

Microsoft.AI-900.v2026-09-24.q146

Exam Code:	AI-900
Exam Name:	Microsoft Azure AI Fundamentals
Certification Provider:	Microsoft
Free Question Number:	146
Version:	v2026-09-24
# of views:	132
# of Questions views:	1460
https://www.freepdfdumps.com/Microsoft.AI-900.v2026-09-24.q146.html	

NEW QUESTION: 1

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

Answer Area

Statements	Yes	No
Labelling is the process of tagging training data with known values.	☐	☐
You should evaluate a model by using the same data used to train the model.	☐	☐
Accuracy is always the primary metric used to measure a model's performance.	☐	☐

Answer:

Answer Area

Statements	Yes	No
Labelling is the process of tagging training data with known values.	☒	☐
You should evaluate a model by using the same data used to train the model.	☐	☒
Accuracy is always the primary metric used to measure a model's performance.	☐	☒

Explanation:

Statements	Yes	No
Labelling is the process of tagging training data with known values.	☒	☐
You should evaluate a model by using the same data used to train the model.	☐	☒
Accuracy is always the primary metric used to measure a model's performance.	☐	☒

This question evaluates understanding of fundamental machine learning concepts as covered in the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Explore the machine learning process." These statements relate to data labeling, model evaluation practices, and performance metrics-three essential parts of building and assessing a machine learning model.

* Labelling is the process of tagging training data with known values # Yes According to Microsoft Learn,

"Labeling is the process of tagging data with the correct output value so the model can learn relationships between inputs and outputs." This is essential for supervised learning, where models require historical data with known outcomes. For example, if training a model to recognize fruit images, each image is labeled as "apple," "banana," or "orange." Hence, this statement is true.

* You should evaluate a model by using the same data used to train the model # No The AI-900 guide stresses that using the same data for both training and evaluation can cause overfitting, where the model performs well on training data but poorly on unseen data. Instead, the dataset is split into training and testing (or validation) subsets. Evaluation must use test data that the model has never seen before to ensure an unbiased measure of performance. Therefore, this statement is false.

* Accuracy is always the primary metric used to measure a model's performance # No Microsoft Learn emphasizes that accuracy is only one metric and not always the best choice. Depending on the problem type, other metrics such as precision, recall, F1-score, or AUC (Area Under the Curve) may be more appropriate-especially in cases with imbalanced datasets. For example, in fraud detection, recall may be more important than accuracy. Thus, this statement is false.

NEW QUESTION: 2


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

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

NOTE: Each correct selection is worth one point.

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

Answer:



Explanation:

Box 1: Reliability and safety

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

Box 2: accountability

Box 3: Privacy and security

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

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

NEW QUESTION: 3

You have a chatbot that answers technical questions by using the Azure OpenAI GPT-3.5 large language model (LLM). Which two statements accurately describe the chatbot? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. Grounding data can be used to constrain the output of the chatbot.
- B. The chatbot will always provide accurate data.
- C. The chatbot might respond with inaccurate data.
- D. The chatbot is suitable for performing medical diagnosis.

Answer: (SHOW ANSWER)

The correct answers are A. Grounding data can be used to constrain the output of the chatbot and C. The chatbot might respond with inaccurate data.

According to the Microsoft Azure AI Fundamentals (AI-900) study material and Microsoft Learn modules on Azure OpenAI, a chatbot built with Azure OpenAI GPT-3.5 is a large language model (LLM) capable of generating natural language responses. However, these models operate based on statistical patterns learned from massive text datasets-they do not inherently guarantee factual accuracy. Hence, while GPT-based models can produce highly coherent text, they may sometimes generate inaccurate, outdated, or fabricated information (commonly referred to as "hallucinations"). This makes C correct.

Grounding data, as described in Microsoft's Responsible AI and Azure OpenAI grounding documentation, refers to integrating trusted external data sources-such as company documents, databases, or knowledge bases-into the prompt context. This helps the model stay aligned with factual or domain-specific content, effectively constraining its output to be relevant and verifiable. Therefore, A is also correct.

Options B and D are incorrect because GPT models do not always provide accurate information, and they are not approved for critical use cases such as medical diagnosis. Microsoft's Responsible AI principles explicitly prohibit unverified use in healthcare or other high-risk domains.

Thus, the verified answers are A and C.

NEW QUESTION: 4

You need to build an app that will read recipe instructions aloud to support users who have reduced vision.

Which version service should you use?

A. Text Analytics

B. Translator Text

C. Speech

D. Language Understanding (LUIS)

Answer: C (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of speech capabilities in Azure Cognitive Services", the Azure Speech service provides functionality for converting text to spoken words (speech synthesis) and speech to text (speech recognition).

In this scenario, the app must read recipe instructions aloud to assist users with visual impairments. This task is achieved through speech synthesis, also known as text-to-speech (TTS). The Azure Speech service uses advanced neural network models to generate natural-sounding voices in many languages and accents, making it ideal for accessibility scenarios such as screen readers, virtual assistants, and educational tools.

Microsoft Learn defines Speech service as a unified offering that includes:

* Speech-to-text (speech recognition): Converts spoken words into text.

* Text-to-speech (speech synthesis): Converts written text into natural-sounding audio output.

* Speech translation: Translates spoken language into another language in real time.

* Speaker recognition: Identifies or verifies a person based on their voice.

The other options do not fit the requirements:

* A. Text Analytics - Performs text-based natural language analysis such as sentiment, key phrase extraction, and entity recognition, but it cannot produce audio output.

* B. Translator Text - Translates text between languages but does not generate speech output.

* D. Language Understanding (LUIS) - Interprets user intent from text or speech for conversational bots but does not read text aloud.

Therefore, based on the AI-900 curriculum and Microsoft Learn documentation, the correct service for converting recipe text to spoken audio is the Azure Speech service.

Final answer: C. Speech

Reference:Microsoft Learn - "What is the Speech service?" (Azure Cognitive Services > Speech)

NEW QUESTION: 5

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

A. Presence penalty

B. Temperature

C. Stop sequence

D. Max responseB

Answer: B (LEAVE A REPLY)

In a chat solution using the Azure OpenAI GPT-3.5 model, the temperature parameter controls the creativity and variability of generated responses. According to the Microsoft Learn documentation for Azure OpenAI Service, temperature is a float value typically between 0 and 2, determining how deterministic or random the model's output is. A lower temperature (e.g., 0-0.3) makes responses more focused and deterministic, while a higher temperature (e.g., 0.8-1.2) produces more verbose, creative, and diverse responses.

When you want the chat model to generate more detailed or expressive output, increasing the temperature encourages the model to explore a broader range of possible continuations, leading to longer and more varied text. This parameter directly affects how "verbose" or elaborate the model's responses can be, which is why it is the correct answer.

The other options are not appropriate for this scenario:

* A. Presence penalty reduces repetition by discouraging reuse of the same phrases but does not control verbosity.

* C. Stop sequence defines tokens where generation should stop, limiting rather than extending response length.

* D. Max response (max tokens) controls the maximum length of the response but does not inherently make answers more verbose or expressive.

Thus, to encourage more elaborate and detailed output from the Azure OpenAI GPT-3.5 model, the correct configuration parameter to adjust is Temperature (B).

NEW QUESTION: 6

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

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

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

Answer: A (LEAVE A REPLY)

Microsoft's Responsible AI Principles, as outlined in the AI-900 certification materials and official Microsoft documentation, emphasize six guiding principles: fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. The principle of transparency means that AI systems should be designed so their decisions and processes are understandable and explainable to users and stakeholders.

In this scenario, the AI system is being developed to decide whether a loan should be approved. Such a decision directly affects people's lives and finances, so it is essential that the system can explain which factors influenced its decision-for example, credit score, income, or payment history. Microsoft's Responsible AI framework stresses that transparency helps ensure trust between humans and AI systems.

When decisions are explainable, users can understand and contest the reasoning if necessary.

The other options do not align precisely with this scenario:

- * B. Inclusiveness focuses on making AI accessible to all people, regardless of ability or background.
- * C. Fairness ensures that AI systems treat all individuals equally and do not discriminate. While fairness is important for loan assessment, the question specifically highlights the need for explainability, not equality.
- * D. Privacy and Security deals with safeguarding user data, which is separate from explaining decisions.

Therefore, the principle demonstrated here is transparency, as it ensures decision-making processes are clear, explainable, and traceable-directly aligning with Microsoft's responsible AI guidance.


Reference:Microsoft Learn: "Identify guiding principles for responsible AI" - explains that transparency ensures AI systems are understandable and their decision logic can be examined and explained.

NEW QUESTION: 7

Select the answer that correctly completes the sentence.



Answer:



Explanation:

Answer Area

Object detection

can be used to identify the location of products on a conveyor belt.

According to the Microsoft Azure AI Fundamentals (AI-900) curriculum and Microsoft Learn's modules on Computer Vision, object detection is the AI technique used to identify and locate multiple objects within an image. Unlike simple image classification, which only labels an entire image with a single category (for example, "This is a product"), object detection not only identifies the type of object but also pinpoints its exact position by providing bounding boxes and coordinates within the image.

In the scenario described - identifying the location of products on a conveyor belt - the system must be able to detect multiple items simultaneously and determine their spatial positions. Object detection algorithms (such as YOLO, Faster R-CNN, or SSD) are specifically designed for this purpose. This allows automation systems, like robotic arms or quality inspection systems, to track product locations in real time for sorting, packaging, or defect detection.

Let's evaluate the other options:

- * Image classification only determines what is in the image, not where it is located. It cannot handle multiple objects or their positions.
- * Image processing involves operations like resizing, filtering, or adjusting contrast, not understanding object placement.
- * Optical character recognition (OCR) extracts text from images and documents, unrelated to locating physical items.

Thus, per Microsoft Learn's AI-900 guidance, object detection is the correct computer vision capability when a task requires both identification and spatial localization of items in an image or video stream.

Final answer Object detection

NEW QUESTION: 8

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.

The screenshot shows the Azure Machine Learning Designer interface. On the left, under the 'Actions' tab, there is a list of actions: 'Evaluate the model against the validation dataset', 'Ingest and prepare a dataset', 'Split the data randomly into training data and validation data', 'Train the model', and 'Evaluate the model against the validation dataset'. On the right, there is an 'Answer Area' with a dashed red border and a list of the same five actions. The actions are arranged in the correct sequence: 'Ingest and prepare a dataset', 'Split the data randomly into training data and validation data', 'Train the model', and 'Evaluate the model against the validation dataset'. The 'Evaluate the model against the validation dataset' action is highlighted with a red dashed box. The 'Answer Area' also has a 'Move' button and a 'Reset' button.

Answer:

Explanation:

Actions

Evaluate the model against the original dataset.

Answer Area

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



Microsoft



According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of common machine learning types", the standard workflow for creating and deploying a machine learning model - especially within Azure Machine Learning Designer - follows a structured sequence of steps to ensure that the model is trained effectively and evaluated correctly. Here's the detailed breakdown of the correct order:

- * Import and prepare a dataset: This is always the first step in the machine learning lifecycle. The dataset is imported into Azure Machine Learning and cleaned or preprocessed. Preparation might include handling missing values, normalizing data, removing outliers, and encoding categorical variables. This ensures the dataset is ready for modeling.
- * Split the data randomly into training data and validation data: The dataset is then divided into two parts - the training set and the validation (or testing) set. Typically, around 70-80% of the data is used for training and 20-30% for validation. This step ensures that the model can be evaluated on unseen data later, preventing overfitting.
- * Train the model: During this stage, the machine learning algorithm learns patterns from the training data. Azure Machine Learning Designer provides multiple algorithms (classification, regression, clustering, etc.) that can be applied using "Train Model" components.
- * Evaluate the model against the validation dataset: Finally, the trained model's performance is tested using the validation dataset. Evaluation metrics such as accuracy, precision, recall, or RMSE (depending on the model type) are calculated to assess how well the model generalizes to new data.

The incorrect option - "Evaluate the model against the original dataset" - is not used in proper ML workflows, because evaluating on the same data used for training would give misleadingly high accuracy due to overfitting.

NEW QUESTION: 9

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

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

Answer:

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

Explanation:

Yes, Yes, and No.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn modules under the topic "Describe features of common AI workloads", conversational AI solutions like chatbots are used to automate and enhance customer interactions. A chatbot is an AI service capable of understanding user inputs (text or voice) and providing appropriate responses, often integrated into websites, mobile apps, or messaging platforms.

* A restaurant can use a chatbot to empower customers to make reservations using a website or an app - Yes. This statement is true because conversational AI is designed to handle structured tasks such as booking, scheduling, and information retrieval. Chatbots built with Azure Bot Service can connect to backend systems (like a reservation database) to let customers make or modify reservations through a chat interface. The AI-900 study guide explicitly notes that chatbots can help businesses "automate processes such as booking or reservations" to improve efficiency and customer experience.

* A restaurant can use a chatbot to answer inquiries about business hours from a webpage - Yes. This is also true. Chatbots can be trained using QnA Maker (now integrated into Azure AI Language) or Azure Cognitive Services for Language to answer common customer questions. FAQs such as opening hours, menu details, and directions are ideal for chatbot automation, as outlined in the AI-900 modules discussing customer support automation.

* A restaurant can use a chatbot to automate responses to customer reviews on an external website - No.

This is not a typical chatbot use case taught in AI-900. Chatbots are meant for direct interactions within controlled channels, such as a company's own website or messaging app. Managing and posting responses to reviews on external platforms (like Yelp or Google Reviews) would involve policy restrictions, authentication issues, and reputational risk. The AI-900 course specifies that responsible AI usage requires maintaining human oversight in public-facing communications that influence brand image.

NEW QUESTION: 10

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: (SHOW ANSWER)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Describe features of computer vision workloads on Azure," the Custom Vision service is a specialized part of Azure Cognitive Services that allows developers to train image classification and object detection models tailored to their own data. It is particularly useful when prebuilt models, such as those in the standard Computer Vision service, cannot accurately recognize domain-specific objects - such as specific product brands or packaging.


In this scenario, the bot must identify brand names of products in images of supermarket shelves. Since brand logos and packaging designs are unique to each company, a general-purpose image analysis model would not perform accurately. The Custom Vision Image Classification capability allows you to upload labeled images (e.g., various brands) and train a model to distinguish between them. Once trained, the model can classify new images and recognize which brand appears on the shelf.

Let's analyze the other options:

- * A. AI enrichment for Azure Search capabilities: Used in knowledge mining to extract information from documents, not image brand identification.
 - * B. Computer Vision Image Analysis capabilities: Provides prebuilt functionality such as detecting objects, describing images, and identifying common items (like "bottle" or "box") but cannot differentiate custom brand names.
 - * D. Language understanding capabilities: Deals with processing and understanding natural language text, not images.
- Therefore, identifying specific brand names from images requires a custom-trained image classification model, making Custom Vision Image Classification capabilities the correct answer.
- # Final Verified answer:
- C). Custom Vision Image Classification capabilities

NEW QUESTION: 11

You are developing a model to predict events by using classification.
You have a confusion matrix for the model scored on test data as shown in the following exhibit.



Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Answer Area



There are [answer choice] correctly predicted positives.

▼

5

11

1,033

13,951

There are [answer choice] false negatives.

▼

5

11

1,033

13,951

Answer:
Answer Area

There are [answer choice] correctly predicted positives.

5

11

1,033

13,951

There are [answer choice] false negatives.

5

11

1,033

13,951

Explanation:

Answer Area

There are [answer choice] correctly predicted positives.

5

11

1,033

13,951

There are [answer choice] false negatives.

5

11

1,033

13,951

Box 1: 11

	Predicted	
	Positive	Negative
Actual True	TP	FN
Actual False	FP	TN

TP = True Positive.

The class labels in the training set can take on only two possible values, which we usually refer to as positive or negative. The positive and negative instances that a classifier predicts correctly are called true positives (TP) and true negatives (TN), respectively. Similarly, the incorrectly classified instances are called false positives (FP) and false negatives (FN).

Box 2: 1,033

FN = False Negative

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio/evaluate-model-performance> Finding TP is easy. It basically means the value where Predicted and True value is 1 and that is 11 in this case. False Negative means where true value was 1 but predicted value was 0 and that is 1033 in this case The confusion matrix shows cases where both the predicted and actual values were 1 (known as true positives) at the top left, and cases where both the predicted and the actual values were 0 (true negatives) at the bottom right. The other cells show cases where the predicted and actual values differ (false positives and false negatives).

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

NEW QUESTION: 12

An app that analyzes social media posts to identify their tone is an example of which type of natural language processing (NLP) workload?

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

Answer: (SHOW ANSWER)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Describe features of natural language processing (NLP) workloads on Azure," sentiment analysis is an NLP workload that determines the emotional tone or opinion expressed in a piece of text. This could be positive, negative, or neutral sentiment.

When an app analyzes social media posts to identify their tone, it is performing sentiment analysis, since it aims to understand the emotional context behind user-generated text such as tweets, reviews, or comments.

Azure provides this functionality through the Azure Cognitive Services - Text Analytics API, which evaluates text and returns sentiment scores.

Other options are not suitable:


- * Key phrase extraction identifies main ideas in text but not tone.
- * Entity recognition identifies names of people, organizations, or locations.
- * Speech recognition converts spoken words into text, not emotional analysis.

Therefore, analyzing social media tone is an example of sentiment analysis, a key NLP workload in Microsoft's AI-900 syllabus.

NEW QUESTION: 13

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

You receive the output shown in the following exhibit.



Which type of natural languages processing was performed?

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

Answer: (SHOW ANSWER)

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/overview> You can provide the Text Analytics service with unstructured text and it will return a list of entities in the text that it recognizes. You can provide the Text Analytics service with unstructured text and it will return a list of entities in the text that it recognizes. The service can also provide links to more information about that entity on the web. An entity is essentially an item of a particular type or a category; and in some cases, subtype, such as those as shown in the following table.

<https://docs.microsoft.com/en-us/learn/modules/analyze-text-with-text-analytics-service/2-get-started-azure>

NEW QUESTION: 14

Which metric can you use to evaluate a classification model?

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

Answer: (SHOW ANSWER)

For evaluating a classification model, the appropriate metric from the options provided is the True Positive Rate (TPR), also known as Sensitivity or Recall. According to the Microsoft Azure AI Fundamentals (AI-

900) official study guide and Microsoft Learn module "Evaluate model performance", classification models are evaluated using metrics that measure how accurately the model predicts categorical outcomes such as "yes

/no," "spam/not spam," or "approved/denied."

The True Positive Rate measures the proportion of correctly identified positive cases out of all actual positive cases. Mathematically, it is expressed as:

True Positive Rate (Recall)= $\frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$

= $\frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$ True Positive Rate (Recall)

= $\frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$

This metric is important when missing positive predictions carries a high cost, such as in medical diagnosis or fraud detection. Microsoft Learn highlights classification evaluation metrics such as accuracy, precision, recall, F1 score, and AUC (Area Under the Curve) as suitable for classification models.

The other options-Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Coefficient of Determination (R^2)-are regression metrics used to evaluate models that predict numeric values rather than categories. For example, they apply to predicting house prices or temperatures, not yes/no decisions.

Therefore, the correct classification evaluation metric among the choices is A. True Positive Rate.

Reference:Microsoft Learn - Evaluate model performance - Understand metrics for classification and regression models

NEW QUESTION: 15

Which type of Azure AI workload should you use to create illustrations based on the text of an article?

- A.** generative AI
- B.** natural language processing
- C.** Azure Document Intelligence in Foundry Tools
- D.** computer vision

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 16

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

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

NOTE: Each correct selection is worth one point.

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

Answer: (SHOW ANSWER)

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and the Microsoft Learn module "Describe features of Natural Language Processing (NLP) workloads on Azure", this scenario combines two major capabilities of AI: speech recognition and natural language understanding.

* Speech-to-Text (E) - This is the first step in processing verbal commands. The Azure Speech service converts the spoken words of a user into textual data that can be understood and processed by downstream components. This workload is commonly referred to as speech recognition, and it falls under the speech capabilities of Azure Cognitive Services. Without this transcription process, the system could not interpret the user's voice input.

* Language Modeling (C) - After the speech input is converted into text, the next step is to interpret the meaning of the text so the system can take appropriate action. Language modeling, also known as language understanding, is responsible for identifying the user's intent (for example, "turn on the lights" or "set the thermostat to 72 degrees") and extracting entities (such as device name or temperature value). In Azure, this function is handled by Language Understanding (LUIS) or Conversational Language Understanding (CLU). These models allow smart systems to process commands and map them to defined actions.

Other options are not correct:

- * A. Text-to-speech converts text output into spoken language, which is not mentioned as a requirement.
- * B. Translation converts text from one language to another, irrelevant to this scenario.
- * D. Key phrase extraction identifies important terms in text but doesn't interpret or execute commands.

Therefore, the solution uses speech-to-text to transcribe verbal commands and language modeling to understand and act upon them - the two key NLP workloads enabling voice-controlled smart devices.

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (**336** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 17

You are creating an app to help employees write emails and reports based on user prompts. What should you use?

- A. Azure AI Speech
- B. Azure OpenAI in Foundry Models
- C. Azure AI Vision
- D. Azure Machine Learning studio

Answer: B (LEAVE A REPLY)

For an app that helps employees write emails and reports based on user prompts, you need a text generation model capable of understanding natural language instructions and producing coherent, contextually appropriate output. Azure OpenAI GPT models-available through Azure AI Foundry (formerly Azure OpenAI Studio)-are specifically designed for such generative tasks.

By integrating GPT-3.5 or GPT-4, the app can analyze prompts like "Write a professional email to a client about project updates" and automatically generate polished text in seconds.

The other options do not fit:

- * A. Azure AI Speech: Converts spoken language to text or text to speech; not suitable for generating written content.
- * C. Azure AI Vision: Processes and analyzes images or video content.
- * D. Azure Machine Learning Studio: Used for training, testing, and deploying custom ML models, not directly for content generation.

Therefore, to create a writing-assistance app for emails and reports, the correct solution is B. Azure OpenAI in Foundry Models using GPT-based language generation.

NEW QUESTION: 18

You need to create a model that labels a collection of your personal digital photographs.

Which Azure AI service should you use?

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

Answer: D (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn module "Describe features of Computer Vision workloads on Azure", the Azure AI Custom Vision service allows users to build, train, and deploy custom image classification or object detection models. It is specifically designed for scenarios where you need a model tailored to your unique image dataset - in this case, personal digital photographs.

Custom Vision lets you upload and label your own images (for example, "family," "friends," "vacations," or

"pets") and then train a model that learns to recognize those categories. The system automatically extracts relevant features from the training images and creates a model capable of classifying new images into the predefined labels. You can iteratively refine your model by adding more images or re-training to improve accuracy.

The other options do not fit this requirement:

- * A. Azure AI Language deals with text-based tasks such as language understanding, sentiment analysis, and key phrase extraction - not image processing.
- * B. Azure AI Computer Vision provides prebuilt image analysis models (e.g., object detection, tag generation, scene description), but it cannot learn custom categories unique to your dataset. It's great for general image recognition but not for specialized labeling tasks.
- * C. Azure AI Document Intelligence (Form Recognizer) is used to extract information from structured or semi-structured documents such as forms, invoices, or receipts - not photographs.

Therefore, when you need to label or categorize personal photos with custom-defined tags, the right service is Azure AI Custom Vision, because it allows you to build a model trained specifically on your own collection of images.

NEW QUESTION: 19

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

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

Answer: C (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Responsible AI Framework, Inclusiveness is one of the six guiding principles for responsible AI. The principle of inclusiveness ensures that AI systems are designed to empower everyone and engage people of all abilities. Microsoft emphasizes that inclusive AI systems must be developed with awareness of potential barriers that could unintentionally exclude certain user groups. This directly aligns with the scenario described-where the company is examining voice recognition technologies in smart home devices to identify barriers that might leave out users, such as those with speech impairments, accents, or language differences.

The official Microsoft Learn module "Identify guiding principles for responsible AI" explains that inclusiveness focuses on creating systems that can understand and serve users with diverse needs. For example, voice recognition models should account for variations in dialect, tone, accent, and speech patterns to ensure equitable access for all. A lack of inclusiveness could cause bias or misrecognition for underrepresented groups, leading to unintentional exclusion.

Microsoft's guidance further stresses that designing for inclusiveness involves involving diverse users in the data collection and testing phases, conducting accessibility assessments, and continuously improving model performance across different demographic groups. In this way, inclusiveness promotes fairness, accessibility, and usability across cultural and physical differences.

In contrast:

- * A. Accountability is about ensuring humans are responsible for AI outcomes.
- * B. Fairness focuses on preventing bias and discrimination in data or algorithms.
- * D. Privacy and security ensure protection of personal data and secure handling of information.

Thus, evaluating potential barriers that could exclude specific user groups exemplifies Inclusiveness, as it demonstrates a proactive approach to making AI accessible and beneficial for all users.

NEW QUESTION: 20

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:

Answer Area



Explanation:

Answer Area



According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure," there are multiple types of computer vision tasks, each designed for different goals such as recognizing, categorizing, or locating objects within an image.

In this question, the application performs two distinct tasks: locating birds within an image and identifying their species. Each of these corresponds to a different type of computer vision capability.

- * Locate the birds # Object detection
- * Object detection is used when an AI system needs to identify and locate multiple objects within a single image.
- * It not only recognizes what the object is but also provides bounding boxes that indicate the exact position of each object.
- * In this scenario, locating the birds (drawing rectangles around each bird) is achieved through object detection models, such as those available in the Azure Custom Vision Object Detection domain.
- * Identify the species of the birds # Image classification
- * Image classification focuses on identifying what is in the image rather than where it is.
- * It assigns a single label (or multiple labels in multilabel classification) to an entire image based on its contents.
- * In this case, determining the species of a bird (e.g., robin, eagle, parrot) is achieved through image classification, where the model compares visual features against learned patterns from training data.

Incorrect options:

- * Automated captioning generates descriptive sentences about an image, not object locations or classifications.
- * Optical character recognition (OCR) extracts text from images, irrelevant in this case.

NEW QUESTION: 21

Select the answer that correctly completes the sentence.



Answer:



Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn module "Describe features of common AI workloads," an anomaly detection workload is designed to identify data points or patterns that deviate significantly from what is expected or normal. These anomalies often indicate irregularities, faults, or potential issues that require attention.

In this scenario, the AI system monitors temperature data from a large machine. Normally, the machine operates within a predictable temperature range. When the AI detects sudden or unexpected temperature spikes or drops - behavior that does not match the historical pattern - it flags these occurrences as anomalies. This type of workload is fundamental in predictive maintenance and industrial monitoring, where it helps detect equipment failures, safety hazards, or energy inefficiencies before they escalate.

Microsoft's AI-900 curriculum emphasizes that anomaly detection workloads are often used in:

- * Industrial IoT systems (detecting abnormal sensor readings or machine behavior)
- * Finance (fraud detection or unusual transaction monitoring)
- * Cybersecurity (detecting irregular network traffic or access patterns)
- * Operations (identifying abnormal variations in production data)

The Azure service used for this purpose is Azure Anomaly Detector, part of Azure Cognitive Services, which uses advanced statistical and machine learning models to automatically detect outliers in time-series data such as temperature, pressure, or transaction logs.

By comparison:

- * Computer vision handles image or video analysis.
- * Knowledge mining extracts insights from large document collections.
- * Natural Language Processing (NLP) interprets human language.

Thus, based on the official Microsoft AI-900 study guide and Microsoft Learn, the correct and verified answer is An anomaly detection workload, since detecting unusual temperature fluctuations precisely fits this AI workload type.

NEW QUESTION: 22

What is an example of a regression model in machine learning?

- A.** dividing the student data in a dataset based on the age of the students and their educational achievements
- B.** identifying subtypes of spam email by examining a large collection of emails that were flagged by users
- C.** predicting the sale price of a house based on historical data, the size of the house, and the number of bedrooms in the house
- D.** identifying population counts of endangered animals by analyzing images

Answer: C (LEAVE A REPLY)

Correct answer is C. Predicting the sale price of a house based on historical data, the size of the house, and the number of bedrooms.

In machine learning, regression is a supervised learning technique used to predict continuous numeric values.

Microsoft's AI-900 study guide defines regression models as those that estimate relationships between variables-predicting a continuous outcome variable from one or more input features.

In this case, the house sale price is a continuous numeric value, and inputs such as size, location, and number of bedrooms are the features. Common regression algorithms include linear regression, decision tree regression, and boosted regression.

Other options represent different ML workloads:

- * A involves segmentation by categories (classification or clustering).
- * B represents clustering, grouping similar items without predefined labels.
- * D represents computer vision, counting animals in images rather than predicting a numeric value.

Hence, the verified answer is C. Regression.

NEW QUESTION: 23

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

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

Answer: B ([LEAVE A REPLY](#))

In Azure Machine Learning, regression is a supervised machine learning technique used to predict continuous numerical values based on input data. According to the Microsoft AI Fundamentals (AI-900) study guide and the Microsoft Learn module "Identify common types of machine learning," regression models are ideal when the goal is to estimate a quantity - such as price, temperature, or, in this case, population size.

In the scenario, the task is to predict the population size of a specific species within a defined area. Population size is a numerical, continuous value that varies depending on multiple factors (like time, environment, and resources). A regression algorithm, such as linear regression or decision tree regression, can be trained on historical data (e.g., species count, area, temperature, food availability) to forecast future population numbers.

Option analysis:

- * A. Clustering: Used for unsupervised learning, where the goal is to group similar data points into clusters without predefined labels (e.g., grouping animals by behavior or habitat).
- * C. Classification: Used to predict discrete categories or labels (e.g., "endangered" vs. "not endangered"), not numerical values.

Therefore, the correct machine learning type for predicting a continuous value such as population size is Regression.

NEW QUESTION: 24

For each of The following statements, select Yes If the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point

Answer Area

Statements	Yes	No
A webchat bot can interact with users visiting a website.	☐	☐
Automatically generating captions for pre-recorded videos is an example of natural language processing.	☐	☐
A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of natural language processing.	☐	☐

Answer:

Answer Area

Statements	Yes	No
A webchat bot can interact with users visiting a website.	☒	☐
Automatically generating captions for pre-recorded videos is an example of natural language processing.	☒	☐
A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of natural language processing.	☒	☐

Explanation:

Statements

Yes

No

A webchat bot can interact with users visiting a website.

Yes

Automatically generating captions for pre-recorded videos is an example of natural language processing.

No

A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of natural language processing.

Yes

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn modules on AI workloads, each of these statements maps to a distinct area of artificial intelligence - namely Conversational AI, Speech AI, and Natural Language Processing (NLP).

* "A webchat bot can interact with users visiting a website." - YesThis is true. A webchat bot represents an example of Conversational AI. It leverages natural language understanding (NLU) to interpret user input and generate appropriate responses. These bots can be created using Azure services such as Azure AI Bot Service and Language Understanding (LUIS). They enable automated interactions with users through text-based communication on websites, applications, or messaging platforms.

* "Automatically generating captions for pre-recorded videos is an example of natural language processing." - NoThis is false. Generating captions from audio involves speech recognition, not NLP.

Specifically, it uses speech-to-text technology to transcribe spoken words into written text. This function is typically performed by Azure's Speech service, which is part of the Speech AI workload, not the language-processing workload.

* "A smart device in the home that responds to questions such as 'What will the weather be like today?' is an example of natural language processing." - YesThis is true. Smart assistants like Alexa or Cortana use NLP to interpret spoken queries, extract meaning, and generate appropriate responses. NLP allows these devices to understand human language, retrieve relevant information, and respond conversationally.

NEW QUESTION: 25

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

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

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

NOTE: Each correct selection is worth one point.

API Features

Entity recognition

Key phrase extraction

Language detection

Sentiment analysis

Answer Area

API Feature

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

API Feature

Summarize important information from the support ticket.

API Feature

Extract key dates from the support ticket.

Answer:

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

Explanation:

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

Box1: Sentiment analysis

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

Box 2: Broad entity extraction

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

Box 3: Entity Recognition

Named Entity Recognition: Identify and categorize entities in your text as people, places, organizations, date

/time, quantities, percentages, currencies, and more. Well-known entities are also recognized and linked to more information on the web.

Reference:

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

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

NEW QUESTION: 26

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
	The Azure AI Language service can identify in which language text is written.	☐	☐
	The Azure AI Language service can detect handwritten signatures in a document.	☐	☐
	The Azure AI Language service can identify companies and organizations mentioned in a document.	☐	☐

Answer:

Answer Area

Statements	Yes	No
The Azure AI Language service can identify in which language text is written.	☒	☐
The Azure AI Language service can detect handwritten signatures in a document.	☐	☒
The Azure AI Language service can identify companies and organizations mentioned in a document.	☒	☐

Explanation:



Statements	Yes	No
The Azure AI Language service can identify in which language text is written.	☒	☐
The Azure AI Language service can detect handwritten signatures in a document.	☐	☒
The Azure AI Language service can identify companies and organizations mentioned in a document.	☒	☐

The Azure AI Language service (part of Azure Cognitive Services) provides a set of natural language processing (NLP) capabilities designed to analyze and interpret text data. Its core features include language detection, key phrase extraction, sentiment analysis, and named entity recognition (NER).

* Language Identification - YES According to the Microsoft Learn module "Analyze text with Azure AI Language," one of the service's built-in capabilities is language detection, which determines the language of a given text string (e.g., English, Spanish, or French). This allows applications to automatically adapt to multilingual input.

* Handwritten Signature Detection - NO The Azure AI Language service only processes text-based data; it does not analyze images or handwriting. Detecting handwritten signatures requires computer vision capabilities, specifically Azure AI Vision or Azure AI Document Intelligence, which can extract and interpret visual content from scanned documents or images.

* Identifying Companies and Organizations - YES The Named Entity Recognition (NER) feature within Azure AI Language can identify entities such as people, locations, dates, organizations, and companies mentioned in text. It tags these entities with categories, enabling structured analysis of unstructured data.

Summary:

* Language detection # Yes (supported by AI Language).

* Handwritten signatures # No (requires Computer Vision).

* Entity recognition for companies/organizations # Yes (supported by AI Language NER).

NEW QUESTION: 27

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

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

Answer: D (LEAVE A REPLY)

The Azure AI Language service provides several NLP features, including language detection, key phrase extraction, sentiment analysis, and named entity recognition (NER).

When you need to extract specific data points such as dates, names, organizations, or locations from unstructured text (for example, social media posts), the correct feature is Entity Recognition.

Entity Recognition identifies and classifies information in text into predefined categories like:

- * Person names (e.g., "John Smith")
- * Organizations (e.g., "Contoso Ltd.")
- * Dates and times (e.g., "October 22, 2025")
- * Locations, events, and quantities

This capability helps transform unstructured textual data into structured data that can be analyzed or stored.

Option analysis:

- * A (Language detection): Determines the language of a text (e.g., English, French).
- * B (Speech recognition): Converts spoken audio to text; not applicable here.
- * C (Key phrase extraction): Identifies important phrases or topics but not specific entities like names or dates.
- * D (Entity recognition): Correctly extracts names, dates, and other specific data from text.

Hence, the accurate feature for this scenario is D. Entity Recognition.

NEW QUESTION: 28

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

NOTE: Each correct answer is worth one point.

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

Answer: B,C,D (LEAVE A REPLY)

A transformer model is the foundational architecture behind many modern natural language processing systems such as GPT and BERT. It processes text data through multiple key stages. According to the Microsoft Azure AI Fundamentals (AI-900) curriculum and Microsoft Learn materials, the major stages of a transformer-based large language model are tokenization, embedding calculation, and next token prediction.

- * Tokenization (C) - The first step converts raw text into smaller units called tokens (words, subwords, or characters). This process allows the model to handle text in a structured numerical form rather than as raw language.
- * Embedding Calculation (B) - After tokenization, the tokens are mapped into high-dimensional numeric vectors, known as embeddings. These embeddings capture semantic relationships between words and phrases so that the model can understand context and meaning.
- * Next Token Prediction (D) - This stage is the heart of transformer operation, where the model predicts the next likely token in a sequence based on prior tokens. Repeated next-token predictions enable text generation, summarization, or translation.

Options A (object detection) and E (anonymization) are incorrect because they relate to vision and privacy workflows, not language modeling.

NEW QUESTION: 29

Select the answer that correctly completes the sentence.

Answer Area

For a vehicle, predicting the miles per gallon based on weight, engine power, and other factors is an example of

regression.

anomaly detection.

classification.

clustering.

regression.

Answer:

Answer Area

For a vehicle, predicting the miles per gallon based on weight, engine power, and other factors is an example of

regression.

anomaly detection.

classification.

clustering.

regression.

Explanation:

Answer Area

For a vehicle, predicting the miles per gallon based on weight, engine power, and other factors is an example of regression.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of common machine learning types", regression is a supervised machine learning technique used to predict continuous numerical values based on one or more input features. In this scenario, the task is to predict a vehicle's miles per gallon (MPG)-a continuous numeric value-based on several measurable factors such as weight, engine power, and other specifications.

Regression models learn the mathematical relationship between input variables (independent features) and a numeric target variable (dependent outcome). Common regression algorithms include linear regression, decision tree regression, and support vector regression. In the example, the model would analyze historical data of vehicles and learn patterns that map characteristics (like engine size, horsepower, and weight) to fuel efficiency. Once trained, it can predict the MPG for a new vehicle configuration.

The other options describe different problem types:

- * Classification predicts discrete categories (for example, whether a car is "fuel efficient" or "not fuel efficient"), not continuous values.
- * Clustering is an unsupervised learning method that groups data points based on similarities without predefined labels, not predictive modeling.
- * Anomaly detection identifies data points that significantly deviate from normal patterns, such as detecting engine sensor failures or fraudulent transactions.

Since predicting MPG involves estimating a numeric value within a continuous range, regression is the most appropriate model type.

In summary, per AI-900 training content, regression models are used when the output variable is numeric, classification for categorical outputs, and clustering for pattern discovery. Therefore, predicting miles per gallon based on vehicle features is a textbook example of a regression problem in Azure Machine Learning.

NEW QUESTION: 30

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation:

You can communicate with a bot by using email # No

You can communicate with a bot by using Microsoft Teams # Yes

You can communicate with a bot by using a webchat interface # Yes

These answers are based on the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore conversational AI in Microsoft Azure." The Azure Bot Service allows developers to build, test, deploy, and manage intelligent chatbots that can interact with users through various channels. Channels are communication platforms or interfaces that connect users to bots. Once a bot is built and published through the Azure Bot Service, it can be connected to multiple channels such as Microsoft Teams, webchat, Skype, Facebook Messenger, Direct Line, Slack, and others.

Let's evaluate each statement:

* You can communicate with a bot by using email # No Azure Bot Service does not support direct interaction via email as a channel. Bots are designed for real-time or conversational interactions through messaging or voice-based platforms, not asynchronous email communication.

* You can communicate with a bot by using Microsoft Teams # Yes Microsoft Teams is one of the primary channels supported by Azure Bot Service. Bots can be integrated directly into Teams to handle chat-based conversations, provide information, automate workflows, or assist users interactively within Teams.

* You can communicate with a bot by using a webchat interface # Yes The Web Chat channel is another core feature of Azure Bot Service. It allows embedding the bot into a website or web application using the Web Chat control or the Direct Line API, enabling users to chat directly from a browser interface.

In summary, Azure Bot Service supports real-time conversational interfaces like Teams and webchat, but not email.

NEW QUESTION: 31

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

Answer Area

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

image classification.

object detection.

optical character recognition (OCR).

facial detection.

Answer:

Answer Area

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

image classification.

object detection.

optical character recognition (OCR).

facial detection.

Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) official study materials, object detection is a type of computer vision workload that not only identifies objects within an image but also determines their location by drawing bounding boxes around them. This functionality is clearly described in the Microsoft Learn module "Identify features of computer vision workloads." In this scenario, the AI system analyzes an image to find a vehicle and then returns a bounding box showing where that vehicle is located within the image frame. That ability - to detect, classify, and localize multiple objects - perfectly defines object detection.

Microsoft's study content contrasts object detection with other computer vision workloads as follows:

- * Image classification: Determines what object or scene is present in an image as a whole but does not locate it (e.g., "this is a car").
- * Object detection: Identifies what objects are present and where they are, usually returning coordinates for bounding boxes (e.g., "car detected at position X, Y").
- * Optical Character Recognition (OCR): Extracts text content from images or scanned documents.
- * Facial detection: Specifically locates human faces within an image or video feed, often as part of face recognition systems.


In Azure, object detection capabilities are available through services such as Azure Computer Vision, Custom Vision, and Azure Cognitive Services for Vision, which can be trained to detect vehicles, products, or other objects in various image datasets.

Therefore, based on the AI-900 study guide and Microsoft Learn materials, the verified and correct answer is Object detection, as it accurately describes the process of returning a bounding box indicating an object's position in an image.

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (336 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 32

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



Answer:

Answer Area



The handling of unusual or missing values provided to an AI system is a consideration for the Microsoft principle for responsible AI.

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

Explanation:

Reliability & Safety

[https://en.wikipedia.org/wiki/Tay_\(bot\)](https://en.wikipedia.org/wiki/Tay_(bot))

"To build trust, it ' s critical that AI systems operate reliably, safely, and consistently under normal circumstances and in unexpected conditions. These systems should be able to operate as they were originally designed, respond safely to unanticipated conditions, and resist harmful manipulation. It ' s also important to be able to verify that these systems are behaving as intended under actual operating conditions. How they behave and the variety of conditions they can handle reliably and safely largely reflects the range of situations and circumstances that developers anticipate during design and testing. We believe that rigorous testing is essential during system development and deployment to ensure AI systems can respond safely in unanticipated situations and edge cases, don ' t have unexpected performance failures, and don ' t evolve in ways that are inconsistent with original expectations"

NEW QUESTION: 33

Which feature of the Azure AI Language service should you use to automate the masking of names and phone numbers in text data?

- A. Personally Identifiable Information (PII) detection
- B. entity linking
- C. custom text classification
- D. custom named entity recognition (NER)

Answer: A (LEAVE A REPLY)

The correct answer is A. Personally Identifiable Information (PII) detection.

In the Azure AI Language service, PII detection is a built-in feature designed to automatically identify and redact sensitive or confidential information from text data. According to the Microsoft Learn module "Identify capabilities of Azure AI Language" and the AI-900 study guide, this capability can detect personal data such as names, phone numbers, email addresses, credit card numbers, and other identifiers. When applied, the service scans input text and either masks or removes these PII elements based on configurable parameters, ensuring compliance with data privacy regulations like GDPR or HIPAA.

For example, if a document contains "John Doe's phone number is 555-123-4567," PII detection can return

"*****'s phone number is *****," thereby preventing exposure of sensitive personal details.

Option analysis:

- * A. Personally Identifiable Information (PII) detection: # Correct. It identifies and masks sensitive data in text.
- * B. Entity linking: Connects recognized entities to known data sources like Wikipedia; not used for redaction.
- * C. Custom text classification: Classifies text into predefined categories; not designed for masking personal data.
- * D. Custom named entity recognition (NER): Detects domain-specific entities you define but doesn't automatically mask them.

Therefore, to automate masking of names and phone numbers, the appropriate Azure AI Language feature is PII detection.

NEW QUESTION: 34

Extracting relationships between data from large volumes of unstructured data is an example of which type of AI workload?

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

Answer: ([SHOW ANSWER](#))

Extracting relationships and insights from large volumes of unstructured data (such as documents, text files, or images) aligns with the Knowledge Mining workload in Microsoft Azure AI. According to the Microsoft AI Fundamentals (AI-900) study guide and Microsoft Learn module "Describe features of common AI workloads," knowledge mining involves using AI to search, extract, and structure information from vast amounts of unstructured or semi-structured content.

In a typical knowledge mining solution, tools like Azure AI Search and Azure AI Document Intelligence work together to index data, apply cognitive skills (such as OCR, key phrase extraction, and entity recognition), and then enable users to discover relationships and patterns through intelligent search. The process transforms raw content into searchable knowledge.

The key characteristics of knowledge mining include:

- * Using AI to extract entities and relationships between data points.
- * Applying cognitive skills to text, images, and documents.
- * Creating searchable knowledge stores from unstructured data.

Hence, B. Knowledge Mining is correct.

The other options-computer vision, NLP, and anomaly detection-deal with image recognition, language understanding, and data irregularities, respectively, not large-scale information extraction.

NEW QUESTION: 35

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

NOTE: Each correct selection is worth one point.

- A. coefficient of determination (R2)
- B. F1 score
- C. root mean squared error (RMSE)
- D. area under curve (AUC)
- E. balanced accuracy

Answer: ([SHOW ANSWER](#))

A: R-squared (R2), or Coefficient of determination represents the predictive power of the model as a value between -inf and 1.00. 1.00 means there is a perfect fit, and the fit can be arbitrarily poor so the scores can be negative.


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

Reference:

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

NEW QUESTION: 36

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



Which type of model is the chart used to evaluate?

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

Answer: B (LEAVE A REPLY)

What is a Predicted vs. True chart?

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

Reference:

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

NEW QUESTION: 37

Select the answer that correctly completes the sentence.

Answer Area

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

computer vision.	▼
forecasting.	
computer vision.	
knowledge mining.	
anomaly detection.	

Answer:

Answer Area

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



computer vision.
forecasting.
computer vision.
knowledge mining.
anomaly detection.

Explanation:

Answer Area

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

computer vision.

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure," computer vision is a field of artificial intelligence that enables computers to interpret and understand visual information from the world - such as images or videos.

In this scenario, the task is to count the number of animals in an area based on a video feed. This requires the system to:

- * Detect the presence of animals in each frame of the video (object detection).
- * Track and count them across multiple frames as they move.

These are classic computer vision tasks, as they involve analyzing visual inputs (video or image data) and identifying objects (in this case, animals). Azure provides services such as Azure Computer Vision, Custom Vision, and Video Indexer, which can perform object detection, counting, and activity recognition using AI models trained on visual datasets.

Why the other options are incorrect:

- * Forecasting: Involves predicting future values based on historical data (e.g., predicting sales or weather), not analyzing video feeds.
- * Knowledge mining: Focuses on extracting insights from large text-based document repositories, not images or videos.
- * Anomaly detection: Identifies unusual patterns in numeric or time-series data, not visual objects.

Therefore, identifying and counting animals in video footage falls under computer vision, since it uses AI to visually detect, classify, and quantify objects in real-time or recorded feeds.

NEW QUESTION: 38

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

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

Which service should you use?

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

Answer: ([SHOW ANSWER](#))

A FAQ PDF file contains structured Q&A content. The QnA Maker (now part of Azure Language Service) can automatically extract questions and answers from such a document and build a knowledge base for conversational bots. This allows users to interact naturally with the content via chat interfaces.

Other options:

- * B. Text Analytics # Extracts insights, not conversational content.
- * C. Computer Vision # Used for image analysis.
- * D. LUIS # Handles intent detection, not static question-answer responses.

NEW QUESTION: 39

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

NOTE Each correct selection is worth one point

Answer Area

Statements	Yes	No
You can fine-tune some Azure OpenAI models by using your own data.	☐	☐
Pretrained generative AI models are a component of Azure OpenAI.	☐	☐
To build a solution that complies with Microsoft responsible AI principles, you must build and train your own model.	☐	☐

Answer:

Answer Area

Statements	Yes	No
You can fine-tune some Azure OpenAI models by using your own data.	☒	☐
Pretrained generative AI models are a component of Azure OpenAI.	☒	☐
To build a solution that complies with Microsoft responsible AI principles, you must build and train your own model.	☐	☒

Explanation:

Answer Area

Statements	Yes	No
You can fine-tune some Azure OpenAI models by using your own data.	☒	☐
Pretrained generative AI models are a component of Azure OpenAI.	☒	☐
To build a solution that complies with Microsoft responsible AI principles, you must build and train your own model.	☐	☒

Full Detailed Explanation (250-300 words):

- * "You can fine-tune some Azure OpenAI models by using your own data." - YESThis statement is true. Azure OpenAI allows customers to fine-tune certain models like GPT-3, GPT-3.5, and some embedding models with their own data. Fine-tuning customizes a model to perform better on specific tasks or match a company's domain terminology, tone, or context. According to Microsoft Learn's AI-900 and Azure OpenAI documentation, fine-tuning is supported for approved use cases while maintaining Microsoft's Responsible AI oversight and compliance process.
- * "Pretrained generative AI models are a component of Azure OpenAI." - YESThis statement is also true. Azure OpenAI provides access to pretrained large language and generative AI models such as GPT-3.5, GPT-4, Codex, and DALL E. These models are pretrained on vast datasets and made available via APIs, allowing developers to generate text, code, and images without needing to train their own models. This is a core feature of Azure OpenAI's service offering.
- * "To build a solution that complies with Microsoft responsible AI principles, you must build and train your own model." - NOThis statement is false. Compliance with Microsoft Responsible AI principles (Fairness, Reliability & Safety, Privacy & Security, Inclusiveness, Transparency, Accountability) does not require building custom models. Prebuilt Azure AI and OpenAI services already align with Responsible AI standards. Developers simply need to use these services responsibly, applying governance and ethical design practices.

NEW QUESTION: 40

What should you implement to identify hateful responses returned by a generative AI solution?

- A. prompt engineering
- B. abuse monitoring
- C. content filtering
- D. fine-tuning

Answer: C (LEAVE A REPLY)

When using generative AI models such as those in Azure OpenAI, it's critical to monitor and control the model's outputs to prevent the generation of harmful, hateful, or unsafe content. According to Microsoft's Responsible AI guidelines and Azure OpenAI safety systems, this is achieved through content filtering mechanisms.

Content filtering automatically analyzes generated text and blocks or flags responses containing hate speech, harassment, sexual content, or violent language. Azure OpenAI applies these filters by default through its built-in content moderation system, which categorizes responses into safety levels (safe, review, block).

Option review:


- * A. Prompt engineering: Helps guide responses but does not detect hate content automatically.
- * B. Abuse monitoring: Refers to manual or post-processing review, not real-time automated filtering.
- * C. Content filtering: # Correct - actively detects and prevents hateful or unsafe responses.
- * D. Fine-tuning: Adjusts model behavior using additional training data but doesn't ensure safety monitoring.

NEW QUESTION: 41


Select the answer that correctly completes the sentence.



Answer:



Explanation:



This question refers to a system that monitors a user's emotions or expressions-in this case, identifying whether a kiosk user is annoyed-through a video feed. According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify Azure services for computer vision," this scenario falls under facial analysis, which is a capability of Azure AI Vision or the Face API.

Facial analysis involves detecting human faces in images or video and analyzing facial features to interpret emotions, expressions, age, gender, or facial landmarks. The AI model does not try to identify who the person is but rather interprets how they appear or feel. For example, facial analysis can detect emotions such as happiness, anger, sadness, or surprise, which allows applications to infer a user's engagement or frustration level while interacting with a system.

Option review:

- * Face detection: Identifies the presence and location of a face in an image but does not interpret expressions or emotions.
- * Facial recognition: Matches a detected face to a known individual's identity (for authentication or security), not for emotion detection.
- * Optical character recognition (OCR): Extracts text from images or scanned documents and has no relation to human emotion or facial features.

Therefore, determining whether a kiosk user is annoyed, happy, or frustrated involves emotion detection within facial analysis, making Facial analysis the correct answer.

This aligns with AI-900's definition of computer vision workloads, where facial analysis provides insights into emotions and expressions, supporting user experience optimization and customer behavior analytics.

NEW QUESTION: 42

Which two components can you drag onto a canvas in Azure Machine Learning designer? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Answer: A,D (LEAVE A REPLY)

In Azure Machine Learning designer, a low-code drag-and-drop interface, users can visually build machine learning workflows. According to the AI-900 study guide and Microsoft Learn module "Create and publish models with Azure Machine Learning designer", two key components that can be dragged onto the designer canvas are datasets and modules.

* Datasets (A): These are collections of data that serve as the input for training or evaluating models.

They can be registered in the workspace and then dragged onto the canvas for use in transformations or model training.

* Modules (D): These are prebuilt processing and modeling components that perform operations such as data cleaning, feature engineering, model training, and evaluation. Examples include "Split Data," "Train Model," and "Evaluate Model."

Compute (B) and Pipeline (C) are not drag-and-drop items within the designer. Compute targets are infrastructure resources used to run the pipeline, while a pipeline represents the overall workflow, not a component that can be added like a dataset or module.

Hence, the correct answers are A. Dataset and D. Module.

Reference: Microsoft Learn - Create a machine learning model with Azure Machine Learning designer

NEW QUESTION: 43

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

Answer:

Statements	Yes	No
Forecasting housing prices based on historical data is an example of anomaly detection.	☐	☒
Identifying suspicious sign-ins by looking for deviations from usual patterns is an example of anomaly detection.	☒	☐
Predicting whether a patient will develop diabetes based on the patient's medical history is an example of anomaly detection.	☒	☐

Explanation:

Statements	Yes	No
Forecasting housing prices based on historical data is an example of anomaly detection.	☐	☒
Identifying suspicious sign-ins by looking for deviations from usual patterns is an example of anomaly detection.	☒	☐
Predicting whether a patient will develop diabetes based on the patient's medical history is an example of anomaly detection.	☐	☒

Box 1: No

Box 2: Yes

Box 3: Yes

Anomaly detection encompasses many important tasks in machine learning:

Identifying transactions that are potentially fraudulent.

Learning patterns that indicate that a network intrusion has occurred.

Finding abnormal clusters of patients.

Checking values entered into a system.

Reference:
<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/anomaly-detection>

NEW QUESTION: 44

Match the principles of responsible AI to appropriate requirements.
To answer, drag the appropriate principles from the column on the left to its requirement on the right. Each principle may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Principles

Fairness

Privacy and security

Reliability and safety

Transparency

Answer Area

The system must not discriminate based on gender, race

Personal data must be visible only to approve

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

Answer:

Principles

Fairness

Privacy and security

Reliability and safety

Transparency

Answer Area

Fairness

Privacy and security

Transparency

The system must not discriminate based on gender, race

Personal data must be visible only to approve

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

Explanation:

Fairness

Privacy and security

Transparency

The system must not discriminate based on gender, race

Personal data must be visible only to approve

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

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify guiding principles for responsible AI", responsible AI is built upon six foundational principles: Fairness, Reliability and Safety, Privacy and Security, Inclusiveness, Transparency, and Accountability. Each principle serves to guide the ethical design, deployment, and management of artificial intelligence systems.

* Fairness - This principle ensures that AI systems treat all people fairly and do not discriminate based on personal attributes such as gender, race, or age. The Microsoft Learn content emphasizes that "AI systems should treat everyone fairly" and that organizations must evaluate datasets and model outputs for bias. In this scenario, "The system must not discriminate based on gender, race" clearly aligns with Fairness because it directly addresses equitable treatment and unbiased decision-making.

* Privacy and Security - Microsoft's responsible AI framework stresses that "AI systems must be secure and respect privacy." This means personal data should be safeguarded, processed lawfully, and visible only to authorized users. The statement "Personal data must be visible only to approved users" reflects the importance of protecting sensitive information and controlling access-precisely the intent of the Privacy and Security principle.

* Transparency - Transparency refers to ensuring that users understand how AI systems operate and make decisions. Microsoft notes that "AI systems should be understandable and users should be able to know why decisions are made." The requirement "Automated decision-making processes must be recorded so that approved users can identify why a decision was made" directly supports this principle. Transparency promotes trust and accountability by documenting the reasoning behind AI outputs.

Reliability and Safety, though another core principle, does not directly relate to any of the provided statements in this question.

NEW QUESTION: 45

Match the facial recognition tasks to the appropriate questions.

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

NOTE: Each correct selection is worth one point.

Tasks

grouping

identification

similarity

verification

Answer Area

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

Task Does this person look like other people?

Task Do all the faces belong together?

Task Who is this person in this group of people?

Answer:

Tasks

grouping

identification

similarity

verification

Answer Area

verification

similarity

grouping

identification

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

Does this person look like other people?

Do all the faces belong together?

Who is this person in this group of people?

Explanation:

Answer Area

verification

similarity

grouping

identification

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

Does this person look like other people?

Do all the faces belong together?

Who is this person in this group of people?

The correct matches are based on the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure." These materials explain that facial recognition tasks can be categorized into four major operations: verification, identification, similarity, and grouping. Each task serves a distinct purpose in facial recognition scenarios.

- * Verification - "Do two images of a face belong to the same person?"The verification task determines whether two facial images represent the same individual. Azure Face API compares the facial features and returns a confidence score indicating the likelihood that the two faces belong to the same person.
- * Similarity - "Does this person look like other people?"The similarity task compares a face against a collection of faces to find visually similar individuals. It does not confirm identity but measures how closely two or more faces resemble each other.
- * Grouping - "Do all the faces belong together?"Grouping organizes a set of unknown faces into clusters based on similar facial features. This is used when identities are not known beforehand, helping discover potential duplicates or visually similar clusters within an image dataset.
- * Identification - "Who is this person in this group of people?"The identification task is used when the system tries to determine who a specific person is by comparing their face against a known collection (face database or gallery). It returns the identity that best matches the input face.

According to Microsoft's AI-900 training, these tasks form the basis of Azure Face API's capabilities. Each helps solve a different type of facial recognition problem-from matching pairs to discovering unknown identities-making them essential components of responsible AI-based vision systems.

NEW QUESTION: 46

You need to use Azure Machine Learning designer to build a model that will predict automobile prices.

Which type of modules should you use to complete the model? To answer, drag the appropriate modules to the correct locations. Each module may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Modules

Convert to CSV

K-Means Clustering


Linear Regression

Split Data


Select Columns in Dataset

Summarize Data


Answer Area



Answer:




Explanation:



Box 1: Select Columns in Dataset

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

Example:



The task is to build a machine learning model in Azure Machine Learning designer to predict automobile prices, which is a regression problem since the output (price) is a continuous numeric value. The pipeline must follow the logical data preparation, training, and evaluation flow as outlined in the Microsoft Azure AI Fundamentals (AI-900) study guide and Microsoft Learn module "Create a machine learning model with Azure Machine Learning designer." Here's the correct sequence and reasoning:

* **Select Columns in Dataset:** The first step after loading the raw automobile dataset is to choose the relevant columns that will be used as features (inputs) and the label (output). This module ensures that only necessary fields (for example, horsepower, engine size, mileage, etc.) are used to train the model while excluding irrelevant columns like vehicle ID or serial number.

* Split Data:Next, the cleaned and filtered dataset must be split into two subsets: training data and testing data (often 70/30 or 80/20). This allows the model to be trained on one portion and evaluated on the other to measure predictive accuracy.

* Linear Regression:Since automobile price prediction is a numeric prediction task, the appropriate learning algorithm is Linear Regression. This supervised algorithm learns relationships between numeric features and the target (price).

Finally, the workflow connects the training data and Linear Regression module to the Train Model module, which outputs a trained regression model. The trained model is then linked to the Score Model module to compare predicted vs. actual prices.

This pipeline fully aligns with Microsoft's recommended process for regression in Azure ML Designer.

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (336 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 47

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation:

Statement

Yes / No

Providing an explanation of the outcome of a credit loan application is an example of the Microsoft transparency principle for responsible AI.

Yes

A triage bot that prioritizes insurance claims based on injuries is an example of the Microsoft reliability and safety principle for responsible AI.

Yes

An AI solution that is offered at different prices for different sales territories is an example of the Microsoft inclusiveness principle for responsible AI.

No

This question is based on the Responsible AI principles defined by Microsoft, which are part of the AI-900 Microsoft Azure AI Fundamentals curriculum. Microsoft's Responsible AI framework consists of six key principles: Fairness, Reliability and Safety, Privacy and Security, Inclusiveness, Transparency, and Accountability. Each principle ensures that AI systems are developed and used in a way that benefits people and society responsibly.

* Transparency Principle - Yes Providing an explanation for a loan decision aligns with the Transparency principle. Microsoft defines transparency as helping users and stakeholders understand how AI systems make decisions. For example, when a credit scoring AI model approves or denies a loan, explaining the factors that influenced that outcome (such as credit history or income level) ensures that customers understand the reasoning process. This builds trust and supports responsible deployment.

* Reliability and Safety Principle - Yes A triage bot that prioritizes insurance claims based on injury severity relates directly to Reliability and Safety. This principle ensures AI systems operate consistently, perform accurately, and produce dependable outcomes. In the case of the triage bot, it must reliably assess the input data (injury descriptions) and rank claims appropriately to avoid harm or misjudgment, aligning with Microsoft's emphasis on designing AI systems that are safe and robust.

* Inclusiveness Principle - No An AI solution priced differently across sales territories is not related to Inclusiveness. Inclusiveness focuses on ensuring accessibility and eliminating bias or exclusion for all users-especially those with disabilities or underrepresented groups. Pricing strategy is a business decision, not an inclusiveness issue. Therefore, this statement is No.

In summary, based on the AI-900 Responsible AI principles, the correct selections are:

NEW QUESTION: 48

Which machine learning technique can be used for anomaly detection?

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

Answer: ([SHOW ANSWER](#))

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore fundamental principles of machine learning," anomaly detection is a specialized machine learning technique used to identify data points, patterns, or events that deviate significantly from normal behavior.

Anomaly detection is widely used for monitoring time-series data and detecting unexpected or rare occurrences that may indicate problems, opportunities, or fraud. For example:

- * Detecting fraudulent transactions in banking systems.
- * Identifying equipment malfunctions in industrial IoT applications.
- * Monitoring network intrusions in cybersecurity.
- * Detecting unexpected spikes or drops in web traffic or sales.

In Azure, this workload is supported by the Azure AI Anomaly Detector service, which uses statistical and machine learning algorithms to learn from historical data and establish a baseline of normal behavior.

When the system detects data points that fall outside expected patterns, it flags them as anomalies.

Let's evaluate the incorrect options:

- * A. A machine learning technique that understands written and spoken language # This describes Natural Language Processing (NLP), not anomaly detection.
- * B. A machine learning technique that classifies objects based on user-supplied images # This refers to image classification, typically using computer vision.
- * D. A machine learning technique that classifies images based on their contents # Also describes computer vision, not anomaly detection.

Therefore, the correct answer is C, since anomaly detection specifically refers to analyzing data over time and identifying unusual or abnormal patterns that differ from the expected trend.

NEW QUESTION: 49

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

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

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

Answer: B ([LEAVE A REPLY](#))

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore natural language processing (NLP) in Azure," sentiment analysis is a core natural language processing (NLP) workload used to determine the emotional tone or attitude expressed in a piece of text. It helps identify whether a statement, review, or comment conveys a positive, negative, neutral, or mixed sentiment.

In this question, the scenario involves analyzing customer reviews and determining how positive or negative each review is. This directly aligns with sentiment analysis, which evaluates subjective text and quantifies the expressed opinion. In Azure, this workload is implemented through the Azure AI Language service (formerly Text Analytics API), where the Sentiment Analysis feature assigns a sentiment score to text inputs and classifies them accordingly.

For example:

- * "I love this product!" # Positive sentiment
- * "It's okay, but could be better." # Neutral or mixed sentiment
- * "I'm disappointed with the service." # Negative sentiment

Let's analyze why the other options are incorrect:

- * A. Language detection: Identifies which language (e.g., English, Spanish, French) the text is written in.

It doesn't measure positivity or negativity.

- * C. Key phrase extraction: Identifies the main topics or keywords in text (e.g., "battery life," "customer support"), not the emotion.
- * D. Entity recognition: Detects and categorizes specific entities such as people, locations, organizations, or dates within the text.

Therefore, based on Microsoft's AI-900 syllabus and Azure AI Language documentation, the workload that analyzes text to determine positive or negative opinions is Sentiment Analysis (Option B). This capability is widely used in customer feedback analysis, brand monitoring, and social media analytics to understand public perception and improve business decisions.

NEW QUESTION: 50

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

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

Answer: A ([LEAVE A REPLY](#))

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn modules under the topic "Describe features of common AI workloads" and "Identify capabilities of Azure AI services", QnA Maker is the service designed specifically to create a knowledge base (KB) or question-and-answer bot from existing content such as FAQ documents, product manuals, support pages, or structured knowledge sources.

QnA Maker enables developers to take semi-structured text (for example, an FAQ document or webpage) and automatically generate a knowledge base of pairs of questions and corresponding answers. This knowledge base can then be connected to a chatbot, typically through the Azure Bot Service, so that users can interact with it conversationally. The key advantage is that the process does not require deep machine learning or programming expertise. The service uses natural language processing (NLP) to match user queries with the most relevant pre-defined answers in the knowledge base.

In the AI-900 curriculum, this falls under the Conversational AI workload-creating intelligent bots that can respond naturally to user questions. Microsoft's training content explains that "QnA Maker extracts pairs of question and answer from your content and builds a knowledge base that can be queried by bots and other applications." The output, as shown in the example diagram, demonstrates how user input (the question) triggers a request to the QnA Maker API, which returns a JSON response containing the best-matched answer.

The other options are not correct because:

- * B. Language Understanding (LUIS) is used to interpret user intent and extract entities, not to create FAQs.

- * C. Text Analytics performs text extraction, sentiment analysis, and key-pharse detection but does not build a Q&A knowledge base.
- * D. Speech handles speech-to-text or text-to-speech, not Q&A matching.

Therefore, per the AI-900 study guide and Microsoft Learn, the verified and correct answer is A. QnA Maker.

NEW QUESTION: 51

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

Answer Area

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

recognition.

analysis.

detection.

recognition.

Answer:

Answer Area

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

recognition.

analysis.

detection.

recognition.

Explanation:

facial analysis.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe features of computer vision workloads on Azure," facial analysis is a computer vision capability that detects faces and extracts attributes such as facial expressions, emotions, pose, occlusion, and image quality factors like exposure and noise. It does not identify or verify individual identities; rather, it interprets facial features and image characteristics to analyze conditions in an image.

In this question, the AI solution helps photographers take better portrait photos by providing feedback on exposure, noise, and occlusion - tasks directly linked to facial analysis. The model analyzes the detected face to determine if the image is well-lit, clear, and unobstructed, thereby improving photo quality. These capabilities are part of the Azure Face service in Azure Cognitive Services, which includes both facial detection and facial analysis functionalities.

Here's how the other options differ:

- * Facial detection only identifies that a face exists in an image and provides its location using bounding boxes, without further interpretation.
- * Facial recognition goes a step further - it attempts to identify or verify a person's identity by comparing the detected face with stored images. This is not what the scenario describes.

Thus, when an AI solution evaluates image quality aspects like exposure, noise, and occlusion, it's performing facial analysis, which focuses on understanding image and facial characteristics rather than identification.

In summary, based on Microsoft's AI-900 study material, this scenario demonstrates facial analysis, a subcategory of computer vision tasks within Azure Cognitive Services.

NEW QUESTION: 52

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

NOTE: Each correct selection is worth one point.

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

D. Identity the retailer from a receipt.

Answer: ([SHOW ANSWER](#))

The correct answers are A and D because both scenarios involve extracting structured information from documents, which is exactly what Azure Form Recognizer is designed to do.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Explore computer vision", Form Recognizer is an Azure Cognitive Service that uses advanced Optical Character Recognition (OCR) and machine learning to extract key-value pairs, tables, and text from structured and semi-structured documents such as receipts, invoices, business cards, and forms. It allows organizations to automate data entry and digitize document processing.

* A. Extract the invoice number from an invoice # CorrectForm Recognizer can identify fields such as invoice number, total amount, date, vendor name, and billing address directly from invoices. It uses prebuilt models for invoices and receipts that automatically detect and extract relevant information without requiring extensive manual labeling. As stated in Microsoft Learn, "Form Recognizer extracts information from documents like receipts and invoices and returns structured data including key-value pairs."

* D. Identify the retailer from a receipt # CorrectThe prebuilt receipt model in Form Recognizer can read printed or scanned receipts and extract data points such as retailer name, transaction date, total amount, and tax information. This makes it ideal for expense reporting, auditing, or financial reconciliation.

The following options are incorrect:

* B. Translate a form from French to English # This task involves language translation, which is performed by Azure Translator, not Form Recognizer.

* C. Find an image of a product in a catalog # This requires object detection or image classification, which are part of Computer Vision, not Form Recognizer.

Therefore, based on Microsoft's AI-900 learning objectives and documentation, the two correct scenarios are:

A. Extract the invoice number from an invoice

D. Identify the retailer from a receipt

NEW QUESTION: 53

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

A. Use features for training and labels for evaluation.

B. Randomly split the data into rows for training and rows for evaluation.

C. Use labels for training and features for evaluation.

D. Randomly split the data into columns for training and columns for evaluation.

Answer: B ([LEAVE A REPLY](#))

<https://docs.microsoft.com/en-us/azure/machine-learning/algorithm-module-reference/split-data> The correct answer is B. Randomly split the data into rows for training and rows for evaluation.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe fundamental principles of machine learning on Azure", the process of developing a machine learning model involves dividing the available dataset into two or more parts-commonly training data and evaluation (or testing) data. The goal is to ensure that the model can learn patterns from one subset of the data (training set) and then be objectively tested on unseen data (evaluation set) to measure how well it generalizes to new situations.

The training dataset contains both features (the measurable inputs) and labels (the target outputs). The model learns from the patterns and relationships between these features and labels. The evaluation dataset also contains features and labels, but it is kept separate during the training phase. Once the model has been trained, it is tested on this unseen evaluation data to calculate metrics like accuracy, precision, recall, or F1 score.

Microsoft emphasizes that the data split should be random and based on rows, not columns. Each row represents a complete observation (for example, one customer record, one transaction, or one image).

Randomly splitting ensures that both subsets represent the same distribution of data, avoiding bias. Splitting by columns would separate features themselves, which would make the model training invalid.

The AI-900 materials often illustrate this using Azure Machine Learning's data preparation workflow, where data is randomly divided (commonly 70% for training and 30% for testing). This ensures the model learns from diverse examples and is fairly evaluated.

Therefore, the verified and correct approach, as per Microsoft's official guidance, is B. Randomly split the data into rows for training and rows for evaluation.

NEW QUESTION: 54

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

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

principle of the

- inclusiveness
- accountability
- reliability and safety
- fairness

Answer:

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

principle of the

- inclusiveness
- accountability
- reliability and safety
- fairness

Explanation:

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

principle of the

- inclusiveness
- accountability
- reliability and safety
- fairness

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

Reference:

<https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding-principles> AI systems should perform reliably and safely. For example, consider an AI-based software system for an autonomous vehicle; or a machine learning model that diagnoses patient symptoms and recommends prescriptions. Unreliability in these kinds of system can result in substantial risk to human life.

<https://docs.microsoft.com/en-us/learn/modules/get-started-ai-fundamentals/7-understand-responsible-ai>

NEW QUESTION: 55

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

Answer Area

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

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

Answer:

Answer Area

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

Custom Vision
Form Recognizer
Ink Recognizer
Text Analytics



Microsoft

Explanation:

Answer Area

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

Custom Vision
Form Recognizer
Ink Recognizer
Text Analytics



Microsoft

Accelerate your business processes by automating information extraction. Form Recognizer applies advanced machine learning to accurately extract text, key/value pairs, and tables from documents. With just a few samples, Form Recognizer tailors its understanding to your documents, both on-premises and in the cloud.

Turn forms into usable data at a fraction of the time and cost, so you can focus more time acting on the information rather than compiling it.

Reference:

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

NEW QUESTION: 56

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

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


anomaly detection.
computer vision.
conversational AI.
forecasting.

Answer:

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

anomaly detection.
computer vision.
conversational AI.
forecasting.

Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) study guide and official Microsoft Learn modules under "Describe features of common AI workloads", Conversational AI refers to technology that enables computers to engage in dialogue or conversation with users through natural language, whether by text or speech. The interactive answering of user-entered questions through a chat interface or virtual assistant is a direct example of a conversational AI workload.

Microsoft defines Conversational AI as systems that use natural language processing (NLP) and language understanding models to interpret what users are asking and respond appropriately. This includes chatbots, virtual assistants (like Cortana or Azure Bot Service), and automated customer service systems that simulate a human-like conversation. In this case, when an application answers questions that a user types interactively, the AI model is processing human language inputs, deriving intent, and generating meaningful replies - precisely what conversational AI is designed to do.

By contrast:

- * Anomaly detection identifies unusual patterns in data, typically used for fraud detection or equipment monitoring - not interactive dialogue.
- * Computer vision deals with interpreting images or video (e.g., object detection, facial recognition), unrelated to answering text-based questions.
- * Forecasting uses historical data to predict future trends or outcomes, often in sales or demand prediction scenarios.

The AI-900 guide emphasizes that Conversational AI helps businesses improve customer interaction efficiency by offering instant, automated, and consistent responses. It enables real-time engagement 24/7 and integrates with tools such as Azure Bot Service, Azure Cognitive Service for Language, and QnA Maker (now part of Azure AI Language Service).

Therefore, based on the Microsoft Learn objectives and definitions from the official AI-900 curriculum, the interactive answering of user questions in an application is best categorized as Conversational AI.

NEW QUESTION: 57

You are processing photos of runners in a race.

You need to read the numbers on the runners' shirts to identify the runners in the photos. Which type of computer vision should you use?

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

Answer: ([SHOW ANSWER](#))

The correct answer is B. Optical Character Recognition (OCR).

Optical Character Recognition (OCR) is a feature of Azure AI Vision that converts printed or handwritten text within images into machine-readable text. In this scenario, the goal is to read runner numbers on shirts from race photos. OCR can identify and extract these numbers, allowing them to be associated with specific participants.

Option analysis:

- * A. Image classification: Categorizes entire images (e.g., "runner," "crowd"), not text.
- * B. Optical Character Recognition (OCR) - # Correct. Extracts alphanumeric text from images.
- * C. Object detection: Identifies and locates objects (e.g., shoes, cars) but doesn't read text.
- * D. Facial recognition: Identifies individuals by matching facial features to known identities, not by reading numbers.

Therefore, to read and extract runner numbers from photos, the correct computer vision technique is Optical Character Recognition (OCR).

NEW QUESTION: 58

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements

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

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

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

Yes

No

☐

☐

☐

☐

☐

Answer:

Answer Area

Statements

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

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

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

Yes

No

☒

☐

☒

☐

☐

Explanation:

Yes, No, Yes.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify capabilities of Azure Cognitive Services for Language", the Azure Translator service is a cloud-based machine translation service used to translate text or entire documents between languages in real time. It uses REST APIs or client libraries to translate text input, detect languages, and support multiple target languages in a single request.

* "The following service call will accept English text as an input and output Italian and French text:

/translate?from=en&to=it,fr - Yes.This URL format is correct because the Translator service API allows multiple target languages to be specified in a single to parameter separated by commas. In this case, from=en defines the source language (English), and to=it,fr requests translations into Italian (it) and French (fr). The API would return results in both target languages simultaneously. This syntax is officially documented in Microsoft Learn as the valid format for multi-language translation.

* "The following service call will accept English text as an input and output Italian and French text:

/translate?from=en&to=fr&to=it - No.This format is incorrect, as the Translator API does not support repeating the to parameter multiple times. Only one to parameter is valid, and multiple target languages must be provided as a comma-separated list within the same to parameter.

* "The Translator service can be used to translate documents from English to French." - Yes.This statement is true. The Translator service supports both text translation and document translation. The document translation capability allows the translation of whole files such as Word, PowerPoint, or PDF documents while preserving formatting and structure. This feature is included in the official Translator API under "Document Translation." In summary, the AI-900 study content clarifies that:

/translate?from=en&to=it,fr # Valid syntax

/translate?from=en&to=fr&to=it # Invalid syntax

Translator can translate full documents between languages

NEW QUESTION: 59

You are building a Language Understanding model for an e-commerce business.

You need to ensure that the model detects when utterances are outside the intended scope of the model.

What should you do?

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

Answer: B (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of conversational AI workloads on Azure", a Language Understanding (LUIS) model is designed to interpret natural language input by identifying intents (the purpose of an utterance) and entities (specific data items in the utterance). Every LUIS model automatically includes a special intent called "None." This intent is used to handle utterances that do not fall into any of the model's defined intents. Adding examples of irrelevant or out-of-scope utterances to the None intent helps the model learn to recognize when a user's input does not match any existing categories. For example, if your e-commerce chatbot handles intents such as "TrackOrder" and "CancelOrder," but a user says "What's your favorite color?", that input should be mapped to the None intent so the bot can respond appropriately, such as "I'm not sure how to answer that." The AI-900 curriculum emphasizes that including diverse None intent examples improves model robustness and prevents false matches, thereby enhancing user experience. Other options are incorrect:

- * A. Test the model by using new utterances: Testing is important but does not define how to detect out- of-scope inputs.
- * C. Create a prebuilt task entity: Entities extract specific data but are unrelated to intent classification.
- * D. Create a new model: Unnecessary; handling out-of-scope utterances is done within the same model via the None intent.

Final answer: B. Add utterances to the None intent

NEW QUESTION: 60

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

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

Answer: B,C (LEAVE A REPLY)

To reduce operator load with a chatbot for predefined answers:

- * QnA Maker provides the knowledge base for answering questions automatically.
- * Azure Bot Service hosts and manages the chatbot interface to interact with users. Text Analytics and Translator Text are not required for basic Q & A chatbots.

NEW QUESTION: 61

Select the answer that correctly completes the sentence.



Answer:



Explanation:



In Azure OpenAI Service, the temperature parameter directly controls the creativity and determinism of responses generated by models such as GPT-3.5. According to the Microsoft Learn documentation for Azure OpenAI models, temperature is a numeric value (typically between 0.0 and 2.0) that determines how "random" or "deterministic" the output should be.

- * A lower temperature value (for example, 0 or 0.2) makes the model's responses more deterministic, meaning the same prompt consistently produces nearly identical outputs.
- * A higher temperature value (for example, 0.8 or 1.0) encourages creativity and variety, causing the model to generate different phrasing or interpretations each time it responds.

When a question specifies the need for more deterministic responses, Microsoft's guidance is to decrease the temperature parameter. This adjustment makes the model focus on the most probable tokens (words) rather than exploring less likely options, improving reliability and consistency-ideal for business or technical applications where consistent answers are essential.

The other parameters serve different purposes:

- * Frequency penalty reduces repetition of the same phrases but does not control randomness.
- * Max response (max tokens) limits the maximum length of the generated output.
- * Stop sequence defines specific tokens that tell the model when to stop generating text.

Thus, the correct and Microsoft-verified completion is:

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

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (336 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 62

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

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

Answer: B (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) learning path and Microsoft Learn documentation on GitHub Copilot, GitHub Copilot is powered by OpenAI Codex, a specialized language model derived from the GPT-3 family but fine-tuned specifically on programming languages and code data.

OpenAI Codex was designed to translate natural language prompts into executable code in multiple programming languages, including JavaScript, Python, C#, TypeScript, and Go. It can understand comments, function names, and code structure to generate relevant code suggestions in real time.

When a developer writes client-side JavaScript, GitHub Copilot uses Codex to analyze the context of the file and generate intelligent suggestions, such as completing functions, writing boilerplate code, or suggesting improvements. Codex can also explain what specific code does and provide inline documentation, which enhances developer productivity.

- * Option A (GPT-4): While some newer versions of GitHub Copilot (Copilot X) may integrate GPT-4 for conversational explanations, the core code completion engine remains based on Codex, as per the AI-900-level content.
- * Option C (DALL-E): Used for image generation, not for programming tasks.
- * Option D (GPT-3): Codex was fine-tuned from GPT-3 but has been further trained specifically for code generation tasks.

Therefore, the verified and official answer from Microsoft's AI-900 curriculum is B. Codex - the OpenAI model used by GitHub Copilot to make suggestions for client-side JavaScript and other programming languages.

NEW QUESTION: 63

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
You can use the Translator service to translate text between languages.	☐	☐
You can use the Translator service to detect the language of a given text.	☐	☐
You can use the Translator service to transcribe audible speech into text.	☐	☐

Answer:

Statements	Yes	No
You can use the Translator service to translate text between languages.	☒	☐
You can use the Translator service to detect the language of a given text.	☒	☐
You can use the Translator service to transcribe audible speech into text.	☐	☒

Explanation:

Statements	Yes	No
You can use the Translator service to translate text between languages.	☐	☐
You can use the Translator service to detect the language of a given text.	☐	☐
You can use the Translator service to transcribe audible speech into text.	☐	☒

The Translator service, part of Microsoft Azure Cognitive Services, is designed specifically for text translation between multiple languages. It is a cloud-based neural machine translation service that supports more than 100 languages. According to Microsoft Learn's module "Translate text with the Translator service", this service provides two main capabilities: text translation and automatic language detection.

* "You can use the Translator service to translate text between languages." # Yes This statement is true.

The primary purpose of the Translator service is to translate text accurately and efficiently between supported languages, such as English to Spanish or French to Japanese. It maintains contextual meaning using neural machine translation models.

* "You can use the Translator service to detect the language of a given text." # Yes This statement is also true. The Translator service includes automatic language detection, which determines the source language before translation. For instance, if a user submits text in an unknown language, the service can identify it automatically before performing translation.

* "You can use the Translator service to transcribe audible speech into text." # No This statement is false.

Transcribing speech (audio) into text is a function of the Azure Speech service, specifically the Speech- to-Text API, not the Translator service.

Therefore, the Translator service is used for text translation and language detection, while speech transcription belongs to the Speech service.

NEW QUESTION: 64

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

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

Answer: B (LEAVE A REPLY)

Azure AI Form Recognizer extracts information from structured and semi-structured documents. A custom model in Form Recognizer allows an organization to train the system on its specific document layouts and data fields.

As per the AI-900 study guide, a key advantage of a custom model is its flexibility. It can be trained with a set of labeled examples (e.g., invoices, purchase orders, receipts) that match the company's format.

Once trained, the model learns where to locate and extract fields such as invoice numbers, dates, or totals-regardless of layout differences between form types.

Option B is correct because a custom model can be trained to recognize a variety of form types, making it adaptable for diverse business processes.

Options A, C, and D are incorrect:

* A: Both prebuilt and custom models are cloud-based; on-premises deployment is not an exclusive feature.

* C: Custom models are not cheaper; they may involve additional training costs.

* D: Custom models do not always guarantee higher accuracy-accuracy depends on the training data quality.

NEW QUESTION: 65

You have an AI-based loan approval system.

During testing, you discover that the system has a gender bias.
Which responsible AI principle does this violate?

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

Answer: C (LEAVE A REPLY)

In Microsoft's Responsible AI principles, Fairness ensures that AI systems treat all individuals and groups equitably and make unbiased decisions. The AI-900 study guide explicitly states that fairness is violated when an AI model produces outcomes that systematically favor one group over another - such as preferring a particular gender, race, or age group.
In this scenario, a loan approval system shows gender bias, meaning it approves or rejects applications differently based on gender. This directly contradicts the fairness principle, as the AI system must make decisions solely based on relevant financial attributes (e.g., credit score, income) rather than personal characteristics.

Other principles explained in the AI-900 course include:

- * Accountability: Ensures human oversight and responsibility.
- * Transparency: Ensures users understand how decisions are made.
- * Reliability and Safety: Ensures consistent and accurate operation.

Since gender bias undermines equitable treatment, the principle violated is Fairness.

NEW QUESTION: 66

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

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

▼

evaluation.

feature engineering

hyperparameter tuning.

labeling.

Answer:

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

▼

evaluation.

feature engineering

hyperparameter tuning.

labeling.

Explanation:

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

▼

evaluation.

feature engineering

hyperparameter tuning.

labeling.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe core concepts of machine learning on Azure", labeling is the process of assigning correct output values (labels) to training data before model training. In supervised learning, every input in the dataset must be paired with its corresponding output so the algorithm can learn the relationship between the two.

In this scenario, the task is to assign classes to images before training a classification model-for example, marking images as "cat," "dog," or "bird." This process defines the target variable (label) that the model will later predict. During training, the classification model uses these labeled examples to learn patterns and distinguish between categories.

Microsoft's official materials clearly define labeling as:

"The process of tagging data with the correct answer so that the model can learn to make predictions." Labeling is a crucial early step in the machine learning lifecycle, especially for image classification and natural language processing (NLP) tasks. Without accurate labels, the model cannot learn correctly and its predictions will be unreliable.

Let's briefly clarify why the other options are incorrect:

- * Evaluation refers to testing the model after training to measure accuracy or performance using metrics like precision, recall, or F1 score.
- * Feature engineering involves creating or selecting the most relevant input features from raw data but does not involve tagging output labels.
- * Hyperparameter tuning adjusts parameters (like learning rate or depth of a tree) to optimize model performance after labeling and training have begun.

Thus, assigning classes to images prior to model training is definitively a Labeling task.

NEW QUESTION: 67

Select the answer that correctly completes the sentence.

Answer Area

Power Virtual Agents

Azure Health Bot

Microsoft Bot Framework

Power Virtual Agents

can be used to build no-code apps that use built-in natural language processing models.

Answer:

Answer Area

Power Virtual Agents

Azure Health Bot

Microsoft Bot Framework

Power Virtual Agents

can be used to build no-code apps that use built-in natural language processing models.

Explanation:

Answer Area

Power Virtual Agents

can be used to build no-code apps that use built-in natural language processing models.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe features of conversational AI workloads," the correct answer is Power Virtual Agents. This service is part of the Microsoft Power Platform and is specifically designed to enable users to create intelligent chatbots without writing any code.

Power Virtual Agents (PVA) provides a no-code/low-code environment where business users and developers can collaboratively design conversational experiences. It integrates built-in natural language processing (NLP) models to understand user intent and respond intelligently to text or speech inputs. The platform's interface allows chatbot creators to design dialogues visually, connect to back-end data via Power Automate, and publish bots on websites, Teams, or other communication channels.

This approach is highlighted in Microsoft Learn as an ideal solution for organizations that want to deploy conversational bots quickly without requiring specialized AI or programming expertise. PVA automatically leverages Microsoft's language understanding models, allowing it to interpret user input and map it to predefined topics or actions.

Let's analyze the other options:

- * Azure Health Bot: A specialized solution for the healthcare industry that provides prebuilt medical compliance and healthcare content. It is not a general-purpose, no-code chatbot builder.
- * Microsoft Bot Framework: A developer-focused SDK for building highly customized bots through code, offering maximum flexibility but not no-code functionality.

Therefore, the most appropriate choice that "can be used to build no-code apps that use built-in natural language processing models" is Power Virtual Agents - the official Microsoft no-code chatbot solution for conversational AI workloads.

NEW QUESTION: 68

Match the Azure Cognitive Services to the appropriate AI workloads.

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

NOTE: Each correct match is worth one point.

Services

Custom Vision

Face

Form Recognizer

Answer Area

Identify objects in an image.

Automatically import data from an invoice to a database.

Identify people in an image.

Answer:

Services

Custom Vision

Face

Form Recognizer

Answer Area

Custom Vision

Form Recognizer

Face

Identify objects in an image.

Automatically import data from an invoice to a database.

Identify people in an image.

Explanation:

Services

Custom Vision

Face

Form Recognizer

Answer Area

Custom Vision

Form Recognizer

Face

Identify objects in an image.

Automatically import data from an invoice to a database.

Identify people in an image.

The correct matches are Custom Vision, Form Recognizer, and Face - each corresponding to a distinct capability under Azure Cognitive Services as described in the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn modules on Computer Vision workloads.

- * Custom Vision # Identify objects in an imageThe Custom Vision service is part of the Azure Cognitive Services suite that enables developers to train custom image classification and object detection models. Unlike the prebuilt Computer Vision API, Custom Vision allows users to upload their own labeled images and teach the model to recognize specific objects relevant to their business context. The AI-900 syllabus explains that Custom Vision is ideal for tasks such as identifying products on a shelf, categorizing images, or detecting defects in manufacturing.
- * Form Recognizer # Automatically import data from an invoice to a databaseForm Recognizer is a document processing AI service that extracts structured data from forms, receipts, and invoices. It uses optical character recognition (OCR) combined with layout and key-value pair detection to automatically capture information such as invoice numbers, amounts, and vendor names. The AI-900 study materials highlight this service under the Document Intelligence category, emphasizing its ability to streamline data entry and business automation workflows by importing extracted data directly into databases or applications.

* Face # Identify people in an imageThe Face service provides advanced facial detection and recognition capabilities. It can locate faces in images, compare similarities between faces, identify known individuals, and even detect facial attributes such as age or emotion. The AI-900 course classifies this under Computer Vision services for person identification and security-related use cases such as access control or identity verification.

Thus, each mapping aligns precisely with the AI-900 official learning outcomes on Cognitive Services capabilities:

* Custom Vision # Object recognition

* Form Recognizer # Data extraction from forms

* Face # People identification

Final verified configuration:

* Custom Vision # Identify objects in an image

* Form Recognizer # Automatically import data from an invoice to a database

* Face # Identify people in an image

NEW QUESTION: 69

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation:

Statement

Yes / No

Providing an explanation of the outcome of a credit loan application is an example of the Microsoft transparency principle for responsible AI.

Yes

A triage bot that prioritizes insurance claims based on injuries is an example of the Microsoft reliability and safety principle for responsible AI.

Yes

An AI solution that is offered at different prices for different sales territories is an example of the Microsoft inclusiveness principle for responsible AI.

No

This question is based on the Responsible AI principles defined by Microsoft, a major topic in the AI-900:

Microsoft Azure AI Fundamentals certification. The goal of Responsible AI is to ensure that artificial intelligence is developed and used ethically, safely, and transparently to benefit people and society. Microsoft's framework defines six core principles: Fairness, Reliability and Safety, Privacy and Security, Inclusiveness, Transparency, and Accountability.

- * Transparency Principle - Yes Providing an explanation for a loan application decision clearly reflects transparency. According to Microsoft's Responsible AI guidelines, transparency involves ensuring that users and stakeholders understand how AI systems make decisions. When a financial AI model explains why a loan was approved or denied, it promotes user trust and confidence in automated decision-making. Transparency helps individuals understand influencing factors (like income or credit score), thereby fostering ethical AI deployment.
- * Reliability and Safety Principle - Yes A triage bot that prioritizes insurance claims based on injury severity demonstrates reliability and safety. This principle ensures that AI systems consistently operate as intended, handle data accurately, and do not cause unintended harm. For a triage bot, safety means it must correctly interpret medical or claim information and consistently provide appropriate prioritization. Microsoft emphasizes that reliable AI systems must be tested rigorously, function correctly in various scenarios, and maintain user safety at all times.
- * Inclusiveness Principle - No An AI solution priced differently for various sales territories is unrelated to inclusiveness. Inclusiveness focuses on designing AI systems that are accessible and fair to all users, including those with disabilities or from different demographic backgrounds. Price variation across territories is a business strategy, not an ethical AI inclusion concern. Hence, this statement does not align with any Responsible AI principle.

NEW QUESTION: 70

Select the answer that correctly completes the sentence.

Answer Area

 Microsoft

Asking a chatbot whether it will rain and being provided with a weather report is an example of

natural language processing.

computer vision.

natural language processing.

prediction and forecasting.

Answer:

Answer Area

 Microsoft

Asking a chatbot whether it will rain and being provided with a weather report is an example of

natural language processing.

computer vision.

natural language processing.

prediction and forecasting.

Answer Area

 Microsoft

Asking a chatbot whether it will rain and being provided with a weather report is an example of

natural language processing.

NEW QUESTION: 71

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

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

Answer: ([SHOW ANSWER](#))

Model Explain Ability.

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

Reference:

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

NEW QUESTION: 72

Providing contextual information to improve the responses quality of a generative AI solution is an example of which prompt engineering technique?

- A.** providing examples
- B.** fine-tuning
- C.** grounding data
- D.** system messages

Answer: **C** ([LEAVE A REPLY](#))

In Microsoft Azure OpenAI Service and the AI-900/AI-102 study materials, grounding data is the correct term used to describe the process of providing contextual or external information to improve the accuracy, relevance, and quality of responses generated by a generative AI model such as GPT-3.5 or GPT-4.

Grounding is a prompt engineering technique where the AI model is supplemented with relevant background data, such as company documents, knowledge bases, or user context, that helps the model generate factually correct and context-aware responses. Microsoft Learn defines grounding as a way to connect the model's general knowledge to specific, real-world information. For example, if you ask a GPT-3.5 model about your organization's HR policies, the base model will not know them unless that policy information is provided (grounded) in the prompt. By embedding this contextual data, the AI becomes "grounded" in the facts it needs to respond reliably.

This technique differs from other prompt engineering concepts:

- * A. Providing examples (few-shot prompting) shows the model sample inputs and outputs to guide formatting or style, not factual context.
- * B. Fine-tuning involves retraining the model with labeled data to permanently adjust its behavior - it's not a prompt-based technique.
- * D. System messages define the model's role, tone, or style (for example, "You are a helpful assistant") but do not add factual context.

Therefore, when you provide contextual information (like product details, policy documents, or reference text) within a prompt to enhance the quality and factual reliability of the model's responses, you are applying the grounding data technique.

NEW QUESTION: 73

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

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

Answer: ([SHOW ANSWER](#))

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe features of computer vision workloads on Azure," the "Describe Image" feature within Azure's Computer Vision service automatically generates textual captions that describe the contents of an image.


This feature uses a deep learning model trained on millions of labeled images to identify objects, people, and scenes, then formulates a natural language sentence summarizing what it sees. For example, given a photograph showing "two people sitting at a table with a laptop," the service might generate the caption: "A man and a woman sitting at a desk using a laptop." Here's how the other options differ:

- * A. Recognize text: Refers to Optical Character Recognition (OCR), which extracts written text from images, not generates descriptive captions.
- * C. Identify the areas of interest: Refers to detecting regions of an image that stand out visually, such as hotspots or significant features, not descriptive captioning.
- * D. Detect objects: Identifies and classifies objects in an image (e.g., cars, chairs, people) but doesn't produce a sentence or caption summarizing them.


Thus, only "Describe the images" generates automatic, human-readable captions that summarize photo content, a core computer vision workload taught in AI-900.

NEW QUESTION: 74

Select the answer that correctly completes the sentence.



Answer:



Explanation:
"features."

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Describe fundamental principles of machine learning on Azure," in a machine learning model, the data used as inputs are known as features, while the data that represents the output or target prediction is known as the label.

Features are measurable attributes or properties of the data used by a model to learn patterns and make predictions. They are also referred to as independent variables because they influence the result that the model tries to predict. For example, in a machine learning model that predicts house prices:

- * Features might include square footage, location, and number of bedrooms, while
- * The label would be the house price (the value being predicted).

In the context of Azure Machine Learning, during model training, features are passed into the algorithm as input variables (X-values), and the label is the corresponding output (Y-value). The model then learns the relationship between the features and the label.

Let's review the incorrect options:

- * Functions: These are mathematical operations or relationships used inside algorithms, not the input data itself.
- * Labels: These are the outputs or results that the model predicts, not the inputs.
- * Instances: These refer to individual data records or rows in the dataset, not the input fields themselves.

Hence, in any supervised or unsupervised learning process, the input data (independent variables) are called features, and the model uses them to predict labels (dependent variables).

NEW QUESTION: 75

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: (SHOW ANSWER)

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure," object detection is a computer vision technique used to locate and identify objects within an image. It not only determines what objects are present but also where they appear in the image by returning bounding box coordinates around each detected item. In this scenario, the goal is to identify the location of widgets within digital photos. This requires both recognition (knowing that the object is a widget) and localization (determining its position). The Custom Vision service in Azure allows you to train a model specifically for your own images, making it ideal for recognizing company-specific products such as widgets. By selecting the Object Detection domain in Custom Vision, you can label regions of interest in your training images. The model then learns to detect and locate those objects in new photos.

Let's examine the other options:

- * A. Computer Vision Spatial Analysis: Used for people tracking, movement detection, and occupancy analytics in video streams - not for locating products in still images.
- * C. Custom Vision classification: This model categorizes an image as a whole (e.g., "contains a widget" or "does not contain a widget") but does not locate objects within the image.
- * D. Computer Vision Image Analysis: Provides general image tagging, description, and OCR capabilities but does not pinpoint object locations.

NEW QUESTION: 76

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

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

classification.

clustering.

regression.

 Microsoft

Answer:

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


classification.

clustering.

regression.

 Microsoft

Explanation:



In the context of Microsoft Azure AI Fundamentals (AI-900) and general machine learning principles, regression refers to a type of supervised learning used to predict continuous numerical values based on historical data. The goal of regression is to model the relationship between input variables (features) and a continuous output variable (target).

In this scenario, the task is to predict how many vehicles will travel across a bridge on a given day. The number of vehicles is a numerical value that can vary continuously depending on factors such as time of day, weather, weekday/weekend, or traffic trends. Because the output is numeric and not categorical, this problem type clearly fits into regression analysis.

Microsoft's official learning content for AI-900, under "Identify features of regression and classification machine learning models," specifies that regression models are used to predict values such as sales forecasts, demand estimation, temperature prediction, or traffic volume-all of which share the same underlying objective: predicting a quantity.

To clarify other options:

- * Classification is used when predicting categories or discrete classes, such as determining whether an email is spam or not spam, or if an image contains a cat or a dog.
- * Clustering is an unsupervised learning technique used to group similar data points without predefined labels (for example, grouping customers by purchasing behavior).


Since predicting the number of vehicles results in a continuous numerical output, it aligns precisely with the regression workload type described in the Microsoft AI-900 study materials.

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (336 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 77

Select the answer that correctly completes the sentence.

Answer Area



Answer:

Answer Area

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

Regression
Classification
Clustering
Regression

Microsoft

Explanation:

Answer Area

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

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore fundamental principles of machine learning," a regression model is used when the goal is to predict a continuous numerical value based on historical data.

In this question, the task is to predict the sale price of auctioned items, which is a numeric output that can take on a wide range of values (for example, \$50.25, \$199.99, etc.). This makes it a regression problem because the output is continuous rather than categorical.

Regression models analyze the relationship between input features (such as item type, condition, age, bidding history, or demand) and a numerical target variable (the sale price). Common regression algorithms include linear regression, decision tree regression, and neural network regression. In Azure Machine Learning, these models are trained using labeled datasets containing known outcomes to learn patterns and make future predictions.

Let's review the incorrect options:

* Classification: Used to predict discrete categories or labels, such as "sold" vs. "unsold" or "low," "medium," "high." It cannot output continuous numeric predictions.

* Clustering: An unsupervised technique used to group similar data points based on shared characteristics, not to predict specific numeric outcomes.

Therefore, because predicting a sale price involves forecasting a continuous numerical value, the correct model type is Regression.

This aligns with Microsoft's AI-900 teaching that regression is used for tasks such as:

- * Predicting house prices
- * Forecasting sales revenue
- * Estimating car values or auction prices

NEW QUESTION: 78

Match the principles of responsible AI to the appropriate descriptions.

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

NOTE: Each correct match is worth one point.

Principles	Answer Area
Fairness	AI systems must consistently operate as intended, even under unexpected conditions.
Inclusiveness	AI systems must protect and secure personal and businesses information.
Privacy and security	
Reliability and safety	

Microsoft

Answer:



Explanation:



The correct answers are derived from the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Identify guiding principles for responsible AI." Microsoft defines six core principles of Responsible AI:

- * Fairness
- * Reliability and safety
- * Privacy and security
- * Inclusiveness
- * Transparency
- * Accountability

Each principle addresses a key ethical and operational requirement for developing and deploying trustworthy AI systems.

* Reliability and safety - "AI systems must consistently operate as intended, even under unexpected conditions."This principle ensures that AI models are dependable, robust, and perform accurately under diverse circumstances. Microsoft emphasizes that systems should be thoroughly tested and monitored to guarantee predictable behavior, prevent harm, and maintain safety. A reliable AI solution should continue to function properly when faced with unusual or noisy inputs, and fail safely when issues arise.

This principle focuses on stability, testing, and dependable performance.

* Privacy and security - "AI systems must protect and secure personal and business information."This principle ensures that AI systems comply with data privacy laws and ethical standards. It protects users' sensitive data against unauthorized access and misuse. Microsoft highlights that organizations must implement strong encryption, data anonymization, and access control mechanisms to maintain confidentiality. Protecting user data is essential to building trust and compliance with global standards like GDPR.

Other principles such as fairness and inclusiveness apply to ensuring equitable and accessible AI, but they do not directly relate to system operation or information protection.

Final Answers:

- * "Operate as intended" # Reliability and safety
- * "Protect and secure information" # Privacy and security

NEW QUESTION: 79

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

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

You have a machine learning model that predicts the value of ColumnE based on the other numeric columns.
Which type of model is this?

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

Answer: A (LEAVE A REPLY)

The dataset described contains numeric columns (ColumnA through ColumnE). The model's task is to predict the value of ColumnE based on the other numeric columns (A-D). This is a classic regression problem.

According to the Microsoft Azure AI Fundamentals (AI-900) study guide and Microsoft Learn module

"Identify common types of machine learning," a regression model is used when the target variable (the value to predict) is continuous and numeric, such as price, temperature, or-in this case-a numerical value in ColumnE.

Regression models analyze relationships between independent variables (inputs: Columns A-D) and a dependent variable (output: ColumnE) to predict a continuous outcome. Common regression algorithms include linear regression, decision tree regression, and neural network regression.

Option analysis:

- * A. Regression: # Correct. Used for predicting numerical, continuous values.
- * B. Analysis: # Incorrect. "Analysis" is a general term, not a machine learning model type.
- * C. Clustering: # Incorrect. Clustering is unsupervised learning, grouping similar data points, not predicting values.

Therefore, the type of machine learning model used to predict ColumnE (a numeric value) from other numeric columns is Regression, which fits perfectly within Azure's supervised learning models.

NEW QUESTION: 80

correctly completes the sentence.

Answer Area



Microsoft

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

features.
functions.
labels.
instances.

Answer:

Answer Area

Microsoft

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

features.
functions.
labels.
instances.

Explanation:

Features

The correct completion of the sentence is:

"In a machine learning model, the data that is used as inputs are called features." According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn module "Identify features of common machine learning types," the term feature refers to an individual measurable property or characteristic of the data that is used by a machine learning model to make predictions or decisions.

In supervised and unsupervised learning, features serve as the inputs to the model. They are the variables that represent the information the algorithm learns from. For example, if a model predicts the price of a house, the features might include square footage, number of bedrooms, location, and age of the property. These features feed into the model so that it can learn the relationship between inputs and outputs.

Microsoft Learn further defines these key concepts:

- * Features: Input variables (independent variables) used by the model to learn patterns.
- * Labels: The desired output or target variable that the model is trained to predict (e.g., price, category).
- * Instances: Individual rows or data records within the dataset (each instance has multiple features).
- * Functions: Algorithms or mathematical operations used during training and prediction - not data inputs.

Therefore, among the provided options - features, functions, labels, instances - only features accurately describe the data elements used as inputs for training or inference in a machine learning model.

In summary, within the AI-900 learning context:

- * Features = inputs to the model.
- * Labels = outputs for supervised learning.
- * Instances = examples (rows) of data.

NEW QUESTION: 81

Match the Azure Cognitive Services service to the appropriate actions.

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

NOTE: Each correct match is worth one point.

Services

Speech

Language service

Translator Text

Answer Area

Convert a user's speech to text.

Identify a user's intent.

Provide a spoken response to the user.



Microsoft

Answer:

Services

Speech

Language service


Translator Text

Answer Area

Convert a user's speech to text.


Identify a user's intent.

Provide a spoken response to the user.



Microsoft

Explanation:



These matches are based on the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore Azure Cognitive Services." Microsoft Azure provides Cognitive Services that enable developers to integrate artificial intelligence capabilities-such as vision, speech, language understanding, and decision-making-into applications without requiring in-depth AI expertise.

* Convert a user's speech to text # Speech ServiceThe Azure Speech Service supports speech-to-text (STT) conversion, which transcribes spoken language into written text. This feature is commonly used in voice assistants, transcription systems, and voice-enabled apps. The service uses advanced speech recognition models to handle different accents, languages, and background noises.

* Identify a user's intent # Language ServiceThe Azure AI Language Service (which includes capabilities from LUIS - Language Understanding) is used to interpret what a user means or wants to achieve based on their words. It identifies intents (the goal or action behind the input) and entities (key pieces of information) from natural language text. This is a key component in conversational AI applications, allowing chatbots and virtual assistants to respond intelligently.

* Provide a spoken response to the user # Speech ServiceThe Speech Service also supports text-to-speech (TTS) functionality, which converts textual responses into natural-sounding speech. This enables applications to communicate audibly with users, completing the conversational loop.

Translator Text is not used here because it's primarily designed for language translation between different languages, not for speech recognition or intent understanding.

NEW QUESTION: 82

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

Statements	Yes	No
A 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:

Answer Area

Statements

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

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

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

Yes

No

☒☐☐☒

Explanation:

< Smart home device answering questions # Yes

Webchat using Azure Bot Service # Yes

Auto-caption generation for videos # No

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe features of common AI workloads", Conversational AI refers to systems designed to engage in human-like dialogue through natural language-either text or speech. These systems include chatbots, virtual assistants, and voice-enabled smart devices, all of which can interpret user intent and respond meaningfully.

* A smart device in the home that responds to questions such as "When is my next appointment?" - Yes.

This is a classic example of Conversational AI. Devices like smart speakers or personal digital assistants use speech recognition, natural language understanding (NLU), and language generation to interpret spoken input and respond conversationally. The AI-900 study materials identify these as examples of voice-based conversational AI systems, which fall under the Speech and Language AI workloads.

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

- Yes. This statement is true. The Azure Bot Service is the primary Microsoft tool for creating and deploying conversational bots across various channels, including websites, Microsoft Teams, and other messaging platforms. The AI-900 syllabus specifically cites this as an example of implementing conversational AI for customer support or information retrieval.

* Automatically generating captions for pre-recorded videos is an example of conversational AI - No.

This is not conversational AI; instead, it falls under the Speech AI workload, specifically speech-to-text transcription. Automatically generating captions involves converting audio from video into written text but does not involve dialogue or interaction between a user and a system.

Thus, based on the official AI-900 guidance, only the first two scenarios describe conversational AI use cases.

NEW QUESTION: 83

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

You clean and transform the clinical data.

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

What should you do next?

A. Train the model by using the clinical data.

B. Split the clinical data into Two datasets.

C. Train the model by using automated machine learning (automated ML).

D. Validate the model by using the clinical data.

Answer: B (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn modules on machine learning concepts, ensuring that the accuracy of a predictive model can be proven requires data partitioning-specifically splitting the available data into training and testing datasets. This is a foundational concept in supervised machine learning.

When you split the data, typically about 70-80% of the dataset is used for training the model, while the remaining 20-30% is used for testing (or validation). The reason behind this approach is to ensure that the model's performance metrics-such as accuracy, precision, recall, and F1-score-are evaluated on data the model has never seen before. This prevents overfitting and allows you to demonstrate that the model generalizes well to new, unseen data.

In the AI-900 Microsoft Learn content under "Describe the machine learning process", it is explained that after cleaning and transforming the data, the next essential step is data splitting to "evaluate model performance objectively." By keeping training and testing data separate, you can prove the reliability and accuracy of the model's predictions, which is particularly crucial in sensitive domains like clinical or healthcare analytics, where decision transparency and validation are vital.

* Option A (Train the model by using the clinical data) is incorrect because you should not train and evaluate on the same data-it would lead to biased results.

* Option C (Train the model using automated ML) is incorrect because automated ML is a method for training and tuning, but it doesn't inherently prove accuracy.

* Option D (Validate the model by using the clinical data) is also incorrect if you use the same dataset for validation and training-it would not prove true accuracy.

Therefore, per Microsoft's official AI-900 study content, the verified correct answer is B. Split the clinical data into two datasets.

NEW QUESTION: 84

What are three Microsoft guiding principles for responsible AI? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. knowledgeability

B. decisiveness

C. inclusiveness

D. fairness

E. opinionatedness

F. reliability and safety

Answer: C,D,F (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe features of common AI workloads and considerations", Microsoft has defined six guiding principles for responsible AI. These principles are intended to ensure that AI systems are developed and deployed in ways that are ethical, transparent, and beneficial to all. The six principles are: Fairness, Reliability and Safety, Privacy and Security, Inclusiveness, Transparency, and Accountability.

Let's break down the three correct options:

* Fairness - Microsoft emphasizes that AI systems should treat all individuals fairly and avoid discrimination against people based on gender, race, age, or other characteristics. Fairness ensures that outcomes and decisions from AI systems are equitable across diverse user groups. In the AI-900 learning materials, fairness is explained as a foundational value that ensures algorithms and models do not introduce or amplify societal bias.

* Reliability and Safety - This principle ensures that AI systems function as intended under all expected conditions and that they can handle unexpected inputs safely. Microsoft states that AI should be tested rigorously and validated for reliability before deployment. AI systems must perform consistently and avoid causing harm due to errors or failures.

* Inclusiveness - Inclusiveness focuses on empowering everyone and engaging people of all backgrounds. Microsoft's responsible AI guidance stresses designing AI systems that understand and respect cultural, linguistic, and ability differences to make technology accessible and beneficial to all users.

Options A (knowledgeability), B (decisiveness), and E (opinionatedness) are not part of Microsoft's Responsible AI principles. These terms do not appear in any Microsoft Learn AI-900 curriculum or official responsible AI documentation.

Thus, based on the verified AI-900 study content and Microsoft's Responsible AI framework, the correct answer is C. Inclusiveness, D. Fairness, and F. Reliability and Safety.

NEW QUESTION: 85

Match the AI solution to the appropriate task.

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

NOTE: Each correct match is worth one point.

Solutions

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

Answer Area

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

Answer:

Solutions

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

Answer Area

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

Explanation:

Solutions

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

Answer Area

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

This question evaluates your understanding of how different Azure AI workloads correspond to specific tasks in image, text, and content generation scenarios, as explained in the Microsoft Azure AI Fundamentals (AI-900) study guide and Microsoft Learn modules covering common AI workloads and Azure services.

- * Generate a caption from a given image # Computer VisionThis is a computer vision task because it involves analyzing the visual elements of an image and producing descriptive text (a caption). Azure AI Vision provides image analysis and captioning capabilities through its Describe Image API, which uses deep learning models to recognize objects, scenes, and actions in an image and automatically generate natural-language descriptions (e.g., "A cat sitting on a sofa").
- * Generate an image from a given caption # Generative AIThis task belongs to Generative AI, which focuses on creating new content such as text, code, or images based on prompts. Tools like Azure OpenAI Service with DALL-E can interpret text descriptions and generate realistic images that match the given caption. Generative AI is capable of creative synthesis, not just analysis, making it the appropriate category.
- * Generate a 200-word summary from a 2,000-word article # Text AnalyticsText analytics (a subset of natural language processing) allows summarization, sentiment analysis, and entity recognition from large text corpora. Azure AI Language includes text summarization capabilities that condense long documents into concise summaries while preserving meaning and key information.

NEW QUESTION: 86

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

Statements	Yes	No
A webchat bot can interact with users visiting a website.	☐	☐
Automatically generating captions for pre-recorded videos is an example of conversational AI.	☐	☐
A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of conversational AI.	☐	☐

Answer:

Statements	Yes	No
A webchat bot can interact with users visiting a website.	☒	☐
Automatically generating captions for pre-recorded videos is an example of conversational AI.	☐	☒
A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of conversational AI.	☒	☐

Explanation:

- < A webchat bot can interact with users visiting a website # Yes
Automatically generating captions for pre-recorded videos is an example of conversational AI # No
A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of conversational AI # Yes
- \ These answers are based on the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore conversational AI in Microsoft Azure."
1. A webchat bot can interact with users visiting a website # Yes
This statement is true. A webchat bot is a key example of conversational AI, which allows users to communicate with an intelligent system through natural language. The Azure Bot Service supports a webchat channel, enabling website visitors to ask questions or get assistance directly through a chat interface embedded on a webpage. This allows businesses to provide 24/7 automated support and interactive engagement without human intervention.
2. Automatically generating captions for pre-recorded videos is an example of conversational AI # No
This is incorrect because automatically generating captions involves speech-to-text transcription, which falls under speech recognition and not conversational AI. While it uses AI to convert audio into text, it does not involve interactive communication or natural language dialogue. This task would be handled by Azure AI's Speech service, not the conversational AI framework.
3. A smart device in the home that responds to questions such as "What will the weather be like today?" is an example of conversational AI # Yes
This is true. Smart assistants like those found in home devices (e.g., voice-activated systems) use conversational AI technologies to process spoken language (using natural language processing and speech recognition) and generate appropriate responses. This interaction represents a classic example of conversational AI, as it allows human-like dialogue between a user and an AI system.
- # Final Answers:
* Webchat bot interacting with users # Yes
* Auto-captioning videos # No
* Smart home device answering questions # Yes

NEW QUESTION: 87

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

Answer Area

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

image classification.

object detection.

optical character recognizer (OCR).

semantic segmentation.

 Microsoft

Answer:

Answer Area

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

image classification.

object detection.

optical character recognizer (OCR).

semantic segmentation.

 Microsoft

Explanation:

Answer Area

 Microsoft

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

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

The correct answer is object detection. According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn module "Explore computer vision", object detection is the process of identifying and locating objects within an image or video. The primary characteristic of object detection, as emphasized in the study guide, is its ability to return a bounding box around each detected object along with a corresponding label or class.

In this question, the task involves returning a bounding box that indicates the location of a vehicle in an image. This is the exact definition of object detection - identifying that the object exists (a vehicle) and determining its position within the frame. Microsoft Learn clearly differentiates this from other computer vision tasks. Image classification, for example, only determines what an image contains as a whole (for instance, "this image contains a vehicle"), but it does not indicate where in the image the object is located.

Optical character recognition (OCR) is specifically used for extracting printed or handwritten text from images, and semantic segmentation involves classifying every pixel in an image to understand boundaries in greater detail, often used in autonomous driving or medical imaging.

The official AI-900 guide highlights object detection as one of the key computer vision workloads supported by Azure Computer Vision, Custom Vision, and Azure Cognitive Services. These services are designed to detect multiple instances of various object types in a single image, outputting bounding boxes and confidence scores for each.

Therefore, based on the AI-900 official curriculum and Microsoft Learn concepts, returning a bounding box that shows the location of a vehicle is a textbook example of object detection, as it involves both recognition and localization of the object within the image frame.

NEW QUESTION: 88

Which action can be performed by using the Azure AI Vision service?

- A. identifying breeds of animals in live video streams
- B. extracting key phrases from documents
- C. extracting data from handwritten letters
- D. creating thumbnails for training videos

Answer: A ([LEAVE A REPLY](#))

The Azure AI Vision service (formerly Computer Vision) is designed to analyze visual content in images and videos. According to Microsoft Learn's "Describe features of computer vision workloads," Azure AI Vision can identify objects, people, text, and scenes, and even classify images or detect objects in real time.

Identifying breeds of animals in live video streams is an example of image classification or object detection- core capabilities of Azure AI Vision. The Vision service can analyze each frame in a video, recognize animals, and classify them according to known categories, making this the correct answer.

The other options are incorrect:

- * B. Extracting key phrases from documents # Done by Azure AI Language (Text Analytics).
- * C. Extracting data from handwritten letters # Done by Azure AI Document Intelligence (Form Recognizer) using OCR.
- * D. Creating thumbnails for training videos # While possible in Azure Media Services, it's not a primary Azure AI Vision function.

Thus, the best answer is A. Identifying breeds of animals in live video streams.

NEW QUESTION: 89

What should you use to extract details from scanned images of contracts?

- A.** Azure AI Document Intelligence
- B.** Azure AI Immersive Reader
- C.** Azure OpenAI
- D.** Azure AI Search

Answer: A ([LEAVE A REPLY](#))

The correct answer is A. Azure AI Document Intelligence (previously known as Form Recognizer). This Azure Cognitive Service is specifically designed to extract structured data and key information from scanned documents, forms, and contracts using advanced Optical Character Recognition (OCR) combined with machine learning models.

According to the Microsoft Learn module "Extract data from documents with Azure AI Document Intelligence", this service enables automated data extraction from unstructured or semi-structured documents such as contracts, invoices, receipts, and purchase orders. It identifies key-value pairs, tables, and fields such as names, dates, amounts, and signatures. This makes it ideal for digitizing legal and business documents like contracts into structured formats that can be easily searched or stored in databases.

Azure AI Document Intelligence offers several model types:

- * Prebuilt models for common documents (invoices, receipts, business cards, etc.).
- * Custom models trained on your specific contract layouts.
- * Layout model for extracting raw text, tables, and structures.

The other options are incorrect:

- * B. Azure AI Immersive Reader enhances reading comprehension and accessibility but does not extract data from documents.
- * C. Azure OpenAI provides natural language generation and understanding but is not used for scanning or data extraction.
- * D. Azure AI Search indexes and searches textual or document content but relies on other services (like Document Intelligence) to extract the data first.

Therefore, to automatically extract details such as contract terms, names, dates, and signatures from scanned contract images, the best Microsoft AI service is A. Azure AI Document Intelligence

NEW QUESTION: 90

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

Which two parameters should you use to consume the pipeline? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A.** the model name
- B.** the training endpoint
- C.** the authentication key
- D.** the REST endpoint

Answer: ([SHOW ANSWER](#))

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore Azure Machine Learning", when you publish an inference pipeline (a deployed web service for real-time predictions) using Azure Machine Learning designer, you make the model accessible as a RESTful endpoint. Consumers-such as applications, scripts, or services-interact with this endpoint to submit data and receive predictions.

To securely access this deployed pipeline, two critical parameters are required:

* REST endpoint (Option D):The REST endpoint is a URL automatically generated when the inference pipeline is deployed. It defines the network location where clients send HTTP POST requests containing input data (usually in JSON format). The endpoint routes these requests to the deployed model, which processes the data and returns prediction results. The REST endpoint acts as the primary access point for consuming the model's inferencing capability programmatically.

* Authentication key (Option C):The authentication key (or API key) is a security token provided by Azure to ensure that only authorized users or systems can access the endpoint. When invoking the REST service, the key must be included in the request header (typically as the value of the Authorization header). This mechanism enforces secure, authenticated access to the deployed model.

The other options are incorrect:

* A. The model name is not required to consume the endpoint; it is used internally within the workspace.

* B. The training endpoint is used for training pipelines, not for inference.

Therefore, according to Microsoft's official AI-900 learning objectives and Azure Machine Learning documentation, when consuming a published inference pipeline, you must use both the REST endpoint (D) and the authentication key (C). These parameters ensure secure, controlled, and programmatic access to the deployed AI model for real-time predictions.

NEW QUESTION: 91

You are building an AI-based app.

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

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

NOTE: Each correct selection is worth one point.

A. Implement an Agile software development methodology

B. Implement a process of AI model validation as part of the software review process

C. Establish a risk governance committee that includes members of the legal team, members of the risk management team, and a privacy officer

D. Prevent the disclosure of the use of AI-based algorithms for automated decision making

Answer: (SHOW ANSWER)

The correct answers are B. Implement a process of AI model validation as part of the software review process and C. Establish a risk governance committee that includes members of the legal team, members of the risk management team, and a privacy officer.

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Responsible AI principles, responsible AI emphasizes six key principles: fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. These principles ensure that AI systems are trustworthy, ethical, and safe for users and society.

Option B aligns with the reliability and safety principle. Model validation ensures that AI models behave as expected, perform accurately across different data conditions, and produce consistent results.

Microsoft teaches that AI models should be validated, tested, and monitored regularly to avoid unintended outcomes, bias, or failures. Validation processes help ensure that the AI behaves responsibly before deployment and continues to perform reliably over time.

Option C aligns with the accountability and governance principle. Establishing a risk governance committee that includes legal, privacy, and risk management experts ensures that AI development and deployment are overseen responsibly. This committee is responsible for reviewing compliance with data protection laws, ensuring ethical practices, and managing risks associated with AI-driven decisions.

Microsoft emphasizes that accountability requires human oversight and governance structures to ensure ethical alignment throughout the AI system's lifecycle.

The incorrect options are:

* A. Implement an Agile software development methodology: Agile is a software project management approach, not a Responsible AI principle.

* D. Prevent the disclosure of the use of AI-based algorithms: This violates the transparency principle, which requires organizations to disclose when and how AI is used.

Therefore, following the official Responsible AI framework taught in AI-900, the correct and verified answers are B and C, as they directly promote reliability, safety, accountability, and governance in AI systems.

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (336 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 92

You deploy the Azure OpenAI service to generate images.
You need to ensure that the service provides the highest level of protection against harmful content.
What should you do?

- A. Configure the Content filters settings.
- B. Customize a large language model (LLM).
- C. Configure the system prompt
- D. Change the model used by the Azure OpenAI service.

Answer: (SHOW ANSWER)

The correct answer is A. Configure the Content filters settings.

When using the Azure OpenAI Service for text or image generation, Microsoft provides built-in content filtering to help detect and block potentially harmful or unsafe outputs. These filters are part of Microsoft's Responsible AI framework and are designed to prevent the generation of offensive, violent, sexual, or otherwise restricted content.

To ensure the highest level of protection, you can configure content filter settings within the Azure OpenAI deployment. This allows you to define stricter policies based on your organization's safety requirements. For image generation models such as DALL E, enabling or strengthening these filters ensures that inappropriate or unsafe images are not generated or returned.

* B (Customize an LLM): Customization affects behavior but not safety filtering.

* C (Configure the system prompt): Adjusts response style but doesn't guarantee content safety.

* D (Change the model): Different models have similar filter systems; protection level depends on filter configuration, not the model itself.

NEW QUESTION: 93

Select the answer that correctly completes the sentence.



Answer:

Answer Area

Microsoft

According to Microsoft's

fairness

accountability

fairness

inclusiveness

transparency

principle of responsible AI,

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

Explanation:

Answer Area

Microsoft

According to Microsoft's

fairness

principle of responsible AI,

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

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Identify guiding principles for responsible AI," Fairness is one of Microsoft's six core principles of Responsible AI. The principle of fairness ensures that AI systems treat all individuals and groups equitably, and that the models do not produce biased or discriminatory outcomes.

Bias in AI systems can occur when training data reflects existing prejudices, inequalities, or imbalances. For example, if a dataset used for a hiring model underrepresents a certain demographic group, the AI system might produce unfair recommendations. Microsoft emphasizes that AI should not reflect or reinforce bias and that developers must actively design, test, and monitor models to mitigate unfairness.

Microsoft's Six Responsible AI Principles:

- * Fairness - AI systems should treat everyone equally and avoid bias.
- * Reliability and safety - AI systems must operate as intended even under unexpected conditions.
- * Privacy and security - AI must protect personal and business data.
- * Inclusiveness - AI should empower all people and be accessible to diverse users.
- * Transparency - AI systems should be understandable and their decisions explainable.
- * Accountability - Humans should be accountable for AI system outcomes.

The other options do not fit this context:

- * Accountability ensures human responsibility for AI decisions.
- * Inclusiveness focuses on accessibility and empowering all users.
- * Transparency relates to making AI systems understandable.

Therefore, the correct answer is fairness, as it directly addresses the principle that AI systems should NOT reflect biases from the datasets used to train them.

NEW QUESTION: 94

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
You train a regression model by using unlabeled data.	☐	☐
The classification technique is used to predict sequential numerical data over time.	☐	☐
Grouping items by their common characteristics is an example of clustering.	☐	☐

Answer:

Statements	Yes	No
You train a regression model by using unlabeled data.	☐	☐
The classification technique is used to predict sequential numerical data over time.	☐	☐
Grouping items by their common characteristics is an example of clustering.	☐	☐

Explanation:

Statements	Yes	No
You train a regression model by using unlabeled data.	☐	☒
The classification technique is used to predict sequential numerical data over time.	☐	☒
Grouping items by their common characteristics is an example of clustering.	☒	☐

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of common machine learning types", there are three main types of machine learning: supervised learning, unsupervised learning, and reinforcement learning. Within supervised learning, two common approaches are regression and classification, while clustering is a primary example of unsupervised learning.

* "You train a regression model by using unlabeled data." - No. Regression models are trained with labeled data, meaning the input data includes both features (independent variables) and target labels (dependent variables) representing continuous numerical values. Examples include predicting house prices or sales forecasts. Unlabeled data (data without target output values) cannot be used to train regression models; such data is used in unsupervised learning tasks like clustering.

* "The classification technique is used to predict sequential numerical data over time." - No.

Classification is used for categorical predictions, where outputs belong to discrete classes, such as spam

/not spam or disease present/absent. Predicting sequential numerical data over time refers to time series forecasting, which is typically a regression or forecasting problem, not classification. The AI-900 syllabus clearly separates classification (categorical prediction) from regression (continuous value prediction) and time series (temporal pattern analysis).

* "Grouping items by their common characteristics is an example of clustering." - Yes. This statement is correct. Clustering is an unsupervised learning technique used to group similar data points based on their features. The AI-900 study materials describe clustering as the process of "discovering natural groupings in data without predefined labels." Common examples include customer segmentation or document grouping.

Therefore, based on Microsoft's AI-900 training objectives and definitions:

* Regression # supervised learning using labeled continuous data (No)

* Classification # categorical prediction, not sequential numeric forecasting (No)

* Clustering # grouping by similarity (Yes)

NEW QUESTION: 95

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

NOTE: Each correct selection is worth one point.

A. a telephone answering service that has a pre-recorder message

B. a chatbot that provides users with the ability to find answers on a website by themselves

C. telephone voice menus to reduce the load on human resources

D. a service that creates frequently asked questions (FAQ) documents by crawling public websites

Answer: B,C ([LEAVE A REPLY](#))

According to the AI-900 official study guide and Microsoft Learn module "Describe features of conversational AI workloads on Azure", conversational AI refers to artificial intelligence systems that interact with users through natural language via text or speech. These systems include chatbots, virtual assistants, and interactive voice response (IVR) systems that simulate human conversation.

* B. Chatbot that provides users with the ability to find answers on a website by themselves This is a classic example of conversational AI. Chatbots use natural language understanding (LUIS) and Azure Bot Service to interpret user input, identify intent, and provide relevant responses automatically. They help users self-serve information without human support, such as retrieving account details or answering FAQs.

* C. Telephone voice menus to reduce the load on human resources Automated telephone systems or IVRs use conversational AI to interpret spoken commands and route calls intelligently. This is often implemented using Azure Cognitive Services Speech (for speech-to-text and text-to-speech) combined with Azure Bot Service for managing dialogue flow.

NEW QUESTION: 96

You need to provide customers with the ability to query the status of orders by using phones, social media, or digital assistants.

What should you use?

A. Azure AI Bot Service

B. the Azure AI Translator service

C. an Azure AI Document Intelligence model

D. an Azure Machine Learning model

Answer: A (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify Azure services for conversational AI," the Azure AI Bot Service is specifically designed to create intelligent conversational agents (chatbots) that can interact with users across multiple communication channels, such as web chat, social media, phone calls, Microsoft Teams, and digital assistants.

In this scenario, customers need the ability to query the status of their orders through various interfaces - including voice and text platforms. Azure AI Bot Service enables this by integrating with Azure AI Language (for understanding natural language), Azure Speech (for speech-to-text and text-to-speech capabilities), and Azure Communication Services (for telephony or chat integration).

The bot can interpret user input like "Where is my order?" or "Check my delivery status," call backend systems (such as an order database or API), and then respond appropriately to the user through the same communication channel.

Let's analyze the incorrect options:

* B. Azure AI Translator Service: Used for real-time text translation between languages; it doesn't handle conversation logic or database queries.

* C. Azure AI Document Intelligence model: Extracts data from structured and semi-structured documents (e.g., invoices, receipts), not user queries.

* D. Azure Machine Learning model: Builds and deploys predictive models, but doesn't provide conversational or multi-channel interaction capabilities.

Thus, for enabling multi-channel conversational experiences where customers can inquire about order statuses using voice, chat, or digital assistants, the most appropriate solution is Azure AI Bot Service, as outlined in Azure's AI conversational workload documentation.

NEW QUESTION: 97

You are evaluating whether to use a basic workspace or an enterprise workspace in Azure Machine Learning.

What are two tasks that require an enterprise workspace? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Use a graphical user interface (GUI) to run automated machine learning experiments.

B. Create a compute instance to use as a workstation.

C. Use a graphical user interface (GUI) to define and run machine learning experiments from Azure Machine Learning designer.

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

Answer: (SHOW ANSWER)

The correct answers are A. Use a graphical user interface (GUI) to run automated machine learning experiments and C. Use a graphical user interface (GUI) to define and run machine learning experiments from Azure Machine Learning designer.

According to the Microsoft Azure AI Fundamentals (AI-900) official documentation and Microsoft Learn module "Create and manage Azure Machine Learning workspaces", there are two workspace tiers: Basic and Enterprise. The Enterprise workspace provides advanced capabilities for automation, visualization, and collaboration that are not available in the Basic tier.

Specifically:

- * Automated machine learning (AutoML) using a GUI is only available in the Enterprise tier. AutoML automatically selects algorithms and tunes hyperparameters through the Azure Machine Learning studio interface.

- * Azure Machine Learning designer, which allows users to visually drag and drop datasets and modules to create machine learning pipelines, also requires the Enterprise workspace.

In contrast:

- * B. Create a compute instance and D. Create a dataset from a CSV file are fundamental actions supported in both Basic and Enterprise workspaces. These do not require the advanced licensing features of the Enterprise edition.

Therefore, tasks involving the graphical, no-code tools-Automated ML (AutoML) and the Designer- require the Enterprise workspace, aligning with AI-900's learning objectives.

Reference:Microsoft Learn - Create and manage Azure Machine Learning workspaces and Automate model training with Azure Machine Learning

NEW QUESTION: 98

You have an Azure Machine Learning pipeline that contains a Split Data module. The Split Data module outputs to a Train Model module and a Score Model module. What is the function of the Split Data module?

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

Answer: B (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of Azure Machine Learning", the Split Data module in an Azure Machine Learning pipeline is used to divide a dataset into two or more subsets-typically a training dataset and a testing (or validation) dataset. This is a fundamental step in the supervised machine learning workflow because it allows for accurate evaluation of the model's performance on data it has not seen during training.

In a typical workflow, the data flows as follows:

- * The dataset is first preprocessed (cleaned, normalized, or transformed).
- * The Split Data module divides this dataset into two parts - one for training the model and another for testing or scoring the model's accuracy.
- * The Train Model module uses the training data output from the Split Data module to learn patterns and build a predictive model.
- * The Score Model module then takes the trained model and applies it to the test data output to measure how well the model performs on unseen data.

The Split Data module typically uses a defined ratio (such as 0.7:0.3 or 70% for training and 30% for testing).

This ensures that the trained model can generalize well to new, real-world data rather than simply memorizing the training examples.

Now, addressing the incorrect options:

- * A. Selecting columns that must be included in the model is done by the Select Columns in Dataset module.
- * C. Diverting records that have missing data is handled by the Clean Missing Data module.
- * D. Scaling numeric variables is done using the Normalize Data or Edit Metadata modules.

Therefore, based on the official AI-900 learning objectives, the verified and most accurate answer is B.

creating training and validation datasets.

NEW QUESTION: 99

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
The Azure OpenAI GPT-3.5 Turbo model can transcribe speech to text.	☐	☐
The Azure OpenAI DALL-E model generates images based on text prompts.	☐	☐
The Azure OpenAI embeddings model can convert text into numerical vectors based on text similarities.	☐	☐

Answer:

Statements	Yes	No
The Azure OpenAI GPT-3.5 Turbo model can transcribe speech to text.	☐	☒
The Azure OpenAI DALL-E model generates images based on text prompts.	☒	☐
The Azure OpenAI embeddings model can convert text into numerical vectors based on text similarities.	☒	☐

Explanation:

Statements	Yes	No
The Azure OpenAI GPT-3.5 Turbo model can transcribe speech to text.	☐	☒
The Azure OpenAI DALL-E model generates images based on text prompts.	☒	☐
The Azure OpenAI embeddings model can convert text into numerical vectors based on text similarities.	☒	☐

- * "The Azure OpenAI GPT-3.5 Turbo model can transcribe speech to text." - NOThis statement is false. The GPT-3.5 Turbo model is a text-based large language model (LLM) designed for natural language understanding and generation, such as answering questions, summarizing text, or writing content. It does not process or transcribe audio input. Speech-to-text capabilities belong to Azure AI Speech Services, specifically the Speech-to-Text API, not Azure OpenAI.
 - * "The Azure OpenAI DALL-E model generates images based on text prompts." - YESThis statement is true. The DALL-E model, available within Azure OpenAI Service, is a generative AI model that creates original images from natural language descriptions (text prompts). For example, given a prompt like "a futuristic city at sunset," DALL-E generates a unique, high-quality image representing that concept. This aligns with generative AI workloads in the AI-900 study guide, where DALL-E is specifically mentioned as an image-generation model.
 - * "The Azure OpenAI embeddings model can convert text into numerical vectors based on text similarities." - YESThis statement is also true. The embeddings model in Azure OpenAI converts text into multi-dimensional numeric vectors that represent semantic meaning. These embeddings enable tasks such as semantic search, recommendations, and text clustering by comparing similarity scores between vectors. Words or phrases with similar meanings have vectors close together in the embedding space.
- In summary:
- * GPT-3.5 Turbo # Text generation (not speech-to-text)
 - * DALL-E # Image generation from text prompts

* Embeddings # Convert text into numerical semantic representations
Correct selections: No, Yes, Yes.

NEW QUESTION: 100

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

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

Answer: B,D (LEAVE A REPLY)

The correct answers are B. a smart device in the home that responds to questions such as, "What will the weather be like today?" and D. a website that uses a knowledge base to interactively respond to users' questions.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of Natural Language Processing (NLP) workloads on Azure", Natural Language Processing (NLP) is a branch of artificial intelligence that enables computers to understand, interpret, and generate human language in a meaningful way. NLP bridges the gap between human communication and machine understanding, allowing systems to process both spoken and written language.

- * Option B - A smart device in the home that responds to questions such as "What will the weather be like today?"This is an example of an NLP workload because the device must process spoken language (speech-to-text), interpret the user's intent (language understanding), and generate a relevant spoken response (text-to-speech). This workflow involves several Azure Cognitive Services, such as Speech Service for recognizing and synthesizing speech, and Language Understanding (LUIS) for interpreting intent. This aligns with conversational AI and NLP tasks in the AI-900 syllabus.
- * Option D - A website that uses a knowledge base to interactively respond to users' questions.This is also an NLP workload because the system interprets text input from users and retrieves appropriate answers from a knowledge base. Microsoft's QnA Maker (now part of the Azure AI Language service) and Azure Bot Service enable such behavior. The model uses NLP to understand the user's question, find the most relevant response, and generate an appropriate reply - key characteristics of natural language processing.

Incorrect options:

- * A (assembly line machinery) represents automation or robotics, not NLP.
- * C (monitoring temperature to activate a fan) is an example of an IoT (Internet of Things) or rule-based system, not related to language processing.

NEW QUESTION: 101

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

According to Microsoft's

▼ principle of responsible AI,

accountability
fairness
inclusiveness
transparency

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

Answer:

According to Microsoft's  principle of responsible AI,

accountability
fairness
inclusiveness
transparency

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

Explanation:

According to Microsoft's  principle of responsible AI,

accountability
fairness
inclusiveness
transparency

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

According to Microsoft's Responsible AI principles, one of the six core principles is fairness, which ensures that AI systems treat all individuals equitably and that their outcomes are not influenced by biases present in the training data or algorithms. The official Microsoft Learn module "Identify the guiding principles for responsible AI" clearly defines fairness as the requirement that AI systems should not amplify or perpetuate existing societal biases.

In this scenario, the statement emphasizes that AI systems should NOT reflect biases from the datasets used to train them, which directly aligns with the fairness principle. Bias in AI models can arise when the data used for training is unbalanced or not representative of the real-world population. For instance, if a facial recognition model is trained mostly on images of one demographic group, it may perform poorly on others- an example of unfair bias. Microsoft advocates building and testing AI systems with diverse, high-quality datasets to ensure fair performance across all groups.

The other principles listed-accountability, inclusiveness, and transparency-are also important but do not directly address bias mitigation:

- * Accountability ensures that people remain responsible for AI systems and their decisions.
- * Inclusiveness promotes accessibility and usability for all people, including those with disabilities.
- * Transparency focuses on explaining how AI systems make decisions.

However, Fairness explicitly deals with avoiding discrimination and bias in AI outcomes and training data.

Thus, in Microsoft's Responsible AI framework, ensuring that systems do not reflect biases from datasets is part of the Fairness principle, which promotes equitable and unbiased treatment for all individuals in AI- driven decisions.

NEW QUESTION: 102

What is an example of a Microsoft responsible AI principle?

- A. AI systems should treat people fairly.
- B. AI systems should NOT reveal the details of their design.
- C. AI systems should use black-box models.
- D. AI systems should protect the interests of developers.

Answer: [\(SHOW ANSWER\)](#)

Full Detailed Explanation (250-300 words):

The correct answer is A. AI systems should treat people fairly.
This statement aligns with one of Microsoft's six Responsible AI principles, which are:

- * Fairness - AI systems should treat all people fairly and avoid bias.
- * Reliability and Safety
- * Privacy and Security
- * Inclusiveness
- * Transparency
- * Accountability

The principle of Fairness ensures that AI models do not discriminate based on factors such as race, gender, age, or socioeconomic background. For example, a loan approval or hiring model must provide equal opportunity to all qualified applicants regardless of demographic differences.

- * B (Not revealing design details) contradicts Transparency, which promotes openness about AI functionality.
- * C (Black-box models) goes against Microsoft's push for Explainable AI.
- * D (Protect developers' interests) is not part of Microsoft's Responsible AI framework.

Therefore, the verified correct answer is A. AI systems should treat people fairly.

NEW QUESTION: 103

Match the Azure OpenAI large language model (LLM) process to the appropriate task.
To answer, drag the appropriate process from the column on the left to its task on the right. Each process may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Processes

Classifying

Generating

Summarizing

Translating

Answer Area

Microsoft

freepdfdumps.com

Detect the genre of a work of fiction.

Create a list of bullet points based on text input.

Create advertising slogans from a product description.

Answer:

Processes

Classifying

Generating

Summarizing

Translating

Answer Area

Classifying

Detect the genre of a work of fiction.

Summarizing

Create a list of bullet points based on text input.

Generating

Create advertising slogans from a product description.



Explanation:

Processes

Classifying

Generating

Summarizing

Translating

Answer Area

Classifying

Detect the genre of a work of fiction.

Summarizing

Create a list of bullet points based on text input.

Generating

Create advertising slogans from a product description.



According to the Microsoft Azure AI Fundamentals (AI-900) study material and Azure OpenAI Service documentation, large language models (LLMs) such as GPT are capable of performing multiple natural language processing (NLP) tasks depending on the intent of the prompt. These tasks generally fall into categories like classification, generation, summarization, and translation, each with a distinct purpose and output type.

* Classifying - This process involves analyzing text and assigning it to a predefined category or label based on its content. The scenario "Detect the genre of a work of fiction" clearly fits this category. The model must evaluate the text and determine whether it belongs to genres like mystery, romance, or science fiction. This is a classic text classification problem, as the output is a discrete category derived from textual features.

* Summarizing - This process means condensing lengthy text into a shorter version that preserves the key information. In the scenario "Create a list of bullet points based on text input," the model extracts essential information and reformats it as concise bullet points, which is an abstraction form of summarization. Summarization models help users quickly understand the main ideas from long documents, meeting efficiency and readability goals.

* Generating - This refers to the LLM's ability to produce new, creative content based on input instructions. The task "Create advertising slogans from a product description" represents generation because it requires the model to construct original text that didn't previously exist. Generation tasks showcase the creativity and contextual fluency of models like GPT in marketing and content creation. Thus, these mappings align directly with the Azure OpenAI LLM capabilities taught in AI-900, linking each NLP process with its most suitable real-world task.

NEW QUESTION: 104

Select the answer that correctly completes the sentence.

Answer Area

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

classification.

clustering.

regression.

classification.

regularization.

Answer:

Answer Area



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

classification.

clustering.

regression.

classification.

regularization.

Explanation:

Clustering.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of common machine learning types", clustering is an unsupervised machine learning technique used to group data points into distinct segments or clusters based on shared characteristics.

Unlike supervised learning (classification or regression), clustering works with unlabeled data, discovering natural groupings without predefined outcomes.

In this question, Recency, Frequency, and Monetary (RFM) values are common marketing metrics used to evaluate customer behavior:

* Recency - how recently a customer made a purchase.

* Frequency - how often they make purchases.

* Monetary - how much money they spend.

Using RFM analysis, a company can segment its customers into groups such as "loyal," "occasional," or "at-risk" buyers. This segmentation process does not rely on predefined labels but rather discovers patterns within the data - which is the defining characteristic of clustering.

In the AI-900 context, clustering is described as a method that "groups items with similar features so that items in the same group are more similar to each other than to those in other groups." Common algorithms used include K-Means, Hierarchical Clustering, and DBSCAN, all available within Azure Machine Learning Designer and other Azure ML environments.

To clarify the incorrect options:

* Classification is supervised learning used to predict discrete categories (e.g., yes/no, spam/not spam).

* Regression predicts continuous numeric values (e.g., house prices).

* Regularization is a model optimization technique, not a type of machine learning.

Therefore, when businesses use RFM values to identify customer segments without labeled outcomes, this is an application of unsupervised learning through clustering.

NEW QUESTION: 105

Capturing text from images is an example of which type of AI capability?

- A. text analysis
- B. optical character recognition (OCR)
- C. image description
- D. object detection

Answer: B ([LEAVE A REPLY](#))

The correct answer is B. Optical character recognition (OCR).

OCR is a key capability within the Computer Vision and Document Intelligence services in Azure AI that enables systems to detect and extract printed or handwritten text from images and scanned documents. When capturing text from images, OCR technology analyzes visual patterns (shapes of letters and numbers) and converts them into machine-readable text. For example, a photo of a receipt, street sign, or printed report can be processed to extract textual content programmatically.

- * A (Text analysis): Applies to NLP tasks such as sentiment detection or key phrase extraction, not image processing.
- * C (Image description): Generates captions describing the scene or objects in an image.
- * D (Object detection): Identifies and locates objects but does not extract text.

NEW QUESTION: 106

Match the machine learning tasks to the appropriate scenarios.

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

NOTE: Each correct selection is worth one point.

Learning Types

Feature engineering

Feature selection

Model deployment

Model evaluation

Model training

Answer Area

TaskExamining the values of a confusion matrix

TaskSplitting a date into month, day, and year fields

TaskPicking temperature and pressure to train a weather model

Answer:

Learning Types

Feature engineering

Feature selection

Model deployment

Model evaluation

Model training

Answer Area

Model evaluation

Feature engineering

Feature selection

Examining the values of a confusion matrix

Splitting a date into month, day, and year fields

Picking temperature and pressure to train a weather model

Explanation:

Answer Area

Model evaluation

Feature engineering

Feature selection

Examining the values of a confusion matrix

Splitting a date into month, day, and year fields

Picking temperature and pressure to train a weather model

This question tests your understanding of machine learning workflow tasks as described in the Microsoft Azure AI Fundamentals (AI-900) study guide and the Microsoft Learn module "Explore the machine learning process." The AI-900 curriculum divides the machine learning lifecycle into key phases: data preparation, feature engineering and selection, model training, model evaluation, and model deployment. Each phase has specific tasks designed to prepare, build, and assess predictive models before deployment.

* Examining the values of a confusion matrix # Model evaluationIn Azure Machine Learning, evaluating a model involves checking its performance using metrics such as accuracy, precision, recall, and F1-score. The confusion matrix is one of the most common tools for this purpose. According to Microsoft Learn, "model evaluation is the process of assessing a trained model's performance against test data to ensure reliability before deployment." Analyzing the confusion matrix helps determine whether predictions align with actual outcomes, making this task part of model evaluation.

* Splitting a date into month, day, and year fields # Feature engineeringFeature engineering refers to transforming raw data into features that better represent the underlying patterns to improve model performance. The study guide describes it as "the process of creating new input features from existing data." Splitting a date field into separate numeric fields (month, day, year) is a classic example of feature engineering because it enables the model to learn from temporal patterns that might otherwise remain hidden.

* Picking temperature and pressure to train a weather model # Feature selectionFeature selection involves identifying the most relevant variables that have predictive power for the model. As defined in Microsoft Learn, "feature selection is the process of choosing the most useful subset of input features for training." In this scenario, selecting temperature and pressure variables as inputs for a weather prediction model fits perfectly within the feature selection stage.

Therefore, the correct matches are:

Examining confusion matrix # Model evaluation

Splitting date field # Feature engineering

Picking temperature & pressure # Feature selection

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (336 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 107

Your website has a chatbot to assist customers.
You need to detect when a customer is upset based on what the customer types in the chatbot.
Which type of AI workload should you use?

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

Answer: D (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) curriculum and Microsoft Learn module "Explore natural language processing", NLP enables computers to understand, interpret, and analyze human language. One of its key capabilities is sentiment analysis, which detects emotional tone (positive, negative, or neutral) in text.

In this scenario, the chatbot must detect when a customer is upset based on what they type. This directly maps to sentiment analysis, a core NLP function. The Text Analytics service within Azure Cognitive Services provides prebuilt sentiment analysis models that return a sentiment score and classification (e.g., positive, neutral, negative). As per Microsoft Learn, "Natural language processing allows systems to understand sentiment and intent within text and speech to derive meaningful insights." Explanation of other options:

- * A. Anomaly detection identifies unusual patterns in data (e.g., fraud detection), not emotions in text.
- * B. Semantic segmentation is a computer vision technique used to label pixels in an image.
- * C. Regression predicts continuous numeric values and is not related to understanding text or emotions.

Therefore, to enable the chatbot to detect when a user is upset based on text input, the correct AI workload is Natural Language Processing (NLP), specifically through Azure Text Analytics sentiment analysis.

NEW QUESTION: 108

You need to predict the income range of a given customer by using the following dataset.

First Name	Last Name	Age	Education Level	Income Range
Orlando	Gee	45	University	25,000-50,000
Keith	Harris	36	High school	25,000-50,000
Donna	Carreras	52	University	50,000-75,000
Janet	Gates	21	University	75,000-100,000
Lucy	Harrington	68	High school	50,000-75,000

Which two fields should you use as features? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

- A. Education Level
- B. Last Name
- C. Age
- D. Income Range
- E. First Name

Answer: A,C (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe core concepts of machine learning on Azure", when building a predictive machine learning model, features are the input variables used by the algorithm to predict the target label. The target label is the output or value the model is trained to predict.

In this dataset, the target variable is clearly the Income Range, since the goal is to predict a customer's income bracket. Therefore, Income Range (D) is the label, not a feature. Features must be other attributes that help the model make this prediction.

The fields Education Level (A) and Age (C) are the most relevant features because both can logically and statistically influence income level.

* Education Level is a categorical variable that often correlates strongly with income. Individuals with higher education levels tend to earn more on average, making this an important predictor.

* Age is a numerical variable that typically affects income level due to factors such as experience and career progression.

By contrast:

* First Name (E) and Last Name (B) are irrelevant as features because they are identifiers, not meaningful predictors of income. Including them could lead to bias or model overfitting without contributing to accurate predictions.

Hence, according to AI-900 principles, the features used to train a model predicting income range would be Education Level and Age.

NEW QUESTION: 109

brectly completes the sentence.

Answer Area

A historian can use

facial analysis

image classification

object detection

optical character recognition (OCR)

to digitize newspaper articles.

Answer:

Answer Area

A historian can use

facial analysis

image classification

object detection

optical character recognition (OCR)

to digitize newspaper articles.

Explanation:

Answer Area

A historian can use

optical character recognition (OCR)

to digitize newspaper articles.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of common AI workloads", OCR (Optical Character Recognition) is a Computer Vision technology that detects and extracts printed or handwritten text from images and scanned documents.

OCR allows organizations and individuals to convert physical or image-based text into machine-readable, editable, and searchable digital text.

In the context of this question, a historian working with old newspaper articles or archival documents would use OCR to digitize printed content. For instance, the historian can scan or photograph old newspaper pages, and then use an OCR tool-such as Azure Computer Vision's OCR API-to automatically recognize and extract the textual content from those images. This process enables the historian to store, edit, and analyze the content digitally without manually typing everything.

OCR works by using deep learning algorithms trained on thousands of text samples. The system analyzes patterns, shapes, and spatial relationships of characters to identify text accurately, even from low-quality or aged paper documents. Once extracted, the digital text can be indexed, translated, or processed further using Natural Language Processing (NLP) tools for content analysis.

Now, addressing the other options:

- * Facial analysis is used to detect emotions, age, or gender from human faces-irrelevant to text digitization.
- * Image classification identifies entire images by categories (e.g., cat, car, flower).
- * Object detection identifies and locates multiple objects within an image but doesn't extract text.

Therefore, per the AI-900 learning objectives under the Computer Vision workload, the correct and verified completion is:

NEW QUESTION: 110

A company employs a team of customer service agents to provide telephone and email support to customers.

The company develops a webchat bot to provide automated answers to common customer queries.

Which business benefit should the company expect as a result of creating the webchat bot solution?

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

Answer: B (LEAVE A REPLY)

Full Detailed Explanation with exact Extract from your Official Study guide and Trained Data at least 250 to 300 words in Explanation:

The correct answer is B. a reduced workload for the customer service agents.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Describe features of common AI workloads", conversational AI solutions such as chatbots are primarily designed to automate repetitive and routine customer interactions. The key business value emphasized in these materials is operational efficiency-chatbots allow organizations to respond to a high volume of customer queries without relying solely on human agents. This results in reduced workload, lower operational costs, and faster response times.

Microsoft's AI-900 learning objectives highlight that AI can be applied to automate tasks that previously required human interaction. In the context of customer support, a webchat bot powered by Azure AI services (such as Azure Bot Service or Azure Cognitive Services for Language) can handle frequently asked questions like order status, password resets, or basic troubleshooting. This allows human agents to focus their time and skills on more complex issues that require empathy, reasoning, or decision-making-tasks that AI cannot yet handle as effectively.

Additionally, the AI-900 course materials explain that one of the measurable business benefits of deploying AI-driven chatbots is improved efficiency and scalability. Chatbots can handle thousands of simultaneous interactions, something that human teams cannot easily do. As a result, the organization experiences reduced operational pressure on support staff, improved customer satisfaction due to quicker responses, and optimized resource utilization.

Options A and C are incorrect because chatbots do not directly influence sales growth or product reliability.

While increased customer satisfaction might indirectly support sales, it is not the primary or guaranteed outcome of implementing a chatbot. Similarly, product reliability is tied to engineering quality, not customer service automation.

Therefore, based on the official AI-900 study materials and Microsoft Learn concepts, the best and verified answer is B. a reduced workload for the customer service agents.

NEW QUESTION: 111

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 (LEAVE A REPLY)

In Natural Language Processing (NLP), entities are pieces of information extracted from text, such as names, locations, or phone numbers. According to the Microsoft Learn module "Explore natural language processing in Azure," Azure's Language Understanding (LUIS) supports several entity types:

- * Machine-learned entities - Automatically learned based on context in training data.
- * List entities - Used for predefined, limited sets of values (e.g., colors or product names).


* Pattern.any entities - Capture flexible, unstructured phrases in user input.

* Regular expression entities - Use regex patterns to match specific data formats such as phone numbers, postal codes, or dates.

A regular expression is ideal for recognizing phone numbers because phone numbers follow specific numeric or symbol-based patterns (e.g., (555)-123-4567 or +1 212 555 0199). By defining a regex pattern, the AI model can accurately extract phone numbers regardless of text context.

NEW QUESTION: 112

Select the answer that correctly completes the sentence.



Answer:



Explanation:



"Optical Character Recognition (OCR) extracts text from handwritten documents." According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of computer vision workloads," Optical Character Recognition (OCR) is a computer vision capability that enables AI systems to detect and extract printed or handwritten text from images, scanned documents, and photographs.

Microsoft Learn explains that OCR uses machine learning algorithms to analyze visual data, locate regions containing text, and then convert that text into machine-readable digital format. This capability is essential for automating processes such as document digitization, form processing, and data extraction.

OCR technology is provided through services such as the Azure Cognitive Services Computer Vision API and Azure Form Recognizer. The Computer Vision API's OCR feature can extract text from both typed and handwritten sources, including receipts, invoices, letters, and forms. Once extracted, this text can be processed, searched, or stored electronically, enabling automation and efficiency in document management systems.

Let's review the incorrect options:

* Object detection identifies and locates objects in an image by drawing bounding boxes (e.g., detecting vehicles or people).

* Facial recognition identifies or verifies individuals by comparing facial features.

* Image classification assigns an image to one or more predefined categories (e.g., "dog," "car," "tree").

None of these perform the task of extracting textual content from images - that is uniquely handled by Optical Character Recognition (OCR).

Therefore, based on the AI-900 official study content, the verified and correct answer is Optical Character Recognition (OCR), as it specifically extracts text (printed or handwritten) from image-based documents.

NEW QUESTION: 113

You are building an AI system.

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

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

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

C. Provide documentation to help developers debug code.

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

Answer: ([SHOW ANSWER](#))

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and Microsoft Learn module "Describe principles of responsible AI", the transparency principle ensures that AI systems are understandable, explainable, and well-documented so that users, developers, and stakeholders can know how the system operates and makes decisions. Transparency involves clear communication, documentation, and interpretability.

Microsoft defines transparency as the responsibility to make sure that people understand how AI systems function, their limitations, and how decisions are made. For developers, this means providing detailed documentation and model interpretability tools so others can inspect, debug, and understand the AI model's behavior. For users, it means ensuring that the purpose, capabilities, and limitations of the AI system are clearly explained.

Providing documentation to help developers debug and understand how a service works directly aligns with this transparency principle. It ensures that the system's logic and behavior are open to inspection and that any unintended consequences can be identified and corrected. Transparency also builds trust in AI solutions by enabling accountability and oversight.

Let's analyze the other options:

* A. Ensure that all visuals have an associated text that can be read by a screen reader - This supports inclusiveness, not transparency, as it focuses on accessibility for all users.

* B. Enable autoscaling to ensure that a service scales based on demand - This is related to system performance and scalability, not responsible AI.

* D. Ensure that a training dataset is representative of the population - This supports fairness, as it prevents bias and ensures equitable outcomes.

Therefore, based on the official AI-900 training content and Microsoft's Responsible AI framework (which includes fairness, reliability, privacy, inclusiveness, transparency, and accountability), the correct answer is C.

Provide documentation to help developers debug code, because this directly promotes transparency in how the AI system operates and communicates its inner workings

NEW QUESTION: 114

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

* Accept customer orders.

* Retrieve support documents.

* Retrieve order status updates.

Which type of NLP should you use?

A. sentiment analysis

B. translation

C. language modeling

D. named entity recognition

Answer: **C** ([LEAVE A REPLY](#))

In the Microsoft Azure AI Fundamentals (AI-900) curriculum, language modeling is described as a core component of Natural Language Processing (NLP) that enables an AI system to understand, interpret, and generate human language in context. The AI-900 Microsoft Learn module "Identify features of Natural Language Processing workloads" explains that language modeling is used to analyze user input and determine intent - essential for conversational systems like chatbots.

In this question, the chatbot must:

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

All these tasks require the bot to understand user intent and context from text input. This understanding process is driven by language modeling, which predicts meaning and structure within sentences, enabling the system to decide what action to take next.

Microsoft Learn distinguishes between various NLP techniques:

- * Sentiment analysis detects emotional tone (positive/negative/neutral).
- * Translation converts text between languages.
- * Named Entity Recognition (NER) identifies specific entities like names or dates. However, none of these individually allow a system to process commands, requests, or user intents - that capability is part of language modeling, which powers LUIS (Language Understanding Intelligent Service) or the modern Azure Cognitive Service for Language.

Therefore, to build a chatbot that can interpret commands and respond contextually - such as processing orders or retrieving documents - you must use language modeling.

NEW QUESTION: 115

Select the answer that correctly completes the sentence.

Answer Area

Azure AI Document Intelligence

Azure AI Bot Service

Azure AI Language

Azure AI Document Intelligence

The Microsoft Bot Framework SDK

can be used to create a knowledge base from frequently asked questions (FAQ).



Answer:

Answer Area

Azure AI Document Intelligence

Azure AI Bot Service

Azure AI Language

Azure AI Document Intelligence

The Microsoft Bot Framework SDK

can be used to create a knowledge base from frequently asked questions (FAQ).



Explanation:

Answer Area

Azure AI Document Intelligence

can be used to create a knowledge base from frequently asked questions (FAQ).



The correct answer is Azure AI Language, which includes the Question Answering capability (previously known as QnA Maker). According to the Microsoft Azure AI Fundamentals (AI-900) study guide and Microsoft Learn documentation, the Azure AI Language service can be used to create a knowledge base from frequently asked questions (FAQ) and other structured or semi-structured text sources.

This service allows developers to build intelligent applications that can understand and respond to user questions in natural language by referencing prebuilt or custom knowledge bases. The Question Answering feature extracts pairs of questions and answers from documents, websites, or manually entered data and uses them to construct a searchable knowledge base. This knowledge base can then be integrated with Azure Bot Service or other conversational platforms to create interactive, self-service chatbots.

Here's how it works:

- * Developers upload FAQ documents, URLs, or structured content.
- * Azure AI Language processes the content and identifies logical question-answer pairs.
- * The model stores these pairs in a knowledge base that can be queried by user input.
- * When users ask questions, the model finds the best matching answer using natural language understanding techniques.

In contrast:

- * Azure AI Document Intelligence (Form Recognizer) is used to extract structured data from forms and documents, not to create FAQ knowledge bases.
 - * Azure AI Bot Service is for managing and deploying conversational bots but does not generate knowledge bases.
 - * Microsoft Bot Framework SDK provides tools for building conversational logic but still requires a knowledge source like Question Answering from Azure AI Language.
- Therefore, the service that can create a knowledge base from FAQ content is Azure AI Language.

NEW QUESTION: 116

What is a form of unsupervised machine learning?

- A.** multiclass classification
- B.** clustering
- C.** binary classification
- D.** regression

Answer: B (LEAVE A REPLY)

As outlined in the AI-900 study guide and Microsoft Learn's "Explore fundamental principles of machine learning" module, clustering is a core example of unsupervised machine learning.

In unsupervised learning, the model is trained on data without labeled outcomes. The goal is to discover patterns or groupings naturally present in the data. Clustering algorithms, such as K-means, DBSCAN, or Hierarchical clustering, analyze similarities among data points and group them into clusters. For example, clustering can group customers by purchasing behavior or segment products by shared characteristics - all without predefined labels.

Supervised learning, by contrast, uses labeled data (input-output pairs) to train a model that predicts outcomes. This includes:

- * A. Multiclass classification - Predicts more than two categories (e.g., classifying images as dog, cat, or bird).
- * C. Binary classification - Predicts two categories (e.g., spam vs. not spam).
- * D. Regression - Predicts continuous numeric values (e.g., price prediction).

Therefore, the only option representing unsupervised learning is clustering, which enables data discovery without predefined labels.

NEW QUESTION: 117

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
The Text Analytics service can identify in which language text is written.		☐	☐
The Text Analytics service can detect handwritten signatures in a document.		☐	☐
The Text Analytics service can identify companies and organizations mentioned in a document.		☐	☐

Answer:

Answer Area			
Statements		Yes	No
The Text Analytics service can identify in which language text is written.		☒	☐
The Text Analytics service can detect handwritten signatures in a document.		☐	☒
The Text Analytics service can identify companies and organizations mentioned in a document.		☒	☐

Explanation:

Answer Area

Statements	Yes	No
The Text Analytics service can identify in which language text is written.	☒	☐
The Text Analytics service can detect handwritten signatures in a document.	☐	☒
The Text Analytics service can identify companies and organizations mentioned in a document.	☒	☐

The Azure Text Analytics service, a component of Azure Cognitive Services, provides natural language processing (NLP) capabilities to analyze and understand text-based data. According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features and uses for Natural Language Processing (NLP)", the Text Analytics service supports multiple text understanding tasks, such as language detection, key phrase extraction, sentiment analysis, and entity recognition.

* Language Identification - Yes:Text Analytics can automatically detect the language in which text is written. This feature analyzes linguistic patterns and assigns a language code (for example, "en" for English, "es" for Spanish). It is one of the primary features described in Microsoft Learn as part of the service's Language Detection API.

* Detect Handwritten Signatures - No: Detecting handwritten signatures is not a text-based NLP task. Instead, it belongs to the computer vision domain, specifically Optical Character Recognition (OCR). The Text Analytics service only processes digital text, not handwritten or image-based data. To detect handwriting or signatures, you would use the Computer Vision OCR API, not Text Analytics.

* Entity Recognition - Yes: The Text Analytics service can identify named entities—such as people, locations, organizations, dates, and quantities—within documents. This is known as Named Entity Recognition (NER), which helps extract structured information from unstructured text.

NEW QUESTION: 118

Select the answer that correctly completes the sentence.

Answer Area

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

classification

classification

clustering

regression

 type of machine learning.

Answer:

Answer Area

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

classification

classification

clustering

regression

 type of machine learning.

Explanation:

Answer Area

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

classification

 type of machine learning.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of common machine learning types", the classification technique is a type of supervised machine learning used to predict which category or class a new observation belongs to, based on patterns learned from labeled training data.

In this scenario, a banking system that predicts whether a loan will be repaid is dealing with a binary outcome—either the loan will be repaid or will not be repaid. These two possible results represent distinct classes, making this problem a classic example of binary classification. During training, the model learns from historical data containing features such as customer income, credit score, loan amount, and repayment history, along with labeled outcomes (repaid or defaulted). After training, it can classify new applications into one of these two categories.

The AI-900 curriculum distinguishes between three key supervised and unsupervised learning approaches:

- * Classification: Predicts discrete categories (e.g., spam/not spam, fraud/not fraud, will repay/won't repay).
- * Regression: Predicts continuous numerical values (e.g., house prices, sales forecast, temperature).
- * Clustering: Groups data based on similarity without predefined labels (e.g., customer segmentation).

Since the banking problem focuses on predicting a categorical outcome rather than a continuous numeric value, it fits squarely into the classification domain. In Azure Machine Learning, such tasks can be performed using algorithms like Logistic Regression, Decision Trees, or Support Vector Machines (SVMs), all configured for categorical prediction.

Therefore, per Microsoft's official AI-900 learning objectives, a banking system predicting whether a loan will be repaid represents a classification type of machine learning problem.

NEW QUESTION: 119

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

Statements	Yes	No
You can use QnA Maker to query an Azure SQL database.	☐	☐
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	☐	☐
The QnA Maker service can determine the intent of a user utterance.	☐	☐

Answer:

Statements	Yes	No
You can use QnA Maker to query an Azure SQL database.	☐	☒
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	☒	☐
The QnA Maker service can determine the intent of a user utterance.	☐	☒

Statements	Yes	No
You can use QnA Maker to query an Azure SQL database.	☐	☒
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	☒	☐
The QnA Maker service can determine the intent of a user utterance.	☐	☒

Reference:

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

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/choose-natural-language-processing-service> QnA maker conversational AI service and has nothing to do with SQL database You can easily create a user support bot solution on Microsoft Azure using a combination of two core technologies:

- QnA Maker. This cognitive service enables you to create and publish a knowledge base with built-in natural language processing capabilities.
- Azure Bot Service. This service provides a framework for developing, publishing, and managing bots on Azure.

<https://docs.microsoft.com/en-us/learn/modules/build-faq-chatbot-qna-maker-azure-bot-service/2-get-started-qna-bot> LUIS is used to understand user intent from utterances. Creating a language understanding application with Language Understanding consists of two main tasks. First you must define entities, intents, and utterances with which to train the language model - referred to as authoring the model. Then you must publish the model so that client applications can use it for intent and entity prediction based on user input.

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

NEW QUESTION: 120

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.

Models

Classification

Clustering

Regression

Answer Area

A supervised machine learning model used to predict numeric values.

A supervised machine learning model used to predict categories.

An unsupervised machine learning model used to group similar entities based on features.

Answer:

Models

Classification

Clustering

Regression

Answer Area

Regression

Classification

Clustering

A supervised machine learning model used to predict numeric values.

A supervised machine learning model used to predict categories.

An unsupervised machine learning model used to group similar entities based on features.

Explanation:

Models

Classification

Clustering

Regression

Answer Area

Regression

Classification

Clustering

A supervised machine learning model used to predict numeric values.

A supervised machine learning model used to predict categories.

An unsupervised machine learning model used to group similar entities based on features.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of common machine learning types", the three main machine learning model types differ by their purpose and the kind of data they use - whether supervised (using labeled data) or unsupervised (using unlabeled data).

* Regression # A supervised machine learning model used to predict numeric values. Regression is a type of supervised learning that predicts continuous numerical outcomes. It learns the relationship between input features (independent variables) and a continuous target variable (dependent variable). Examples include predicting house prices, sales revenue, or temperature. The AI-900 curriculum emphasizes regression for "predicting numeric values based on known data," using algorithms such as linear regression or decision tree regression.

* Classification # A supervised machine learning model used to predict categories. Classification is also a supervised learning technique, but it predicts discrete outcomes (classes) instead of continuous values.

It assigns input data to one or more categories based on learned patterns. Typical examples include spam detection (spam vs. not spam), sentiment analysis (positive, neutral, negative), or predicting loan approval (approved/denied). The AI-900 study materials describe classification as "predicting a category or label for new observations."

* Clustering # An unsupervised machine learning model used to group similar entities based on features.

Clustering is an unsupervised learning approach, meaning it works on unlabeled data. It automatically identifies patterns and groups similar data points into clusters based on shared characteristics. Examples include customer segmentation (grouping customers by behavior) and grouping similar documents. The AI-900 learning module explains clustering as "discovering natural groupings in data without predefined labels." Thus, per Microsoft's official AI-900 learning objectives:

- * Regression # Predicts numeric/continuous values.
 - * Classification # Predicts categories/labels.
 - * Clustering # Groups similar entities (unsupervised).
- # Final Verified Configuration:
- * Regression # Predict numeric values
 - * Classification # Predict categories
 - * Clustering # Group similar entities based on features

NEW QUESTION: 121

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

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

Answer:

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

Explanation:

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

This question tests understanding of Microsoft's six guiding principles for Responsible AI, which are:

fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. These principles, as described in the Microsoft Azure AI Fundamentals (AI-900) study guide and Microsoft Learn Responsible AI module, help ensure that AI systems are developed and used ethically and responsibly.

* Transparency - Yes: Transparency means users should understand how and why an AI system makes certain decisions. Providing an explanation of the outcome of a credit loan application clearly supports transparency because it helps customers know the reasoning behind approval or rejection. According to Microsoft Learn, transparency ensures that "AI systems are understandable by users and stakeholders," especially in sensitive applications such as finance and credit scoring. Thus, the first statement is Yes.

* Reliability and Safety - Yes: The reliability and safety principle ensures AI systems perform consistently, safely, and as intended, even in complex or high-risk environments. A triage bot that prioritizes insurance claims based on injury type aligns with this principle-it must be accurate, dependable, and safe to ensure claims are processed correctly and not influenced by errors or faulty algorithms. Microsoft teaches that AI should be "reliable under expected and unexpected conditions" to prevent harm or misjudgment. Therefore, this statement is Yes.

* Inclusiveness - No: Inclusiveness focuses on ensuring AI systems empower and benefit all users, especially those with different abilities or backgrounds. Offering an AI solution at different prices across sales territories is a business decision, not an ethical or inclusiveness principle issue. It does not relate to accessibility or equal participation of diverse users. Therefore, this final statement is No.

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (336 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 122

Which Azure OpenAI model should you use to summarize the text from a document?

- A. Whisper
- B. DALL-E
- C. Codex
- D. GPT

Answer: (SHOW ANSWER)

According to the Microsoft Learn documentation and the Azure AI Fundamentals (AI-900) study guide, the GPT (Generative Pre-trained Transformer) family of models within Azure OpenAI Service is used for text-based natural language tasks, including summarization, content generation, and text completion.

When you need to summarize text from a document, GPT models (such as GPT-3.5 or GPT-4) can process large sections of text, extract the most relevant details, and generate concise summaries that retain the key meaning. The summarization task uses the model's natural language understanding capabilities to identify core concepts and generate human-like, coherent text.

Other options are incorrect:

- * A. Whisper # Used for speech-to-text transcription, not text summarization.
- * B. DALL-E # Generates images from text prompts, not text summaries.
- * C. Codex # Specializes in code generation and completion, not document summarization.

NEW QUESTION: 123

You have an Azure subscription that uses Azure OpenAI.

You need to create an original image of a rural scene to use on a website.

What should you do?

- A.** From Azure AI Foundry, deploy a GPT-3.5 Turbo model and provide instructions to create the image of the rural scene.
- B.** From Microsoft Bing, search the term " rural scene " and download the results.
- C.** From GitHub Copilot, provide instructions to create the image of the rural scene.
- D.** From Azure AI Foundry, deploy a DALL-E model, and provide instructions to create the rural scene.

Answer: ([SHOW ANSWER](#))

The Azure OpenAI DALL-E model is specifically designed for generating original images from natural language prompts. If you want to create an image of a "rural scene," you can use Azure AI Foundry (formerly Azure OpenAI Studio) to deploy the DALL-E model and provide descriptive instructions such as "create an image of a peaceful rural village with trees and a sunset."

- * A. GPT-3.5 Turbo # Handles text generation, not image creation.
- * B. Bing search # Finds existing images, not generate original ones.
- * C. GitHub Copilot # Assists with writing code, not generating images.

NEW QUESTION: 124

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

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

NOTE: Each correct match is worth one point.

Workload Types

Image classification

Object detection

Optical character recognition (OCR)

Answer Area

Generate captions for images.

Extract movie title names from movie poster images.

Locate vehicles in images.

Answer:

Workload Types

Image classification

Object detection

Optical character recognition (OCR)

Answer Area

Image classification

Generate captions for images.

Optical character recognition (OCR)

Extract movie title names from movie poster images.

Object detection

Locate vehicles in images.

Explanation:



In the Microsoft Azure AI Fundamentals (AI-900) curriculum, computer vision workloads are grouped into distinct types, each serving a specific purpose. The three major workloads illustrated here are image classification, object detection, and optical character recognition (OCR). Understanding their use cases is essential for correctly mapping them to real-world scenarios.

* **Generate captions for images # Image classification** The image classification workload is used to identify the main subject or context of an image and assign descriptive labels. In Microsoft Learn's "Describe features of computer vision workloads," image classification models are trained to recognize content (e.g., a cat, a beach, or a city). Caption generation expands on classification results by describing the image's contents in human-readable language-based on what the model identifies as key visual features.

* **Extract movie title names from movie poster images # Optical character recognition (OCR)** OCR is a vision workload that detects and extracts text from images. Azure AI Vision's Read API or Document Intelligence OCR models can identify printed or handwritten text within posters, signs, or documents.

In this case, the movie title text from a poster is best extracted using OCR.

* **Locate vehicles in images # Object detection** The object detection workload identifies multiple objects within an image and provides their locations using bounding boxes. It's ideal for tasks like counting cars in a parking lot or tracking objects in traffic images.

NEW QUESTION: 125

Which Azure Machine Learning capability should you use to quickly build and deploy a predictive model without extensive coding?

- A. ML pipelines
- B. Copilot
- C. DALL-E
- D. automated machine learning (automated ML)

Answer: (SHOW ANSWER)

According to the Microsoft Azure AI Fundamentals (AI-900) curriculum and Microsoft Learn's "Explore Automated Machine Learning in Azure Machine Learning" module, Automated ML (AutoML) is the Azure Machine Learning capability that allows users to quickly build, train, and deploy predictive models with minimal or no coding experience.

Automated ML automatically performs tasks that would normally require expert data science knowledge, such as:

- * Selecting appropriate algorithms (e.g., decision trees, logistic regression, random forests)
- * Performing hyperparameter tuning to optimize model accuracy
- * Handling missing data and feature scaling automatically
- * Generating performance metrics and best model recommendations

This feature is especially useful for business analysts, developers, or beginners who want to leverage machine learning for predictions (like sales forecasting, churn analysis, or demand prediction) without having to write complex Python code.

Other options explained:

- * A. ML pipelines automate and organize workflows for model training and deployment but still require pre-built models.
- * B. Copilot is an AI-powered assistant embedded in Microsoft tools for productivity, not a model training feature.
- * C. DALL-E is an image generation model under Azure OpenAI, not a predictive modeling tool.

Thus, per official Microsoft Learn content, Automated Machine Learning is the correct capability to quickly build and deploy predictive models with minimal coding.

NEW QUESTION: 126

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

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

Answer: ([SHOW ANSWER](#))

When interacting with the Azure OpenAI REST API, requests and responses are structured using the JSON format. According to the Microsoft Learn documentation on the Azure OpenAI REST API, JSON is the required format for sending prompts, configuration parameters (such as temperature, max_tokens, or top_p), and metadata to the API endpoint.

A typical request to an Azure OpenAI endpoint (for example, /openai/deployments/{model-name}/chat/completions?api-version=2024-02-01) includes a JSON body similar to the following:

```
{
  "messages": [
    {"role": "system", "content": "You are an assistant."},
    {"role": "user", "content": "Write a JavaScript function to add two numbers."}
  ],
  "max_tokens": 200,
  "temperature": 0.7
}
```

This JSON format allows for structured data exchange between the client and the server, ensuring that key- value pairs are properly parsed by the API.

* Option A (CSV): Used for storing tabular data, not suitable for API communication.

* Option C (XML): Though historically used in web services, Microsoft's modern APIs, including Azure OpenAI, rely on JSON.

* Option D (YAML): Common in configuration files but not for REST API payloads.

Therefore, per official Microsoft documentation and the AI-900 study guide, the correct and verified answer is B. JSON, as it is the standard format required when sending requests to the Azure OpenAI REST API endpoint.

NEW QUESTION: 127

You are building a Conversational Language Understanding model for an e-commerce business.

You need to ensure that the model detects when utterances are outside the intended scope of the model.

What should you do?

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

Answer: C ([LEAVE A REPLY](#))

In Conversational Language Understanding (CLU), a core service within Azure AI Language, intents represent the goals or purposes behind user utterances (for example, "Track my order" or "Cancel my subscription"). However, in real-world scenarios, users often provide inputs that do not match any defined intent. To handle such cases gracefully, Microsoft recommends including a "None" intent that captures out-of- scope utterances - text that doesn't belong to any other intent in your model.

According to the Microsoft Learn module: "Build a Conversational Language Understanding app", the None intent serves as a catch-all or fallback category for utterances that the model should ignore or respond to with a default message (e.g., "I'm sorry, I don't understand that."). By training the model with multiple examples of irrelevant or unrelated utterances in this intent, you improve its ability to distinguish between valid and invalid user inputs.

The other options are incorrect:

* A. Export the model: Exporting only saves or transfers the model; it does not influence how the model detects irrelevant utterances.

* B. Create a new model: A new model would not inherently solve out-of-scope detection unless properly trained with a None intent.

* D. Create a prebuilt task entity: Entities identify specific data (like dates or products) within valid intents, not irrelevant utterances.

Thus, the correct approach to ensure that your CLU model can detect utterances outside its intended scope is to add examples of unrelated or off-topic phrases to the None intent. This improves classification accuracy and prevents incorrect intent matches.

Correct answer: C. Add utterances to the None intent

NEW QUESTION: 128

Match the computer vision service to the appropriate AI workload.

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

NOTE: Each correct match is worth one point.

Services

Azure AI Custom Vision

Azure AI Document Intelligence

Azure AI Vision

Answer Area

Extract information from scanned forms and invoices.

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

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

Services

Azure AI Custom Vision

Azure AI Document Intelligence

Azure AI Vision

Answer Area

Azure AI Custom Vision

Extract information from scanned forms and invoices.

Azure AI Document Intelligence

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

Azure AI Vision

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

Explanation:

Answer Area

Azure AI Custom Vision

Extract information from scanned forms and invoices.

Azure AI Document Intelligence

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

Azure AI Vision

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

This question evaluates understanding of the different Azure AI Computer Vision services and their distinct functionalities, as covered in the Microsoft AI-900 study guide and Microsoft Learn modules under "Describe features of common AI workloads" and "Identify Azure services for computer vision."

- * Azure AI Document Intelligence (formerly known as Form Recognizer): This service is designed to extract structured information from documents, such as forms, receipts, and invoices. It uses optical character recognition (OCR) combined with AI models to detect key-value pairs, tables, and handwritten text. This makes it ideal for automating data entry and digitizing scanned documents. Hence, it matches "Extract information from scanned forms and invoices."
- * Azure AI Vision (formerly Computer Vision): This service provides image and video analysis capabilities. It can detect objects, people, text, and scenes; generate image captions; and extract descriptive tags. It also supports OCR for printed and handwritten text within images. Therefore, it matches "Analyze images and video, and extract descriptions, tags, objects, and text."
- * Azure AI Custom Vision: Custom Vision allows you to train your own image classification and object detection models using your own labeled images. Unlike the general Vision service, Custom Vision lets you build domain-specific models-for example, detecting your company's products or identifying manufacturing defects. Hence, it matches "Train custom image classification and object detection models by using your own images." These three services complement each other within Azure's computer vision ecosystem, collectively supporting both general-purpose and specialized AI solutions for visual data analysis.

NEW QUESTION: 129

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

Answer Area

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

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


Answer:

Answer Area

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

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

Explanation:



In the Microsoft Azure AI Fundamentals (AI-900) and Azure Machine Learning (AML) learning paths, deploying a real-time inference pipeline refers to making a trained machine learning model available as a web service that can process incoming data and return predictions instantly. To achieve this, the model must be deployed to an infrastructure capable of handling continuous, low-latency requests with high reliability and scalability.

Microsoft's official guidance from Azure Machine Learning documentation specifies that:

- * For testing or development, you can deploy to Azure Container Instances (ACI) because it provides a lightweight, temporary environment suitable for small-scale or non-production workloads.

- * For production-grade, real-time inference, the deployment should be made to Azure Kubernetes Service (AKS).

AKS provides enterprise-level scalability, load balancing, and high availability, which are critical for serving real-time predictions to multiple consumers simultaneously. It manages containerized applications using Kubernetes orchestration, allowing the model to scale automatically based on traffic demands.

Azure Machine Learning Compute is mainly used for model training and batch inference pipelines, not real-time endpoints. A local web service is typically used only for debugging or offline testing on a developer machine and cannot be shared for external consumption.

Therefore, when deploying a real-time inference pipeline as a service for others to consume, the correct and Microsoft-verified option is Azure Kubernetes Service (AKS). This environment ensures production readiness, secure endpoint management, and scalability for live AI applications, fully aligning with best practices outlined in the Azure Machine Learning designer documentation and AI-900 exam objectives.

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

NEW QUESTION: 130

What should you implement to prevent hateful responses from being returned by a generative AI solution?

- A. abuse monitoring
- B. content filtering
- C. fine-tuning
- D. prompt engineering

Answer: B (LEAVE A REPLY)

To prevent hateful or harmful responses from being returned by a generative AI solution, the correct approach is to implement content filtering. According to the Microsoft Learn documentation for Azure OpenAI Service and the Responsible AI principles, content filtering is a built-in safety mechanism that automatically screens both user prompts (inputs) and model outputs (responses) for inappropriate, harmful, or policy-violating material.

Content filters are designed to detect and block content such as:

- * Hate speech or harassment
- * Sexual or explicit material
- * Self-harm or violent content
- * Personally identifiable information (PII) misuse

In Azure OpenAI, the content filtering system is part of Microsoft's Responsible AI standard and cannot be disabled. It ensures that generative AI models such as GPT-3.5 or GPT-4 operate safely and ethically, reducing the risk of producing offensive or discriminatory text. The filter evaluates model responses in real time and can modify, block, or flag inappropriate outputs before they reach the user.

Let's review the other options:

- * A. Abuse monitoring tracks misuse after deployment but does not actively prevent hateful responses.
- * C. Fine-tuning customizes a model's style or domain knowledge but does not guarantee filtering of offensive content.
- * D. Prompt engineering helps steer model behavior but cannot fully prevent harmful outputs.

Therefore, to proactively prevent hateful, unsafe, or offensive responses in a generative AI system built on Azure OpenAI, the correct and Microsoft-verified approach is B. Content filtering.

NEW QUESTION: 131

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

Which service should you use?

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

Answer: B (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals official study guide and Microsoft Learn module

"Describe features of conversational AI workloads on Azure", Azure Bot Service is the core Azure platform for building, testing, deploying, and managing conversational agents or chatbots. These bots can communicate with users across multiple channels, including email, Microsoft Teams, Slack, Facebook Messenger, and webchat.

Azure Bot Service integrates deeply with the Bot Framework SDK and Azure Cognitive Services such as Language Understanding (LUIS) or Azure AI Language, enabling natural language processing and multi- channel message delivery. The service abstracts away channel management, meaning that developers can build one bot logic that connects seamlessly to several communication platforms.

Option analysis:

- * A. Text Analytics is a Cognitive Service used for text mining tasks like key phrase extraction, language detection, and sentiment analysis - not for building chatbots.
- * C. Translator provides language translation but cannot manage conversations or multi-channel delivery.
- * D. Form Recognizer extracts structured information from documents and forms - unrelated to conversational interaction.

The AI-900 course explicitly defines Azure Bot Service as "a managed platform that enables intelligent, multi- channel conversational experiences between users and bots." This service allows businesses to unify chat experiences across multiple digital communication channels.

Thus, based on the official Microsoft Learn content and AI-900 syllabus, the best and verified answer is B.

Azure Bot Service, as it is the designated Azure solution for deploying a single conversational AI experience accessible from multiple platforms such as email, Teams, and webchat.

NEW QUESTION: 132

What are two tasks that can be performed by using computer vision? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Answer: B,C (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals study guide and Microsoft Learn module "Identify features of computer vision workloads", computer vision is an AI workload that allows systems to interpret and understand visual information from the world, such as images and videos.

Computer vision tasks typically include:

- * Object detection and image classification (e.g., detecting brands, logos, or items in images)
- * Image analysis (e.g., identifying colors, patterns, or visual features)
- * Face detection and recognition
- * Optical Character Recognition (OCR) for reading text in images

Therefore, both detecting brands and detecting color schemes in an image are clear examples of computer vision tasks because they involve analyzing visual content.

In contrast:

- * A. Predict stock prices # Regression task, not vision-based.
- * D. Translate text between languages # Natural language processing (NLP).
- * E. Extract key phrases # NLP as well.

Thus, the correct computer vision tasks are B and C.

Reference:Microsoft Learn - Identify features and uses of computer vision in Azure AI services

NEW QUESTION: 133

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 (LEAVE A REPLY)

The correct answers are B. assigning the color pixels in an image to object names and C. describing the contents of an image.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Describe features of computer vision workloads on Azure," computer vision is a branch of AI that enables systems to analyze, interpret, and understand visual data from images and videos. It allows machines to identify objects, people, text, and even describe scenes automatically.

* Option B: Assigning color pixels in an image to object names represents image classification or object detection, which are key computer vision workloads. In these tasks, AI analyzes pixel patterns to determine which pixels correspond to specific objects (for example, classifying pixels as "car," "tree," or "road").

* Option C: Describing the contents of an image corresponds to image captioning, another computer vision workload. It involves using AI models trained to generate natural language descriptions of what is visible in an image, such as "A group of people sitting at a dining table." Azure's Computer Vision service provides this functionality through its "Describe Image" API.

Incorrect options:

- * A. Creating photorealistic images involves generative AI and 3D modeling, not traditional computer vision.
- * D. Detecting inconsistencies and anomalies in a data stream relates to anomaly detection, not computer vision.

* E. Creating visual representations of numerical data involves data visualization, not AI-driven image analysis.

NEW QUESTION: 134

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

A. classification

B. regression

C. clustering

Answer: C ([LEAVE A REPLY](#))

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe features of common AI workloads", clustering is a type of unsupervised machine learning used to group data points that share similar characteristics. In unsupervised learning, the data provided to the model does not have predefined labels or outcomes. Instead, the algorithm identifies inherent patterns or groupings within the dataset based on similarities in input features.

In this scenario, the task is to identify groups of people who have similar purchasing habits. There is no predefined label such as "buyer type" or "purchase category." The goal is to discover hidden patterns- such as grouping customers by spending behavior, preferred products, or frequency of purchases. This is precisely what clustering algorithms are designed to do.

Clustering is commonly used in:

* Customer segmentation for marketing analytics.

* Market basket analysis to find associations in purchasing patterns.

* Recommender systems to identify similar user profiles.

* Anomaly detection when outliers deviate from natural clusters.

Typical algorithms for clustering include K-means, Hierarchical clustering, and DBSCAN. These models analyze multidimensional data to form clusters that maximize intra-group similarity and minimize inter-group similarity.

By contrast:

* Classification (A) is a supervised learning method that predicts a categorical label (e.g., whether a customer will churn or not). It requires labeled training data.

* Regression (B) is used to predict continuous numeric values (e.g., sales revenue, temperature).

Since the question focuses on discovering groups of similar customers without prior labels, the correct type of machine learning is Clustering.

NEW QUESTION: 135

You have a webchat bot that provides responses from a QnA Maker knowledge base.

You need to ensure that the bot uses user feedback to improve the relevance of the responses over time.

What should you use?

A. key phrase extraction

B. sentiment analysis

C. business logic

D. active learning

Answer: D ([LEAVE A REPLY](#))

According to the Microsoft Azure AI Fundamentals (AI-900) study guide and the official Microsoft Learn module "Describe features of common AI workloads", QnA Maker (now part of Azure AI Language services) allows developers to build, train, and publish a knowledge base that provides natural-language answers to user queries. A key capability of this service is active learning, which enables the knowledge base to automatically suggest improvements by analyzing user feedback and usage patterns.

Active learning is an iterative process in which the service observes real user interactions and identifies ambiguous questions or pairs of similar questions that produce uncertain or multiple answers. The system then recommends updates or refinements to the knowledge base to improve the accuracy and relevance of responses. This feedback loop helps ensure that over time, the chatbot's responses align more closely with actual user expectations and language variations.

In contrast:

- * A. Key phrase extraction identifies main ideas in text and is used in content summarization, not in response optimization.
- * B. Sentiment analysis detects emotional tone (positive, negative, neutral), but it doesn't refine QnA responses.
- * C. Business logic defines operational rules in an application, not machine learning-driven feedback.

The AI-900 guide specifically emphasizes that QnA Maker supports active learning to improve the quality of answers based on end-user feedback, making this the verified and official Microsoft answer.

Reference (from Microsoft Learn AI-900 content):

"Active learning uses feedback from end users to automatically suggest improvements to a knowledge base, helping improve the accuracy of answers over time."

NEW QUESTION: 136

You need to convert receipts into transactions in a spreadsheet. The spreadsheet must include the date of the transaction, the merchant the total spent and any taxes paid.

Which Azure AI service should you use?

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

Answer: ([SHOW ANSWER](#))

To extract structured data such as transaction date, merchant name, total amount, and taxes from receipts, the best service is Azure AI Document Intelligence (formerly known as Form Recognizer). As described in the Microsoft Learn module: "Extract data from documents with Azure AI Document Intelligence", this service applies optical character recognition (OCR) combined with machine learning models to identify and extract key-value pairs and tabular data from semi-structured documents like invoices, receipts, and forms.

The prebuilt receipt model of Document Intelligence can automatically recognize common receipt fields, including:

- * Merchant Name
- * Transaction Date
- * Total Amount
- * Taxes
- * Items Purchased

It outputs structured JSON that can easily be converted into spreadsheet or database entries. This capability eliminates the need for manual data entry, ensuring accuracy and efficiency in digitizing financial documents.

The other options are incorrect:

- * A. Face detects and verifies human faces but does not extract text or numerical data.
- * B. Azure AI Language analyzes text sentiment, key phrases, and entities but does not interpret scanned documents.
- * D. Azure AI Custom Vision is for training image classification or object detection models, not document data extraction.

Therefore, the most accurate and Microsoft-verified service for converting receipts into structured transactions in a spreadsheet is C. Azure AI Document Intelligence.

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (**336** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 137

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

NOTE: Each correct selection is worth one point.

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

Answer: A,D (LEAVE A REPLY)

According to the Microsoft Learn module "Explore speech capabilities of Azure AI" and the AI-900 Official Study Guide, speech synthesis (also known as text-to-speech) is the process of converting written text into spoken audio output. Azure's Speech service provides this functionality, allowing applications to produce human-like voices dynamically.

Let's evaluate each scenario:

- * A. Automated voice that reads back a credit card number entered into a telephone keypad # YesThis is a classic text-to-speech (TTS) use case. The application converts numeric or textual input (such as a credit card number) into audio output that the caller hears. Azure Speech service can handle such voice responses in automated phone systems or IVR (Interactive Voice Response) setups.
- * B. Generating live captions for a news broadcast # NoThis is a speech-to-text scenario (speech recognition), not speech synthesis. It involves converting audio speech into written text.
- * C. Extracting key phrases from an audio recording of a meeting # NoThis involves speech-to-text followed by text analytics, not speech synthesis.
- * D. An AI character in a computer game that speaks audibly to a player # YesThis is a direct example of speech synthesis, where the character's dialog text is converted into realistic spoken output for immersive interaction.

Therefore, based on Microsoft's AI-900 curriculum, speech synthesis is used in applications that convert text into audible speech, such as automated voice systems or interactive digital characters.

NEW QUESTION: 138

Select the answer that correctly completes the sentence.

Answer Area

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

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



Answer selections

Answer:



Explanation:

Answer Area

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

a natural language processing (NLP) workload.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe features of natural language processing (NLP) workloads on Azure," Natural Language Processing refers to the branch of AI that enables computers to interpret, understand, and generate human language. One of the main NLP workloads identified by Microsoft is speech-to-text conversion, which transforms spoken words into written text. Creating a text transcript of a voice recording perfectly fits this definition because it involves converting audio language data into text form - a process handled by speech recognition models. These models analyze the acoustic features of human speech, segment phonemes, identify words, and produce a text transcript. On Azure, this function is implemented using the Azure Cognitive Services Speech-to-Text API, part of the Language and Speech services.

Let's examine the other options to clarify why they are incorrect:

- * Computer vision workload: Involves interpreting and analyzing visual data such as images and videos (e.g., object detection, facial recognition). It does not deal with speech or audio.
- * Knowledge mining workload: Refers to extracting useful information from large amounts of structured and unstructured data using services like Azure Cognitive Search, not transcribing audio.
- * Anomaly detection workload: Involves identifying unusual patterns in data (e.g., fraud detection or sensor anomalies), unrelated to language or speech.

In summary, when a system creates a text transcript from spoken audio, it is performing a speech recognition task-classified under Natural Language Processing (NLP). This workload helps make spoken content searchable, analyzable, and accessible, aligning with Microsoft's Responsible AI goal of enhancing accessibility through language understanding.

NEW QUESTION: 139

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

▼

Accuracy

Confidence

Root Mean Square Error

Sentiment

is the calculated probability of a correct image classification.



Answer:

▼

Accuracy

Confidence

Root Mean Square Error

Sentiment

is the calculated probability of a correct image classification.



Explanation:

Confidence.

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure," the confidence score represents the calculated probability that a model's prediction is correct. In image classification, when an AI model analyzes an image and assigns it to a specific category, it also produces a confidence value-a numerical probability (usually between 0 and 1) indicating how certain the model is about its prediction.

For example, if an image classification model identifies an image as a "cat" with a confidence of 0.92, it means the model is 92% certain that the image depicts a cat. The confidence value helps developers and users understand the model's certainty level about its classification output.

Microsoft Learn emphasizes that in Azure Cognitive Services-such as the Custom Vision Service-each prediction result includes both the predicted label (class) and a confidence score. These confidence scores are essential for evaluating model performance and determining thresholds for automated decisions (e.g., accepting predictions only above a 0.8 probability).

Let's evaluate the other options:

- * Accuracy: This is an overall performance metric measuring the percentage of correct predictions across the dataset, not a probability for a single prediction.
- * Root Mean Square Error (RMSE): This is a metric for regression models, not classification tasks. It measures average error magnitude between predicted and actual values.
- * Sentiment: This is a type of prediction (positive, negative, neutral) in text analysis, not a probability metric.

Therefore, based on Microsoft's AI-900 training materials and Azure Cognitive Services documentation, the calculated probability of a correct image classification is called Confidence, which expresses how sure the model is about its prediction for a specific input.

NEW QUESTION: 140

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

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

Answer: D (LEAVE A REPLY)

According to the Microsoft Azure AI Fundamentals (AI-900) learning content, Language Understanding Intelligent Service (LUIS) is part of Azure Cognitive Services used to interpret the meaning or intent behind a user's input in natural language. When a user says, "Call me back later," the system must recognize that the user intends for a call to be scheduled or delayed-this is not just about translating or analyzing text but understanding intent and relevant entities.

LUIS allows developers to train models to identify intents (such as ScheduleCall, CancelMeeting, etc.) and extract key entities (like names, times, or actions) from text inputs. It is typically integrated with conversational agents such as Azure Bot Service, enabling more natural, human-like interactions.

Other options do not fit the scenario:

- * Translator Text (A) translates text between languages but does not interpret meaning.
- * Text Analytics (B) performs sentiment analysis, key phrase extraction, and named entity recognition, but it doesn't identify intent.
- * Speech (C) converts spoken language to text or text to speech but doesn't interpret what the words mean.

Therefore, for understanding user intent such as "Call me back later," the correct AI service is D. Language Understanding (LUIS).

NEW QUESTION: 141

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

Answer Area

_____ is used to generate additional features

Feature engineering

Feature selection

Model evaluation

Model training



Answer:

Answer Area

_____ is used to generate additional features

Feature engineering


Feature selection

Model evaluation

Model training



Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Explore fundamental principles of machine learning", feature engineering is the process used to generate additional features or transform existing data into forms that improve model performance. Features are individual measurable properties or characteristics used as input for machine learning algorithms. The goal of feature engineering is to create new informative variables that better represent the underlying patterns in the data.

Feature engineering may include tasks such as:

- * Combining or transforming raw data columns (e.g., creating a "total purchase amount" from price × quantity).
- * Extracting time-based components (e.g., year, month, day, hour) from datetime values.
- * Encoding categorical variables (e.g., one-hot encoding or label encoding).
- * Scaling or normalizing numerical features.
- * Creating polynomial or interaction terms to capture complex relationships.

Microsoft's AI-900 learning material emphasizes that the process of preparing data for machine learning involves data cleaning, feature engineering, and feature selection. While feature selection is about choosing the most relevant features from the existing dataset, feature engineering focuses on creating or generating new features to enhance model accuracy and generalization.

The other options do not fit this definition:

- * Feature selection is about removing redundant or irrelevant features, not generating new ones.
- * Model evaluation involves assessing the model's performance using metrics like accuracy or F1 score.
- * Model training is the phase where the algorithm learns patterns from the data, not when features are created.

Therefore, based on the AI-900 official concepts and Microsoft's documentation, the correct answer is Feature engineering, as it is the process specifically used to generate additional features that improve machine learning model performance and predictive capability.

NEW QUESTION: 142

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Mass (kg) is a feature.	☐	☐
Quality Test is a label.	☐	☐
Temperature (C) is a label.	☐	☐



Answer:

Answer Area

Statements	Yes	No
Mass (kg) is a feature.	☒	☐
Quality Test is a label.	☒	☐
Temperature (C) is a label.	☐	☒



Explanation:

Answer Area

Statements	Yes	No
Mass (kg) is a feature.	☒	☐
Quality Test is a label.	☒	☐
Temperature (C) is a label.	☐	☒



This question tests the understanding of features and labels in machine learning, a core concept covered in the Microsoft Azure AI Fundamentals (AI-900) syllabus under "Describe fundamental principles of machine learning on Azure." In supervised machine learning, data is divided into features (inputs) and labels (outputs).

- * Features are the independent variables - measurable properties or characteristics used by the model to make predictions.
- * Labels are the dependent variables - the target outcome the model is trained to predict.

From the provided dataset, the goal of the Azure Machine Learning model is to predict product quality (Pass or Fail). Therefore:

- * Mass (kg) is a feature - Yes"Mass (kg)" represents an input variable used by the model to learn patterns that influence product quality. It helps the algorithm understand how variations in mass might correlate with passing or failing the quality test. Thus, it is correctly classified as a feature.
- * Quality Test is a label - YesThe "Quality Test" column indicates the outcome of the manufacturing process, marked as either Pass or Fail. This is the target the model tries to predict during training. In Azure ML terminology, this column is the label, as it represents the dependent variable.
- * Temperature (C) is a label - No"Temperature (C)" is an input that helps the model determine quality outcomes, not the outcome itself. It influences the quality result but is not the value being predicted. Therefore, temperature is another feature, not a label.

In conclusion, per Microsoft Learn and AI-900 study materials, features are measurable inputs (like mass and temperature), while the label is the target output (like the quality test result).

NEW QUESTION: 143

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

Answer Area		Statements	Yes	No
		You can use Language Service's question answering to query an Azure SQL database.	☐	☐
		You should use Language Service's question answering when you want a knowledge base to provide the same answer to different users who submit similar questions.	☐	☐
		Language Service's question answering can determine the intent of a user utterance.	☐	☐

Answer:

Answer Area		Statements	Yes	No
		You can use Language Service's question answering to query an Azure SQL database.	☐	☒
		You should use Language Service's question answering when you want a knowledge base to provide the same answer to different users who submit similar questions.	☒	☐
		Language Service's question answering can determine the intent of a user utterance.	☐	☒

Explanation:

Answer Area		Statements	Yes	No
		You can use Language Service's question answering to query an Azure SQL database.	☐	☒
		You should use Language Service's question answering when you want a knowledge base to provide the same answer to different users who submit similar questions.	☒	☐
		Language Service's question answering can determine the intent of a user utterance.	☐	☒

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of Natural Language Processing (NLP) workloads and services," the Azure Cognitive Service for Language - Question Answering capability is designed to allow applications to respond to user questions using information from a prebuilt or custom knowledge base. It relies on Natural Language Processing (NLP) to match user queries to the most relevant answers but does not directly execute queries against databases or infer user intent.

* "You can use Language Service's question answering to query an Azure SQL database." # NOThe Question Answering feature does not connect directly to or query structured databases such as Azure SQL. Instead, it retrieves answers from unstructured or semi-structured content (FAQs, manuals, documents). Querying SQL databases would require traditional database access, not a cognitive service.

* "You should use Language Service's question answering when you want a knowledge base to provide the same answer to different users who submit similar questions." # YESThis statement is correct and aligns exactly with Microsoft's official documentation. Question Answering enables organizations to create a knowledge base that can automatically answer repeated or similar customer queries using natural language understanding. For instance, two users asking "How do I reset my password?" and "Can you help me change my password?" would receive the same predefined response.

* "Language Service's question answering can determine the intent of a user utterance." # NODetermining user intent is handled by Language Understanding (LUIS) or Conversational Language Understanding, not by Question Answering. While both belong to the Language Service, Question Answering focuses on retrieving relevant answers, whereas LUIS focuses on intent detection and entity extraction.

NEW QUESTION: 144

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

Answer Area

Data values that influence the prediction of a model are called

dependant variables.

features.

identifiers.

labels.



Answer:

Answer Area

Data values that influence the prediction of a model are called

dependant variables.

features.

identifiers.

labels.



Explanation:

Features

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Explore fundamental principles of machine learning," data values that influence the prediction of a model are called features. In the context of machine learning, a feature is an individual measurable property, attribute, or input variable used by the model to make predictions.

Features are the independent variables that describe the characteristics of the data. For example, in a housing price prediction model, features might include square footage, location, number of bedrooms, and year built.

These inputs help the model understand relationships in the data so it can predict the target outcome (the house price).

Microsoft Learn explains that features are the input variables that the algorithm uses to identify patterns and relationships in the training data. During training, the model learns how changes in these features influence the label (also known as the dependent variable or target variable). The label is the value the model tries to predict-such as "price," "category," or "yes/no." Here's how the other options differ:

- * Dependent variables (labels): These are the outcomes or target values the model predicts, not the inputs.
- * Identifiers: These are unique keys (like customer ID or transaction ID) used to distinguish records but not to influence predictions.
- * Labels: As mentioned, labels are the results the model tries to predict.

Therefore, based on the AI-900 learning objectives and Microsoft's official explanation, the data values that influence the prediction of a model-that is, the input variables that guide the model's learning-are called features. These features form the foundation of the model's predictive capabilities and directly impact its accuracy and performance.

NEW QUESTION: 145

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
Azure Machine Learning designer provides a drag-and-drop visual canvas to build, test, and deploy machine learning models.	☐	☐
Azure Machine Learning designer enables you to save your progress as a pipeline draft.	☐	☐
Azure Machine Learning designer enables you to include custom JavaScript functions.	☐	☐

Answer:

Statements	Yes	No
Azure Machine Learning designer provides a drag-and-drop visual canvas to build, test, and deploy machine learning models.	☒	☐
Azure Machine Learning designer enables you to save your progress as a pipeline draft.	☒	☐
Azure Machine Learning designer enables you to include custom JavaScript functions.	☐	☒

Explanation:

Answer Area			
	Statements	Yes	No
	Azure Machine Learning designer provides a drag-and-drop visual canvas to build, test, and deploy machine learning models.	☒	☐
	Azure Machine Learning designer enables you to save your progress as a pipeline draft.	☒	☐
	Azure Machine Learning designer enables you to include custom JavaScript functions	☐	☒

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Describe core concepts of machine learning on Azure", the Azure Machine Learning Designer is a drag-and-drop, no-code/low-code interface that allows users to build, test, and deploy machine learning models visually without needing to write extensive code.

* Drag-and-Drop Visual Canvas # YES The Azure Machine Learning Designer indeed provides a graphical interface where users can connect prebuilt modules for data preprocessing, training, evaluation, and deployment. Microsoft documentation describes it as a "drag-and-drop visual environment that simplifies machine learning model creation." This allows beginners and business users to construct machine learning pipelines intuitively, confirming this statement as True.

* Save Progress as a Pipeline Draft # YES The designer lets users save their current work as a pipeline draft, enabling them to pause and return later. Microsoft Learn explicitly states that you can "save and publish pipeline drafts before running or deploying them." This functionality ensures workflow continuity, collaboration, and version management-making this statement also True.

* Include Custom JavaScript Functions # NO The Azure Machine Learning Designer allows the integration of Python scripts through the "Execute Python Script" module for custom logic, but it does not support JavaScript. Custom code in the designer environment is limited to Python, as the platform is built for data science and machine learning tasks typically handled in Python-based environments. Therefore, this statement is False.

NEW QUESTION: 146

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 ([**LEAVE A REPLY**](#)**)**

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn module "Identify features of common AI workloads," the Azure Form Recognizer service is part of Azure Cognitive Services for Document Intelligence. It enables organizations to extract, analyze, and identify information from structured and unstructured documents, including sensitive or confidential data such as names, addresses, financial figures, and identification numbers.

Form Recognizer uses optical character recognition (OCR) combined with machine learning to automatically extract key-value pairs, tables, and text fields from documents like invoices, receipts, contracts, and forms. It can be customized to identify and classify documents that contain specific sensitive data, allowing businesses to automate compliance and data governance tasks.

By contrast:

* A. Custom Vision is used for image classification and object detection - it analyzes visual data, not document content.

* B. Conversational Language Understanding (formerly LUIS) identifies intent and entities in text conversations, not document structure or sensitive data.

Form Recognizer is explicitly mentioned in the AI-900 course as the tool for document analysis and extraction. It can even integrate with Azure Cognitive Search or Azure Purview for further data management and compliance workflows.

Therefore, the verified and correct answer, aligned with Microsoft's official training content, is C. Form Recognizer, as it is the Azure Cognitive Service capable of identifying and processing documents containing sensitive information.

Valid AI-900 Dumps shared by Actual4test.com for Helping Passing AI-900 Exam! Actual4test.com now offer the **newest AI-900 exam dumps**, the Actual4test.com AI-900 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-900 dumps with Test Engine here: https://www.actual4test.com/AI-900_examcollection.html (**336** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)