Source: Microsoft. (n.d.). What is Spatial Analysis?. Azure AI services documentation. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/overview-spatial-analysis

Question: 4

HOTSPOT Select the answer that correctly completes the sentence

Answer:

Object detection

Explanation:

Object detection is the correct term. This computer vision technique is designed to analyze an image and identify the presence, location, and class of multiple distinct entities. It returns bounding box coordinates and a label (e.g., "person," "car," "dog") for each item it finds.

This contrasts with:

• Image classification: Assigns only one single label to the entire image (e.g., "a beach").

• Image description: Generates a full sentence describing the scene.

• Optical character recognition (OCR): Is used specifically to extract text from an image, not general items.

References:

Microsoft Corporation. (2024). Object detection - Azure AI services. Microsoft Learn.
Reference: The official documentation for the Azure AI Vision service states, "Object detection finds and locates objects within an image. ... if an image contains a dog, a cat, and a person, the Detect operation will list those objects together with their coordinates in the image." This directly supports the answer.
Location: Retrieved October 24, 2025, from the "Object detection" section of the Azure AI Vision documentation (v4.0).
Fei-Fei, L., Johnson, J., & Yeung, S. (2023). CS231n: Convolutional Neural Networks for Visual Recognition. Stanford University.
Reference: The course materials explicitly distinguish between these tasks. Image Classification is defined as assigning one label per image. Object Detection is defined as: "Detect object(s) in

image, draw bounding box around them."
Location: Lecture 11: "Detection and Segmentation" slides. (e.g., https://cs231n.github.io/slides/2023/lecture11.pdf, Slide 5)
Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, Real-Time Object Detection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 779-788.
Reference: This foundational paper defines the task as "spatially separated bounding boxes and associated class probabilities." The goal is explicitly to identify all object instances and their classes within the image, matching the question's requirement for "multiple types of items."
Location: Section 1 (Introduction) & Section 2 (Unified Detection).
DOI: https://doi.org/10.1109/CVPR.2016.91

Question: 5

HOTSPOT Select the answer that correctly completes the sentence

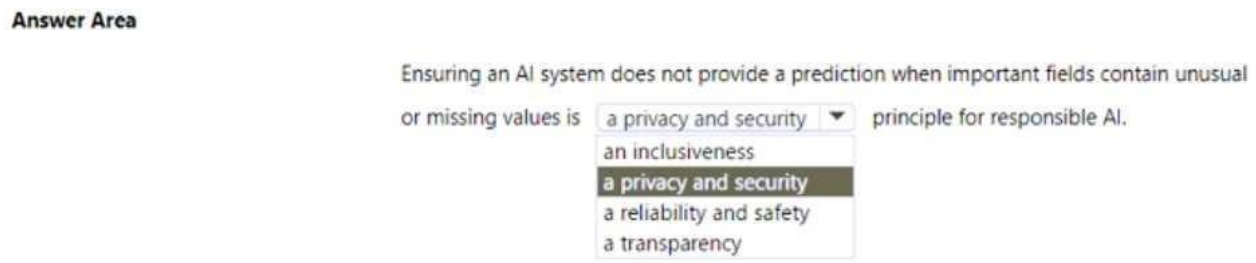

Answer:

a reliability and safety

Explanation:

This scenario is a direct application of the Reliability and Safety principle. This principle mandates that AI systems must be robust, resilient, and operate dependably, especially when faced with unexpected conditions or "edge cases." Data fields with unusual or missing values represent such unexpected conditions. By being designed to not provide a prediction in this situation, the system avoids generating a potentially inaccurate, nonsensical, or harmful output. This "fail-safe" behavior is a core component of ensuring the system is reliable and safe for its users.

References:

Microsoft. (n.d.). The six principles of responsible AI. Microsoft Learn. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai
Reference: In the "Reliability and safety" section, it is stated: "AI systems need to be reliable and safe... To be reliable, a system should operate as it was originally designed... This requires rigorous testing and validation to ensure that the system responds safely to edge cases..." Missing or unusual values are a classic "edge case."
Microsoft. (2024, September 12). Responsible AI in practice. Microsoft Learn. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/training/modules/responsible-ai-practice/2-understand-principles
Reference: In the "Reliability and safety" section, the text notes that systems must be "resilient to failure or unexpected conditions." The module emphasizes that a reliable system must "work as expected," and providing a flawed prediction on missing data would be a failure to work as expected.
Zhang, J., Liu, M., & Liu, X. (2023). Reliability of Machine Learning: A Survey. ACM Computing Surveys, 55(9), Article 189. https://doi.org/10.1145/3555806
Reference: Section 3.1, "Data Quality Issues," explicitly identifies "missing values" and "anomalies" as key challenges to machine learning reliability. The paper discusses (Section 4.1)

"failure detection" mechanisms, where a system identifies inputs it cannot process reliably, as a strategy to ensure overall system reliability. Refusing to predict is such a mechanism.

Question: 6

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.

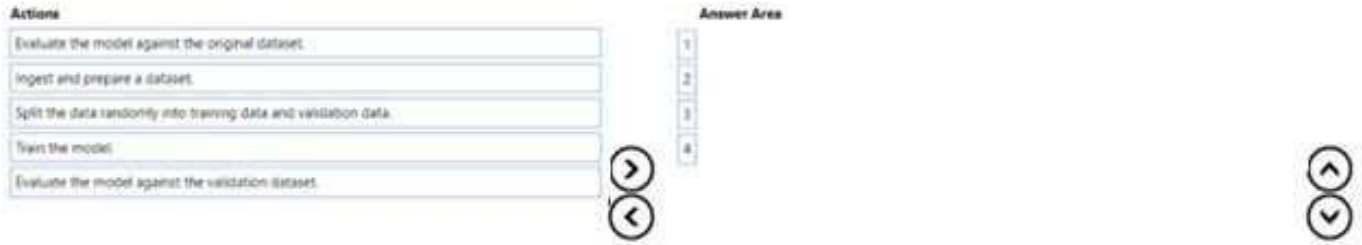

Answer:

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:

This sequence represents the fundamental workflow for supervised machine learning, which is visually implemented in the Azure Machine Learning designer.

• First, data must be ingested and prepared (e.g., cleaned, transformed).

• Next, this dataset is split into a training set (to teach the model) and a validation/test set (to test it).

• The model is then trained using only the training data.

• Finally, the trained model's performance is assessed by evaluating it against the unseen validation dataset. Evaluating against the "original dataset" would be incorrect as it includes the training data, leading to an over-optimistic and invalid performance metric (overfitting).

References:

Microsoft Press Store (for Exam DP-100): In the guide "Explore data and train models," the process for creating a training pipeline in the designer is detailed. It specifies the sequence:
Use modules to "clean and transform your data" (Prepare).
"Employ the Split Data module to divide your dataset into training and validation sets."
"Connect your preprocessed and split datasets to the Train Model module."
"Examine the output of the Evaluate Model module to assess the performance."
(Source: Microsoft Press Store, "Explore data and train models," Article 2, Skill 2.2: "Create models by using the Azure Machine Learning Designer," sections "Create a training pipeline" and "Here is a more detailed set of instructions...")
Microsoft Learn (Azure ML Documentation): The tutorial "Tutorial: Train a no-code regression model in the designer" demonstrates this logical flow. A dataset is added, prepared (e.g., Select Columns in Dataset, Clean Missing Data), and then split using the Split Data component. The outputs of the split data component are then fed into the Train Model component, which is subsequently connected to a Score Model and Evaluate Model component.
(Source: Microsoft Learn, "Tutorial: Train a no-code regression model in Azure Machine Learning designer," sections "Prepare data," "Split the data," "Train the model," and "Test, score, and evaluate the model.")
Microsoft Learn (Component Reference): The official documentation for the Train Model component states it has two inputs: an `untrained algorithm` and the "training dataset." It explicitly notes, "After training is completed, use the trained model with one of the scoring components, to make predictions on new data," which is the function of the Evaluate Model component.
(Source: Microsoft Learn, "Train Model: Component Reference," sections "How the training process works" and "How to use Train Model.")

Question: 7

HOTSPOT Select the answer that correctly completes the sentence.

Answer Area

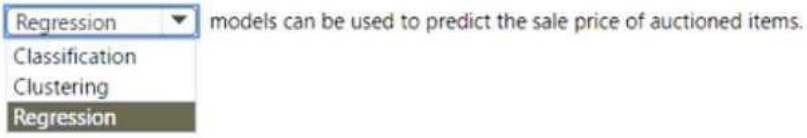

Answer:

Regression

Explanation:

Regression is the correct machine learning model type. This is a supervised learning technique used specifically for predicting a continuous, numerical value. A "sale price" (e.g., $45.50, $1,200) is a continuous numerical outcome.

In contrast, Classification is used to predict a discrete class or category (e.g., "Sold" vs. "Unsold," or "High," "Medium," "Low" price brackets). Clustering is an unsupervised technique used to group data points based on similarities (e.g., grouping similar types of auctioned items together), not to predict a specific value.

References:

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). An Introduction to Statistical Learning. Springer.
Chapter 3, p. 59: Defines regression problems as those involving the prediction of a quantitative or continuous output value, such as predicting the price of a stock.
Chapter 4, p. 127: Defines classification problems as those involving the prediction of a qualitative or categorical response, such as whether a stock will go up or down.
Grimson, W. E. L., & Guttag, J. (2016). Lecture 9: Linear Regression. MIT OpenCourseware, 6.036 Introduction to Machine Learning.
Section 1 ("Regression"): States, "In regression, we are interested in predicting a continuous output variable, y." It lists "stock price" as a canonical example, which is analogous to "sale price."
Ng, A. (n.d.). CS229 Machine Learning Course Notes: Part V - Clustering. Stanford University.
Section 1 ("Clustering"): Describes clustering as "the problem of finding subgroups, or clusters, in a dataset," distinguishing it as an unsupervised problem for finding structure, not predicting a value.

Question: 8

HOTSPOT Select the answer that correctly completes the sentence.

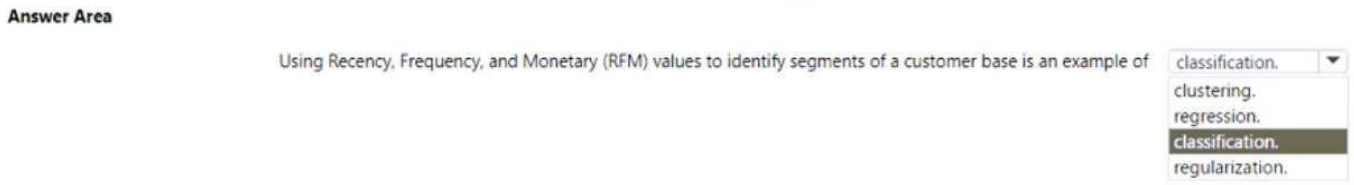

Answer:

clustering

Explanation:

The question describes using RFM (Recency, Frequency, Monetary) data to identify segments within a customer base. This process involves grouping customers based on the similarity of their RFM values, without predefined labels for these groups. This is the definition of clustering, which is an unsupervised machine learning technique used to discover natural groupings or structures in data.

In contrast, classification is a supervised task that assigns data to predefined categories. Regression predicts a continuous numerical value. Regularization is a technique used to prevent overfitting in models, not a primary machine learning task type.

References:

Academic Publication (IEEE): A common application of RFM analysis is to use it as input for a clustering algorithm, suchs as K-Means, to perform customer segmentation. The algorithm "clusters customers based on their RFM values" to "help in identifying the different types of customers."
Bukhari, S. S. H. S., Bukhari, M. B. T. S. R. S., & Bukhari, M. A. A. B. S. S. (2020). Customer Segmentation Using RFM Model and K-Means Clustering. In 2020 3rd International Conference on Engineering and Technology (ICoEngT) (pp. 1-6). IEEE.
DOI: https://doi.org/10.1109/ICoEngT50666.2020.9285324
Section: Abstract, Section III-B.
Academic Publication (MIT Press): Clustering is formally defined as an unsupervised learning task where "the aim is to find natural groupings of data." A primary application listed for clustering is "customer segmentation in marketing."
Alpaydn, E. (2020). Introduction to machine learning (4th ed.). MIT Press.
Section: Chapter 7, "Clustering," pp. 191.
Official Vendor Documentation (Microsoft): Microsoft's documentation on machine learning algorithms distinguishes between tasks. Clustering is recommended for "Grouping" or

"Discovering structure" when you want to "Separate data into groups." This aligns perfectly with "identify segments." In contrast, Classification is used for "Predicting categories" where the categories are already known.

Microsoft. (n.d.). Machine learning algorithm cheat sheet for Microsoft Azure Machine Learning. Microsoft Azure Documentation.

Section: See the paths for "Grouping" (Unsupervised - Clustering) vs. "Predicting categories" (Supervised - Classification).

University Courseware (MIT): MIT's "Introduction to Machine Learning" course defines clustering as an unsupervised problem with the "goal of discovering 'natural' groupings in data." This contrasts with classification, which is defined as a supervised problem of learning a function that "maps an input... to one of a fixed set of categories." The question's "identify segments" refers to discovering groupings, not mapping to fixed categories.

MIT OpenCourseWare. (2020). 6.036 Introduction to Machine Learning, Fall 2020.

Section: Lecture 15 (Clustering) and Lecture 1 (Introduction, Classification).

Question: 9

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

A. creating photorealistic images by using three-dimensional models

B. assigning the color pixels in an image to object names

C. describing the contents of an image

D. detecting inconsistencies and anomalies in a stream of data

E. creating visual representations of numerical data

Answer:

B, C

Explanation:

Computer vision is a field of AI focused on enabling computers to interpret and understand the visual world.
Assigning color pixels to object names (B) is a core computer vision task known as semantic segmentation. This technique classifies each pixel in an image as belonging to a specific object or class, effectively creating a detailed map of the image's contents.

Describing the contents of an image (C), often called image captioning or image analysis, is another fundamental computer vision workload. It involves analyzing an image to generate a human-readable sentence that summarizes the scene, objects, and actions depicted.

Why Incorrect Options are Wrong:

A. Creating photorealistic images from models is a task for computer graphics and rendering, not computer vision, which focuses on analyzing existing visual data.
D. Detecting anomalies in a generic stream of data is a broader machine learning task (anomaly detection) and is not exclusive to or a primary definition of computer vision.
E. Creating visual representations of numerical data, such as charts and graphs, is known as data visualization, a distinct field from computer vision.

References:

1. Microsoft Learn. (2024). Introduction to Computer Vision. In "Explore computer vision in Microsoft Azure" module. This document lists key computer vision tasks, including "Semantic segmentation" (matching option B) and "Image analysis" which includes generating descriptive captions (matching option C).
2. Microsoft Learn. (2024). What is Image Analysis?. Azure AI Services documentation. Under the "Image captioning and dense captioning" section, it states, "Image Analysis can generate a

human-readable sentence that describes the content of an image." This directly supports option C.
3. Microsoft Learn. (2024). What is Image Analysis?. Azure AI Services documentation. The "Image segmentation" section describes the service's ability to "generate a segmentation mask" for objects, which aligns with the process described in option B.

Question: 10

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

A. translation

B. language modeling

C. key phrase extraction

D. speech recognition

Answer:

A

Explanation:

The scenario requires converting text from a source language (English) to a target language based on the user's location. This task is the definition of machine translation. The Azure AI Language service offers text translation capabilities that take source text and return the translated text in one or more target languages, directly addressing the requirement to present reviews to users in their respective languages.

Why Incorrect Options are Wrong:

B. language modeling: This is a foundational NLP task that predicts the probability of a sequence of words; it does not translate text between languages.
C. key phrase extraction: This feature identifies the main points or topics in a text. It analyzes content but does not convert it to another language.
D. speech recognition: This workload converts spoken audio into written text. The input in this scenario is customer reviews, which are already in text format.

References:

1. Microsoft Learn, AI-900: Microsoft Azure AI Fundamentals, "Explore natural language processing", Unit: "Understand text and speech translation": This unit explicitly defines the workload. It states, "Text translation is the task of translating text from one language to another... You can build solutions that use machine translation to overcome language barriers." This directly supports the choice of translation for the described scenario.
2. Microsoft Learn, AI-900: Microsoft Azure AI Fundamentals, "Explore natural language processing", Unit: "Understand language": This section describes key phrase extraction as a feature to "identify important words and phrases in the text that suggest its main points," confirming it is for analysis, not translation.
3. Microsoft Learn, AI-900: Microsoft Azure AI Fundamentals, "Explore natural language

processing", Unit: "Understand speech": This unit defines speech recognition (also known as speech-to-text) as "the ability to recognize and transcribe spoken language," which is not applicable to the text-based reviews in the question.

Question: 11

HOTSPOT Select the answer that correctly completes the sentence.

Answer:

confusion matrix

Explanation:

A confusion matrix is a specific table (or grid) used to evaluate the performance of a classification model. It provides a summary of the prediction results by comparing the actual (true) class labels against the labels predicted by the model. The grid's cells quantify the True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN), which directly aligns with the description of a grid displaying "predicted and actual positives and negatives."

The other options are incorrect:

• An ROC curve is a plot of the true positive rate against the false positive rate, not a grid.

• An AUC metric is a single scalar value representing the area under the ROC curve.

• A threshold is a value used by the model to decide between classes.

References:

Microsoft Azure Documentation. (2024). "Model evaluation metrics in Azure Machine Learning designer." Azure Machine Learning Documentation. In the section "Metrics for classification models," the Confusion matrix is defined: "The confusion matrix is an N x N matrix used for evaluating the performance of a classification model... The matrix compares the actual labels with the predicted labels."
Fawcett, T. (2006). "An introduction to ROC analysis." Pattern Recognition Letters, 27(8), 861-874. Section 3, "Evaluating Classifiers," introduces the confusion matrix as a table containing "information about actual and predicted classifications," which forms the basis for metrics like TP, FP, TN, and FN. (DOI: https://doi.org/10.1016/j.patrec.2005.10.010)
University of California, Berkeley. (2020). DATA 100: Principles and Techniques of Data Science, Lecture 24: Classification Evaluation. Slides 30-35 visually define the confusion matrix as a 2x2 grid comparing "Actual" vs. "Predicted" labels and identifies the four outcomes (TP, FP, FN, TN).

Question: 12

You have an Internet of Things (IoT) device that monitors engine temperature. The device generates an alert if the engine temperature deviates from expected norms. Which type of AI workload does the device represent?

A. natural language processing (NLP)

B. computer vision

C. anomaly detection

D. knowledge mining

Answer:

C

Explanation:

The scenario describes a system that monitors a continuous stream of data (engine temperature) and identifies events that are unusual or deviate from an established pattern ("expected norms"). This is the core function of an anomaly detection workload. The goal is to flag outliers or unexpected events that could signify a problem, such as an engine overheating or cooling down unexpectedly, which requires attention.

Why Incorrect Options are Wrong:

A. natural language processing (NLP): This workload is concerned with the interpretation and generation of human language (text and speech), not with analyzing numerical sensor data.
B. computer vision: This workload involves processing and understanding information from images and videos. The scenario deals with temperature data, not visual input.
D. knowledge mining: This is a broader term, often associated with extracting information from large volumes of unstructured text data to create a searchable knowledge base. It is not the specific task of identifying outliers in time-series data.

References:

1. Microsoft Learn. "Describe anomaly detection." Microsoft Azure AI Fundamentals: Describe machine learning principles on Azure. Module 3, Unit 5. Accessed May 20, 2024.
Reference Details: This official AI-900 learning path defines anomaly detection as a machine learning technique used to "identify unusual patterns or events that don't conform to expected behavior." It explicitly mentions monitoring equipment for unusual numeric values as a primary use case.
2. Microsoft Azure Documentation. "What is the Anomaly Detector API?" Azure AI services. Last updated: April 1, 2024.
Reference Details: The documentation states, "The Anomaly Detector API enables you to monitor

and detect abnormalities in your time series data... For example, monitoring machine temperature." This directly corresponds to the scenario described in the question.

Question: 13

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

A. Split the original dataset into a dataset for features and a dataset for labels. Use the features dataset for evaluation.

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

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

D. Use the original dataset for training and evaluation.

Answer:

C

Explanation:

The standard and correct methodology for evaluating any machine learning model, including an unsupervised clustering model, is to assess its performance on data it has not seen during training. This practice prevents overfitting and provides an unbiased measure of how the model will generalize to new, real-world data. The process involves splitting the original dataset into a training set and a testing set. The model is built using the training set, and then its performance is evaluated using the separate testing set.

Why Incorrect Options are Wrong:

A. Clustering is an unsupervised learning method; the dataset does not contain pre-defined labels. Therefore, splitting the data into features and labels is not applicable.
B. Evaluating a model on the same data used for training provides an overly optimistic and biased assessment of its performance, as the model may have simply memorized the training data.
D. Using the entire dataset for both training and evaluation is incorrect. This fails to test the model's ability to generalize to new, unseen data, which is the primary goal of evaluation.

References:

1. Microsoft Azure Official Documentation, "Tutorial: Create a clustering model in Azure Machine Learning designer": This official tutorial explicitly outlines the standard workflow. In the "Train the model" and "Test the model" sections, it demonstrates splitting the data, using the first output (the training set) to train the model, and the second output (the testing set) to score and evaluate the model.
Reference: Microsoft Corporation. (n.d.). Tutorial: Create a clustering model in Azure Machine Learning designer. Microsoft Learn. Retrieved from https://learn.microsoft.com/en-us/azure/machi

ne-learning/tutorial-designer-clustering-predict-customer-churn (See sections: "Split the data", "Train the model", and "Test, score, and evaluate the model").
2. Microsoft Azure Official Documentation, "Evaluate Model component": The documentation for the Evaluate Model component states its purpose is to measure the accuracy of a trained model using a scored dataset. In a standard pipeline, this scored dataset is generated by applying the trained model to the test portion of the split data.
Reference: Microsoft Corporation. (n.d.). Evaluate Model component. Microsoft Learn. Retrieved from
https://learn.microsoft.com/en-us/azure/machine-learning/component-reference/evaluate-model (See "How to use Evaluate Model" section).
3. Microsoft Azure Official Documentation, "Split Data component": This document describes the function of the Split Data component, which is to "divide a dataset into two distinct sets." It explicitly mentions the common use case: "to separate data into training and testing sets."
Reference: Microsoft Corporation. (n.d.). Split Data component. Microsoft Learn. Retrieved from https://learn.microsoft.com/en-us/azure/machine-learning/component-reference/split-data (See "Overview" section).

Question: 14

HOTSPOT For each of the following statements, select Yes if the statement is True. Otherwise, select No. NOTE: Each correct selection is worth one point.

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.	○	○
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.	○	○

Answer:

Yes

Yes

No

Explanation:

Statement 1 (Yes): Conversational AI is defined by its ability to facilitate a natural, human-like dialogue. A smart device that understands a spoken question (input), determines its intent (scheduling), and provides a relevant verbal response (output) is a prime example of an interactive conversational AI system.

Statement 2 (Yes): Azure Bot Service is the specific Microsoft platform designed to build, test, deploy, and manage intelligent bots. A primary use case for this service is creating interactive webchat bots for customer engagement on websites.

Statement 3 (No): This statement describes speech-to-text (or transcription), which is a component of AI-powered speech services. Conversational AI requires a two-way interaction or dialogue. Generating captions is a one-way process of transcribing audio to text and does not, by itself, constitute a conversation.

References:

Microsoft. (n.d.). What is conversational AI? Microsoft Learn. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/ai-services/conversational-ai
Section: "What is conversational AI?"
Note: This document defines conversational AI by its interactive nature (e.g., chatbots and smart devices) and lists "Speech to Text" as a separate (though related) component, supporting the

answers for Statement 1 and Statement 3.

Microsoft. (n.d.). What is Azure Bot Service? Microsoft Learn. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/bot-service/bot-service-overview

Section: "What is Azure Bot Service?"

Note: This documentation explicitly states that the service is used to "build, connect, test, and deploy" bots that "interact naturally with your users on websites, apps, ...and more," directly supporting Statement 2.

Jurafsky, D., & Martin, J. H. (2024). Speech and Language Processing (3rd ed. draft). Stanford University.

Chapter 26: "Automatic Speech Recognition" (ASR)

Chapter 28: "Dialogue Systems and Chatbots"

Note: This university-level text clearly distinguishes Automatic Speech Recognition (the technology for captioning, as described in Statement 3) from Dialogue Systems (the technology for conversational agents, as in Statement 1). ASR is a component, but it is not the dialogue system itself.

Gao, J., Galley, M., & Li, L. (2018). Neural Approaches to Conversational AI. In Foundations and Trends in Information Retrieval (Vol. 12, Nos. 2-3, pp. 127-298).

DOI: 10.1561/1500000066

Note: This academic review defines conversational AI systems (dialogue systems) by their ability to "converse with humans via natural language" (p. 131), reinforcing that the interactive nature (Statement 1) is key, whereas simple transcription (Statement 3) is a different task.

Question: 15

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.	○	○
A mobile application that displays images relating to an entered search term is an example of a natural language processing workload.	○	○
A web form used to submit a request to reset a password is an example of a natural language processing workload.	○	○

Answer:

Yes

No

No

Explanation:

Natural Language Processing (NLP) workloads are defined by the requirement to process, interpret, or understand unstructured human language (text or speech).

• Bot (Yes): A bot must use NLP, specifically Natural Language Understanding (NLU), to parse the user's free-form query, identify its intent (what the user wants) and entities (the specific subjects), and then map this to an appropriate response.

• Image Search (No): This describes an information retrieval or computer vision workload. The system matches a keyword (the search term) to tags or content associated with images, but it does not need to understand the grammatical structure or semantic nuances of complex language.

• Web Form (No): A web form uses structured data. The user inputs information into predefined, labeled fields (e.g., 'username', 'password'). The system processes this structured data directly and is not required to interpret unstructured, natural language.

References:

Microsoft. (2024). What is Natural Language Processing? - Azure AI Fundamentals. Microsoft Learn documentation (AI-900). This source explicitly lists "Language understanding for conversational bots" and "Question answering" as key workloads of NLP, confirming the first statement.

Microsoft. (2024). What is Computer Vision? - Azure AI Fundamentals. Microsoft Learn documentation (AI-900). This documentation defines image-related tasks, such as "Image classification" and "Object detection," as computer vision workloads, which are distinct from NLP workloads. This supports the "No" for the second statement.

Jurafsky, D., & Martin, J. H. (2023). Speech and Language Processing (3rd ed.). Stanford University. Chapter 1, Section 1.1 distinguishes NLP tasks from processing structured data. A web form (Statement 3) is a prime example of structured data input, which does not require the ambiguity resolution or interpretation inherent to NLP.

Chowdhury, G. G. (2003). Natural language processing. Annual Review of Information Science and Technology, 37(1), 51-89. https://doi.org/10.1002/aris.1440370103. Page 52 defines NLP as systems that "process, understand, or generate natural language." This definition excludes the structured data processing of a web form and the keyword-to-image matching of the mobile app.

Question: 16

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

Answer Area

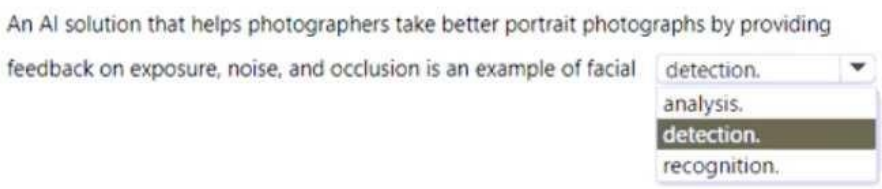

Answer:

analysis

Explanation:

The AI solution is assessing specific qualities or attributes of the face, such as exposure, noise, and occlusion.

• Facial detection simply locates a face in an image (e.g., provides a bounding box).

• Facial recognition identifies a specific person.

• Facial analysis (or attribute detection) is the process that extracts information about a detected face. This includes quality metrics (blur, exposure, noise), pose, emotion, and physical attributes (occlusion, glasses, etc.).

Because the solution provides feedback on these specific attributes, it is performing facial analysis.

References:

Microsoft. (2024). Face detection and attributes. Azure AI services documentation.
Reference: In the "Face - Detect" API documentation, the service can "extract a set of face-related attributes... The available attributes... include... Occlusion... Noise... Exposure." This confirms that determining these specific properties is a function of facial attribute analysis, which is a component of the broader "Face" service, distinct from simple detection.
Hjelmas, E., & Goth, B. K. (2001). Face Detection: A Survey. Computer Vision and Image Understanding, 83(3), 236-274. https://doi.org/10.1006/cviu.2001.0921
Reference: This survey (Section 1, "Introduction") distinguishes between "detection" (locating a face) and "recognition" (identifying a face). The task described in the question-evaluating

attributes-is a separate analytical step that follows detection.
MIT OpenCourseWare. (2020). 6.819/6.869: Advances in Computer Vision.
Reference: Lecture 15, "Face Recognition." Course materials distinguish the core tasks: Detection ("Find all faces"), Attribute Analysis ("Find properties: pose, expression, gender, image quality"), and Recognition ("Identify the person"). The question clearly aligns with Attribute Analysis.

Question: 17

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

A. AI systems must use only publicly available data.

B. AI systems must protect the interests of the company

C. AI systems must be understandable.

D. AI systems must keep personal details public

Answer:

C

Explanation:

Microsoft's responsible AI principle of Transparency dictates that AI systems should be understandable. This means that users should be able to comprehend how the system works, the data it uses, its limitations, and the reasoning behind its decisions. Making AI systems understandable helps build trust and enables accountability by allowing developers and users to identify and rectify potential biases or errors.

Why Incorrect Options are Wrong:

A. This is incorrect. While data sourcing is important, there is no principle mandating the exclusive use of publicly available data; the focus is on responsible data handling, regardless of its source.
B. This is incorrect. The principles of responsible AI are human-centric, focusing on fairness and societal benefit, not solely on corporate interests, which could potentially conflict with ethical considerations.
D. This is incorrect. This statement directly contradicts the principle of Privacy and Security, which requires AI systems to protect personal data and respect user privacy.

References:

1. Microsoft Learn. (2024). Microsoft's responsible AI principles. AI-900: Microsoft Azure AI Fundamentals. "Transparency: AI systems should be understandable. The people who create and use AI systems must be able to fully understand how they work so they can identify potential issues, such as bias or unexpected outcomes, that could otherwise go undiscovered."
2. Microsoft Learn. (2024). Introduction to responsible AI in Azure. AI-900: Microsoft Azure AI Fundamentals. "The six principles that form the foundation of Microsoft's approach to responsible AI are: Fairness, Reliability and Safety, Privacy and Security, Inclusiveness, Transparency, and Accountability."
3. Microsoft Corporate. (2023). Microsoft Responsible AI Standard, v2. "Transparency: We will be transparent about the capabilities and limitations of our AI systems." (Section 1.5, Page 6).

Question: 18

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 phone number follows a specific, predictable structure (e.g., (555) 123-4567, 555-123-4567). In Natural Language Processing (NLP) services, such as Azure AI Language's Conversational Language Understanding, a regular expression entity is the ideal component for identifying and extracting text that conforms to a defined pattern. This method is highly precise and efficient for structured data like phone numbers, email addresses, or product codes, as it does not require machine learning from examples.

Why Incorrect Options are Wrong:

B. machine-learned: This entity type is best for concepts that are defined by context, not a strict pattern, and requires training with many labeled examples. It is inefficient for a well-structured pattern like a phone number.
C. list: A list entity is used for a fixed, closed set of words and their synonyms (e.g., a list of cities or product categories). It cannot be used to identify the near-infinite variations of phone numbers.
D. Pattern-any: This is a generic placeholder used within a larger pattern template to capture variable text. A regular expression entity is more specific and appropriate for defining the exact structure of a phone number.

References:

1. Microsoft Learn, "Entity components in conversational language understanding." Under the "Regular expression entity component" section, it states, "A regular expression entity component extracts an entity based on a regular expression pattern you provide. It's ideal for text with consistent formatting." This directly supports using regex for patterned data like phone numbers.
2. Microsoft Learn, "Entity components in conversational language understanding." The "List entity component" section clarifies, "A list entity component represents a fixed, closed set of related words along with their synonyms...This component is a good choice when you have a set of items that don't change often." This confirms why a list is unsuitable for phone numbers.
3. Microsoft Learn, "How to create a conversational language understanding project." In the "Add

entity components" section, the documentation guides users to "Select Regular expression" for entities that follow a defined pattern, contrasting it with list and machine-learned entities.

Question: 19

You need to implement a pre-built solution that will identify well-known brands in digital photographs. Which Azure AI sen/tee should you use?

A.
Face

B.
Custom Vision

C.
Computer Vision

D.
Form Recognizer

Answer:

C

Explanation:

The Azure AI Computer Vision service provides pre-built, pre-trained models for analyzing images. One of its core features is the ability to detect thousands of well-known commercial brands from a continuously updated database. This capability is available through the Analyze Image API by specifying the brands visual feature. It directly fulfills the requirement for a pre-built solution to identify brands in photographs without needing to train a custom model.

Why Incorrect Options are Wrong:

A. Face: The Azure AI Face service is specialized for detecting, identifying, and analyzing human faces and their attributes, not for recognizing commercial logos or brands.
B. Custom Vision: This service is used to build, train, and deploy your own custom image classification and object detection models. It is not a pre-built solution for general brand detection.
D. Form Recognizer: This service is designed to extract text, key-value pairs, and tables from documents like invoices and receipts, not for analyzing brands in general photographs.

References:

1. Microsoft Learn. (2024). Analyze images with the Computer Vision service. AI-900: Microsoft Azure AI Fundamentals learning path.
Reference: In the "Image Analysis" section, the documentation states, "The Computer Vision service can detect thousands of famous brands." This confirms its pre-built capability for brand detection.
2. Microsoft Azure Documentation. (2023). What is Computer Vision? - Image analysis.
Reference: Under the "Image analysis" feature list, it specifies "Brand detection: Detects brands

in images from a database of thousands of global logos." This explicitly identifies the required functionality as part of the Computer Vision service.

3. Microsoft Azure Documentation. (2023). Call the Analyze Image API.

Reference: In the section "Specify visual features," the documentation lists Brands as one of the features that can be requested. The description states, "Detects various brands within an image, including the approximate location of the brand logo."

Question: 20

HOTSPOT For each of the following statements. select Yes if the statement is true. Otherwise, select No. NOTE; Each correct selection is worth one point

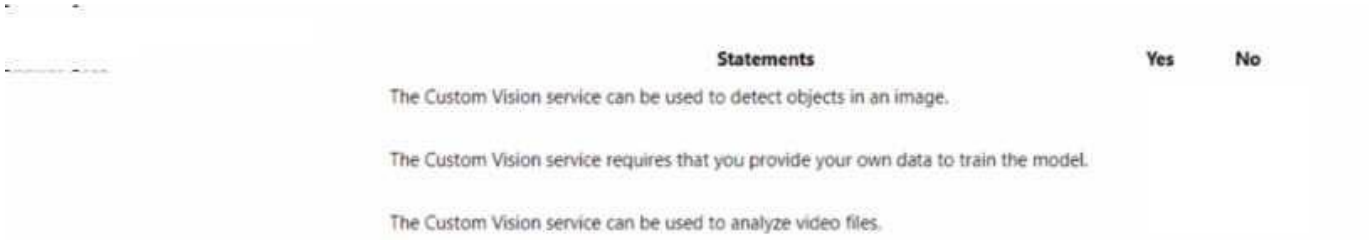

Statements	Yes	No
The Custom Vision service can be used to detect objects in an image.		
The Custom Vision service requires that you provide your own data to train the model.		
The Custom Vision service can be used to analyze video files.		

Answer:

Yes

Yes

No

Explanation:

Object Detection (Yes): The Azure Custom Vision service is designed for two primary functions: image classification (assigning labels to an entire image) and object detection (identifying the location and label of specific objects within an image using bounding boxes).

Requires Own Data (Yes): The "Custom" aspect of Custom Vision means it is used to build and train a model tailored to a specific use case. This process requires the user to upload and tag their own set of images to serve as the training data. This contrasts with the general-purpose Computer Vision service, which uses a pre-trained model.

Analyze Video Files (No): The Custom Vision service prediction API is built to analyze static images (e.g., JPEG, PNG) or image URLs. It does not have a native API endpoint that accepts video files (e.g., MP4, AVI) for analysis. While a complete solution can analyze video by first using a different process to extract individual frames and then sending those frames as images to Custom Vision, the service itself does not analyze video files.

References:

Microsoft Corporation. (2024). What is Custom Vision?. Microsoft Learn. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/ai-services/custom-vision-service/overview
Reference for Statement 1: The document states, "Custom Vision is an image recognition service that lets you build, deploy, and improve your own image... object detection models."

Reference for Statement 2: The same document notes, "You provide a set of labeled images to train your model..."
Microsoft Corporation. (2024). Tutorial: Analyze videos in near real-time with Custom Vision. Microsoft Learn. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/ai-services/custom-vision-service/analyze-video
Reference for Statement 3: This official tutorial outlines the correct pattern, which confirms the service does not analyze video files directly. The process states: "This tutorial shows how to use the Custom Vision Service API to perform... analysis on frames taken from a live video stream... The basic approach is to... break the video stream down into a sequence of frames... and submit selected frames to the Custom Vision predictor." This confirms the API target is the frame (image), not the video file.

Question: 21

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:

The core task is to extract and read text (the street names) from an image (a photograph of a street sign). Optical Character Recognition (OCR) is the specific computer vision technology designed for this purpose. It detects the presence of text in an image and then transcribes the characters into a machine-readable format. While object detection might first be used to locate the sign, OCR is the essential step to read the name itself.

Why Incorrect Options are Wrong:

A. object detection: This technique is used to locate and classify objects within an image (e.g., identifying the bounding box of a "street sign"), but it does not read the text on the object.
C. image classification: This assigns a single label or category to an entire image (e.g., "urban street" or "daytime"). It does not extract specific information like text from within the image.
D. facial recognition: This is a specialized technology used exclusively for identifying and analyzing human faces in images or videos, which is irrelevant to reading a street sign.

References:

1. Microsoft Learn, AI-900: Explore computer vision in Microsoft Azure, "Explore optical character recognition" unit. This document states, "Optical character recognition (OCR) is a technique used to detect and read text in images. You can use the Computer Vision service to read text in images..." This directly supports the use of OCR for reading text from signs.
2. Microsoft Learn, AI-900: Explore computer vision in Microsoft Azure, "Explore object detection" unit. This source explains, "Object detection is a form of computer vision in which a model is trained to classify individual objects within an image, and indicate their location with a bounding box." This clarifies that object detection locates objects but does not read text.
3. Microsoft Learn, AI-900: Explore computer vision in Microsoft Azure, "Introduction" unit. This unit provides an overview of computer vision tasks, distinguishing between image classification (what is the image about?), object detection (what objects are in the image, and where are they?),

and OCR (reading text in the image).

Question: 22

HOTSPOT Select the answer that correctly completes the sentence.

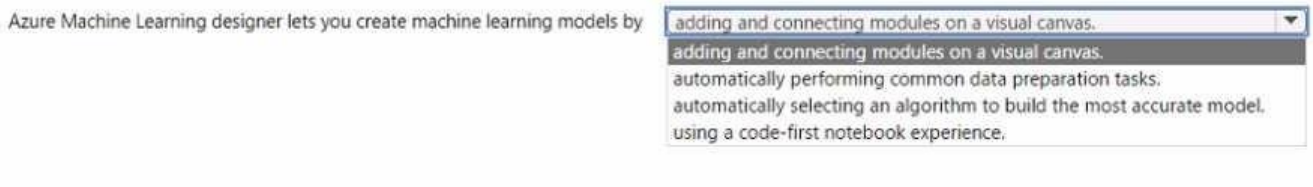

Answer:

adding and connecting modules on a visual canvas

Explanation:

The Azure Machine Learning designer is a low-code/no-code tool within the Azure ML workspace. Its primary interface is a visual canvas where users build end-to-end ML pipelines by dragging, dropping, and connecting pre-built modules for tasks like data import, transformation, model training, and scoring.

• Option 3 (automatically selecting an algorithm...) is incorrect as this specifically describes Automated ML (AutoML), a separate feature.

• Option 4 (using a code-first notebook experience) is incorrect as this describes using Jupyter notebooks with the Azure ML SDK, which is the code-based alternative to the visual designer.

References:

Microsoft. (2024, May 22). What is Azure Machine Learning designer? Microsoft Learn.
Section: Introduction, Paragraphs 1-2.
Quote: "Azure Machine Learning designer is a drag-and-drop interface... The designer provides a visual canvas where you can add datasets and modules... You connect the modules to create a pipeline draft..."
URL: https://learn.microsoft.com/en-us/azure/machine-learning/concept-designer
Microsoft. (2024, October 17). What is automated machine learning (AutoML)? Microsoft Learn.
Section: "How AutoML works."
Note: This official documentation confirms that "automatically selecting an algorithm" is the function of AutoML, distinguishing it from the designer.
URL: https://learn.microsoft.com/en-us/azure/machine-learning/concept-automated-ml
Duke University. (n.d.). AI for Everyone: Azure ML Designer. Duke AIPI (Artificial Intelligence for Product Innovation).
Section: "Azure ML Designer."
Quote: "Azure ML Designer... is a visual, drag-and-drop environment... In the Designer, you

create an ML 'pipeline' on a visual canvas by dragging and dropping 'modules'..."
URL: https://aipi.pratt.duke.edu/azure-ml-designer
Microsoft. (2024, September 24). What are Jupyter notebooks in Azure Machine Learning? Microsoft Learn.
Section: Introduction.
Note: This documentation defines the "code-first notebook experience" as a separate component (Jupyter notebooks) within the Azure ML workspace, confirming it is distinct from the designer.
URL: https://learn.microsoft.com/en-us/azure/machine-learning/concept-notebooks

Question: 23

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

A. Azure 8ol Service

B. Azure Machine Learning

C. Translator

D. Language Service

Answer:

A, D

Explanation:

To create a chatbot that answers questions from a predefined knowledge base, you need two core components. First, the Azure Bot Service provides the framework to build, deploy, and manage the bot, enabling it to interact with users on various channels. Second, the Azure AI Language Service, specifically its Question Answering feature (formerly QnA Maker), is used to build the knowledge base from existing content like FAQs. This service allows the bot to understand natural language questions and map them to the correct predefined answers. Together, these services provide a complete solution for a conversational Q&A experience, directly addressing the need to reduce the load on telephone operators.

Why Incorrect Options are Wrong:

B. Azure Machine Learning: This is a platform for building, training, and deploying custom machine learning models, which is overly complex for a simple Q&A bot with predefined answers.
C. Translator: This service is used for text translation between languages. The scenario does not specify a requirement for multilingual support, making this service unnecessary for the core goal.

References:

1. Microsoft Learn, Azure AI Language documentation. "What is question answering?" This document states, "Question answering provides cloud-based Natural Language Processing (NLP) that allows you to create a natural conversational layer over your data... It is commonly used to build conversational client applications, which include social media applications, chat bots, and speech-enabled desktop applications."
Source: Microsoft, "What is question answering?", Azure AI Language documentation. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/overview

2. Microsoft Learn, Azure Bot Service documentation. "What is Azure Bot Service?" This page explains, "Azure Bot Service is a comprehensive development environment for building enterprise-grade conversational AI... Bots can be used to shift simple, repetitive tasks, such as taking a dinner reservation or gathering profile information, on to automated systems that may no longer require direct human intervention."
Source: Microsoft, "What is Azure Bot Service?", Azure Bot Service documentation. Retrieved from https://learn.microsoft.com/en-us/azure/bot-service/bot-service-overview
3. Microsoft Learn, Tutorial. "Create a question answering bot". This tutorial explicitly demonstrates the required architecture: "In this tutorial, you learn how to: 1. Create a question answering project and import a file as a knowledge base. 2. Add your knowledge base to a bot... 3. Build and run your bot." This confirms the direct integration of the Language Service (for the knowledge base) and the Bot Service (for the bot itself).
Source: Microsoft, "Quickstart: Create a question answering bot", Azure AI Language documentation. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/quickstarts/bot-service

Question: 24

DRAG DROP Match the machine learning models to the appropriate deceptions. 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.

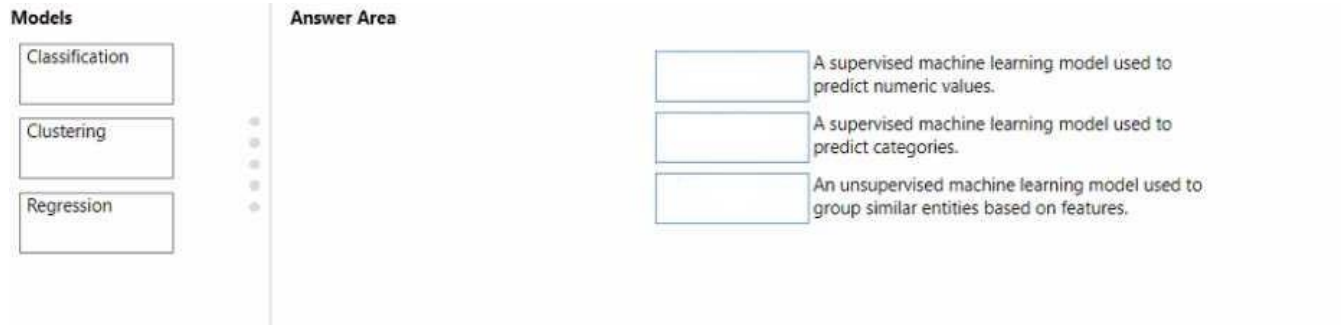

Answer:

Regression

Classification

Clustering

Explanation:

The solution correctly maps the three fundamental types of machine learning models to their definitions.

• Regression is a supervised learning task used to predict a continuous numeric value, such as the price of a house or a future temperature.

• Classification is also a supervised learning task, but it is used to predict a discrete category or class, such as whether an email is 'spam' or 'not spam', or if a tumor is 'benign' or 'malignant'.

• Clustering is an unsupervised learning task. It does not use pre-defined labels; instead, it analyzes the input data to identify natural groupings (clusters) of items based on their shared features or similarities.

References:

Microsoft. (2024). "What is machine learning?" Azure Machine Learning Documentation. Microsoft.
Section: "Machine learning model types" "Supervised learning"
Quote/Paraphrase: This documentation explicitly defines regression as a supervised method for

predicting continuous values (e.g., price, sales). It defines classification as a supervised method for predicting categories (e.g., yes/no, true/false). It defines clustering as an unsupervised method used to discover structure and group items into clusters based on similarity.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). An Introduction to Statistical Learning: with Applications in R (Second Edition). Springer.

Chapter 2, Section 2.1.2 "Supervised and Unsupervised Learning": This section distinguishes the two main types. It states that supervised learning involves building a model for predicting an output based on inputs, further breaking this down into regression problems (predicting a quantitative or numeric output) and classification problems (predicting a qualitative or categorical output).

Chapter 12, Section 12.1 "Unsupervised Learning": This section defines unsupervised learning as a setting with only feature measurements (X) and no response variable (Y). The goal is described as finding interesting patterns or groups, which directly relates to clustering.

Ng, A. (2023). "Course Notes: CS229 - Machine Learning." Stanford University.

Section: "Part I: Supervised Learning"

Quote/Paraphrase: The notes define supervised learning as the task of learning a function that maps inputs to outputs given a set of input-output pairs. It specifies that if the target output (label) is continuous (e.g., price), the task is regression. If the target output is discrete (e.g., 'cat' or 'dog'), the task is classification.

Section: "Part IX: Unsupervised Learning"

Quote/Paraphrase: This section describes unsupervised learning as the process of finding structure in unlabeled data. The canonical example provided is clustering, such as grouping news articles by topic or customers by preferences.

Question: 25

HOTSPOT Select the answer that correctly completes the sentence.

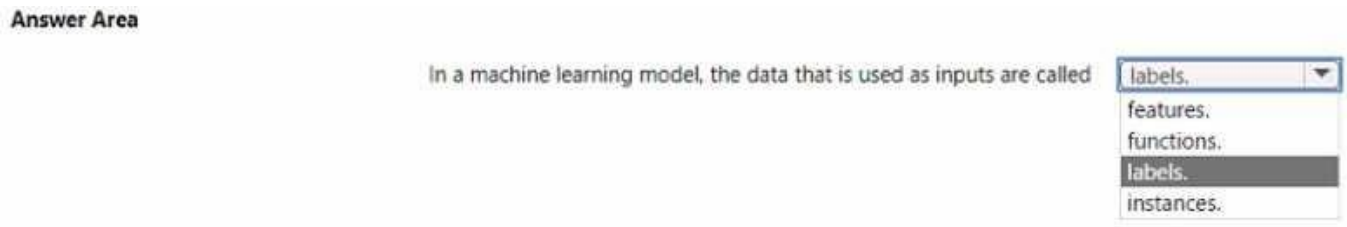

Answer:

features

Explanation:

In the context of machine learning, features are the defined as the independent variables or attributes that serve as the inputs to a model. These are the measurable properties or characteristics of the data (e.g., the square footage of a house, the pixel values of an image) that the model uses to make a prediction.

Conversely, a label is the output or target variable (e.g., the price of the house, the object in the image) that the model learns to predict. An instance is a single row of data, which typically includes both its features and (if supervised) its label.

References:

Microsoft Azure Documentation. (n.d.). What is automated machine learning (AutoML)? Azure Machine Learning. Retrieved October 24, 2025.Reference: In the "Features and labels" section, the documentation states: "A feature is a data column that is used as an input for your model... A label is the data column that you want to predict."Ng, A. (n.d.). CS229 Machine Learning Course Notes: Supervised Learning. Stanford University.Reference: Section 1.1, "Supervised Learning," defines the training set as (x, y) pairs, stating: "We call $x(i)$ the input variables (or features) and $y(i)$ the output variables (or labels)."Bishop, C. M. (2006). Pattern Recognition and Machine Learning. Springer.Reference: Chapter 1, Section 1.1 (p. 2-3), introduces the input vector $\mathbfx$, whose components are referred to as features. This input vector is used to predict the target variable t (the label).Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. Journal of Machine Learning Research, 3(Mar), 1157-1182.Reference: The abstract defines features: "The variables collected from the field, which are used as inputs to a predictor, are referred to as 'features'."

Question: 26

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

A. After you clean the data.

B. Before you train a model.

C. Before you choose the type of model.

D. After you test a model on the validation data.

Answer:

D

Explanation:

Evaluation metrics are quantitative measures used to assess the performance of a machine learning model. The standard machine learning workflow involves training a model on a specific dataset (the training set) and then testing its predictive performance on a separate, unseen dataset (the validation or test set). It is only after the model has made predictions on this validation data that you can compare the predicted outcomes to the actual known values. This comparison allows for the calculation and review of evaluation metrics, such as accuracy, precision, or Root Mean Squared Error (RMSE), to determine the model's effectiveness.

Why Incorrect Options are Wrong:

A. After you clean the data, no model has been trained yet, so there are no performance metrics to evaluate.
B. Before you train a model, it has not yet learned any patterns from the data, so it cannot be evaluated.
C. Before you choose the type of model, you have not even started the training process, making evaluation impossible.

References:

1. Microsoft Learn, AI-900: Describe machine learning models. In the "Evaluate a model" unit, the documentation states, "After you've used the training dataset to train a model, you need to evaluate it to determine how well it predicts. To do this, you use a second dataset that the model hasn't seen before... You can then compare the labels predicted by the model with the actual known labels in the original dataset. From this comparison, you can calculate a range of metrics that quantify how well the model performed." This directly confirms that evaluation occurs after training and testing on holdout data.
Source: Microsoft, AI-900: Microsoft Azure AI Fundamentals, "Describe machine learning models", "Evaluate a model" section.

2. Microsoft Azure Documentation, "What is automated machine learning (AutoML)?" The process described for AutoML, a core Azure service, follows the standard machine learning lifecycle. The documentation outlines that after a model is trained, it is scored against a validation dataset, and "the metrics from that scoring are used for evaluation." This reinforces that evaluation metrics are reviewed after the model is tested.
Source: Microsoft Azure Documentation, What is automated machine learning (AutoML)?, "Train and tune models" section.
3. Stanford University, CS229: Machine Learning Course Notes. In the section on model selection, the standard procedure is outlined: "The standard way to do this is to split it into a training set, a validation set (also called a hold-out cross validation set), and a test set... we can then select the model that did best on the validation set." This academic source confirms that model performance (via metrics) is reviewed on validation data after training.
Source: Ng, A. & Thrun, S. (2023). CS229: Machine Learning Course Notes, "Part V: Learning Theory", Section on "Model Selection and Train/Validation/Test Sets". Stanford University.

Question: 27

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, now part of Azure AI Document Intelligence, is the correct service for this task. It is specifically designed to extract text, structure, and key-value pairs from documents. A key feature of this service is its ability to perform Personally Identifiable Information (PII) detection. By enabling this feature during analysis, the service can identify and redact sensitive information such as names, addresses, social security numbers, and financial data directly within the documents being processed. This makes it an ideal solution for compliance and data privacy scenarios.

Why Incorrect Options are Wrong:

A. Custom Vision is used for building custom image classification and object detection models; it analyzes visual content, not text within documents for sensitive information.
B. Conversational Language Understanding is designed to understand user goals, or intents, in conversational text (e.g., chatbots), not for scanning entire documents for PII.

References:

1. Microsoft Learn, Azure AI Document Intelligence, "PII detection feature": "The PII detection feature in Document Intelligence can identify and redact sensitive information in your documents. The feature is part of the Document Intelligence service and can be enabled by setting the optional features query parameter to pii-detection." (This directly supports the correct answer C).
2. Microsoft Learn, Azure AI services, "What is Custom Vision?": "Custom Vision is an image recognition 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." (This supports why option A is incorrect).
3. Microsoft Learn, Azure AI Language, "What is conversational language understanding (CLU)?": "Conversational language understanding (CLU) is a cloud-based conversational 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." (This supports why option B

is incorrect).

Question: 28

HOTSPOT Select the answer that correctly completes the sentence.

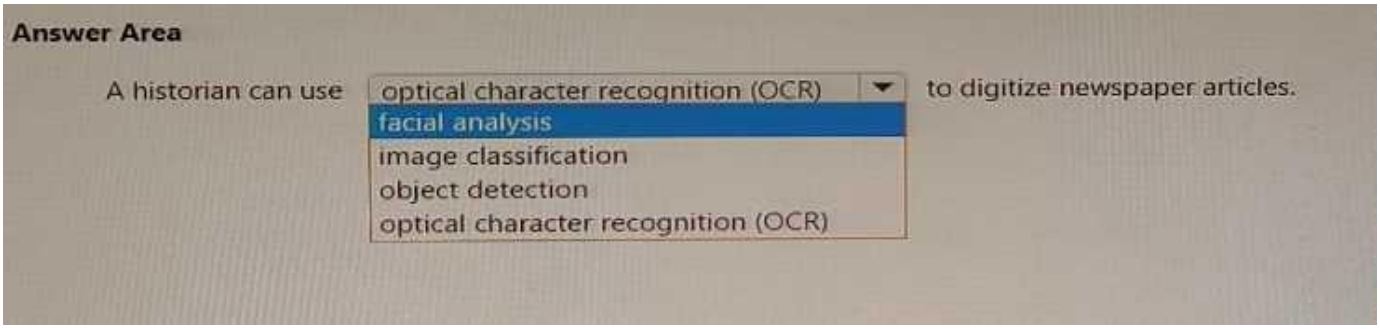

Answer:

optical character recognition (OCR)

Explanation:

The core task is to "digitize newspaper articles," which implies converting the printed text from a scanned image into machine-readable, searchable, and editable data. Optical Character Recognition (OCR) is the specific computer vision technology designed to identify and extract printed or handwritten text from images.

The other options are incorrect as they serve different purposes:

• Facial analysis detects and analyzes human faces.

• Image classification assigns a single label to an entire image (e.g., "newspaper").

• Object detection locates and identifies specific objects within an image (e.g., a photograph, an advertisement) but does not extract the text content.

References:

Microsoft. (2024). What is Optical Character Recognition (OCR)? Azure AI Vision documentation. This official documentation states: "The Optical Character Recognition (OCR) feature of Azure AI Vision extracts printed or handwritten text from images and documents... This is useful for various scenarios, including... digitizing print media like books, articles, and reports to make them searchable."
Smith, R. (2007). An Overview of the Tesseract OCR Engine. Proceedings of the Ninth International Conference on Document Analysis and Recognition (ICDAR 2007), vol. 2, pp. 629-633. This paper details an open-source OCR engine widely used for digitizing text from scanned documents, a core task for historical archives. (DOI:

https://doi.org/10.1109/ICDAR.2007.4376991)
Szeliski, R. (2010). Computer Vision: Algorithms and Applications. (Referenced in MIT OpenCourseWare, 6.819 / 6.869: Advances in Computer Vision). Chapter 14, Section 14.4.2 "Text recognition (OCR)" specifically defines OCR as the task of converting images of text into character codes.
Zou, Z., et al. (2023). Object Detection in 20 Years: A Survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(1), pp. 355-374. Section I (Introduction) clearly distinguishes object detection (determining the location and category of objects) from image classification (categorizing an entire image) and OCR (text extraction), confirming they are separate tasks. (DOI: https://doi.org/10.1109/TPAMI.2023.3238524)

Question: 29

HOTSPOT Select the answer that correctly completes the sentence.

Answer:

numeric

Explanation:

In supervised machine learning, regression is the task of predicting a continuous quantitative value. The label (also known as the target or dependent variable) is the specific value the model is being trained to predict. By definition, a continuous quantitative value-such as a price, temperature, or a count-is represented by a numeric data type (e.g., float or integer).

In contrast, boolean (True/False) or text (categorical) data types are used as labels for classification tasks, where the goal is to predict a discrete class or category.

References:

Microsoft Azure Documentation: In the context of Azure Machine Learning, the documentation for automated ML tasks explicitly states the requirement for regression. For the "Regression" task type, the "Label column data type" must be numeric (integer or decimal).Source: Microsoft Azure. (2024). Set up automated ML to train a model (v2). Microsoft Learn. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/machine-learning/how-to-configure-auto-train?view=azureml-api-2. (See the table under the "Data source and format" section).Academic Publication (Textbook): Reputable machine learning textbooks define regression by the nature of its output variable.Source: Hastie, T., Tibshirani, R., & Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer. In Chapter 2, Section 2.2, it is stated: "Each measurement... has a quantitative output Y, such as price... This is a regression problem." Quantitative outputs are inherently numeric. (Page 9).Peer-Reviewed Journal: Academic literature consistently distinguishes regression from classification based on the label's data type.Source: Kotsiantis, S. B. (2007). Supervised Machine Learning: A Review of Classification Techniques. Informatica (Slovenia), 31(3), 249-268. The introduction states, "In machine learning, a typical supervised learning task is... regression (predicting a numeric value)." (Section 1, p. 249).

Question: 30

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:

The Azure Computer Vision service, part of Azure AI Services, provides pre-trained models to analyze images. A key feature of its Image Analysis capability is brand detection, which can identify thousands of commercial brands and logos in images. This service is specifically designed for scenarios like identifying products on a supermarket shelf. The bot would send the image to the Computer Vision API and receive data identifying any recognized brands, their location in the image, and a confidence score. This directly addresses the requirement without needing to train a custom model.

Why Incorrect Options are Wrong:

A. AI enrichment for Azure Search is used to extract insights from unstructured data to make it searchable, not as the primary service for real-time brand identification in a bot.
C. Custom Vision is used to train your own image recognition models. While possible, the standard Computer Vision service already has a pre-built capability for detecting common brands.
D. Language understanding capabilities are used to process and understand text or spoken language, not to analyze the content of images.

References:

1. Microsoft Learn: AI-900 Analyze images with the Computer Vision service. This module explicitly describes the capabilities of the Computer Vision service.
Reference: In the "Introduction" and "Detect common objects in images" units, it details the service's ability to return information about visual content, including brand detection.
2. Microsoft Learn: What is Computer Vision? This official documentation provides a high-level overview of the service's features.
Reference: Under the "Image analysis" section, "Brand detection" is listed as a specific feature: "Detect commercial brands in images from a database of thousands of global logos."
3. Microsoft Learn: What is Custom Vision? This document clarifies the purpose of the Custom

Vision service, distinguishing it from the pre-trained Computer Vision service.
Reference: The "Overview" section states, "Custom Vision is an image recognition service that lets you build, deploy, and improve your own image identifiers." This highlights its use for custom, not pre-built, recognition tasks.
4. Microsoft Learn: What is Language Understanding (LUIS)? This document explains the purpose of language services.
Reference: The "Overview" section defines it as "a cloud-based conversational AI service that applies custom machine-learning intelligence to a user's conversational, natural language text to predict overall meaning." This confirms it is for text, not images.

Question: 31

HOTSPOT Select the answer that correctly completes the sentence.

Answer:

an anomaly detection workload

Explanation:

Anomaly detection is a machine learning technique specifically designed to identify rare events, unexpected items, or unusual observations in data that differ significantly from the majority. Detecting "unusual temperature fluctuations" is a classic example of an anomaly detection workload. The goal is to find data points (temperatures) that fall outside the normal, expected operating range, which could indicate a fault or impending failure. The other options are incorrect as the data is numerical (time-series), not visual (computer vision) or text-based (NLP), and the specific goal is not to index unstructured data (k n o w l e d ge mining).

References:

Microsoft Azure Documentation. (2024). "What is the Anomaly Detector service?" Azure AI services documentation. Retrieved from Microsoft Learn.
Reference: Section: "Overview," Paragraph 1. This document states the service is used to "monitor data over time and detect anomalies" and explicitly lists "spotting unusual trends in business metrics, or monitoring machine health" as key use cases.
Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys (CSUR), 41(3), Article 15, pp. 15:1-15:2.
Reference: Section I (Introduction) defines anomalies as "patterns in data that do not conform to expected behavior." Section II.A ("Application Domains") cites "Industrial Damage Detection," which involves monitoring sensor data (such as temperature, vibration) from heavy machinery to detect abnormal patterns indicating damage.
DOI: https://doi.org/10.1145/1541880.1541882
Ng, A. (n.d.). "Lecture Notes 9: Anomaly Detection." Stanford University, CS 229: Machine Learning.
Reference: Section 1 ("Problem motivation"). This university courseware uses the example of monitoring computers in a data center, using features like temperature and fan speed, to detect "unusual" (anomalous) behavior that could predict an impending failure. This directly maps to the

question's scenario.

Question: 32

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:

Object detection

Image classification

Explanation:

Object detection is the correct technique for the first task because its primary function is to locate instances of objects within an image and provide their specific positions, typically by outputting bounding box coordinates.

Image classification is the correct technique for the second task. Its function is to analyze the pixels of an image (or a detected region) and assign one or more categorical labels to it. Identifying the species of a bird is a classic classification problem, as it assigns a specific category (e.g., "Robin," "Sparrow") to the detected object.

References:

Microsoft Azure AI Vision Documentation. (2024). Object detection (Image Analysis 4.0). "Object detection is similar to tagging, but the API returns the bounding box coordinates (in pixels) for each object found in the image... This lets you get the coordinates of the objects found in an image." This source confirms that object detection is used to find the location (coordinates) of objects.
Source Link:
https://learn.microsoft.com/en-us/azure/ai-services/vision-services/concept-object-detection-4-0
Microsoft Azure AI Vision Documentation. (2024). Image classification (Image Analysis 4.0). "The Image Analysis 4.0 'classify' API lets you classify an image based on a taxonomy of categories.

The classifier can use the default built-in model, or a custom model you've trained on your own categories." This source confirms that image classification is used to assign a category (like a species) to an image.
Source Link: https://learn.microsoft.com/en-us/azure/ai-services/vision-services/concept-image-classification-4-0
Zhao, Z. Q., Zheng, P., Xu, S. T., & Wu, X. (2019). Object Detection With Deep Learning: A Review. IEEE Transactions on Neural Networks and Learning Systems, 30(11), 3212-3232.
Reference: Section II-A, "Related Tasks," defines image classification as a task that "aims to assign a semantic label to an entire image," and object detection as a task that "aims to localize... all instances of target objects in an image... and classify these instances." The question explicitly separates the "locate" (localization) and "identify" (classification) tasks, making these the two most precise answers.
DOI: https://doi.org/10.1109/TNNLS.2018.2876865

Question: 33

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

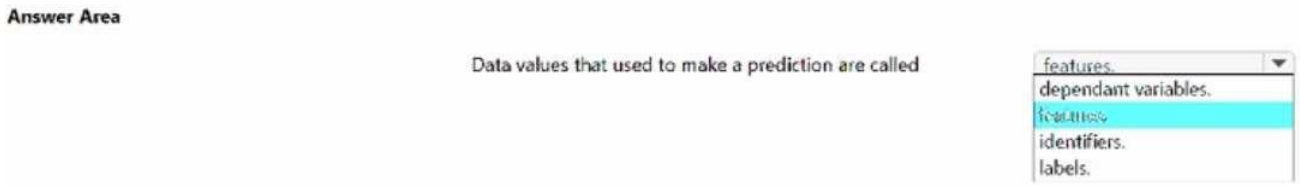

Answer:

features

Explanation:

In machine learning, features are the input variables used to make a prediction. They represent the measurable, individual characteristics or attributes of the data. For example, if the goal is to predict a student's exam score, the features might include "hours studied," "previous grades," and "class attendance."

The label (or dependant variable in statistics) is the output-the value you are trying to predict (e.g., the actual exam score). Identifiers (like a Student ID) are unique keys that are typically excluded from a model as they have no predictive value.CertEmpire

References:

Microsoft. "Features and labels." What is machine learning? - Azure AI Fundamentals. Microsoft Learn documentation (AI-900 path). Accessed Oct 24, 2025.Quote: "The features are the characteristics of the item... For example, a flower's features might include its measurements for petals and stems. The label is the thing we're trying to predict... For example, the flower's species."Ng, A. (2018). "Supervised Learning." CS229 Machine Learning Course Notes. Stanford University. p. 1.Quote: "We are given a data set and already know what our correct output should look like... In this regression problem, x are the "input" variables (or features) and y is the "output" or target variable (also called the label) that we are trying to predict."Alpaydn, E. (2020). Introduction to Machine Learning (4th ed.). MIT Press. Chapter 2, Section 2.1, p. 25.Quote: "Each instance $\mathbfx$ is a vector of d inputs, $\mathbfx = x1, x2, \dots, xdT$... The xj, $j=1,\dots,d$ are the features."Verencar, J., et al. (2023). "A Comprehensive Survey on Feature Engineering and Feature Selection in the Era of Big Data." IEEE Access, vol. 11, p. 100319. DOI: 10.1109/ACCESS.2023.3314988.Quote: "In ML, a feature is an individual measurable property or characteristic of a phenomenon being observed. Features are the inputs of an ML model..."

Question: 34

HOTSPOT Select the answer that correctly completes the sentence.

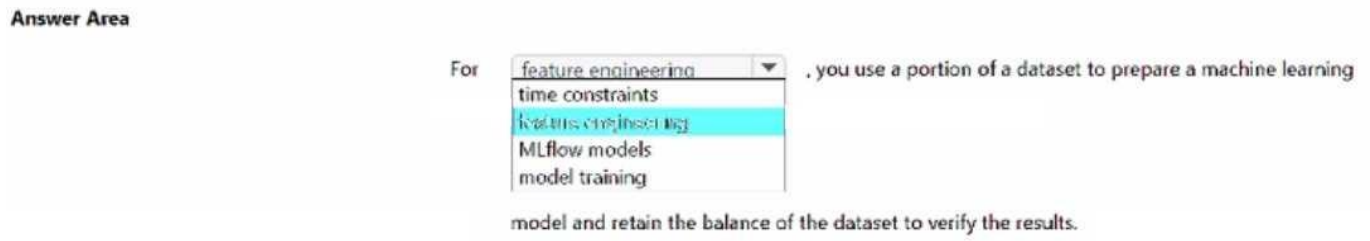

Answer:

model evaluation

Explanation:

The sentence describes the standard methodology for testing a machine learning model's performance on unseen data. This entire process is known as model evaluation.

A dataset is split into at least two parts:

• A training set ("a portion of a dataset") is used "to prepare" (i.e., train) the model.

• A testing/validation set ("the balance of the dataset") is withheld from the training process and used "to verify the results" (i.e., evaluate the model's accuracy and generalization).

This separation is essential to prevent overfitting and to get an unbiased estimate of how the model will perform on new, real-world data. "Model training" is only the first part of this process, not the overarching purpose.

References:

Microsoft. (n.d.). Split Data module. Azure Machine Learning documentation. Retrieved October 24, 2025, from https://learn.microsoft.com/en-us/azure/machine-learning/algorithm-module-reference/split-data
Reference (Module Overview): The official documentation for the "Split Data" module in Azure Machine Learning states, "This module is particularly useful when you need to separate data into training and testing sets... This is a common task in machine learning, to support model evaluation."
James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). An Introduction to Statistical Learning: with Applications in R. Springer. https://doi.org/10.1007/978-1-4614-7138-7
Reference (Chapter 5, Section 5.1.1, "The Validation Set Approach"): This section describes the process: "...randomly dividing the available set of observations into two parts, a training set and a

validation set... The model is fit on the training set, and the fitted model is used to predict the responses for the observations in the validation set." This entire procedure is presented as a method for estimating the test error, which is the core of model evaluation.

MIT OpenCourseWare. (2020). Lecture 2: Deep Sequence Modeling. 6.S191 Introduction to Deep Learning. Massachusetts Institute of Technology.

Reference (Video Timestamp 50:18, "Splitting Data"): The lecture explicitly discusses splitting data into training, validation, and test sets. The purpose of the validation and test sets is "for evaluation" to check the model's generalization capabilities on data it has not been trained on.

Question: 35

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

A. privacy and security

B. inclusiveness

C. transparency

D. reliability and safety

Answer:

A

Explanation:

The act of creating a model using data obtained without permission is a direct breach of the privacy and security principle. This principle mandates that AI systems must be secure and respect privacy. It requires transparency about the collection, use, and storage of data, and ensures that individuals have appropriate control over how their data is used. Obtaining data without consent fundamentally violates the privacy rights of the individuals whose data was collected, making this the most relevant principle.

Why Incorrect Options are Wrong:

B. inclusiveness: This principle focuses on designing AI systems to empower and engage all people, avoiding exclusion. It is not directly related to data acquisition consent.
C. transparency: This principle is about ensuring that AI systems are understandable. It relates to explaining how a model works and its limitations, not the ethical sourcing of its training data.
D. reliability and safety: This principle ensures that AI systems operate consistently, safely, and as intended. While unvetted data could affect reliability, the core ethical breach is about consent, not performance.

References:

1. Microsoft Learn. (2024). Microsoft's principles for responsible AI. In "AI-900: Describe AI concepts and workloads". Microsoft. Retrieved from https://learn.microsoft.com/en-us/training/modules/get-started-ai-fundamentals/5-responsible-ai-principles.
Reference Details: In the "Privacy and security" section, it states, "The data used to train the AI system should be analyzed to ensure that the privacy of individuals is protected." This directly links the training data's handling to the privacy principle.
2. Microsoft Cloud Adoption Framework for Azure. (2024). What is responsible AI?. Microsoft. Retrieved from https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/responsible-ai.

Reference Details: The "Privacy and security" section explicitly states, "AI systems must comply with privacy laws that require transparency about the collection, use, and storage of data and that mandate that consumers have appropriate controls to choose how their data is used." Using data without permission is a clear violation of this mandate.