

Microsoft.AI-900J.v2026-08-28.q153

試験コード :	AI-900J
試験名称 :	Microsoft Azure AI Fundamentals (AI-900日本語版)
認証ベンダー :	Microsoft
無料問題の数 :	153
バージョン :	v2026-08-28
ページの閲覧量 :	108
問題集の閲覧量 :	2041
https://www.jpnsiken.com/shiken/Microsoft.AI-900J.v2026-08-28.q153.html	

質問: 1

モデルをトレーニングするときに、行をランダムに別々のサブセットに分割する必要があるのはなぜですか？

- A. モデルを2回トレーニングして、精度を向上させる
- B. 複数のモデルを同時にトレーニングしてパフォーマンスを向上させる
- C. モデルのトレーニングに使用されなかったデータを使用してモデルをテストする

正解: ([正解を表示します](#))

When training a machine learning model, it is standard practice to randomly split the dataset into training and testing subsets. The purpose of this is to evaluate how well the model generalizes to unseen data. According to the AI-900 study guide and Microsoft Learn module "Split data for training and evaluation", this ensures that the model is trained on one portion of the data (training set) and evaluated on another (test or validation set).

The correct answer is C. to test the model by using data that was not used to train the model.

Random splitting prevents data leakage and overfitting, which occur when a model memorizes patterns from the training data instead of learning generalizable relationships. By testing on unseen data, developers can assess true performance, ensuring that predictions will be accurate on future, real-world data.

Options A and B are incorrect because:

- * A. Train the model twice does not improve accuracy; model accuracy depends on data quality, feature engineering, and algorithm choice.
- * B. Train multiple models simultaneously refers to model comparison, not the purpose of splitting data.

Thus, the correct reasoning is that random splitting provides a reliable estimate of the model's predictive power on new data.

Reference:Microsoft Learn - Split data for training and evaluation in Azure Machine Learning designer

質問: 2

文を正しく完成させる答えを選択してください。



正解:



Explanation:

Azure Kubernetes Service (AKS).

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn documentation on Azure Machine Learning, the Azure Kubernetes Service is commonly used to host and deploy machine learning models, including Automated ML models, into production environments. Once a model is trained using Azure Machine Learning (Azure ML), it must be deployed as a web service endpoint so it can receive data and return predictions.

Azure ML offers two primary options for hosting and deploying models:

- * Azure Kubernetes Service (AKS) - for high-scale, production-grade deployments that require fast response times, high availability, and scalability.
- * Azure Container Instances (ACI) - for testing or low-scale workloads where cost and simplicity are more important than performance.

AKS provides a managed Kubernetes cluster that allows for automated container orchestration, load balancing, scaling, and monitoring of deployed machine learning models. When you use Automated ML in Azure ML Studio, the generated model can be containerized and deployed directly to AKS, making it accessible as a REST API endpoint. This enables applications, systems, or users to send data and receive predictions in real time.

The other options serve different purposes:

- * Azure Data Factory is used for data integration and pipeline orchestration, not model hosting.
- * Azure Automation focuses on automating administrative tasks and runbooks, not ML deployment.
- * Azure Logic Apps is used to automate workflows and integrate services, not to serve ML models.

Therefore, the correct service to host automated machine learning (AutoML) models in production is Azure Kubernetes Service (AKS), as it provides a reliable, scalable, and secure environment for real-time inference and enterprise AI workloads.

質問: 3

音声合成ソリューションを使用できるシナリオは2つありますか？それぞれの正解は完全な解決策を提示します。

注：正しい選択はそれぞれ1ポイントの価値があります。

- A. テンキーを使用して電話に入力されたクレジットカード番号を読み戻す自動音声
- B. ニュース放送のライブキャプションを生成する
- C. 会議の録音からキーフレーズを抽出する
- D. プレイヤーに音声で話すコンピューターゲームのAIキャラクター

正解: ([正解を表示します](#))

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.

質問: 4

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area

Microsoft

Statements

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

正解:

Answer Area

Microsoft

Statements

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

Explanation:

Answer Area

Microsoft

Statements

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

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn documentation for Azure AI Custom Vision, this service is a specialized part of the Azure AI Vision family that enables developers to train custom image classification and object detection models. It allows organizations to build tailored computer vision models that recognize images or specific objects relevant to their business needs.

* Detect objects in an image # YesThe Azure AI Custom Vision service supports both image classification (assigning an image to one or more categories) and object detection (identifying and locating objects within an image using bounding boxes). This means it can indeed detect and differentiate multiple objects in a single image, making this statement true.

* Requires your own data to train the model # YesThe Custom Vision service is designed to be customizable. Unlike prebuilt Azure AI Vision models that work out of the box, Custom Vision requires you to upload and label your own dataset for training. The model then learns from your examples to perform specialized image recognition tasks relevant to your domain. Thus, this statement is also true.

* Analyze video files # NoWhile Custom Vision can analyze images, it does not directly process or analyze video files. Video analysis is handled by a different service-Azure Video Indexer-which can extract insights such as spoken words, scenes, and faces from videos.

In summary:

Yes - Detect objects in images
Yes - Requires your own data
No - Does not analyze video files.

質問: 5

コードを分析し、コード関数とコードコメントの説明を生成するために使用できる 2 つのリソースはどれですか？それぞれの正解は完全なソリューションを示します。

注意: 正解ごとに 1 ポイントが付与されます。

- A. Azure OpenAI DALL-E モデル
- B. Azure OpenAI Whisper モデル
- C. Azure OpenAI GPT-4 モデル
- D. GitHub Copilot サービス

正解: ([正解を表示します](#))

The correct answers are C. the Azure OpenAI GPT-4 model and D. the GitHub Copilot service.

According to the Microsoft Azure AI Fundamentals (AI-900) curriculum and Microsoft Learn documentation on Azure OpenAI and GitHub Copilot, both GPT-4 and GitHub Copilot can be used to analyze and generate explanations for code functionality, as well as produce or refine code comments.

* Azure OpenAI GPT-4 model (C):The GPT-4 model is a large language model (LLM) developed by OpenAI and available through the Azure OpenAI Service. It is trained on vast amounts of text, including programming languages, documentation, and natural language instructions. This enables it to interpret source code, explain what it does, suggest optimizations, and automatically generate detailed code comments. When prompted with code snippets, GPT-4 can provide structured natural language explanations describing the logic and intent of the code. In enterprise scenarios, developers use Azure OpenAI GPT models for code understanding, review automation, and documentation generation.

* GitHub Copilot service (D):GitHub Copilot, powered by OpenAI Codex, is an AI coding assistant integrated into IDEs such as Visual Studio Code. It can analyze code context and generate inline comments and explanations in real time. GitHub Copilot understands the syntax and intent of numerous programming languages and provides intelligent suggestions or explanations directly in the developer's environment.

The other options are not suitable:

* A. DALL-E is a generative image model for creating visual content, not text or code analysis.

* B. Whisper is an automatic speech recognition (ASR) model used for converting speech to text, unrelated to code interpretation.

Therefore, based on the official Azure AI and GitHub documentation, the correct and verified answers are C.

Azure OpenAI GPT-4 model and D. GitHub Copilot service.

質問: 6

GitHub Copilot は、クライアント側 JavaScript の提案を行うためにどの OpenAI モデルを使用しますか？

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

正解: B ([コメントを発表する](#))

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.

質問: 7

Azure Machine Learning Studio と自動機械学習 (自動 ML) を使用してモデルを構築し、トレーニングする予定です。最初に何を作成する必要がありますか？

A. Jupyter ノートブック

B. 機械学習ワークスペース

C. 登録されたデータセット

D. 機械学習デザイナー パイプライン

正解: ([正解を表示します](#))

Before building or training any model in Azure Machine Learning Studio-including when using Automated ML (AutoML)-you must first create a Machine Learning workspace.

A workspace serves as the central environment for all machine learning assets such as datasets, compute targets, models, pipelines, and experiments. According to the AI-900 study guide and Microsoft Learn module

"Describe features and tools for machine learning in Azure," a workspace is the foundational setup required to organize and manage all ML-related resources.

The sequence typically follows these steps:

* Create a Machine Learning workspace.

* Configure compute resources (e.g., compute instance or cluster).

* Upload or register datasets.

* Use Automated ML or Designer to train models.

* Deploy and manage the trained models.

Option A (Jupyter notebook) is an optional tool for coding experiments.

Option C (Registered dataset) is created after the workspace exists.

Option D (Designer pipeline) is a visual tool used within the workspace.

Hence, B. a Machine Learning workspace is the correct answer because it is the first and mandatory step before using Automated ML or any training component in Azure Machine Learning Studio.

質問: 8

文を正しく完成させる答えを選択してください。



正解:

Answer Area

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

Computer Vision
Conversational Language Understanding
Custom Vision
Form Recognizer

Explanation:

Answer Area

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

The correct completion of the sentence is:

"The Form Recognizer service can be used to extract information from a driver's license to populate a database." According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of computer vision workloads," Azure Form Recognizer (part of Azure AI Document Intelligence) is a document processing service that uses machine learning and optical character recognition (OCR) to extract structured data, key-value pairs, and text from documents such as invoices, receipts, identity cards, and driver's licenses.

This service allows businesses to automate data entry and document processing workflows by converting physical or scanned documents into machine-readable formats. For example, with a driver's license, Form Recognizer can extract structured data fields such as Name, Date of Birth, License Number, and Expiration Date, and automatically populate those values into a database or CRM system. The AI-900 study materials emphasize that Form Recognizer is designed to handle both structured and unstructured document layouts. It includes prebuilt models for common document types (like invoices, receipts, and identity documents) and supports custom models for domain-specific forms.

By comparison:

- * Computer Vision extracts general text or image content but doesn't structure or label extracted fields.
- * Custom Vision is used for training image classification or object detection models.
- * Conversational Language Understanding is for processing text or speech to determine intent, not extracting document data.

Therefore, based on the Microsoft Learn AI-900 official study content, the Form Recognizer service is the correct choice, as it is explicitly designed to extract and structure data from documents like driver's licenses, forms, and receipts - making it ideal for automatically populating a database.

質問: 9

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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.	☐	☐

正解:

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).

質問: 10

文を正しく完成させる答えを選択してください。

Answer Area

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

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

正解:

Answer Area

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

AUC metric
confusion matrix
ROC curve
threshold



Explanation:

Answer Area

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



"When evaluating the performance of a model, the confusion matrix displays the predicted and actual positives and negatives by using a grid of 0 and 1 values." According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn module "Identify features of common machine learning types", a confusion matrix is a tool used to evaluate the performance of classification models. It visually summarizes how many predictions were correctly or incorrectly classified by comparing the predicted labels to the actual (true) labels.

A confusion matrix is a table, typically 2×2 for binary classification, with the following components:

- * True Positives (TP): The model correctly predicted the positive class.
- * True Negatives (TN): The model correctly predicted the negative class.
- * False Positives (FP): The model incorrectly predicted the positive class.
- * False Negatives (FN): The model incorrectly predicted the negative class.

The confusion matrix allows data scientists and analysts to derive important performance metrics such as accuracy, precision, recall, and F1-score, which together provide a more complete understanding of how well a model performs beyond a single number.

In Microsoft Learn's AI-900 curriculum, the confusion matrix is highlighted as a key visualization tool that

"compares actual values to predicted values to evaluate classification performance." The grid format (using 0s and 1s for predicted classes) helps identify where misclassifications occur.

By contrast:

- * AUC metric (Area Under Curve) and ROC curve evaluate model discrimination ability.
- * Threshold defines decision cutoffs but doesn't display classifications.

Therefore, based on the official Microsoft AI-900 study guide and Microsoft Learn resources, the correct answer is Confusion Matrix, as it provides a grid view comparing actual versus predicted values in classification models.

質問: 11

従業員とその写真のリストを含むデータベースがあります。

従業員の新しい写真にタグを付けています。

次の各ステートメントについて、ステートメントが真の場合は[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area

Statements	Yes	No
The Face service can be used to group all the employees who have similar facial characteristics.	☐	☐
The Face service will be more accurate if you provide more sample photos of each employee from different angles.	☐	☐
If an employee is wearing sunglasses, the Face service will always fail to recognize the employee.	☐	☐

正解:

Answer Area

Statements	Yes	No
The Face service can be used to group all the employees who have similar facial characteristics.	☒	☐
The Face service will be more accurate if you provide more sample photos of each employee from different angles.	☒	☐
If an employee is wearing sunglasses, the Face service will always fail to recognize the employee.	☐	☒

Explanation:

Answer Area

Statements

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

☒☐

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

☒☐

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

☐☒

These answers are derived from the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure." The Azure Face service, part of Azure Cognitive Services, provides advanced facial recognition capabilities including detection, verification, identification, grouping, and similarity analysis.

Let's analyze each statement:

* "The Face service can be used to group all the employees who have similar facial characteristics." # Yes
The Face service supports a grouping function that automatically organizes a collection of unknown faces into groups based on visual similarity. It doesn't require labeled data; instead, it identifies clusters of similar-looking faces. This is particularly useful when building or validating datasets of people.

* "The Face service will be more accurate if you provide more sample photos of each employee from different angles." # Yes
According to Microsoft documentation, model accuracy improves when you provide multiple high-quality images of each person under different conditions-such as varying lighting, poses, and angles. This diversity allows the service to better learn unique facial characteristics and improves recognition reliability, especially for identification and verification tasks.

* "If an employee is wearing sunglasses, the Face service will always fail to recognize the employee." # No
This is incorrect. While occlusions (like sunglasses or hats) can reduce accuracy, the service may still recognize the person depending on how much of the face remains visible. Microsoft Learn explicitly notes that partial occlusion affects recognition confidence but does not guarantee failure.

In conclusion, the Face service can group similar faces (Yes), become more accurate with diverse samples (Yes), and still recognize partially covered faces though with lower confidence (No). These principles align directly with the Face API's core functions and AI-900 learning objectives regarding computer vision and responsible AI-based facial recognition.

質問: 12

教師なし機械学習の例は何ですか？

A. 分類

B. クラスタリング

C. 回帰

正解: ([正解を表示します](#))

In unsupervised machine learning, the algorithm learns patterns or structure within data without pre-labeled outputs or target values. The primary goal is to discover hidden relationships or group similar data points automatically. The Microsoft Azure AI Fundamentals (AI-900) study materials identify clustering as the key example of unsupervised learning.

In clustering, algorithms such as K-means, hierarchical clustering, or DBSCAN group data based on feature similarity. For example, a business may cluster customers by purchase behavior to discover natural customer segments without prior category labels. The model finds inherent patterns within the data rather than being told what to predict.

By contrast, classification and regression are supervised learning techniques. In supervised learning, the algorithm is trained using labeled data where correct outputs are already known. Therefore, the correct answer is B. Clustering, as it best represents unsupervised learning in Azure AI-900 principles.

質問: 13

視力が低下しているユーザーをサポートするために、レシピの説明を声に出して読むアプリを作成する必要があります。

どのバージョンのサービスを使用する必要がありますか？

A. テキスト分析

B. 翻訳者のテキスト

C. スピーチ

D. 言語理解 (UIS)

正解: ([正解を表示します](#))

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)

質問: 14

Azure AI Document Intelligence の事前構築済み領収書モデルを使用できる Azure サービスはどれですか？

A. Azure AI コンピューター ビジョン

B. Azure 機械学習

C. Azure AI サービス

D. アズール アル カスタム ビジョン

正解: ([正解を表示します](#))

The prebuilt receipt model is part of Azure AI Document Intelligence (formerly Form Recognizer), which belongs to the broader Azure AI Services family. The prebuilt receipt model is designed to automatically extract key information such as merchant names, dates, totals, and tax amounts from receipts without requiring custom training.

Among the given options, C. Azure AI Services is correct because it encompasses all cognitive AI capabilities-vision, language, speech, and document processing. Specifically, Azure AI Document Intelligence is included within Azure AI Services and provides both prebuilt and custom models for processing invoices, receipts, business cards, and identity documents.

Options A (Computer Vision) and D (Custom Vision) are image-based services, not form-processing tools.

Option B (Azure Machine Learning) focuses on building custom predictive models, not using prebuilt document models.

Therefore, the correct answer is C. Azure AI Services, which includes the prebuilt receipt model in Document Intelligence.

質問: 15

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Statements	Yes	No
Chatbots can support voice input.	☐	☐
A separate chatbot is required for each communication channel.	☐	☐
Chatbots manage conversation flows by using a combination of natural language and constrained option responses.	☐	☐

正解:

Statements	Yes	No
Chatbots can support voice input.	☐	☒
A separate chatbot is required for each communication channel.	☐	☒
Chatbots manage conversation flows by using a combination of natural language and constrained option responses.	☒	☐

Explanation:

Statements	Yes	No
Chatbots can support voice input.	☒	☐
A separate chatbot is required for each communication channel.	☐	☒
Chatbots manage conversation flows by using a combination of natural language and constrained option responses.	☒	☐

This question is derived from the Microsoft Azure AI Fundamentals (AI-900) learning module, particularly under "Describe features of conversational AI workloads on Azure." It tests understanding of chatbot capabilities and design principles within the context of Azure Bot Service and Conversational AI.

* Chatbots can support voice input - YesAccording to the AI-900 official materials, conversational AI systems such as chatbots can interact with users through text or voice. Using speech recognition services like Azure Cognitive Services Speech-to-Text, bots can interpret spoken input, and with Text- to-Speech, they can respond verbally. This enables voice-based chatbots used in virtual assistants, call centers, and customer support. Hence, voice input is fully supported by conversational AI solutions in Azure.

* A separate chatbot is required for each communication channel - NoThe Azure Bot Service is designed to provide multi-channel communication from a single bot instance. A single chatbot can communicate across several channels such as Microsoft Teams, Web Chat, Slack, Facebook Messenger, and email without needing separate bots for each platform. This centralized design allows developers to create, deploy, and manage one bot while configuring multiple channel connections through the Azure portal.

Therefore, the statement is false.

* Chatbots manage conversation flows by using a combination of natural language and constrained option responses - YesIn Microsoft's AI-900 training, chatbots are described as using Natural Language Processing (NLP) to understand free-form user input while also guiding interactions with predefined options such as buttons or quick replies. This hybrid approach ensures both flexibility and control, improving user experience and accuracy. Bots can interpret natural language via services like Language Understanding (LUIS) and also present structured options to guide conversations efficiently.

質問: 16

マイクロソフトの責任あるAI原則の例として適切な記述はどれですか？

- A. AIシステムは公開されているデータのみを使用しなければなりません。
- B. 人工知能システムは企業の利益を保護しなければならない
- C. すべてのシステムは理解可能でなければならない。
- D. すべてのシステムは個人情報を公開しなければならない

正解: C ([コメントを发表する](#))

The correct answer is C. AI systems must be understandable, which corresponds to the Transparency principle of Microsoft's Responsible AI framework.

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Identify guiding principles for responsible AI", Microsoft defines six key principles for responsible AI:

- * Fairness - AI systems should treat everyone equitably.
- * Reliability and safety - AI should function as intended, even under unexpected conditions.
- * Privacy and security - AI must protect personal and business data.
- * Inclusiveness - AI should empower everyone and engage diverse users.
- * Transparency - AI systems should be understandable.
- * Accountability - People should be accountable for AI systems.

The statement "AI systems must be understandable" aligns directly with the Transparency principle, ensuring that AI decisions and behaviors can be explained and interpreted by developers, users, and stakeholders.

Microsoft emphasizes that transparent AI builds trust, allows debugging, and ensures ethical usage.

Other options are incorrect:

- * A. Use only publicly available data - Not a principle of Responsible AI.
- * B. Protect the interests of the company - Focused on business goals, not ethical AI.
- * D. Keep personal details public - Violates the Privacy and Security principle.

Final Answer (Q179): C. AI systems must be understandable.

有効的な**AI-900J**問題集はJPNTTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AI-900J**試験問題集を提供します。JPNTTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **336**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」

質問: 17

Text AnalyticsAPI機能をテクニカルサポートチケットシステムに適用することを計画しています。
Text AnalyticsAPI機能を適切な自然言語処理シナリオに一致させます。
答えるには、適切な機能を左側の列から右側のシナリオにドラッグします。各機能は、1回使用することも、複数回使用することも、まったく使用しないこともできます。
注：正しい選択はそれぞれ1ポイントの価値があります。

API Features

Entity recognition

Key phrase extraction

Language detection

Sentiment analysis

Answer Area

API Feature

API Feature

API Feature

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

Summarize important information from the support ticket.

Extract key dates from the support ticket.

正解:

API Features

Entity recognition

Key phrase extraction

Language detection

Sentiment analysis

Answer Area

Sentiment analysis

Key phrase extraction

Entity recognition

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

Summarize important information from the support ticket.

Extract key dates from the support ticket.

Explanation:

Answer Area

Sentiment analysis

Key phrase extraction

Entity recognition

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

Summarize important information from the support ticket.

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>

質問: 18

顧客サポートシステム向けのWebベースのAIソリューションを開発する必要があります。ユーザーは、最適なリソースや回答へと誘導するWebアプリを操作できなければなりません。

目標を達成するには、どのサービスを Web アプリと統合する必要がありますか？

A. Azure AI 言語サービス

B. 顔

C. Azure AI Translator

D. Azure AI カスタムビジョン

正解: D ([コメントを发表する](#))

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

Reference:

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

質問: 19

文を完成させるには、回答領域で適切なオプションを選択します。

Answer Area

Data values that influence the prediction of a model are called

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

正解:

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.

質問: 20

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。
注：正しい選択はそれぞれ1ポイントの価値があります。

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.	☐	☐

正解:

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.

質問: 21

スキャンしたドキュメントからテキスト、キーと値のペア、テーブルデータを自動的に抽出するには、どのサービスを使用すればよいですか？

- A. Azure AI カスタムビジョン
- B. Azure AI ドキュメント インテリジェンス
- C. Azure AI 言語
- D. アズールアルフェイス

正解: ([正解を表示します](#))

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/>

質問: 22

文を正しく完成させる答えを選択してください。



正解:



Explanation:



The question describes a process where an AI system generates text that describes an image - for example,

"A dog playing with a ball in the park." This process is an example of image classification, a core workload in computer vision that allows a system to recognize and categorize the content of an image.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify Azure services for computer vision," image classification involves analyzing the pixels of an image and assigning one or more predefined categories or labels to it. In more advanced implementations, image classification models are combined with caption generation algorithms to produce descriptive text. For example, Azure AI Vision can generate captions and tags that describe an image's content, such as "outdoor scene," "a person riding a bicycle," or "a group of people smiling." Let's review the other options to clarify why they are incorrect:

* Facial detection: Identifies the presence and location of human faces in an image, but does not generate descriptive text.

* Object detection: Identifies and locates multiple objects within an image by drawing bounding boxes, not by describing the overall scene.

* Optical character recognition (OCR): Extracts text from images or scanned documents (for example, reading a street sign), but it doesn't create descriptive language about what's depicted.

Therefore, the correct answer is Image classification, as it aligns with the AI-900 learning objective that describes this task as recognizing and categorizing the main content of an image, often leading to caption generation in modern vision models such as those in Azure AI Vision.

質問: 23

Azure OpenAI GPT-3.5 モデルを使用するチャット ソリューションからより詳細な応答を生成するには、どのパラメーターを構成する必要がありますか？

A. プレゼンスペナルティ

B. 温度

C. シーケンスを停止

D. 最大レスポンスB

正解: B ([コメントを发表する](#))

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).

質問: 24

Azureで自然言語処理ソリューションを開発しています。このソリューションは、顧客のレビューを分析し、各レビューがどれほど肯定的か否定的かを判断します。

これは、どのタイプの自然言語処理ワークロードの例ですか？

A. 言語検出

B. 感情分析

C. キーフレーズ抽出

D. エンティティの認識

正解: **B** ([コメントを発表する](#))

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.

質問: 25

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area			
Statements		Yes	No
Monitoring online service reviews for profanities is an example of natural language processing.		☐	☐
Identifying brand logos in an image is an example of natural languages processing.		☐	☐
Monitoring public news sites for negative mentions of a product is an example of natural language processing.		☐	☐

正解:

Answer Area			
Statements		Yes	No
Monitoring online service reviews for profanities is an example of natural language processing.		☒	☐
Identifying brand logos in an image is an example of natural languages processing.		☐	☒
Monitoring public news sites for negative mentions of a product is an example of natural language processing.		☒	☐

Explanation:

Answer Area

Statements	Yes	No
Monitoring online service reviews for profanities is an example of natural language processing.	☒	☐
Identifying brand logos in an image is an example of natural languages processing.	☐	☒
Monitoring public news sites for negative mentions of a product is an example of natural language processing.	☒	☐

This question examines your understanding of Natural Language Processing (NLP) as described in the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Explore natural language processing." NLP is a branch of artificial intelligence that enables computers to analyze, understand, and generate human language - both written and spoken. Typical NLP tasks include text analytics, language understanding, sentiment analysis, key phrase extraction, and profanity detection.

- * Monitoring online service reviews for profanities # Yes This is a classic example of NLP. Detecting profane or inappropriate words in customer reviews requires analyzing text content. Azure Cognitive Services offers Content Moderator and Text Analytics APIs that can detect and filter profanity, sentiment, and offensive language automatically. Microsoft Learn states: "Natural language processing is used to process and analyze text to detect sentiment, key phrases, and inappropriate content." Hence, this task is correctly classified as NLP.
- * Identifying brand logos in an image # No This task belongs to Computer Vision, not NLP. The Computer Vision API and Custom Vision service in Azure are designed to detect and classify visual elements like logos, objects, or scenes. Since it involves images, not text, it is unrelated to natural language processing.
- * Monitoring public news sites for negative mentions of a product # Yes This is another valid example of NLP. The process involves analyzing the sentiment of text from online articles to determine whether mentions of a product are positive, neutral, or negative. Azure Text Analytics provides prebuilt sentiment analysis and entity recognition capabilities that help automate such monitoring.

質問: 26

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area

Statements	Yes	No
A transformer model architecture uses self-attention.	☐	☐
A transformer model architecture includes an encoder block and a decoder block.	☐	☐
A transformer model architecture includes an encryption block or a decryption block.	☐	☐

正解:

Answer Area

Statements	Yes	No
A transformer model architecture uses self-attention.	☐	☒
A transformer model architecture includes an encoder block and a decoder block.	☒	☐
A transformer model architecture includes an encryption block or a decryption block.	☐	☒

Explanation:

Answer Area

Microsoft

Statements

A transformer model architecture uses self-attention.	Yes ☒	No ☐
A transformer model architecture includes an encoder block and a decoder block.	Yes ☒	No ☐
A transformer model architecture includes an encryption block or a decryption block.	☐	No ☒

The Transformer model architecture is a foundational deep learning model introduced in the 2017 research paper "Attention Is All You Need." It serves as the core architecture for modern large language models such as GPT, BERT, and T5, all of which are used in Azure OpenAI Service.

* "A transformer model architecture uses self-attention." - Yes The self-attention mechanism is the defining feature of transformer models. It allows the model to evaluate the relationships between words (tokens) in a sequence and assign weights based on contextual relevance. This means that each word in an input sentence can "attend" to every other word, capturing dependencies regardless of their position in the text. This mechanism replaced older recurrent (RNN) and convolutional (CNN) architectures for sequence processing because it provides parallelization and better context understanding.

* "A transformer model architecture includes an encoder block and a decoder block." - Yes The original Transformer architecture includes both an encoder and a decoder. The encoder processes the input sequence into contextual representations, and the decoder generates the output sequence based on both the encoder's output and previously generated tokens. Models like BERT use only the encoder stack, while GPT models use only the decoder stack, but the full Transformer design conceptually includes both.

* "A transformer model architecture includes an encryption block or a decryption block." - No Transformers are not related to cryptography. They perform encoding and decoding of language data for representation learning-not encryption or decryption for data security. The terms "encoder" and "decoder" here refer to neural network components, not cryptographic processes.

質問: 27

文を正しく完成させる答えを選択してください。

Answer Area

Natural language processing can be used to

Microsoft

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

正解:

Answer Area

Microsoft

Natural language processing can be used to

Microsoft

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

Explanation:

Answer Area



Microsoft

Natural language processing can be used to

classify email messages as work-related or personal.

This question is drawn from the Microsoft Azure AI Fundamentals (AI-900) syllabus section "Describe features of natural language processing (NLP) workloads on Azure." According to the Microsoft Learn materials, Natural Language Processing (NLP) is a branch of artificial intelligence that allows computers to analyze, understand, and generate human language. NLP enables machines to work with text or speech data in a way that extracts meaning, sentiment, and intent.

Microsoft defines NLP as enabling scenarios such as language detection, text classification, key phrase extraction, sentiment analysis, and named entity recognition. The example given-classifying emails as

"work-related" or "personal"-is a text classification task, which falls under NLP capabilities. The AI model processes the textual content of emails, identifies linguistic patterns, and categorizes them based on the detected topic or context.

Let's analyze the other options:

* Predict the number of future car rentals # This is a forecasting task, handled by machine learning regression models, not NLP.

* Predict which website visitors will make a transaction # This is a classification or prediction problem in machine learning, not NLP, since it deals with behavioral or numerical data rather than language.

* Stop a process in a factory when extremely high temperatures are registered # This is an IoT or anomaly detection scenario, focusing on sensor data, not language understanding.

Therefore, only classifying email messages as work-related or personal correctly represents an NLP use case.

It illustrates how NLP can analyze written text and make intelligent categorizations-a key capability covered in AI-900's natural language workloads section.

質問: 28

会社のプレスリリースをさまざまな言語で利用できるようにする必要があります。

どのサービスを使うべきですか？

A. 翻訳者のテキスト

B. テキスト分析

C. スピーチ

D. 言語理解 (UIS)

正解: ([正解を表示します](#))

The Translator Text service (part of Azure Cognitive Services) provides real-time text translation across multiple languages. According to Microsoft Learn's AI-900 module on "Identify features of Natural Language Processing (NLP) workloads", translation is one of the four main NLP tasks, alongside key phrase extraction, sentiment analysis, and language understanding.

In this scenario, the company wants to make press releases available in a range of languages, which requires converting text from one language to another while preserving meaning and tone. The Translator Text API supports more than 100 languages and can be integrated into web apps, chatbots, or content management systems for automatic multilingual publishing.

The other options perform different functions:

* Text Analytics (B) extracts insights such as key phrases or sentiment but does not translate.

* Speech (C) focuses on converting between speech and text, not text translation.

* Language Understanding (LUIS) (D) identifies user intent but does not perform translation.

Therefore, to provide multilingual press releases, the appropriate service is A. Translator Text, which ensures accurate, fast, and scalable translation across global audiences.

質問: 29

Azure Machine Learning Designerのキャンバスにドラッグできる2つのコンポーネントはどれですか？それぞれの正解は完全な解決策を提示します。

注：正しい選択はそれぞれ1ポイントの価値があります。

- A. データセット
- B. 計算
- C. パイプライン
- D. モジュール

正解: [\(正解を表示します\)](#)

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

質問: 30

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area

Statements	Yes	No
A smart device in the home that responds to questions such as "When is my next appointment?" is an example of conversational AI.	☐	☐
An interactive webchat feature on a company website can be implemented by using Azure Bot Service.	☐	☐
Automatically generating captions for pre-recorded videos is an example of conversational AI.	☐	☐

正解:

Answer Area

Statements	Yes	No
A smart device in the home that responds to questions such as "When is my next appointment?" is an example of conversational AI.	☒	☐
An interactive webchat feature on a company website can be implemented by using Azure Bot Service.	☒	☐
Automatically generating captions for pre-recorded videos is an example of conversational AI.	☐	☒

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.

質問: 31

ある地域の動物の個体数を予測する必要があります。

どの Azure Machine Learning タイプを使用する必要がありますか？

A. クラスタリング

B. 分類

C. 回帰

正解: C ([コメントを发表する](#))

According to the AI-900 official study materials, regression is a type of supervised machine learning used to predict continuous numeric values. Predicting the animal population of an area involves estimating a numeric quantity, which makes regression the appropriate model type.

Microsoft Learn defines regression workloads as predicting real-valued outputs, such as:

* Forecasting sales or demand.

* Predicting housing prices.

* Estimating resource usage or population sizes.

In contrast:

* Classification predicts discrete categories (e.g., "cat" or "dog").

* Clustering groups data into similar clusters but doesn't produce numeric predictions.

Therefore, because the task requires predicting a numerical population size, the verified answer is C.

Regression, as per Microsoft's AI-900 official guidelines.

有効的なAI-900J問題集はJPNTTest.com提供され、AI-900J試験に合格することに役に立ちます！JPNTTest.comは今最新AI-900J試験問題集を提供します。JPNTTest.com AI-900J試験問題集はもう更新されました。ここでAI-900J問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> 336問、30%ディスカウント、特別な割引コード: **JPNshiken**」

質問: 32

ツールをAzureMachineLearningタスクに一致させます。

答えるには、適切なツールを左側の列から右側のタスクにドラッグします。各ツールは、1回使用することも、複数回使用することも、まったく使用しないこともできます。注：正しい一致はそれぞれ1ポイントの価値があります。

Tools

Automated machine learning (automated ML)

The Azure portal

Machine Learning designer

Answer Area

Tool

Tool

Tool

Create a Machine Learning workspace

Use a drag-and-drop interface used to train and deploy models

Use a wizard to select configurations for a machine learning run

正解:

Tools

Automated machine learning (automated ML)

The Azure portal

Machine Learning designer

Answer Area

The Azure portal

Machine Learning designer

Automated machine learning (automated ML)

Create a Machine Learning workspace

Use a drag-and-drop interface used to train and deploy models

Use a wizard to select configurations for a machine learning run

Explanation:

Answer Area

The Azure portal

Machine Learning designer

Automated machine learning (automated ML)

Create a Machine Learning workspace

Use a drag-and-drop interface used to train and deploy models

Use a wizard to select configurations for a machine learning run

The correct matching aligns directly with the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn modules under "Identify features of Azure Machine Learning". Azure Machine Learning provides a suite of tools that serve different functions within the model development lifecycle - from creating workspaces, to training models, to automating experimentation.

* The Azure portal # Create a Machine Learning workspace. The Azure portal is a web-based graphical interface for managing all Azure resources. According to Microsoft Learn, you use the portal to create and configure the Azure Machine Learning workspace, which acts as the central environment where datasets, experiments, models, and compute resources are organized. Creating a workspace through the portal involves specifying a subscription, resource group, and region - tasks that are part of the setup stage rather than model development.

* Machine Learning designer # Use a drag-and-drop interface used to train and deploy models. The Machine Learning designer (formerly "Azure ML Studio (classic)") provides a visual, no-code/low-code interface for building, training, and deploying machine learning pipelines. The designer uses a drag-and-drop workflow where users connect modules representing data transformations, model training, and evaluation. This tool is ideal for beginners and those who want to quickly experiment with machine learning concepts without writing code.

* Automated machine learning (Automated ML) # Use a wizard to select configurations for a machine learning run. Automated ML simplifies model creation by automatically selecting algorithms, hyperparameters, and data preprocessing options. Users interact through a guided wizard (within the Azure Machine Learning studio) that walks them through configuration steps such as selecting datasets, target columns, and performance metrics. The system then iteratively trains and evaluates multiple models to recommend the best-performing one.

Together, these tools streamline the machine learning workflow:

* Azure portal for setup and resource management,

* Machine Learning designer for visual model creation, and

* Automated ML for guided, automated model selection and tuning.

質問: 33

を選択。



正解:



Explanation:



The correct completion of the sentence is:

"You can use the Custom Vision service to train an object detection model by using your own images." According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of computer vision workloads," the Azure Custom Vision service is a specialized component of Azure Cognitive Services for Vision that enables developers to train custom image classification or object detection models using their own labeled image datasets.

The Custom Vision service differs from the Computer Vision service in that it allows full customization - meaning you can upload your own images, tag them manually, and train the model to recognize objects specific to your use case (for example, detecting your company's products, tools, or vehicles). Once trained, the model can identify and localize these objects in new images by returning bounding boxes and confidence scores, which is precisely what defines an object detection workload.

Microsoft's AI-900 materials describe object detection as the process of identifying objects in an image and determining their position, typically represented by bounding boxes. Custom Vision supports two main project types:

* Image Classification: Determines what is present in the image (e.g., "dog," "cat," "car").

* Object Detection: Identifies what is present and where it is located in the image.

In contrast:

- * Computer Vision provides prebuilt models for general image analysis but doesn't allow custom model training.
- * Form Recognizer is used for extracting text and data from structured or semi-structured documents.
- * Azure Video Analyzer for Media focuses on video content analysis, not custom object detection.

Therefore, based on the official Microsoft AI-900 study guide and Microsoft Learn content, the verified and correct answer is Custom Vision, as it specifically allows training of a custom object detection model using your own images.

質問: 34

エンジン温度を監視するモノのインターネット (IoT) デバイスがあります。
エンジン温度が予想される基準から逸脱した場合、デバイスはアラートを生成します。
デバイスはどのタイプの AI ワークロードを表していますか？

- A. 自然言語処理 (NLP)**
- B. コンピュータビジョン**
- C. 異常検出**
- D. ナレッジマイニング**

正解: ([正解を表示します](#))

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 machine learning workload used to identify data points or patterns that deviate significantly from expected behavior.

In this scenario, the IoT device monitors engine temperature and generates alerts when the readings deviate from normal operating ranges. This directly matches the definition of anomaly detection, where the AI system learns what "normal" looks like and identifies outliers or abnormal conditions that may indicate potential issues.

Common real-world uses of anomaly detection include:

- * Detecting equipment malfunctions or overheating in IoT systems.
- * Identifying fraudulent transactions in finance.
- * Detecting unusual spikes or drops in system metrics (e.g., temperature, traffic, or pressure).

Other options are incorrect:

- * A. NLP (Natural Language Processing): Focuses on understanding and interpreting human language, not sensor data.
- * B. Computer Vision: Involves analyzing images or videos, which is unrelated to temperature data.
- * D. Knowledge Mining: Refers to extracting information from large document stores, not identifying abnormal readings.

質問: 35

友達の画像に自動的にタグを付けるソーシャルメディア用の画像タグ付けソリューションを構築する必要があります。
どのAzureCognitive Servicesサービスを使用する必要がありますか？

- A. コンピュータビジョン**
- B. 顔**
- C. テキスト分析**
- D. フォーム認識機能**

正解: ([正解を表示します](#))

The correct answer is B. Face because the Azure Face service, part of Azure Cognitive Services, provides capabilities for detecting, recognizing, and analyzing human faces in images. According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Explore computer vision", the Face service can automatically detect faces, determine attributes such as age, gender, emotion, and identify or verify individuals based on facial features.

In this scenario, the goal is to build an image tagging solution that automatically tags images of friends on social media. This requires the ability to recognize and match faces of known individuals, which is a key feature of the Face API. The API uses facial recognition technology to create face IDs for detected faces and can compare new faces against stored face data to identify known individuals.

Microsoft Learn notes: "The Face service can detect, recognize, and identify people in images, enabling automated tagging or authentication scenarios." Other options explained:

- * A. Computer Vision identifies objects, landmarks, and general content but does not specialize in identifying individual people.
- * C. Text Analytics analyzes textual data for sentiment, key phrases, or entities, not image data.
- * D. Form Recognizer extracts structured data from forms or receipts, unrelated to face detection.

Thus, for automatically tagging people in photos, Azure Face is the correct and most suitable service.

質問: 36

Azure AI Vision サービスを使用して実行できるアクションはどれですか？

- A. ライブビデオストリームで動物の品種を識別する
- B. 文書からキーフレーズを抽出する
- C. 手書きの文字からデータを抽出する
- D. トレーニングビデオのサムネイルを作成する

正解: ([正解を表示します](#))

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.

質問: 37

あなたはAIシステムを構築しています。

サービスが責任あるAIに関するマイクロソフトの透明性の原則を確実に満たすために、どのタスクを含める必要がありますか？

- A. すべてのビジュアルに、スクリーンリーダーで読み込めるテキストが関連付けられていることを確認します。
- B. 自動スケーリングを有効にして、サービスが需要に基づいてスケーリングされるようにします。
- C. 開発者がコードをデバッグするのに役立つドキュメントを提供します。
- D. トレーニングデータセットが母集団を代表していることを確認します。

正解: C ([コメントを發表する](#))

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

質問: 38

画像からテキストをキャプチャすることは、どのタイプの AI 機能の例ですか？

- A. テキスト分析
- B. 光学文字認識 (OCR)
- C. 画像の説明
- D. 物体検出

正解: B ([コメントを発表する](#))

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.

質問: 39

生成型AIモデルにはどのような2種類がありますか？正解はそれぞれ完全な解答を示しています。注：正解つにつき1ポイントです。

- A. トランスフォーマー
- B. 自然言語処理
- C. コンピュータビジョン
- D. GPT
- E. DALL-E

正解: D,E ([コメントを発表する](#))

質問: 40

画像セット内の類似した顔を識別するには何を使用すればよいですか？

- A. アズール・アル・ビジョン
- B. Azure AI カスタムビジョン
- C. Azure AI 言語

D. Azure OpenAI サービス

正解: ([正解を表示します](#))

The correct service to identify similar faces in a set of images is Azure AI Vision, which includes the Face API capability. According to the Microsoft Learn module "Analyze images with Azure AI Vision", this service provides prebuilt models for face detection, facial recognition, and similarity matching.

The Face API can detect individual faces in images and extract unique facial features to create a face embedding (a numerical representation of the face). It then compares these embeddings across multiple images to determine whether faces are similar or belong to the same person. This functionality is commonly used in identity verification, photo management systems, and security solutions.

The other options are incorrect:

* B. Azure AI Custom Vision is used for custom image classification or object detection but does not provide face similarity or recognition features.

* C. Azure AI Language processes text-based data (sentiment, entities, key phrases) - not visual content.

* D. Azure OpenAI Service focuses on text generation, summarization, and conversation, not facial analysis.

Therefore, the Microsoft-verified service for identifying similar faces across images is A. Azure AI Vision.

質問: 41

会話型 AI ワークロードの例となるシナリオを 2 つ挙げてください。それぞれの正解は完全なソリューションを示しています。

注意: 正しい選択ごとに 1 ポイントが加算されます。

A. 今日の天気はどうか?」などの質問に答える家庭内のスマート デバイス。

B. ナレッジベースを使用してユーザーの質問に対話的に応答するウェブサイト

C. ヘッドランプを車に自動で取り付ける組立ラインの機械

D. 機械の温度を監視し、温度が特定のしきい値に達したときにファンをオンにします。

正解: ([正解を表示します](#))

Conversational AI workloads involve human-like dialogue with AI systems.

* A: A smart assistant (e.g., smart speaker) uses voice-based conversational AI.

* B: A knowledge-based chatbot interacts with users via natural language. Options C and D describe automation/IoT workloads, not conversational AI.

Final Answer (Q110): A and B

質問: 42

文を完成させるには、回答領域で適切なオプションを選択します。

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

classification.

clustering.

regression.

 Microsoft

正解:

Answer Area

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

classification.
clustering.
regression.

Explanation:

Answer Area

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

classification.
clustering.
regression.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of regression machine learning", regression is a type of supervised machine learning used when the target variable (the value you want to predict) is a continuous numeric value.

In this scenario, the task is to predict how many hours of overtime a delivery person will work based on the number of orders received. Both the input (number of orders) and the output (hours of overtime) are numeric variables. Since the goal is to estimate a quantitative value rather than categorize or group data, this is a classic example of a regression problem.

Regression models analyze the relationship between variables to make numerical predictions. For example, the model might learn that each additional 20 orders increases overtime by about two hours. Common algorithms used for regression include linear regression, decision tree regression, and boosted regression models. These models produce outputs such as "expected overtime = 5.6 hours," which are continuous numeric results.

To contrast with the other options:

* Classification is used for predicting categories or labels, such as "overtime required" vs. "no overtime," or "high-risk" vs. "low-risk." It deals with discrete outputs rather than continuous numbers.

* Clustering is an unsupervised learning approach used to group similar data points based on shared characteristics, such as grouping delivery staff by performance patterns or customer types.

As emphasized in Microsoft's Responsible AI and Machine Learning Fundamentals learning paths, regression models are ideal for numeric forecasting problems such as predicting sales, revenue, demand, or working hours.

Therefore, the correct answer is: Regression.
Reference:Microsoft Learn - Identify the types of machine learning models: Regression, Classification, and Clustering (AI-900 Learning Path)

質問: 43
次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。
注：正しい選択はそれぞれ1ポイントの価値があります。

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.	☐	☐

正解:

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.

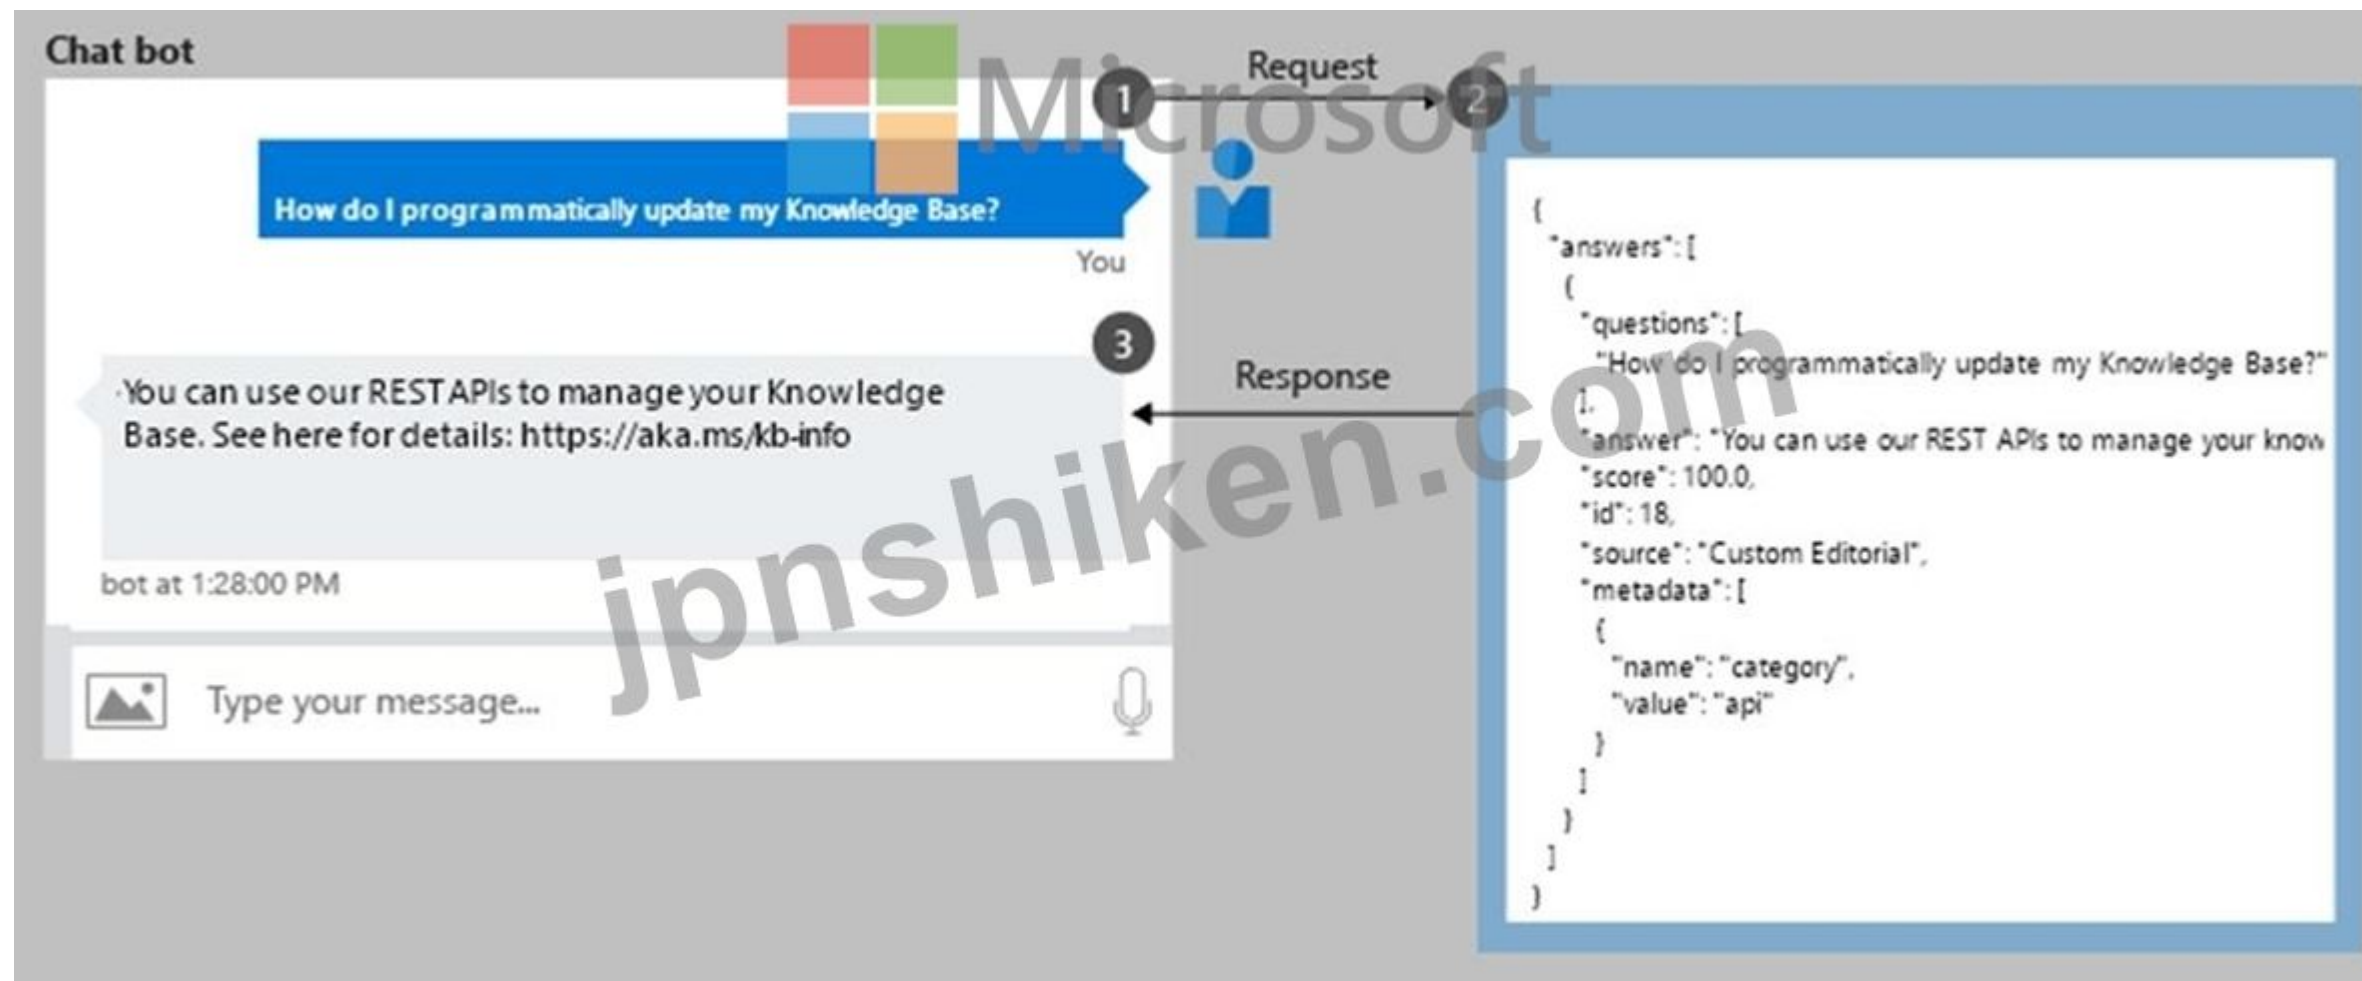
* Drag-and-Drop Visual Canvas # YESThe Azure Machine Learning Designer indeed provides a graphical interface where users can connect prebuilt modules for data preprocessing, training, evaluation, and deployment. Microsoft documentation describes it as a "drag-and-drop visual environment that simplifies 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 # YESThe 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 # NOThe 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.

質問: 44

次の展示に示すプロセスがあります。



図に示されているタイプのAIソリューションはどれですか？

- A. 機械学習モデル
- B. チャットボット
- C. 感情分析ソリューション
- D. コンピュータービジョンアプリケーション

正解: ([正解を表示します](#))

質問: 45

次の表に示す列を含むデータセットがあります。

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

他の数値列に基づいて ColumnE の値を予測する機械学習モデルがあります。

これはどのタイプのモデルですか？

- A. 回帰
- B. 分析
- C. クラスタリング

正解: ([正解を表示します](#))

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.

質問: 46

小売店からの画像を処理し、競合他社の製品を識別するツールを構築しています。

このソリューションでは、カスタムモデルを使用します。

どのAzureCognitive Servicesサービスを使用する必要がありますか？

- A. カスタムビジョン
- B. フォーム認識機能
- C. 顔
- D. コンピュータビジョン

正解: ([正解を表示します](#))

to 300 words in Explanation:

The Custom Vision service under Azure Cognitive Services is specifically designed for image classification and object detection tasks that require a custom-trained model. According to the AI-900 official study materials, Custom Vision enables developers to "build, deploy, and improve image classifiers that recognize specific objects in images based on custom data." In this question, the goal is to build a system that processes images from retail stores and identifies products of competitors. Since these are unique products that may not be part of Microsoft's pre-trained models, a custom model must be created. The Custom Vision service allows you to upload your own labeled images (e.

g., product pictures), train a model to recognize those products, and then deploy it as an API for image recognition tasks.

Other options explained:

- * B. Form Recognizer is used to extract text, key-value pairs, and tables from structured or semi- structured documents like invoices or receipts. It is not suitable for object identification.
- * C. Face service detects and analyzes human faces, providing attributes like age, emotion, and facial landmarks, but cannot recognize general objects like products.
- * D. Computer Vision is a general-purpose image analysis service used for tagging, OCR, and scene recognition, but it uses pre-trained models. It doesn't allow for custom product identification.

Thus, based on Microsoft's guidance, the best fit for recognizing competitor products from images using a custom-trained model is A. Custom Vision.

有効的な**AI-900J**問題集はJPNTTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AI-900J**試験問題集を提供します。JPNTTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **336**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」

質問: 47

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。
注：正しい選択はそれぞれ1ポイントの価値があります。

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

正解:

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

Explanation:

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

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of Computer Vision workloads on Azure", the Custom Vision service is a part of Azure Cognitive Services that allows users to build, train, and deploy custom image classification and object detection models. It is primarily designed for still-image analysis, not video processing.

* "The Custom Vision service can be used to detect objects in an image." - Yes. This is correct. The Custom Vision service supports two major model types: classification (categorizing entire images) and object detection (identifying and locating multiple objects within a single image). In object detection mode, the model outputs both the object's category and its position in the image using bounding boxes. This capability is emphasized in the AI-900 curriculum as an example of applying computer vision to real-world scenarios, such as identifying products on shelves or detecting equipment parts in manufacturing.

* "The Custom Vision service requires that you provide your own data to train the model." - Yes. This statement is also true. Unlike prebuilt computer vision models, Custom Vision is a trainable model that requires users to upload their own labeled images to create a domain-specific AI model. The model's accuracy depends on the quality and quantity of this user-provided data. The AI-900 study materials explain that Custom Vision is used when prebuilt models do not meet specific needs, enabling businesses to train models tailored to unique image sets.

* "The Custom Vision service can be used to analyze video files." - No. This is incorrect. Custom Vision is limited to image-based analysis. To analyze video content (detecting objects or motion in moving frames), Azure provides Video Indexer, which is a separate service designed for extracting insights from video files, including speech, objects, faces, and emotions.

質問: 48

ローンを承認する必要があるかどうかを評価するAIシステムを設計するときは、決定を下すために使用される要素を説明する必要があります。

これは、責任あるAIに関するマイクロソフトの指針となる原則の例です。

- A. 透明性
- B. 包括性
- C. 公平性
- D. プライバシーとセキュリティ

正解: ([正解を表示します](#))

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.

質問: 49

文を正しく完成させる答えを選択してください。

Answer Area

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

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

Microsoft

正解:

Answer Area

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

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

Microsoft

Explanation:

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



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.

質問: 50

あなたの会社は、ボトルのリサイクル機を作りたいと考えています。リサイクル機は、正しい形状のボトルを自動的に識別し、他のすべてのアイテムを拒否する必要があります。会社はどのタイプのAIワークロードを使用する必要がありますか？

- A. 異常検出
- B. 会話型AI
- C. コンピュータービジョン
- D. 自然言語処理

正解: ([正解を表示します](#))

According to Microsoft's AI-900 study guide and the Microsoft Learn module "Identify common types of AI workloads", computer vision is the branch of AI that enables machines to interpret and act upon visual inputs such as photos, videos, or camera feeds.

In this scenario, the recycling machine must identify bottles based on their shape and reject other objects. This requires the AI system to visually analyze an image or a video stream and make a decision based on shape recognition-a classic computer vision application.

Computer vision models are trained with labeled image data (e.g., images of correct bottle shapes vs. incorrect items). Using techniques such as object detection or image classification, the model learns to detect patterns and identify whether an image matches the desired object.

The other options are not suitable:

- * A. Anomaly detection identifies unusual data patterns, not visual features.

- * B. Conversational AI handles human-computer interactions via text or voice.
- * D. Natural language processing (NLP) deals with understanding and generating human language.

Therefore, because the task requires visual recognition of bottle shapes, the appropriate AI workload is Computer Vision.

Reference: Microsoft Learn - Identify common types of AI workloads: Computer Vision, NLP, Conversational AI, and Anomaly Detection (AI-900 Learning Path)

質問: 51

文を完成させるには、回答領域で適切なオプションを選択します。

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

inclusiveness
accountability
reliability and safety
fairness

正解:

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

inclusiveness
accountability
reliability and safety
fairness

Explanation:

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

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>

質問: 52

ドローンを使って作物の列の間に生えている雑草の位置を特定し、雑草除去の指示を出す。これはどのタイプのコンピュータビジョンの一例でしょうか？

- A. 光学文字認識(OCR)
- B. シーン分割
- C. 物体検出

正解: (正解を表示します)

質問: 53

文を正しく完成させる答えを選択してください。



正解:



Explanation:



In the Microsoft Azure AI Fundamentals (AI-900) curriculum, computer vision capabilities refer to artificial intelligence systems that can analyze and interpret visual content such as images and videos. The Azure AI Vision and Face API services provide pretrained models for detecting, recognizing, and analyzing visual information, enabling developers to build intelligent applications that understand what they "see." When asked how computer vision capabilities can be deployed, the correct answer is to integrate a face detection feature into an app. This aligns with Microsoft Learn's module "Describe features of computer vision workloads," which explains that computer vision can identify objects, classify images, detect faces, and extract text (OCR). The Face API, a part of Azure AI Vision, specifically provides face detection, verification, and emotion recognition capabilities.

Integrating these services into an application allows it to perform actions such as:

- * Detecting human faces in photos or video streams.
- * Recognizing facial attributes like age, emotion, or head pose.
- * Enabling secure authentication based on face recognition.

The other options are incorrect because they relate to different AI workloads:

- * Develop a text-based chatbot for a website: This falls under Conversational AI, implemented with Azure Bot Service or Conversational Language Understanding (CLU).
- * Identify anomalous customer behavior on an online store: This task relates to machine learning and anomaly detection models, not computer vision.
- * Suggest automated responses to incoming email: This uses Natural Language Processing (NLP) capabilities, not visual analysis.

Therefore, the correct and Microsoft-verified completion of the statement is:

"Computer vision capabilities can be deployed to integrate a face detection feature into an app."

質問: 54

手書きのメモをデジタルテキストに変換する必要があります。
どのタイプのコンピューター ビジョンを使用すればよいですか？

- A. 光学文字認識 (OCR)
- B. 物体検出
- C. 画像分類
- D. 顔検出

正解: ([正解を表示します](#))

According to the Microsoft Azure AI Fundamentals (AI-900) study guide and Microsoft Learn documentation on Azure AI Vision, OCR is a computer vision technology that detects and extracts printed or handwritten text from images, scanned documents, or photographs. The OCR feature in Azure AI Vision can analyze images containing handwritten notes, recognize the characters, and convert them into machine-readable digital text.

This process is ideal for digitizing handwritten meeting notes, forms, or classroom materials. OCR works by identifying text regions in an image, segmenting characters or words, and then applying language models to interpret them correctly. Azure's OCR capabilities support multiple languages and can handle varied handwriting styles.

Other options are incorrect because:

- * B. Object detection identifies and locates objects (like cars, animals, or furniture) within an image, not text.
- * C. Image classification assigns an image to a predefined category (e.g., "dog" or "cat") rather than extracting text.
- * D. Facial detection detects or recognizes human faces, not written text.

Therefore, to convert handwritten notes into digital text, the correct computer vision technique is Optical Character Recognition (OCR).

質問: 55

文を正しく完成させる答えを選択してください。

Answer Area

You can use the

Azure AI Custom Vision

Azure AI Computer Vision

Azure AI Custom Vision

Azure AI Document Intelligence

Azure Video Analyzer for Media

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

正解:

Answer Area

You can use the

Azure AI Custom Vision

Azure AI Computer Vision

Azure AI Custom Vision

Azure AI Document Intelligence

Azure Video Analyzer for Media

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



Explanation:

Answer Area

You can use the

Azure AI Custom Vision

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

The correct answer is Azure AI Custom Vision.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Explore computer vision in Azure", the Azure AI Custom Vision service is specifically designed to allow users to train, test, and deploy custom image classification and object detection models using their own images.

The Azure AI Custom Vision service extends the capabilities of Azure's general-purpose Computer Vision API by enabling organizations to upload their own labeled datasets, define custom tags (labels), and train models that are optimized for their specific use cases. This makes it ideal for object detection scenarios-such as detecting equipment in a manufacturing line, identifying products on store shelves, or recognizing medical images-where general-purpose models may not suffice.

By contrast, the Azure AI Computer Vision service provides pre-built models for tasks such as image description, tagging, face detection, and OCR (optical character recognition). It does not allow users to train their own models with custom data. Similarly, Azure AI Document Intelligence is used for extracting structured information from documents (forms, receipts, invoices), and Azure Video Analyzer for Media focuses on analyzing video content for insights and metadata extraction.

The AI-900 study guide emphasizes that the Custom Vision service supports two key model types:

- * Image Classification - categorizing entire images based on predefined tags.
- * Object Detection - identifying and locating multiple objects within an image by drawing bounding boxes.

Therefore, when the question specifies "train an object detection model by using your own images," the correct Azure service is Azure AI Custom Vision, as it provides the necessary tools for training, evaluating, and deploying custom computer vision models tailored to a user's dataset.

Hence, the verified correct answer is: Azure AI Custom Vision.

質問: 56

Azure AI Document Intelligence の事前構築済み領収書モデルを使用して処理できる画像の最大サイズはどれくらいですか？

- A. 50 MB
- B. 10MB
- C. 100MB
- D. 5 MB

正解: ([正解を表示します](#))

質問: 57

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Statements	Yes	No
Named entity recognition can be used to retrieve dates and times in a text string.	☐	☐
Key phrase extraction can be used to retrieve important phrases in a text string.	☐	☐
Key phrase extraction can be used to retrieve all the city names in a text string.	☐	☐

正解:

Statements	Yes	No
Named entity recognition can be used to retrieve dates and times in a text string.	☒	☐
Key phrase extraction can be used to retrieve important phrases in a text string.	☒	☐
Key phrase extraction can be used to retrieve all the city names in a text string.	☐	☒

Explanation:

Answer Area

Statements

Named entity recognition can be used to retrieve dates and times in a text string.

Yes



No



Key phrase extraction can be used to retrieve important phrases in a text string.



Key phrase extraction can be used to retrieve all the city names in a text string.



In Microsoft Azure AI Language Service, both Named Entity Recognition (NER) and Key Phrase Extraction are core features for text analytics. They serve distinct purposes in analyzing and structuring unstructured text data.

* Named Entity Recognition (NER):NER is used to identify and categorize specific entities within text, such as people, organizations, locations, dates, times, and quantities. According to Microsoft Learn's "Analyze text with Azure AI Language" module, NER scans text to extract these entities along with their types. Therefore, the statement "Named entity recognition can be used to retrieve dates and times in a text string" is True (Yes).

* Key Phrase Extraction:This feature identifies the most important phrases or main topics in a block of text. It is useful for summarization or highlighting central ideas without classifying them into specific categories. Therefore, the statement "Key phrase extraction can be used to retrieve important phrases in a text string" is also True (Yes).

* City Name Retrieval:While key phrase extraction highlights major phrases, it does not extract specific entities like cities or dates. Extracting such details requires Named Entity Recognition, which is designed to find named entities such as city names, people, or organizations. Hence, the statement "Key phrase extraction can be used to retrieve all the city names in a text string" is False (No).

質問: 58

AI ベースのローン承認システムがあります。

テスト中に、システムに性別バイアスがあることがわかりました。

これはどの責任ある AI 原則に違反していますか？

- A. 説明責任
- B. 透明度
- C. 公平性
- D. 信頼性と安全性

正解: C ([コメントを发表する](#))

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.

質問: 59

次の展示に示す予測対真のチャートがあります。



チャートはどのタイプのモデルを評価に使用しますか？

- A. 分類
- B. 回帰
- C. クラスタリング

正解: B ([コメントを发表する](#))

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>

質問: 60

文を正しく完成させる答えを選択してください。

Answer Area



正解:



Explanation:

Safety system.

According to the Microsoft Learn documentation and the AI-900: Microsoft Azure AI Fundamentals official study guide, the safety system layer in generative AI architecture plays a crucial role in monitoring, filtering, and mitigating harmful or unsafe model outputs. This layer works alongside the model and user experience layers to ensure that generative AI systems-such as those powered by Azure OpenAI-produce responses that are safe, aligned, and responsible.

The safety system layer uses various techniques including content filtering, prompt moderation, and policy enforcement to prevent outputs that could be harmful, biased, misleading, or inappropriate. It evaluates both user inputs (prompts) and model-generated outputs to identify and block unsafe or unethical content. The system might use predefined rules, classifiers, or human feedback signals to decide whether to allow, modify, or stop a response.

In contrast, the other layers serve different purposes:

- * The model layer contains the core large language or generative model (e.g., GPT or DALL-E) that processes inputs and produces outputs.
- * The metaprompt and grounding layer ensures the model's responses are contextually relevant and factually supported, often linking to organizational data sources or system prompts.
- * The user experience layer defines how users interact with the AI system, including the interface and conversational flow, but does not manage safety enforcement.

Therefore, the layer that uses system inputs and context to mitigate harmful outputs from a generative AI model is the Safety system layer.

This aligns with Microsoft's responsible AI principles-Fairness, Reliability and Safety, Privacy and Security, Inclusiveness, Transparency, and Accountability-ensuring generative AI operates ethically and safely.

質問: 61

文を完成させるには、回答領域で適切なオプションを選択します。



正解:

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

classification.
clustering.
regression.

Explanation:

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

classification.
clustering.
regression.

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.

有効的な**AI-900J**問題集はJPNTTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AI-900J**試験問題集を提供します。JPNTTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **336**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」

質問: 62

コンピューター ビジョンのワークロードの種類を適切なシナリオに合わせます。

答えるには、適切なワークロード タイプを左側の列から右側のシナリオにドラッグします。各ワークロード タイプは複数回使用されることも、まったく使用されないこともあります。

注: 正しく一致するたびに 1 ポイントの価値があります。

Workload Types	Answer Area
Image classification	Generate captions for images.
Object detection	Extract movie title names from movie poster images.
Optical character recognition (OCR)	Locate vehicles in images.

正解:

Workload Types	Answer Area
Image classification	Generate captions for images.
Object detection	Extract movie title names from movie poster images.
Optical character recognition (OCR)	Locate vehicles in images.

Explanation:

Workload Types	Answer Area
Image classification	Generate captions for images.
Object detection	Extract movie title names from movie poster images.
Optical character recognition (OCR)	Locate vehicles in images.

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.

質問: 63

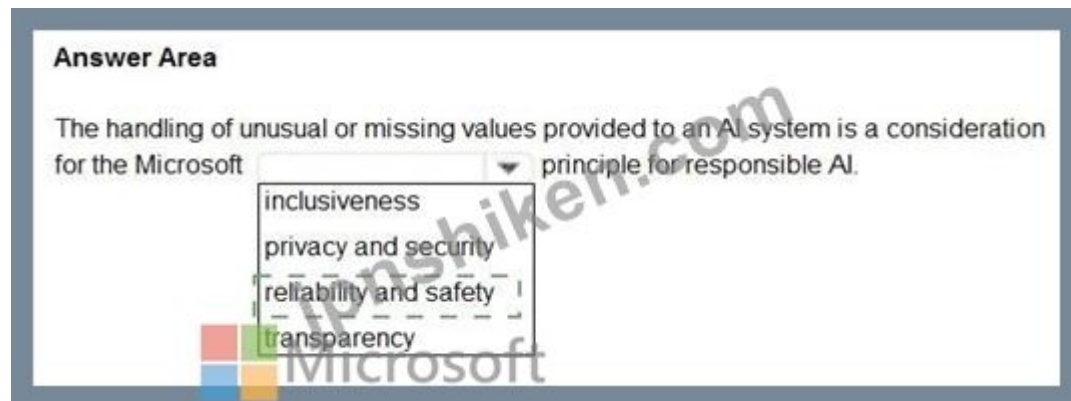
文を完成させるには、回答領域で適切なオプションを選択します。

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

正解:



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"

質問: 64

よくある質問 (FAQ) ドキュメントからボットを作成するには、どのAIサービスを使用する必要がありますか？

- A. QnAメーカー
- B. 言語理解 (UIS)
- C. テキスト分析
- D. スピーチ

正解: ([正解を表示します](#))

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-phrase 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.

質問: 65

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area

Statements	Yes	No
Azure Bot Service and Azure Cognitive Services can be integrated.	☐	☐
Azure Bot Service engages with customers in a conversational manner.	☐	☐
Azure Bot Service can import frequently asked questions (FAQ) to question and answer sets.	☐	☐

正解:

Answer Area

Statements	Yes	No
Azure Bot Service and Azure Cognitive Services can be integrated.	☒	☐
Azure Bot Service engages with customers in a conversational manner.	☒	☐
Azure Bot Service can import frequently asked questions (FAQ) to question and answer sets.	☒	☐

Explanation:

Azure Bot Service and Azure Cognitive Services can be integrated. # Yes

* Azure Bot Service engages with customers in a conversational manner. # Yes

* Azure Bot Service can import frequently asked questions (FAQ) to question and answer sets. # Yes

\

All three statements are true, as confirmed by the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Explore conversational AI." The Azure Bot Service is Microsoft's platform for building, deploying, and managing intelligent bots that can communicate naturally with users across various channels (web, Teams, Facebook Messenger, etc.).

* Azure Bot Service and Azure Cognitive Services can be integrated # YesMicrosoft Learn specifies that Azure Bot Service can be enhanced with Azure Cognitive Services such as Language Understanding (LUIS), QnA Maker, and Speech Services to add intelligence. For example, integration with LUIS allows bots to understand user intent and context, while QnA Maker helps them respond accurately to FAQs. As stated in the official documentation: "The Azure Bot Service can be combined with Cognitive Services to create bots that understand language, speech, and meaning."

* Azure Bot Service engages with customers in a conversational manner # YesThe primary function of Azure Bot Service is to create conversational AI agents that interact naturally with users. These bots simulate human-like dialogue using text or speech. According to Microsoft Learn, "Bots created using Azure Bot Service communicate with users in a conversational format through natural language."

* Azure Bot Service can import frequently asked questions (FAQ) to question and answer sets # YesAzure Bot Service can integrate with the QnA Maker (now part of Azure Cognitive Service for Language) to automatically import FAQs from existing documents or web pages and generate a knowledge base of question-answer pairs. This allows the bot to respond intelligently to customer queries. In conclusion, Azure Bot Service supports intelligent, conversational interaction, integrates seamlessly with Cognitive Services, and can use QnA Maker to import and manage FAQ-based knowledge sets-making all three statements true.

質問: 66

Azure Cognitive Services を使用して、音声制御のパーソナル アシスタント アプリを開発する予定です。

Azure Cognitive Services を適切なタスクに合わせます。

回答するには、左側の列から適切なサービスを右側の説明にドラッグします。各サービスは、1 回だけ使用することも、複数回使用することも、まったく使用しないこともできます。

注: 正しく選択するたびに 1 ポイントの価値があります。

正解:

Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn Cognitive Services documentation, developing a voice-controlled personal assistant app involves integrating multiple Azure AI services that specialize in different aspects of language and speech processing. The three services in focus-Azure AI Speech, Azure AI Language Service, and Azure AI Translator Text-perform unique but complementary roles in conversational AI systems.

- * Convert a user's speech to text # Azure AI SpeechThe Azure AI Speech service provides speech-to-text (STT) capabilities. It enables applications to recognize spoken language and convert it into written text in real time. This is often the first step in voice-enabled applications, transforming audio input into a machine-readable format that can be analyzed further.
- * Identify a user's intent # Azure AI Language serviceOnce speech has been transcribed, the Azure AI Language service (which includes capabilities like Conversational Language Understanding and Text Analytics) interprets the meaning of the text. It detects the user's intent (what the user wants to accomplish) and extracts entities (key data points) from the input. This service helps the assistant understand commands like "Book a flight" or "Set a reminder."
- * Provide a spoken response to the user # Azure AI SpeechAfter determining an appropriate response, the system uses the text-to-speech (TTS) feature of Azure AI Speech to convert the assistant's text-based reply back into natural-sounding spoken language, allowing the user to hear the response.

Together, these services form the backbone of a conversational AI system: Speech-to-Text # Language Understanding # Text-to-Speech, aligning precisely with the AI-900 curriculum's explanation of how Azure Cognitive Services enable intelligent voice-based interactions.

質問: 67

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area		Microsoft	
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.		☐	☐

正解:

Answer Area

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:

Answer Area

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>

質問: 68

ソーシャル メディアの投稿から日付や人の名前などのデータを取得するために使用できる Azure AI Language 機能はどれですか？

A. 言語検出

B. 音声認識

C. キーフレーズ抽出

D. エンティティ認識

正解: (正解を表示します)

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.

質問: 69

次のアプリがあります:

- * アプリ1: 画像と写真のセットを使用してブランド名を抽出します
- * アプリ2: 建物のタッチレスアクセス制御を可能にする

各アプリはどの Azure AI Vision サービスを使用していますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが付与されます。



正解:

Answer Area



App1:

App2:

Explanation:

Answer Area

App1:

App2:

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn documentation on Azure AI Vision services, different Azure AI Vision capabilities are suited for different use cases such as object detection, brand recognition, facial recognition, and spatial analysis.

* App1: Uses a set of images and photos to extract brand names # Image AnalysisThe Azure AI Vision - Image Analysis service (formerly part of Computer Vision) can detect and extract brands, objects, text, and other visual features from images. It uses advanced image classification and object detection models to recognize logos and identify brand names (for example, "Microsoft" or "Coca-Cola") in photos. The Image Analysis API can also return descriptive tags, scene descriptions, and confidence scores. Therefore, since App1 analyzes static images to extract brand names, it specifically relies on the Image Analysis feature of Azure AI Vision.

* App2: Enables touchless access control for buildings # FaceThe Azure AI Face service is designed for facial detection, verification, and identification. It can recognize and match faces in real time, making it ideal for access control, identity verification, and attendance tracking systems. A "touchless access control" system uses a camera to detect a person's face and verify identity against a stored profile, allowing or denying entry without physical interaction.

The other options are not suitable:

* Optical Character Recognition (OCR) extracts text, not brand logos.

* Spatial Analysis is for detecting movement or presence in video feeds.

* Video Analysis is for analyzing dynamic video content rather than still images.

質問: 70

あなたはレースのランナーの写真を処理しています。

写真のランナーを識別するには、ランナーのシャツの番号を読む必要があります。

どのタイプのコンピュータビジョンを使用する必要がありますか？

A. 顔認識

B. 光学式文字認識 (OCR)

C. セマンティックセグメンテーション

D. オブジェクト検出

正解: ([正解を表示します](#))

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features of computer vision workloads", Optical Character Recognition (OCR) is a core capability within the computer vision domain that enables systems to detect and extract text from images or documents. OCR technology can identify printed or handwritten characters in photographs, scanned documents, or camera feeds, and convert them into machine-readable text.

In this scenario, the task is to read the numbers on runners' shirts in race photos. These numbers are textual or numeric characters embedded within images. OCR is specifically designed for this purpose - to locate and recognize characters within visual data and convert them into usable text. Once extracted, those numbers can be cross-referenced with a database to identify each runner.

Let's analyze why the other options are incorrect:

* A. Facial recognition focuses on identifying individuals based on unique facial features, not reading text or numbers.

* C. Semantic segmentation classifies each pixel of an image into categories (for example, separating road, sky, and people), but it doesn't read text.

* D. Object detection identifies and locates objects within an image (such as detecting people or vehicles) but does not extract readable text or numbers.

Therefore, since the task involves reading textual or numeric content from an image, the appropriate type of computer vision to use is Optical Character Recognition (OCR).

Reference: Microsoft Learn - Identify features and uses of computer vision in Azure AI services (Cognitive Services - Optical Character Recognition).

質問: 71

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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 like today?" is an example of conversational AI	☐	☐

正解:

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 like today?" is an example of conversational AI	☒	☐

Explanation:

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 conversational AI	☐	☒
A smart device in the home that responds to questions such as "What will the weather like today?" is an example of conversational AI	☒	☐

These answers align with 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. The Azure Bot Service allows developers to create intelligent chatbots that can be integrated into a webchat interface. This enables visitors to interact with the bot directly from a website, asking questions and receiving automated responses. This is a typical use case of conversational AI, where natural language processing (NLP) is used to interpret and respond to user input conversationally.

2. Automatically generating captions for pre-recorded videos is an example of conversational AI # No This statement is false. Automatically generating captions from video content is an example of speech-to-text (speech recognition) technology, not conversational AI. While it uses AI to convert spoken words into text, it lacks the two-way interactive communication characteristic of conversational AI. This task is typically handled by the Azure AI Speech service, which transcribes spoken content.

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 statement is true. Smart home assistants that engage in dialogue with users are powered by conversational AI. These devices use speech recognition to understand spoken input, natural language understanding (NLU) to determine intent, and speech synthesis (text-to-speech) to respond audibly. This represents the full conversational AI loop, where machines communicate naturally with humans.

質問: 72

Azure OpenAI サービスを使用する生成 AI ソリューション内の有害なコンテンツを識別する必要があります。

何をすべきでしょうか？

- A. 顔
- B. ビデオ分析
- C. 言語
- D. コンテンツの安全性

正解: ([正解を表示します](#))

According to the Microsoft Azure AI Fundamentals (AI-900) curriculum and Azure OpenAI documentation, the appropriate service for detecting and managing harmful, unsafe, or inappropriate content in text, images, or other generative AI outputs is Azure AI Content Safety.

Azure AI Content Safety is designed to automatically detect potentially harmful material such as hate speech, violence, self-harm, sexual content, or profanity. It ensures that generative AI applications like chatbots, image generators, and content creation tools comply with Microsoft's Responsible AI principles - specifically Reliability & Safety and Accountability.

This service integrates directly with the Azure OpenAI Service, meaning that when developers build AI solutions using models like GPT-4 or DALL E, they can use Content Safety to filter and moderate both input prompts and model outputs. This protects users from unsafe or offensive content generation.

Let's analyze why the other options are incorrect:

* A. Face - The Face service detects and analyzes human faces in images or videos. It is unrelated to moderating harmful textual or generative content.

* B. Video Analysis - This service analyzes video streams to detect objects, actions, or events but not inappropriate or harmful text or imagery from AI models.

* C. Language - The Azure AI Language service focuses on text understanding tasks like sentiment analysis, entity recognition, and translation, not content safety filtering. Therefore, per Microsoft Learn's official AI-900 guidance, when identifying or filtering harmful content in a generative AI solution built with Azure OpenAI, the correct and verified service to use is Azure AI Content Safety.

質問: 73

文を正しく完成させる答えを選択してください。

Answer Area

A randomly extracted subset of data from a dataset is commonly used for

features
algorithms
features
labels
validation

of the model.

正解:

Answer Area

A randomly extracted subset of data from a dataset is commonly used for

features
algorithms
features
labels
validation

of the model.

Explanation:

validation.

In the Microsoft Azure AI Fundamentals (AI-900) study materials, a key concept in machine learning model development is splitting data into subsets for training, validation, and testing. A randomly extracted subset of data from a dataset is most commonly used for validation - that is, for evaluating the performance of the model during or after training.

Here's how this process works:

- * Training set - This portion of the dataset is used to train the machine learning model. The model learns patterns, relationships, and parameters from this data.
- * Validation set - This is a randomly selected subset (separate from training data) used to fine-tune model hyperparameters and evaluate how well the model generalizes to unseen data. It helps detect overfitting
 - when the model performs well on training data but poorly on new data.
- * Test set - A final, untouched dataset used to measure the model's real-world performance after all training and tuning are complete.

By reserving a random subset for validation, data scientists ensure that the model's performance metrics reflect generalization, not memorization of the training data.

Let's review the incorrect options:

- * Algorithms - These are the mathematical frameworks or methods used to build models (e.g., decision trees, neural networks). They are not data subsets.
- * Features - These are input variables (attributes) used by the model, not randomly selected data subsets.
- * Labels - These are target values or outcomes the model predicts; again, not data subsets.

Therefore, in alignment with Azure AI-900's machine learning fundamentals, the correct completion is:

"A randomly extracted subset of data from a dataset is commonly used for validation of the model."

質問: 74

文を正しく完成させる答えを選択してください。



正解:

Answer Area



Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn documentation on GitHub Copilot, the GitHub Copilot extension for Microsoft Visual Studio Code (VS Code) is powered by the OpenAI Codex model, a specialized descendant of the GPT (Generative Pre-trained Transformer) family of models. The OpenAI Codex model is specifically fine-tuned for programming tasks- it understands and generates code in multiple programming languages such as Python, JavaScript, C#, and more.

GitHub Copilot functions as an AI-powered coding assistant that suggests complete lines or blocks of code, helps write comments, and can even generate functions based on natural language prompts entered by the developer. When a user types a comment like "// sort a list of numbers in ascending order," Copilot uses the Codex model to understand the intent and generate an appropriate code implementation.

The integration with Visual Studio Code allows developers to work efficiently without needing to switch between documentation and their coding environment. GitHub Copilot leverages context-aware suggestions, meaning it understands the surrounding code, function definitions, and variable names to provide relevant completions.

The other options are incorrect because:

- * GitHub source control manages code repositories but doesn't use AI models.
- * IntelliSense is a built-in VS Code feature for code completion but doesn't rely on OpenAI models.
- * Microsoft 365 Copilot uses large language models (like GPT-4) for productivity tools such as Word, Excel, and Outlook-not for code generation.

Therefore, the correct choice that uses the OpenAI Codex model in Visual Studio Code is GitHub Copilot, aligning with Microsoft's AI-900 learning objectives under "Identify common use cases for Azure OpenAI and GitHub Copilot."

質問: 75

Azure Machine Learningを使用してトレーニングされたモデルの複数のバージョンを追跡する必要があります。あなたは何をするべきか？

- A. 推論ダスターをプロビジョニングします。
- B. モデルを説明してください。
- C. モデルを登録します。
- D. トレーニングデータを登録します。

正解: (正解を表示します)

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore Azure Machine Learning," registering a model is the correct way to track multiple versions of models in Azure Machine Learning.

When you train models in Azure Machine Learning, each trained version can be registered in the workspace's Model Registry. Registration stores the model's metadata, including version, training environment, parameters, and lineage. Each registration automatically increments the version number, enabling you to manage, deploy, and compare multiple model iterations efficiently.

The other options are incorrect:

- * A. Provision an inference cluster - Used for model deployment, not version tracking.
- * B. Explain the model - Provides interpretability but does not track versions.
- * D. Register the training data - Registers data assets, not models.

質問: 76

音声を含まない説明ビデオが 100 本あります。各説明ビデオにはスクリプトがあります。スクリプトに基づいて、各ビデオのナレーション オーディオ ファイルを生成する必要があります。どのタイプのワークロードを使用する必要がありますか？

- A. 音声認識
- B. 言語モデル
- C. 音声合成
- D. 翻訳

正解: C (コメントを発表する)

Speech synthesis, also known as text-to-speech (TTS), is the AI workload that converts written text into spoken words. In this case, the task is to generate narration audio from provided scripts for silent instructional videos.

Speech recognition performs the opposite function - it converts speech into text. Language modeling is for text understanding and prediction (e.g., GPT). Translation converts text between languages, not from text to audio.


Therefore, the most appropriate workload, according to Microsoft's AI-900 study material under the "Speech AI capabilities" section, is speech synthesis, which enables natural voice narration generation.

有効的な**AI-900J**問題集はJPNTTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AI-900J**試験問題集を提供します。JPNTTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **336**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」

質問: 77

文を正しく完成させる答えを選択してください。

Answer Area



正解:



Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of Computer Vision workloads on Azure", Object Detection is a specific computer vision capability used to identify and locate multiple types of objects within a single image. Unlike image classification, which assigns one label to an entire image, object detection identifies individual objects, their categories, and their positions using bounding boxes or polygons.

In practical terms, Object Detection combines two key outputs:

- * Classification - recognizing what the object is (for example, "car", "person", "dog").
- * Localization - determining where the object appears in the image by drawing bounding boxes around it.

This technology is commonly used in scenarios such as traffic monitoring (detecting vehicles and pedestrians), retail shelf analysis (detecting products and inventory levels), and manufacturing quality control (identifying defective parts).

Microsoft's Azure Cognitive Services - Custom Vision includes a dedicated Object Detection domain, which allows developers to train custom models to recognize multiple object types within a single image. The service uses deep learning techniques, particularly convolutional neural networks (CNNs), to process pixel patterns and spatial relationships for accurate detection.

For contrast:

- * Image Classification identifies only the overall category of an image (e.g., "This is a cat").
- * Image Description generates captions summarizing the visual content (e.g., "A cat sitting on a couch").
- * Optical Character Recognition (OCR) detects and extracts text from images, not physical objects.

Therefore, per the official AI-900 learning content and Azure documentation, when the goal is to identify multiple types of items within a single image, the correct AI workload is Object Detection.

質問: 78

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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

正解:

Statements	Yes	No
Organizing documents into groups based on similarities of the text contained in the documents is an example of clustering.	☒	☐
Grouping similar patients based on symptoms and diagnostic test results is an example of clustering.	☒	☐
Predicting whether a person will develop mild, moderate, or severe allergy symptoms based on pollen count is an example of clustering.	☐	☒

Explanation:

Statements	Yes	No
Organizing documents into groups based on similarities of the text contained in the documents is an example of clustering.	☒	☐
Grouping similar patients based on symptoms and diagnostic test results is an example of clustering.	☒	☐
Predicting whether a person will develop mild, moderate, or severe allergy symptoms based on pollen count is an example of clustering.	☐	☒

This question evaluates understanding of clustering-an unsupervised learning technique explained in the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Explore fundamental principles of machine learning." Clustering involves finding natural groupings within data without prior knowledge of output labels. The algorithm identifies similarities among data points and groups them accordingly, with each group (or cluster) containing items that are more similar to each other than to those in other groups.

- * Organizing documents into groups based on similarities of the text contained in the documents # YesThis is a classic clustering application. In text analytics or natural language processing (NLP), clustering algorithms such as K-means or hierarchical clustering are used to group documents with similar content or topics. According to Microsoft Learn, "clustering identifies relationships in data and groups items that share common characteristics." Therefore, organizing text documents based on content similarity is a textbook example of clustering.
- * Grouping similar patients based on symptoms and diagnostic test results # YesThis is another example of clustering. In healthcare analytics, clustering can be used to segment patients with similar health patterns or risks. The study guide emphasizes that clustering can "discover natural groupings in data such as customers with similar buying patterns or patients with similar clinical results." Thus, this task correctly describes unsupervised clustering because it does not involve predicting a known outcome but grouping based on similarity.
- * Predicting whether a person will develop mild, moderate, or severe allergy symptoms based on pollen count # NoThis is a classification problem, not clustering. Classification is a supervised learning technique where the model is trained with labeled data to predict predefined categories (in this case, mild, moderate, or severe). Microsoft Learn clearly distinguishes between clustering (discovering hidden patterns) and classification (predicting predefined categories).

質問: 79

Azure OpenAI モデルが最近のイベントを含む正確な応答を生成するようにするには、どうすればよいでしょうか？

- A. システム メッセージを変更します。
- B. 接地データを追加します。
- C. 少数ショット学習を追加します。
- D. トレーニングデータを追加します。

正解: ([正解を表示します](#))

In Azure OpenAI, grounding refers to the process of connecting the model to external data sources (for example, a database, search index, or API) so that it can retrieve accurate and up-to-date information before generating a response. This is particularly important for scenarios requiring current facts or events, since OpenAI models like GPT-3.5 and GPT-4 are trained on data available only up to a certain cutoff date.

By adding grounding data, the model's responses are "anchored" to factual sources retrieved at runtime, improving reliability and factual accuracy. Grounding is commonly implemented in Azure OpenAI + Azure Cognitive Search solutions (Retrieval-Augmented Generation or RAG).

Option review:

- * A. Modify the system message: Changes model tone or behavior but doesn't supply real-time data.
- * B. Add grounding data: # Correct - allows access to recent and domain-specific information.
- * C. Add few-shot learning: Provides examples in the prompt to improve context understanding but not factual accuracy.
- * D. Add training data: Refers to fine-tuning; this requires retraining and doesn't update the model's awareness of current events.

Hence, the best method to ensure accurate and current responses from an Azure OpenAI model is to add grounding data, enabling the model to reference real, updated sources dynamically.

質問: 80

単純なユーザーのクエリに回答するのに役立つビジネス チャットボットのコンテンツを提供する必要があります。

Azure AI Language Service の質問応答機能を使用して質問と回答のテキストを作成する 3 つの方法は何ですか？それぞれの正しい回答は完全なソリューションを示します。

注意: 正解および選択による質問はそれぞれ 1 ポイントになります。

- A. Cortana を使用してボットを Cortana チャネルに接続します。
- B. 定義済みのデータ ソースから雑談コンテンツをインポートします。
- C. 質問と回答を手動で入力します。
- D. Azure Machine Learning Automated ML を使用して、質問と回答のペアを含むファイルに基づいてモデルをトレーニングします。
- E. 既存の Web ページから質問と回答を生成します。

正解: ([正解を表示します](#))

The correct answers are B. Import chat content from a predefined data source, C. Manually enter the questions and answers, and E. Generate the questions and answers from an existing webpage.

According to Microsoft Learn and the Azure AI Fundamentals (AI-900) study guide, the Question Answering feature of the Azure AI Language Service (formerly part of QnA Maker) allows developers to create a knowledge base (KB) that enables a chatbot to answer common questions automatically. This knowledge base can be built in three main ways:

- * Import chat content (B): Azure provides predefined chat datasets that can be imported to make a bot more conversational and natural. This includes small talk such as greetings, acknowledgments, and polite responses (for example, "How are you?" # "I'm doing great, thanks!"). Importing this content enriches the bot's personality and improves user engagement.
- * Manually enter questions and answers (C): Developers can manually add pairs of questions and answers directly into the question answering knowledge base. This approach is suitable for custom FAQs or domain-specific content. It gives complete control over how each question is phrased and what answer is returned, ensuring high precision and clarity.
- * Generate questions and answers from an existing webpage (E): Azure AI Language can automatically extract Q & A pairs from a website's FAQ or support page. This is done by providing the webpage URL to the service, which scans the page and builds a knowledge base from the detected questions and corresponding answers.

The other options are incorrect:

- * A (Cortana channel) relates to bot deployment, not knowledge creation.
- * D (Automated ML) is used for predictive modeling, not for building Q & A datasets.

Thus, the verified correct answers are B, C, and E.

質問: 81

文を正しく完成させる答えを選択してください。



正解:



Explanation:



The correct answer is "An embedding."

In the context of large language models (LLMs) such as GPT-3, GPT-3.5, or GPT-4, an embedding refers to a multi-dimensional numeric vector representation assigned to each word, token, or phrase. According to the Microsoft Azure AI Fundamentals (AI-900) study guide and Microsoft Learn documentation for Azure OpenAI embeddings, embeddings are used to represent textual or semantic meaning in a numerical form that a machine learning model can process mathematically.

Each embedding captures the semantic relationships between words. Words or tokens with similar meanings (for example, "car" and "automobile") are represented by vectors that are close together in the multi-dimensional space, while unrelated words (like "tree" and "laptop") are farther apart. This vector representation enables the model to understand context, similarity, and relationships between different pieces of text.

Embeddings are fundamental in tasks such as:

- * Semantic search: Finding documents or sentences with similar meaning.
- * Clustering: Grouping related concepts together.
- * Recommendation systems: Suggesting similar content based on text meaning.
- * Contextual understanding: Helping generative models produce coherent and context-aware text.

Option review:

- * Attention: A mechanism used within transformers to focus on relevant parts of input sequences but not a representation of words.
- * A completion: Refers to the generated text output from a model, not the internal representation.
- * A transformer: The architecture that powers models like GPT, not the vector representation of tokens.

Therefore, the correct term for a multi-dimensional vector assigned to each word or token in a large language model (LLM) is An embedding, which represents how meaning is numerically encoded and compared within language models.

質問: 82

カスタマー レビューを含む Web サイトがあります。

レビューを英語で保存し、各ユーザーの地理的な場所を認識して、それぞれの言語でレビューをユーザーに提示する必要があります。

どのタイプの自然言語処理ワークロードを使用する必要がありますか？

- A. 翻訳
- B. 言語モデリング
- C. キーフレーズ抽出
- D. 音声認識

正解: ([正解を表示します](#))

According to the Microsoft Azure AI Fundamentals (AI-900) syllabus and Microsoft Learn module "Describe features of natural language processing (NLP) workloads on Azure," translation is a core NLP workload that converts text from one language into another while maintaining meaning and context.

In this scenario, the website stores reviews in English and must present them in the user's native language based on geographical location. This directly requires a translation workload, which uses Azure Cognitive Services - specifically, the Translator service - to automatically translate content dynamically for each user.

Other options explained:

- * B. Language modeling involves predicting the next word in a sentence or understanding linguistic patterns; it's used in model training, not translation.
- * C. Key phrase extraction identifies main ideas in text, not language conversion.
- * D. Speech recognition converts spoken words into written text but does not perform translation or handle geographic adaptation.

Microsoft's Translator service supports real-time text translation, multi-language detection, and context preservation, making it ideal for global websites. The AI-900 study guide emphasizes translation as one of the most common NLP workloads, enabling applications to break language barriers and enhance accessibility for diverse audiences.

Therefore, based on official Microsoft Learn material, the correct answer is:

A. translation.

質問: 83

あなたは、小売店からの画像を処理し、競合他社の製品を識別するツールを構築しています。

ソリューションは、会社から提供された画像を使用してトレーニングする必要があります。

どの Azure AI サービスを使用する必要がありますか？

- A. アズール アル カスタム ビジョン
- B. Azure AI コンピューター ビジョン
- C. 顔
- D. Azure AI Document Intelligence

正解: ([正解を表示します](#))

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and Microsoft Learn documentation, Azure AI Custom Vision is specifically designed for training custom image classification and object detection models using images that a company provides. In this scenario, the company wants to identify competitor products from images captured in retail stores - a classic use case for custom image classification or object detection, depending on whether you are labeling entire images or identifying multiple items within an image.

Azure AI Custom Vision allows users to:

- * Upload their own labeled training images.
- * Train a model that learns to recognize specific objects (in this case, competitor products).
- * Evaluate, iterate, and deploy the model as an API endpoint for real-time inference.

This fits perfectly with the requirement that the solution "must be trained on images provided by your company." The key phrase here indicates the need for a custom-trained model rather than a prebuilt one.


The other options are not suitable for this scenario:

- * B. Azure AI Computer Vision provides prebuilt models for general-purpose image understanding (e.g., detecting common objects, reading text, describing scenes). It is not intended for training on custom datasets.
- * C. Face service is limited to detecting and recognizing human faces; it cannot be trained to identify products.
- * D. Azure AI Document Intelligence (formerly Form Recognizer) is focused on extracting structured data from documents and forms, not analyzing retail images.

Therefore, per Microsoft's official AI-900 training content, when a solution must be trained on custom company images to recognize specific products, the appropriate service is Azure AI Custom Vision.

質問: **84**

画像をComputerVision APIに送信し、展示に示されている注釈付きの画像を受け取ります。



どのタイプのコンピュータービジョンが使用されましたか？

- A. 物体検出
- B. セマンティックセグメンテーション
- C. 光学式文字認識 (OCR)
- D. 画像分類

正解: ([正解を表示します](#))

Object detection is similar to tagging, but the API returns the bounding box coordinates (in pixels) for each object found. For example, if an image contains a dog, cat and person, the Detect operation will list those objects together with their coordinates in the image. You can use this functionality to process the relationships between the objects in an image. It also lets you determine whether there are multiple instances of the same tag in an image.

The Detect API applies tags based on the objects or living things identified in the image. There is currently no formal relationship between the tagging taxonomy and the object detection taxonomy. At a conceptual level, the Detect API only finds objects and living things, while the Tag API can also include contextual terms like "indoor", which can't be localized with bounding boxes.

Reference:

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

質問: 85

スプレッドシートで領収書をトランザクションに変換する必要があります。スプレッドシートには、取引の日付、販売者が支払った総額、および支払った税金を含める必要があります。どの Azure AI サービスを使用する必要がありますか？

- A. Face
- B. アズールアル言語
- C. Azure AI Document Intelligence
- D. アズール アル カスタム ビジョン

正解: ([正解を表示します](#))

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.

質問: 86

文を完成させるには、回答領域で適切なオプションを選択します。



正解:

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:

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.



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>

質問: 87

記事のテキストに基づいてイラストを作成するには、どのタイプの Azure AI ワークロードを使用すればよいですか？

A. 自然言語処理

B. コンピュータビジョン

C. Foundry Tools の Azure Document Intelligence

D. 生成AI

正解: ([正解を表示します](#))

質問: 88

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area

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

正解:

Answer Area

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

Explanation:

Answer Area

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

Box 1: Yes

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

Box 2: No

Box 3: Yes

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

Box 4: No

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

The label is the column you want to predict.

Reference:

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

質問: 89

法的文書から当事者と管轄区域を抽出するには、どの Azure AI Document Intelligence 事前構築済みモデルを使用すればよいですか？

- A. 契約
- B. レイアウト
- C. 一般的な文書
- D. 読む

正解: A ([コメントを発表する](#))

Within Azure AI Document Intelligence (formerly Form Recognizer), the Contract prebuilt model is designed to extract key information from legal and business contracts, including parties, jurisdictions, dates, and terms.

According to Microsoft Learn, this prebuilt model identifies structured entities such as contracting parties, effective dates, governing jurisdictions, and termination clauses.

* Layout (B) extracts text, tables, and structure but does not identify semantic information such as parties or jurisdictions.

* General document (C) extracts key-value pairs and entities but lacks domain-specific contract analysis.

* Read (D) performs OCR (optical character recognition) to extract raw text but not contextual metadata.

Thus, when the requirement is to extract parties and jurisdictions from a legal document, the Contract model is the correct Azure AI Document Intelligence choice.

質問: 90

ある会社は、顧客サービスエージェントのチームを雇用して、顧客に電話と電子メールのサポートを提供しています。

同社は、一般的な顧客の質問に自動回答を提供するWebチャットボットを開発しています。

Webチャットボットソリューションを作成した結果、会社はどのビジネス上のメリットを期待する必要がありますか？

- A. 売上の増加
- B. カスタマーサービスエージェントの作業負荷の軽減
- C. 製品の信頼性の向上

正解: ([正解を表示します](#))

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.

質問: 91

電子メール、Microsoft Teams、Webチャットなどの複数のチャネルを介してユーザーと通信する会話型AIソリューションを開発しています。

どのサービスを使うべきですか？

A. テキスト分析

B. Azureボットサービス

C. Translator

D. フォーム認識機能

正解: ([正解を表示します](#))

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.

有効的な**AI-900J**問題集はJPNTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTest.comは今最新**AI-900J**試験問題集を提供します。JPNTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **836**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」

質問: 92

生成 AI ソリューションによって憎悪的な応答が返されるのを防ぐには、何を実装する必要がありますか？

- A. 不正使用監視
- B. コンテンツフィルタリング
- C. 微調整
- D. 迅速なエンジニアリング

正解: **B** ([コメントを发表する](#))

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.

質問: **93**

ユーザーのプロンプトに基づいてイメージを生成する必要があります。どの Azure OpenAI モデルを使用する必要がありますか？

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

正解: ([正解を表示します](#))

According to the Microsoft Azure OpenAI Service documentation and AI-900 official study materials, the DALL-E model is specifically designed to generate and edit images from natural language prompts. When a user provides a descriptive text input such as "a futuristic city skyline at sunset", DALL-E interprets the textual prompt and produces an image that visually represents the content described. This functionality is known as text-to-image generation and is one of the creative AI capabilities supported by Azure OpenAI.

DALL-E belongs to the family of generative models that can create new visual content, expand existing images, or apply transformations to images based on textual instructions. Within Azure OpenAI, the DALL-E API enables developers to integrate image creation directly into applications-useful for design assistance, marketing content generation, or visualization tools. The model learns from vast datasets of text-image pairs and is optimized to ensure alignment, diversity, and accuracy in the produced visuals.

By contrast, the other options serve different purposes:

- * A. GPT-4 is a large language model for text-based generation, reasoning, and conversation, not for creating images.
- * C. GPT-3 is an earlier text generation model, primarily used for language tasks like summarization, classification, and question answering.
- * D. Whisper is an automatic speech recognition (ASR) model used to convert spoken language into written text; it has no image-generation capability.

Therefore, when the requirement is to generate images based on user prompts, the only Azure OpenAI model that fulfills this purpose is DALL-E. This aligns directly with the AI-900 learning objective covering Azure OpenAI generative capabilities for text, code, and image creation.

質問: 94

文を完成させるには、回答領域で適切なオプションを選択します。

Answer Area

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

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

正解:

Answer Area

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

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

Explanation:

Answer Area

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

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

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Identify features and uses of speech capabilities", speech recognition refers to the process of converting spoken words into written text. When a speaker's voice is transcribed into subtitles during a presentation, the system listens to the audio input, identifies the spoken words, and generates corresponding text in real time. This is precisely what speech recognition technology accomplishes.

Azure provides this functionality through the Azure Speech Service, which supports multiple speech-related features:

- * Speech-to-Text (Speech Recognition) - Converts spoken audio into text.
- * Text-to-Speech (Speech Synthesis) - Converts written text into spoken audio.
- * Speech Translation - Translates spoken words into another language.

In this case, the session is transcribed into subtitles in the same language, not translated or spoken aloud, so the correct feature is Speech Recognition.

Let's review the other options:

- * Sentiment Analysis: This belongs to the Text Analytics service under natural language processing (NLP) and is used to determine the emotional tone of text, not to convert speech to text.
- * Speech Synthesis: Converts text into audible speech (Text-to-Speech), the reverse of what is happening in this scenario.
- * Translation: Converts spoken or written words from one language to another. Here, no translation is mentioned-only transcription.

Therefore, the described process-turning live spoken language into readable subtitles-is an example of Speech Recognition, a speech-to-text AI capability provided by Azure Cognitive Services.

Final answer: Speech recognition

Reference:Microsoft Learn - Identify speech capabilities of Azure AI services (AI-900 Learning Path)

質問: 95

文を完成させるには、回答領域で適切なオプションを選択します。

Answer Area

Natural language processing can be used to

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

正解:

Answer Area

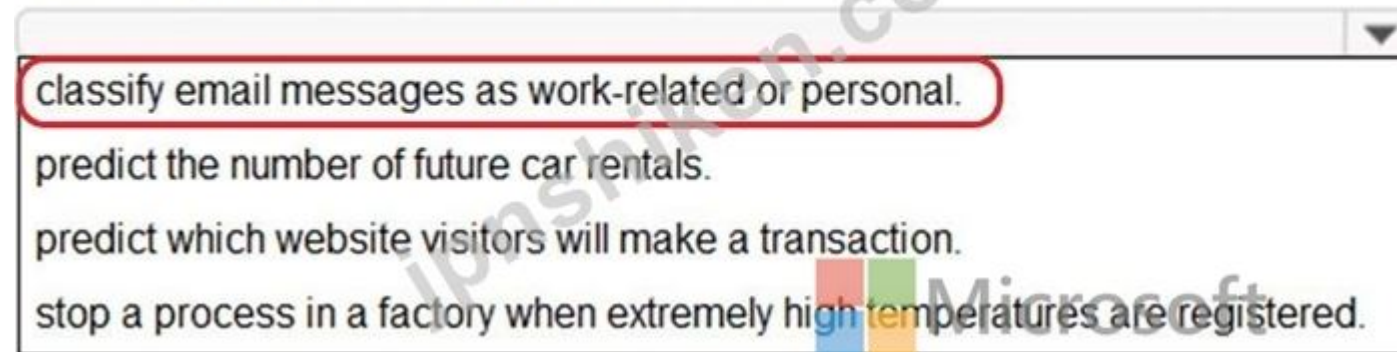
Natural language processing can be used to

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

Explanation:

Answer Area

Natural language processing can be used to



classify email messages as work-related or personal.

predict the number of future car rentals.

predict which website visitors will make a transaction.

stop a process in a factory when extremely high temperatures are registered.

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore natural language processing (NLP) in Azure", Natural Language Processing (NLP) is a branch of artificial intelligence that focuses on enabling computers to understand, interpret, and generate human language. NLP is used to extract meaning and intent from text or speech, perform sentiment analysis, identify entities, and classify content based on context.

One of the primary applications of NLP is text classification, where an AI model automatically categorizes text documents or messages into predefined classes. Classifying emails as work-related or personal is a textbook example of this NLP capability. It involves analyzing the words, phrases, and structure of the text to determine the email's category. Microsoft Learn highlights this type of problem as document classification, an essential NLP use case often implemented through Azure Cognitive Services such as Text Analytics or Language Studio.

Let's examine why the other options are incorrect:

- * Predict the number of future car rentals - This is a time series forecasting or regression task, not NLP.
- * Predict which website visitors will make a transaction - This is a predictive analytics or machine learning classification problem based on behavioral data, not language understanding.
- * Stop a process in a factory when extremely high temperatures are registered - This relates to IoT automation or sensor-based anomaly detection, not NLP.

Therefore, based on Microsoft's AI-900 materials, Natural Language Processing is best used for tasks involving understanding and classifying text, such as classifying email messages as work-related or personal.

This example perfectly aligns with NLP's goal-to enable machines to process and derive insights from human language inputs.

質問: 96

口頭のコマンドを使用してスマートデバイスを制御する機能をユーザーに提供するAIソリューションがあります。

ソリューションが使用する2種類の自然言語処理 (NLP) ワークロードはどれですか？それぞれの正解は、解決策の一部を示しています。

注：正しい選択はそれぞれ1ポイントの価値があります。

- A. テキスト読み上げ
- B. 翻訳
- C. 言語モデリング
- D. キーフレーズ抽出
- E. 音声認識

正解: C,E ([コメントを發表する](#))

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.

質問: **97**

質問に答えるスマートデバイス。Contoso, Ltd.の株価はいくらですか?」はどの AI ワークロードの一例ですか?

A. コンピュータ ビジョン

B. 異常検出

C. ナレッジマイニング

D. 自然言語処理

正解: **D** ([コメントを发表する](#))

The question describes a smart device that can understand and respond to a spoken or written question such as, "What is the stock price of Contoso, Ltd.?" This scenario directly maps to the Natural Language Processing (NLP) workload in Microsoft Azure AI.

According to the Microsoft AI Fundamentals (AI-900) study guide and the Microsoft Learn module "Describe features of common AI workloads," NLP enables systems to understand, interpret, and generate human language. Azure AI Language and Azure Speech services are examples of NLP-based solutions.

In this case, the smart device performs several NLP tasks:

* Speech recognition - converts spoken input into text.

* Language understanding - interprets the user's intent, i.e., retrieving the stock price of a specific company.

* Response generation - formulates a meaningful answer that can be presented back as text or speech.

This process shows a full pipeline of natural language understanding (NLU) and conversational AI. It does not involve visual data (computer vision), data pattern analysis (anomaly detection), or document search (knowledge mining).

Hence, the correct AI workload is D. Natural Language Processing.

質問: **98**

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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.	☐	☐

正解:

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.

質問: 99

文を正しく完成させる答えを選択してください。



正解:



Explanation:



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.

質問: 100

Azure OpenAI サービスを呼び出すために使用できる 2 つのツールはどれですか? それぞれの正解は完全な解決策を示します。

注: 各正解は 1 ポイントの価値があります。

A. Azure コマンドライン インターフェイス (CLI)

- B. Azure REST API
- C. Python 用 Azure SDK
- D. JavaScript 用 Azure SDK

正解: ([正解を表示します](#))

The correct answers are B. Azure REST API and C. Azure SDK for Python.

The Azure OpenAI Service can be accessed using multiple development interfaces. According to Microsoft Learn documentation, developers can call the service via the Azure REST API, which provides direct HTTPS- based access to the model endpoints for tasks like completions, chat, embeddings, and image generation. This interface is platform-independent and supports integration with any system capable of making HTTP requests.

Additionally, Azure SDKs offer higher-level libraries for convenient integration into applications. The Azure SDK for Python and Azure SDK for JavaScript are both supported for Azure OpenAI interaction, allowing developers to authenticate with Azure credentials, send prompts, and receive model responses programmatically.

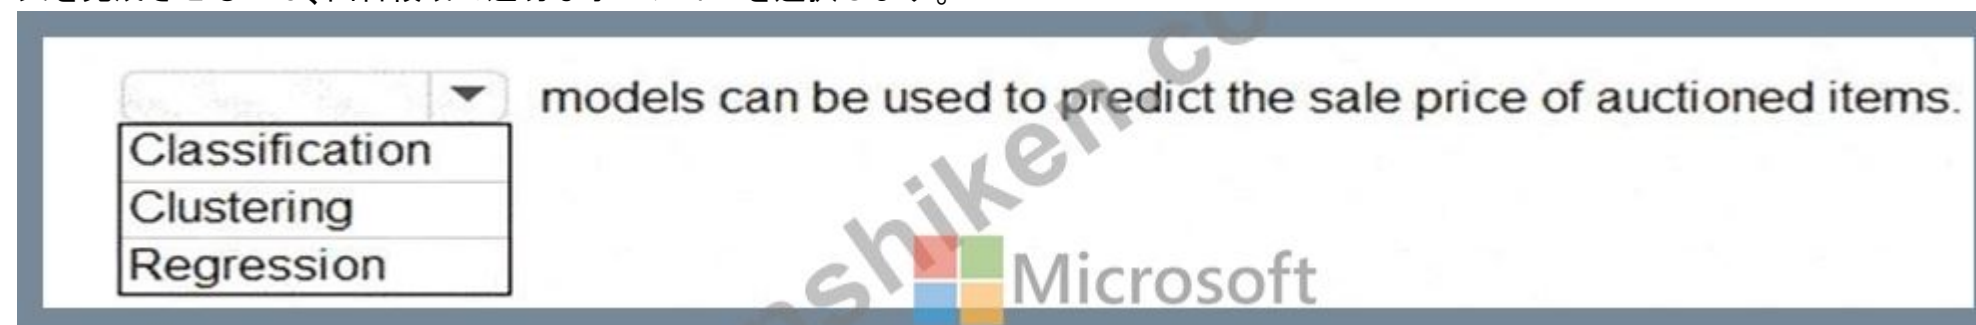
However, among the listed options, the REST API (B) and SDK for Python (C) are most explicitly referenced in the AI-900 learning modules and Microsoft documentation as standard tools to call Azure OpenAI services.

Option A (Azure CLI) is incorrect because the CLI is used primarily for provisioning and managing Azure resources, not for directly calling OpenAI model endpoints.


Therefore, based on the Azure AI-900 and OpenAI integration guidance, the correct answers are B. Azure REST API and C. Azure SDK for Python.

質問: 101

文を完成させるには、回答領域で適切なオプションを選択します。



正解:



Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Explore fundamental principles of machine learning", regression models are used to predict numerical or continuous values based on patterns found in historical data. When the goal is to forecast or estimate a real-valued outcome-such as price, temperature, sales, or age-the appropriate model type is regression.

In this question, the task is to predict the sale price of auctioned items. Since price is a continuous numeric value that can vary within a range (for example, \$100.50, \$105.75, \$120.00, etc.), it fits perfectly into a regression problem. Microsoft Learn defines regression as "a supervised machine learning technique that predicts a numeric value based on relationships found in input features." Common regression algorithms include linear regression, decision tree regression, and neural network regression.

By contrast:

- * Classification is used when the output variable represents categories or classes, such as predicting whether an email is spam or not spam, or whether a transaction is fraudulent or legitimate. Classification predicts discrete labels, not continuous values.
- * Clustering, on the other hand, is an unsupervised learning method used to group similar data points together without predefined labels. Examples include grouping customers by purchasing behavior or grouping images by visual similarity.

In a predictive business scenario, like estimating the price of an auctioned item based on features such as age, condition, and demand, regression models are most appropriate. Azure Machine Learning supports regression experiments using built-in algorithms and AutoML to automatically choose the best-performing model for continuous output prediction.

質問: 102
文を完成させるには、回答領域で適切なオプションを選択します。

Answer Area

Feature engineering

Feature selection

Model evaluation

Model training

is used to generate additional features.

正解:

Answer Area

Feature engineering


Feature selection

Model evaluation

Model training

is used to generate additional features

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.

質問: 103

文を正しく完成させる答えを選択してください。



正解:



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

質問: 104

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area		Statements	Yes	No
		A restaurant can use a chatbot to answer queries through Cortana.	☐	☐
		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.	☐	☐

正解:

Answer Area		Statements	Yes	No
		A restaurant can use a chatbot to answer queries through Cortana.	☒	☐
		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, No.

According to the Microsoft Azure AI Fundamentals (AI-900) study materials, conversational AI enables applications, websites, and digital assistants to interact with users via natural language. A chatbot is a key conversational AI workload and can be integrated into multiple channels such as web pages, Microsoft Teams, Facebook Messenger, and Cortana using Azure Bot Service and Bot Framework.

* "A restaurant can use a chatbot to answer queries through Cortana" - Yes. Azure Bot Service supports multi-channel deployment, which includes Cortana integration. This means the same bot can respond to voice or text input via Cortana, making it a valid use case for a restaurant to provide menu details, reservations, or order tracking through voice-based AI assistants.

* "A restaurant can use a chatbot to answer inquiries about business hours from a webpage" - Yes. This is a standard scenario for chatbots embedded on a company website. As per Microsoft Learn's Describe features of conversational AI module, a chatbot can be added to a website to handle FAQs such as business hours, location, or menu details, thereby improving response time and reducing repetitive human workload.

* "A restaurant can use a chatbot to automate responses to customer reviews on an external website" - No. Azure bots and other conversational AI tools cannot automatically interact with or post on external third-party platforms where the business does not control the data or API integration. Automated posting or replying to reviews on external review sites (e.g., Yelp or Google Reviews) would violate both ethical and technical boundaries of responsible AI usage outlined by Microsoft.

質問: 105

文を正しく完成させる答えを選択してください。



正解:



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).

質問: 106

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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.	☐	☐

正解:

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:

Yes, Yes, and No.

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and the Microsoft Learn module "Identify features of natural language processing (NLP) workloads on Azure", the Azure Translator service is a cloud-based AI service within Azure Cognitive Services that provides real-time text translation across multiple languages.

* "You can use the Translator service to translate text between languages." - Yes. This is the core function of the Translator service. It takes text as input in one language and returns it in another using advanced neural machine translation models. This aligns with the AI-900 learning objective: "Describe the capabilities of Azure Cognitive Services for language", which specifically names Azure Translator as the service used to perform automatic text translation. The service supports over 100 languages and dialects, offering both single-sentence and document-level translations.

* "You can use the Translator service to detect the language of a given text." - Yes. This statement is also true. The Translator service automatically detects the source language if it is not specified in the request. This feature is documented in the Azure Translator API, where the system identifies the input language before performing translation. The AI-900 exam content emphasizes this as one of the Translator service's built-in capabilities-language detection for untagged text.

* "You can use the Translator service to transcribe audible speech into text." - No. This is not a function of Translator. Transcription (converting speech to text) is a speech AI workload, handled by the Azure Speech Service, not Translator. The Speech-to-Text capability in Azure Cognitive Services processes spoken audio input and returns the text transcription. The Translator service only works with text input, not direct audio.

Therefore, based on official AI-900 guidance, the verified configuration is:

Yes - for text translation

Yes - for language detection

No - for speech transcription.

This aligns precisely with the AI-900 learning outcomes describing Text Translation and Language Detection as Translator capabilities, and Speech Transcription as part of the separate Speech service.

有効的な**AI-900J**問題集はJPNTTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AI-900J**試験問題集を提供します。JPNTTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **336**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」

質問: 107

Azure OpenAI DALL-E モデルを使用して実行できる 2 つのアクションはどれですか。それぞれの正解は完全なソリューションを示します。

注意: 正解ごとに 1 ポイントが付与されます。

- A. イメージを作成します。
- B. 光学式文字認識 (OCR) を使用します。
- C. 画像内のオブジェクトを検出します。
- D. 画像を変更します。

E. 画像のキャプションを生成します。

正解: [\(正解を表示します\)](#)

The correct answers are A. Create images and D. Modify images.

The Azure OpenAI DALL-E model is a text-to-image generative AI model that can create original images and modify existing ones based on text prompts. According to Microsoft Learn and Azure OpenAI documentation, DALL-E interprets natural language descriptions to produce unique and creative visual content, making it useful for design, illustration, marketing, and educational applications.

* Create images (A) - DALL-E can generate new images entirely from textual input. For example, the prompt "a futuristic city skyline at sunrise" would result in a custom-generated artwork that visually represents that description.

* Modify images (D) - DALL-E also supports inpainting and outpainting, allowing users to edit or expand existing images. You can replace parts of an image (for example, changing a background or object) or add new elements consistent with the visual style of the original.

The remaining options are incorrect:

* B. OCR is performed by Azure AI Vision, not DALL-E.

* C. Detect objects in images is also an Azure AI Vision (Image Analysis) feature.

* E. Generate captions for images is handled by Azure AI Vision, not DALL-E, since DALL-E generates-not interprets-visuals.

質問: 108

次のデータセットがあります。

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

データセットを使用して、住宅の住宅価格カテゴリを予測するモデルをトレーニングすることを計画しています。

世帯収入と住宅価格のカテゴリとは何ですか？回答するには、回答領域で適切なオプションを選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

Answer Area

Household Income:

A feature

A label

House Price Category:

A feature

A label

正解:

Answer Area

Household Income:

House Price Category:

Explanation:



Answer Area

Household Income:

House Price Category:

In machine learning, especially within the Microsoft Azure AI Fundamentals (AI-900) framework, datasets used for supervised learning are composed of features (inputs) and labels (outputs). According to the Microsoft Learn module "Explore the machine learning process", a feature is any measurable property or attribute used by the model to make predictions, whereas a label is the actual value or category the model is trying to predict.

* Household Income # Feature A feature (also known as an independent variable) represents the input data that the machine learning algorithm uses to detect patterns or correlations. In this dataset, Household Income is a numeric value that influences the prediction of house price categories. During training, the model learns how variations in household income correlate with changes in the house price category.

Microsoft Learn defines features as "the attributes or measurable inputs that are used to train the model." Thus, Household Income serves as a predictive input or feature.

* House Price Category # LabelThe label (or dependent variable) represents the output the model aims to predict. It is the known result during training that helps the algorithm learn correct mappings between features and outcomes. In this scenario, House Price Category-which can take values such as "Low," "Middle," or "High"-is the classification outcome that the model will predict based on household income (and possibly other variables). According to Microsoft Learn, "the label is the variable that contains the known values that the model is trained to predict." In summary, the dataset defines a supervised learning classification problem, where Household Income is the feature (input) and House Price Category is the label (output) that the model will learn to predict.

質問: 109

クライアントアプリケーションで使用されるサービスとして、Azure MachineLearningモデルをデプロイすることを計画しています。
モデルをデプロイする前に、どの3つのプロセスを順番に実行する必要がありますか？回答するには、適切なプロセスをプロセスのリストから回答領域に移動し、正しい順序で配置します。

Processes

data encryption

model retraining

model training


data preparation

model evaluation


Answer Area

正解:

Processes



Answer Area



Explanation:



The correct order of processes before deploying a model as a service is:

(1) Data preparation # (2) Model training # (3) Model evaluation.

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Explore the machine learning process", machine learning follows a structured lifecycle that involves several sequential stages. Before a model can be deployed, the data must be properly prepared, the model must be trained, and then its performance must be evaluated to ensure accuracy and reliability.

* **Data Preparation:** The first stage involves collecting, cleaning, and transforming raw data into a usable format. Azure Machine Learning provides tools like Data Wrangler, Data Labeling, and Data Transformation pipelines to ensure the dataset is accurate and consistent. As per Microsoft Learn, "data preparation is essential to remove noise, handle missing values, and split the dataset into training and testing sets." This step ensures the model learns from quality input.

* **Model Training:** In this step, algorithms are applied to the prepared training data to create a predictive model. The system learns patterns and relationships from the data. Azure Machine Learning allows model training using AutoML, custom code, or designer pipelines. The training process produces a model that can make predictions, but it still needs to be tested before deployment.

* **Model Evaluation:** Once trained, the model's performance is tested against unseen (test) data.

Evaluation metrics like accuracy, precision, recall, and F1-score are analyzed to verify if the model meets business and performance requirements. Microsoft Learn defines this stage as "assessing the model's performance to determine its readiness for deployment." After these three processes, the model can then be deployed as a web service using Azure Machine Learning endpoints. Model retraining happens later when new data becomes available, and data encryption is a security process, not part of model development steps.

質問: 110

データ分割モジュールを含むAzureMachineLearningパイプラインがあります。Split Dataモジュールは、TrainModelモジュールとScoreModelモジュールに出力します。データ分割モジュールの機能は何ですか？

- A. モデルに含める必要のある列を選択する
- B. トレーニングおよび検証データセットの作成
- C. データが欠落しているレコードを迂回させる
- D. 一貫した数値範囲内に収まるように数値変数をスケーリングする

正解: ([正解を表示します](#))

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.

質問: 111

モデルのトレーニングに使用されたデータのソースを明記することは、どの責任ある AI 原則の例ですか？

- A. 公平性
- B. 透明性
- C. 信頼性と安全性
- D. プライバシーとセキュリティ

正解: ([正解を表示します](#))

According to Microsoft's Responsible AI Principles, Transparency means that AI systems should clearly communicate how they operate, including data sources, limitations, and decision-making processes. Stating the source of data used to train a model helps users understand where the model's knowledge comes from, enabling informed trust and accountability.

Transparency ensures that organizations disclose relevant details about data collection and model design, especially for compliance, fairness, and reproducibility.

Other options are incorrect:

- * A. Fairness: Focuses on avoiding bias and ensuring equitable outcomes.

- * C. Reliability and safety: Ensures AI performs consistently and safely.
 - * D. Privacy and security: Protects user data and maintains confidentiality.
- Thus, the principle illustrated by disclosing training data sources is Transparency.

質問: 112

コンピュータビジョンの種類を適切なシナリオに一致させます。

答えるには、適切なワークロードタイプを左側の列から右側のシナリオにドラッグします。各ワークロードタイプは、1回使用することも、複数回使用することも、まったく使用しないこともできます。

注：正しい選択はそれぞれ1ポイントの価値があります。

Workloads Types

Facial recognition

Image classification

Object detection

Optical character recognition (OCR)

Answer Area

Workload TypeIdentify celebrities in images.

Workload TypeExtract movie title names from movie poster images.

Workload TypeLocate vehicles in images.

正解:

Workloads Types

Facial recognition

Image classification

Object detection

Optical character recognition (OCR)

Answer Area

Facial recognitionIdentify celebrities in images.

Optical character recognition (OCR)Extract movie title names from movie poster images.

Object detectionLocate vehicles in images.

Explanation:

The correct mappings are based on the Microsoft Azure AI Fundamentals (AI-900) curriculum topic: "Describe features of computer vision workloads." Microsoft divides computer vision tasks into key workload types - image classification, object detection, facial recognition, and optical character recognition (OCR) - each designed for specific visual analysis objectives.

* Identify celebrities in images # Facial recognitionFacial recognition goes beyond simple face detection; it can identify or verify specific individuals by comparing facial features with known profiles. According to Microsoft Learn, the Face service in Azure Cognitive Services can detect, recognize, and identify people in photos or videos. Recognizing celebrities or known individuals is a prime example of facial recognition.

* Extract movie title names from movie poster images # Optical Character Recognition (OCR)OCR is used to detect and extract text content from images, such as printed or handwritten words. Azure's Computer Vision API uses OCR technology to read text in various languages from photos, scanned documents, or posters. Therefore, extracting movie titles or actor names from a poster image is a perfect use case for OCR.

* Locate vehicles in images # Object detectionObject detection identifies and locates specific objects within an image, returning bounding boxes that indicate their positions. In Azure, the Custom Vision service or Computer Vision object detection models are used to detect multiple objects like vehicles, pedestrians, or animals in a single image.

Summary:

- * Facial recognition # Identifies specific people (celebrities)
- * OCR # Extracts text (movie titles)
- * Object detection # Finds and locates physical items (vehicles)

Thus, the verified and official answer is:

質問: 113

電子商取引ビジネス用の会話型言語理解モデルを構築しました。

発話がモデルの意図した範囲外にある場合に、それをモデルが検出できるようにする必要があります。

何をすべきでしょうか？

- A. モデルをエクスポートします。
- B. 新しいモデルを作成します。
- C. None インテントに発話を追加します。
- D. 事前に構築されたタスク エンティティを作成します。

正解: ([正解を表示します](#))

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

質問: 114


AI ソリューションを適切なタスクに一致させます。

回答するには、左側の列から適切なソリューションを右側のタスクにドラッグします。各ソリューションは、1 回、複数回、またはまったく使用されない場合があります。


注意: 正解ごとに 1 ポイントが付与されます。

The screenshot shows a matching exercise interface. On the left, under the heading "Solutions", there are four boxes: "Computer vision", "Data mining", "Generative AI", and "Text analytics". On the right, under the heading "Answer Area", there are three tasks: "Generate a caption from a given image.", "Generate an image from a given caption.", and "Generate a 200-word summary from a 2,000-word article." A Microsoft logo is visible in the background.

正解:



Explanation:



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 Vision This 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 AI This 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 Analytics Text 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.

質問: 115

コンピューター ビジョン サービスを適切な AI ワークロードに合わせます。

答えるには、左側の列から適切なサービスを右側のワークロードにドラッグします。各サービスは、1 回だけ使用することも、複数回使用することも、まったく使用しないこともできます。

注: 正しく一致するたびに 1 ポイントの価値があります。

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.

正解:

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.

質問: 116

Azure MachineLearningでベーシックワークスペースとエンタープライズワークスペースのどちらを使用するかを評価しています。

エンタープライズワークスペースを必要とする2つのタスクは何ですか？それぞれの正解は完全な解決策を提示します。

注：正しい選択はそれぞれ1ポイントの価値があります。

- A. グラフィカルユーザーインターフェイス (GUI)を使用して、自動機械学習実験を実行します。
- B. ワークステーションとして使用するコンピューティングインスタンスを作成します。
- C. グラフィカルユーザーインターフェイス (GUI)を使用して、Azure MachineLearningデザイナーから機械学習実験を定義および実行します。
- D. コンマ区切り値 (CSV) ファイルからデータセットを作成します。

正解: ([正解を表示します](#))

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

質問: 117

テキストとして保存されている保険金請求レポートがあります。

要約を生成するには、レポートから重要な用語を抽出する必要があります。

どのタイプのAIワークロードを使用する必要がありますか？

- A. 会話型アル
- B. 異常検出
- C. 自然言語処理
- D. コンピュータービジョン

正解: ([正解を表示します](#))

According to the AI-900 study guide and Microsoft Learn module "Identify features of natural language processing workloads", Natural Language Processing (NLP) enables computers to understand, interpret, and generate human language. Tasks such as extracting key terms, summarizing documents, identifying topics, or determining sentiment fall under NLP workloads.

In this question, you have insurance claim reports stored as text, and you need to extract key terms to generate summaries. This matches the Text Analytics service in Azure Cognitive Services, which uses NLP techniques such as key phrase extraction to identify important concepts within textual data.

The other options are incorrect because:

* A. Conversational AI focuses on chatbots or dialogue systems.

* B. Anomaly detection identifies unusual data patterns, not textual meaning.

* D. Computer vision processes image or video content, not text.

Therefore, extracting key terms from documents is a clear example of Natural Language Processing.

Reference:Microsoft Learn - Identify features of Natural Language Processing (NLP) workloads

質問: 118

スーパーマーケットの棚の画像に含まれる製品のブランド名を識別するボットがあります。

ボットはどのサービスを使用しますか？

A. Azure Search 機能の強化

B. コンピュータビジョン画像解析機能

C. カスタム ビジョン画像分類機能

D. 言語理解能力

正解: ([正解を表示します](#))

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

質問: 119

Azure OpenAI 大規模言語モデル (LLM) プロセスを適切なタスクに一致させます。

答えるには、適切なプロセスを左側の列から右側のタスクにドラッグします。各プロセスは 1 回だけ使用することも、複数回使用することも、まったく使用しないこともできます。

注: 正しく一致するたびに 1 ポイントの価値があります。

Processes

Classifying

Generating

Summarizing

Translating

Answer Area

Detect the genre of a work of fiction.

Create a list of bullet points based on text input.

Create advertising slogans from a product description.

正解:

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:



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.

質問: 120

ある地域における特定の動物種の個体数を予測する必要があります。

どのタイプの Azure Machine Learning を使用すればよいですか？

A. クラスタリング

B. 回帰

C. 分類

正解: ([正解を表示します](#))

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.

質問: 121

Azure AI Foundry を使用して AI アプリケーションを作成する予定です。このソリューションは専用の仮想マシンにデプロイされます。どのデプロイ オプションを使用すべきでしょうか？

- A. マネージドコンピューティング
- B. Azure仮想マシン
- C. Azure Kubernetes Service (AKS) クラスタ
- D. サーバーレスAPI

正解: (正解を表示します)

有効的なAI-900J問題集はJPNTest.com提供され、AI-900J試験に合格することに役に立ちます！JPNTest.comは今最新AI-900J試験問題集を提供します。JPNTest.com AI-900J試験問題集はもう更新されました。ここでAI-900J問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> 836問、30%ディスカウント、特別な割引コード: JPNshiken」

質問: 122

Azure AI サービスを適切な生成 AI 機能と一致させます。

回答するには、左側の列から適切なサービスを右側の機能にドラッグしてください。各サービスは、1回、複数回、またはまったく使用されない場合があります。

注意: 正解ごとに 1 ポイントが加算されます。

Services

Azure OpenAI Service

Azure AI Speech

Azure AI Vision

Microsoft

Answer Area

Classify and label images.

Generate conversational responses.

Convert speech to text in real time.

正解:

Services

Azure OpenAI Service

Azure AI Speech

Azure AI Vision

Microsoft

Answer Area

Azure AI Vision

Azure OpenAI Service

Azure AI Speech

Classify and label images.

Generate conversational responses.

Convert speech to text in real time.

Explanation:



This question maps each Azure AI service to its correct capability based on the Microsoft Azure AI Fundamentals (AI-900) syllabus and Microsoft Learn documentation on Azure Cognitive Services.

* Classify and label images # Azure AI Vision Azure AI Vision (formerly Computer Vision) provides capabilities to analyze visual content, detect objects, classify images, and extract information from pictures. It includes object detection, image classification, and tagging, which are core vision tasks.

This service enables businesses to build solutions that understand visual input, such as identifying products, reading signs, or detecting faces in images.

* Generate conversational responses # Azure OpenAI Service Azure OpenAI Service integrates powerful large language models such as GPT-3.5 and GPT-4, capable of generating human-like text responses, summarizations, translations, and dialogues. These models are designed for natural language generation (NLG) and conversational AI, making them ideal for chatbots, virtual agents, and intelligent assistants that produce dynamic, context-aware replies.

* Convert speech to text in real time # Azure AI Speech Azure AI Speech provides speech-to-text capabilities (speech recognition) that convert spoken language into written text instantly. It is commonly used in transcription services, voice command systems, and live captioning applications.

Additionally, the Speech service supports text-to-speech (speech synthesis) and speech translation, making it versatile for voice-based AI applications.

By understanding each service's specialization-Vision for visual data, OpenAI for generative text, and Speech for audio processing-you can correctly match the capabilities.

質問: 123

医学研究プロジェクトでは、事前定義された脳出血タイプに分類された脳スキャン画像の匿名化された大規模なデータセットを使用します。

機械学習を使用して、画像を人がレビューする前に、画像内のさまざまな脳出血タイプの早期検出をサポートする必要があります。

これは、どのタイプの機械学習の例ですか？

A. クラスタリング

B. 回帰

C. 分類

正解: C ([コメントを发表する](#))

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Identify features of classification machine learning", classification is a supervised machine learning technique used when the output variable represents discrete categories or classes. In this case, the brain scan images are already labeled into predefined haemorrhage types, such as "subarachnoid," "epidural," or

"intraventricular." The model's goal is to learn patterns from labeled examples and then predict the correct class for new, unseen images.

The use of categorized brain scan images clearly indicates a supervised learning setup because both the input (image data) and output (haemorrhage type) are known during training. This aligns with Microsoft's definition: classification problems "predict which category or class an item belongs to," often using algorithms such as logistic regression, decision trees, neural networks, or convolutional neural networks (CNNs) for image-based data.

In contrast:

* A. Clustering is an unsupervised learning approach that groups data into clusters based on similarity when no predefined labels exist.

* B. Regression predicts continuous numeric values (e.g., predicting age or temperature), not categories.

Because this project aims to automatically classify medical images into known diagnostic categories, it is a textbook example of classification.
Reference:Microsoft Learn - Identify common types of machine learning models: Regression, Classification, and Clustering (AI-900 Learning Path)

質問: 124

AI ワークロードを適切なタスクに一致させます。

回答するには、適切な AI ワークロードを左側の列から右側のタスクにドラッグします。各ワークロードは、1 回、複数回、またはまったく使用されない場合があります。

注意: 正解ごとに 1 ポイントが付与されます。

AI workloads

Azure AI Document Intelligence

Computer vision

Generative AI

Knowledge mining

Natural language processing

Answer Area

Extract data from medical admission forms for import into a patient tracking database.

Automatically create drafts for a monthly newsletter.

Analyze aerial photos to identify flooded areas.

正解:

AI workloads

Azure AI Document Intelligence

Computer vision

Generative AI

Knowledge mining

Natural language processing

Answer Area

Azure AI Document Intelligence

Generative AI


Computer vision

Extract data from medical admission forms for import into a patient tracking database.

Automatically create drafts for a monthly newsletter.

Analyze aerial photos to identify flooded areas.

Explanation:



This question tests your understanding of AI workloads as described in the Microsoft Azure AI Fundamentals (AI-900) study guide. Each Azure AI workload is designed to handle specific types of data and tasks: text, images, documents, or content generation.

* Extract data from medical admission forms for import into a patient tracking database # Azure AI Document Intelligence Formerly known as Form Recognizer, this service belongs to the Azure AI Document Intelligence workload. It extracts key-value pairs, tables, and textual information from structured and semi-structured documents such as forms, invoices, and admission sheets. For medical forms, Document Intelligence can identify fields like patient name, admission date, and diagnosis and export them into structured formats for database import.

* Automatically create drafts for a monthly newsletter # Generative AI This task involves creating original written content, which is a capability of Generative AI. Microsoft's Azure OpenAI Service uses large language models (like GPT-4) to generate human-like text, summaries, or articles. Generative AI workloads are ideal for automating creative writing, drafting newsletters, producing blogs, or summarizing reports.

* Analyze aerial photos to identify flooded areas # Computer Vision Computer Vision workloads involve analyzing and interpreting visual data from images or videos. This includes detecting objects, classifying scenes, and identifying patterns such as flooded regions in aerial imagery. Azure's Computer Vision or Custom Vision services can be trained to detect water coverage or terrain changes using image recognition techniques.

Thus, the correct matches are:

* Azure AI Document Intelligence # Extract medical form data

* Generative AI # Create newsletter drafts

* Computer Vision # Identify flooded areas from aerial photos

質問: 125

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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.	☐	☐

正解:

The screenshot shows a Microsoft quiz interface with the following content:

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

質問: 126

文を完成させるには、回答領域で適切なオプションを選択します。

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.

正解:

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.

質問: 127

文を完成させるには、回答領域で適切なオプションを選択します。

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.
--	---

正解:



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:

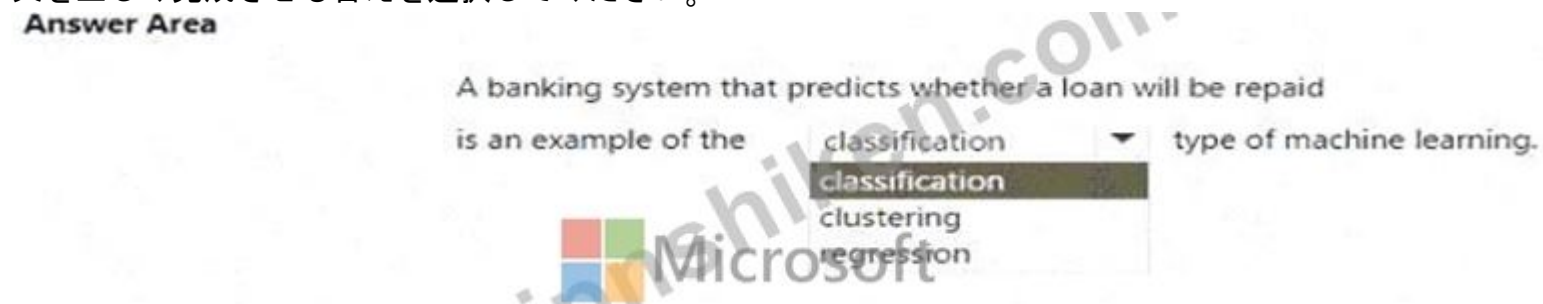
- * Image classification: Determines what object or scene is present in an image as a whole but does not locate it (e.g., "this is a car").
- * Object detection: Identifies what objects are present and where they are, usually returning coordinates for bounding boxes (e.g., "car detected at position X, 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.

質問: 128

文を正しく完成させる答えを選択してください。



正解:



Explanation:



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.

質問: 129


文を正しく完成させる答えを選択してください。



正解:



Explanation:



According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure," Optical Character Recognition (OCR) is a computer vision capability that detects and extracts printed or handwritten text from images or scanned documents and converts it into machine-readable digital text.

In this scenario, a historian wants to digitize newspaper articles - which means converting physical or scanned images of printed text into digital text for easier searching, archiving, and analysis. This is exactly the function of OCR. By using OCR, the historian can take photos or scans of old newspapers and extract the words into editable digital documents, preserving valuable historical information.

OCR is a key feature of the Azure Computer Vision service, which provides capabilities such as:

- * Extracting text from images or PDFs.
- * Reading both printed and handwritten text in multiple languages.
- * Converting physical documents into searchable digital files.

Let's examine the incorrect options:

- * Facial analysis: Detects facial features, age, gender, and emotions - unrelated to text extraction.
- * Image classification: Identifies what an image contains (e.g., "dog," "car," or "building") but doesn't extract text.
- * Object detection: Identifies and locates objects within an image using bounding boxes, not suitable for text recognition.

Therefore, to digitize newspaper articles and convert printed words into editable digital text, the correct technology to use is Optical Character Recognition (OCR), provided by the Azure Computer Vision API.

Final answer: optical character recognition (OCR)

質問: 130

パンフレットで使用する漫画を生成する必要があります。各漫画はテキストの説明に基づいています。

どの Azure OpenAI モデルを使用する必要がありますか？

- A. コーデックス
- B. ダルイー
- C. GPT-3.5
- D. GPT-4

正解: ([正解を表示します](#))

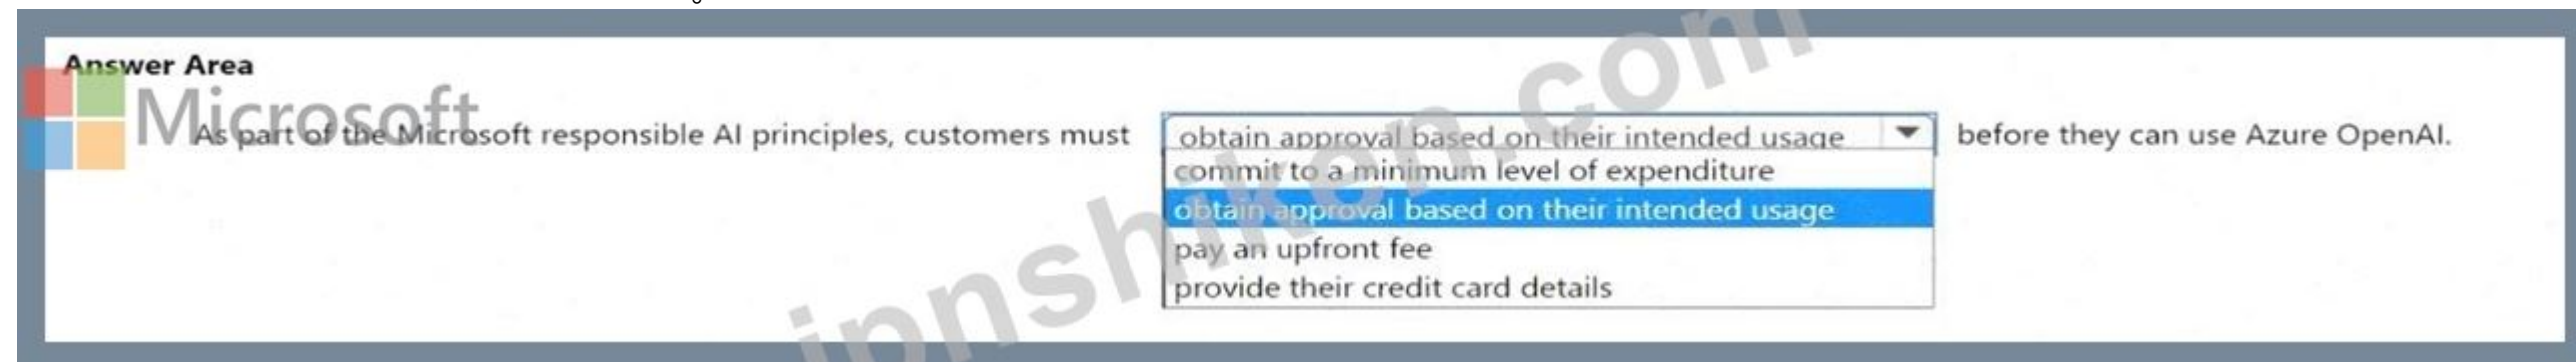
To generate cartoons or images from text descriptions, the correct Azure OpenAI model is DALL-E. As described in Microsoft's OpenAI integration documentation, DALL-E is a generative image model that converts natural language prompts into images, illustrations, and artwork.

Codex is for code generation, GPT-3.5 and GPT-4 are for text and reasoning tasks, not image creation.

Therefore, B. DALL-E is correct.

質問: 131

文を正しく完成させる答えを選択してください。

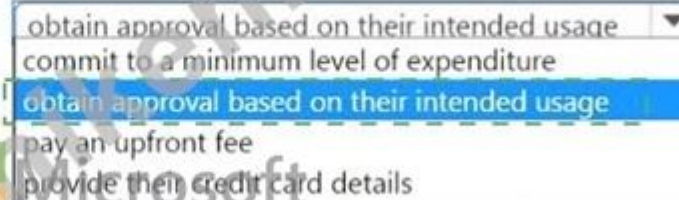


The screenshot shows a quiz interface with the Microsoft logo and the text "As part of the Microsoft responsible AI principles, customers must". A dropdown menu is open, showing five options: "obtain approval based on their intended usage", "commit to a minimum level of expenditure", "obtain approval based on their intended usage" (highlighted in blue), "pay an upfront fee", and "provide their credit card details". The text "before they can use Azure OpenAI." is visible to the right of the dropdown.

正解:

Answer Area

As part of the Microsoft responsible AI principles, customers must




before they can use Azure OpenAI.

Explanation:

Answer Area

As part of the Microsoft responsible AI principles, customers must



before they can use Azure OpenAI.

According to the Microsoft Azure AI Fundamentals (AI-900) study materials and Microsoft's Responsible AI guidelines, customers must obtain approval based on their intended usage before accessing and deploying Azure OpenAI Service. This requirement ensures that Microsoft upholds its commitment to Responsible AI principles, which include fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability.

The Azure OpenAI Service provides access to powerful language models such as GPT series and Codex, which can generate, summarize, and understand natural language and code. Because of the potential for misuse-such as generating harmful content, misinformation, or unethical automation-Microsoft enforces a use case review and approval process before granting customers access to the service. This process involves submitting an application describing the intended purpose, deployment method, and compliance measures.

Only after Microsoft validates that the proposed use aligns with responsible AI practices will access be approved.

This aligns with Microsoft's documented commitment that "customers are required to submit an application that describes their intended use of the Azure OpenAI Service," ensuring that all deployments follow ethical and legal standards. This approval step helps maintain transparency and prevent harmful or non-compliant use cases such as deepfake generation, biased automation, or malicious chatbot deployment.

Other options listed in the question are incorrect:

* Commit to a minimum level of expenditure - Microsoft does not require financial commitments for ethical approval.

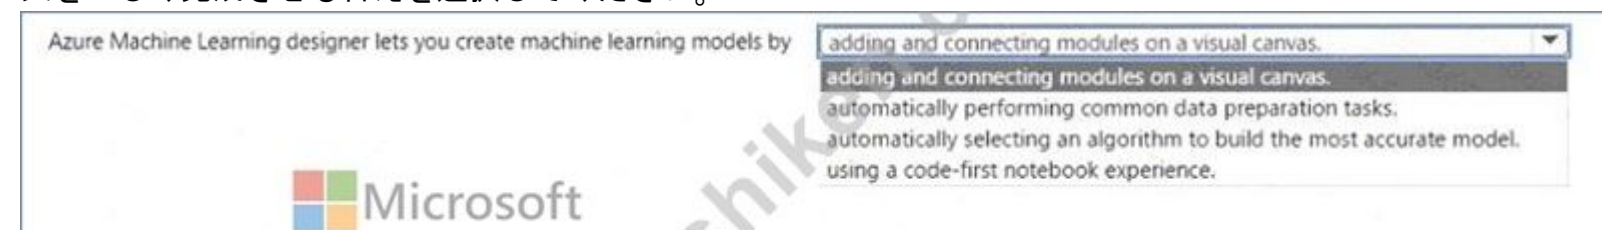
* Pay an upfront fee - Payment is handled through normal Azure billing, not a special fee.

* Provide credit card details - Not a responsible AI requirement; this is standard for any Azure subscription.

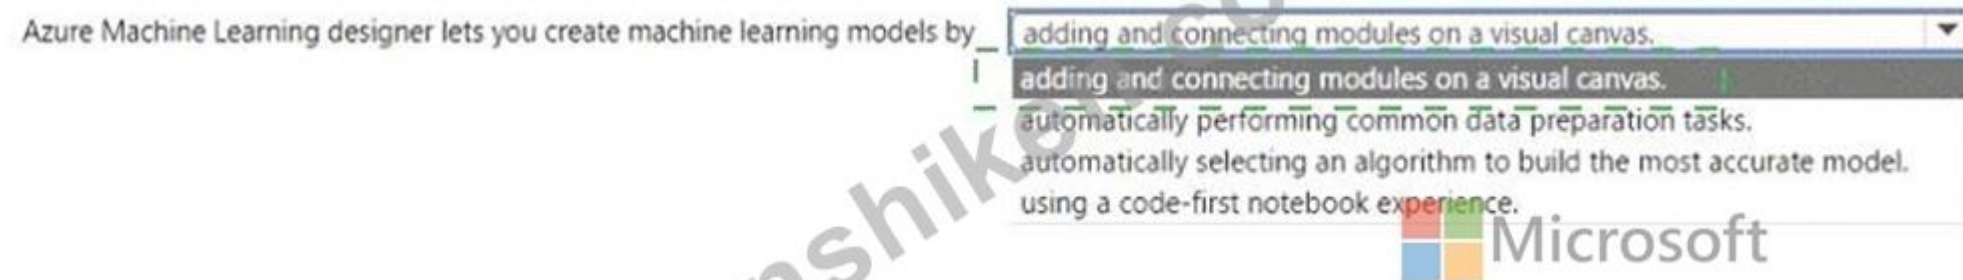
Therefore, the correct and verified answer per Microsoft's Responsible AI framework and Azure AI-900 study

質問: 132

文を正しく完成させる答えを選択してください。



正解:



Explanation:

Azure Machine Learning designer lets you create machine learning models by

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and the Microsoft Learn module "Describe Azure Machine Learning and Automated ML," Azure Machine Learning designer (formerly known as Azure Machine Learning Studio) is a drag-and-drop, low-code/no-code environment that allows users to create, train, and evaluate machine learning models visually - without the need for extensive programming knowledge.

The designer provides a visual interface, known as the canvas, where users can:

- * Import and prepare data using modules for data transformation and cleaning.
- * Split data into training and testing datasets.
- * Select and configure algorithms (classification, regression, or clustering).
- * Train and evaluate the model.
- * Deploy the model as a web service directly from the designer.

The official Microsoft Learn content emphasizes that "Azure Machine Learning designer enables users to build, test, and deploy models by adding and connecting prebuilt modules on a visual interface."

This allows business analysts, data professionals, and beginners to experiment with machine learning workflows without writing code.

By comparison:

- * Automatically performing common data preparation tasks refers to Automated ML, not the designer.
- * Automatically selecting an algorithm is also part of Automated ML, which optimizes models algorithmically.
- * Using a code-first notebook experience applies to Azure Machine Learning notebooks, intended for data scientists familiar with Python and SDKs.

Therefore, as per the AI-900 study guide and Microsoft Learn documentation, the verified and correct answer is:

Adding and connecting modules on a visual canvas, which accurately describes how Azure Machine Learning designer operates.

質問: 133

文を正しく完成させる答えを選択してください。



正解:



Explanation:



In Microsoft's Responsible AI framework, the Reliability and Safety principle ensures that AI systems perform consistently, safely, and as intended across diverse conditions - even when faced with incomplete, unusual, or unexpected data. Correctly handling unusual or missing values in a dataset directly demonstrates this principle, as it helps prevent faulty predictions, biased results, or unsafe system behaviors.

According to the Microsoft Learn Responsible AI module (from the AI-900 and AI-102 study paths), a reliable AI model should maintain its performance when encountering data anomalies. This includes validating inputs, managing missing or extreme values, and testing models to ensure they behave as expected in real-world scenarios. Such practices make AI systems robust and trustworthy, which aligns exactly with the Reliability and Safety principle.

The other Responsible AI principles address different concerns:

- * Inclusiveness: Ensures AI empowers and serves all users equitably.
- * Privacy and Security: Focuses on safeguarding personal data and preventing unauthorized access.
- * Transparency: Ensures that AI decisions are understandable and explainable to users.

While all principles are essential, managing data integrity and system stability-including how a model responds to missing or anomalous values-is primarily a matter of reliability and safety. It ensures the AI behaves predictably and minimizes risks of errors or unintended harm.

Therefore, the correct completion of the sentence is:

"Correctly handling unusual or missing values is an example of the application of the Reliability and Safety principle for Responsible AI."

質問: 134

文を正しく完成させます。



正解:

Answer Area

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



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.

質問: 135

機械学習の進歩のために、トレーニングと評価のためにデータをどのように分割する必要がありますか？

- A. トレーニング用の機能と評価用のラベルを使用します。
- B. データをトレーニング用の行と評価用の行にランダムに分割します。
- C. トレーニングにはラベルを使用し、評価には機能を使用します。
- D. データをトレーニング用の列と評価用の列にランダムに分割します。

正解: ([正解を表示します](#))

<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.

質問: 136

文を正しく完成させる答えを選択してください。



正解:



Explanation:



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

有効的な**AI-900J**問題集はJPNTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTest.comは今最新**AI-900J**試験問題集を提供します。JPNTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **336**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」

質問: 137

機械学習における回帰モデルの例は何ですか？

- A. データセット内の生徒データを生徒の年齢と学業成績に基づいて分割する
- B. ユーザーによってフラグが付けられた大量のメールを調べて、スパムメールのサブタイプを識別する
- C. 過去のデータ、家の大きさ、寝室の数に基づいて家の販売価格を予測する
- D. 画像解析による絶滅危惧動物の個体数の特定

正解: C (コメントを発表する)

rrrect 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.

質問: 138

Microsoftやサードパーティプロバイダーから提供されている、事前学習済みの生成型AIモデルを探索するには、何を使用すればよいでしょうか？

- A. 言語スタジオ
- B. Azure Machine Learning デザイナー
- C. Azure Synapse Analytics
- D. Azure AI Foundry

正解: (正解を表示します)

質問: 139

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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.	☐	☐

正解:

Microsoft	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:

Microsoft	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)

質問: 140

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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

正解:

Statements	Yes	No
A validation set includes the set of input examples that will be used to train a mode.	☐	☒
A validation set can be used to determine how well a model predicts labels.	☒	☐
A validation set can be used to verify that all the training data was used to train the model.	☐	☒

Explanation:

Statements	Yes	No
A validation set includes the set of input examples that will be used to train a mode.	☐	☒
A validation set can be used to determine how well a model predicts labels.	☒	☐
A validation set can be used to verify that all the training data was used to train the model.	☐	☒

The correct answers are based on the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore fundamental principles of machine learning." In supervised machine learning, data is typically divided into three main subsets:

- * Training set - used to train the model, i.e., to teach the algorithm the patterns and relationships between input features and output labels.
- * Validation set - used to evaluate the model during training to tune hyperparameters and prevent overfitting.
- * Test set - used after training to assess the final model's performance on unseen data.

Let's analyze each statement in light of these definitions:

* "A validation set includes the set of input examples that will be used to train a model." # NoThis is incorrect because the training set, not the validation set, contains the input examples used for model training. The validation set is separate from the training data to ensure unbiased evaluation.

* "A validation set can be used to determine how well a model predicts labels." # YesThis is correct. The validation set helps assess how effectively the model generalizes during training. It measures performance and helps tune model parameters for optimal results.

* "A validation set can be used to verify that all the training data was used to train the model." # NoThis is false. The validation set is not used to verify the completeness of training data usage. It exists independently to evaluate the model's performance during training cycles.

According to Microsoft Learn, using a validation set helps ensure that a model generalizes well and avoids overfitting to the training data. It plays a crucial role in refining and optimizing models before final testing.

質問: 141

文を完成させるには、回答領域で適切なオプションを選択します。

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.

正解:

Answer Area

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

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

Explanation:

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.



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.

質問: 142

Webチャットボットの例はどのシナリオですか？

- A. コンサートのウェブサイトに入力されたレビューが肯定的か否定的かを判断し、レビューに賛成または反対の絵文字を追加します。
- B. キオスクで顧客が入力した英語の質問に翻訳して、適切な担当者が顧客に折り返し電話できるようにします。
- C. 電子メールで質問を受け入れ、メッセージの内容に基づいて電子メールメッセージを正しい人にルーティングします。
- D. ウェブサイトのインターフェースから、音楽祭の予定されているイベントやチケット購入に関する一般的な質問に答えます。

正解: ([正解を表示します](#))

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials, a webchat bot is defined as a conversational AI application that interacts with users through a web-based chat interface. It simulates human conversation using text (and sometimes voice) to answer questions, assist with transactions, or provide information automatically. Microsoft Learn's "Describe features of

common AI workloads" module highlights conversational AI as a major AI workload, where bots and virtual agents are used to provide automated, intelligent responses in real time through web, mobile, or messaging platforms.

In this scenario, the chatbot on the festival website provides immediate answers about scheduled events and ticket purchases. This aligns exactly with how a webchat bot operates - interacting with users through a website, handling repetitive inquiries, and providing consistent information without human intervention. This type of solution is commonly built using Azure Bot Service integrated with Azure Cognitive Services for Language, which allows the bot to understand user intent and respond naturally.

Let's examine the other options to reinforce why D is correct:

* A describes a text analytics or sentiment analysis scenario, not a conversational bot, because it classifies text sentiment but doesn't "chat" with a user.

* B is an example of machine translation using the Translator service, not a chatbot.

* C is an email classification or natural language processing task, not a webchat interaction.

The AI-900 exam objectives clearly distinguish conversational AI from other cognitive services such as translation or sentiment analysis. Conversational AI focuses on dialogue and interaction through natural language conversation channels like websites or messaging apps.

Therefore, the verified and officially aligned answer is D. From a website interface, answer common questions about scheduled events and ticket purchases for a music festival.

質問: 143

スキャンした請求書を分析し、請求先住所や合計金額などのデータを抽出するには何を使用すればよいですか？

A. Azure AI 検索

B. Azure AI ドキュメント インテリジェンス

C. Azure AI カスタムビジョン

D. Azure OpenAI

正解: ([正解を表示します](#))

The correct answer is B. Azure AI Document Intelligence (formerly Form Recognizer).

This Azure service uses AI and OCR technologies to analyze and extract structured data from documents such as invoices, receipts, and purchase orders. It identifies key fields like billing address, invoice number, total amount due, and line items. The service supports prebuilt models for common document types and custom models for specialized layouts.

Option review:

* A. Azure AI Search: Used for knowledge mining and semantic search, not document data extraction.

* B. Azure AI Document Intelligence - # Correct. Designed for form and invoice extraction.

* C. Azure AI Custom Vision: Used for image classification and object detection, not text extraction.

* D. Azure OpenAI: Generates or processes language but not structured document data.

Therefore, Azure AI Document Intelligence is the right service to extract data from scanned invoices.

質問: 144

売上データと顧客タイプを示すラベルが定義されたデータセットがあるとします。売上データに基づいて顧客タイプを分類するモデルを作成する必要があります。どのタイプの機械学習を使用すべきでしょうか？

A. 回帰

B. 分類

C. クラスタリング

正解: ([正解を表示します](#))

質問: 145

自動車の販売データを含む大規模なデータセットがあります。

車両の種類に基づいて車両の販売価格を予測するには、自動機械学習 (Automated ML) モデルをトレーニングする必要があります。
どのタスクを選択すればよいですか? 回答するには、回答エリアで適切なタスクを選択してください。
注意: 正しい選択ごとに 1 ポイントが付与されます。

Answer Area 

Select task and settings
Select the machine learning task type for the experiment. To fine tune the experiment, choose additional configuration or featurization settings.

-  **Classification**
To predict one of several categories in the target column: yes/no, blue, red, green.
-  **Regression**
To predict continuous numeric values.
-  **Time series forecasting**
To predict values based on time.
-  **Natural Language Processing (preview)**
Predict based on text-only data types using multi-class or multi-label classification.
-  **Computer Vision (preview)**
Multi-class or multi-label image classification, object detection, and instance segmentation.

[View additional configuration settings.](#) [View featurization settings.](#)

正解:

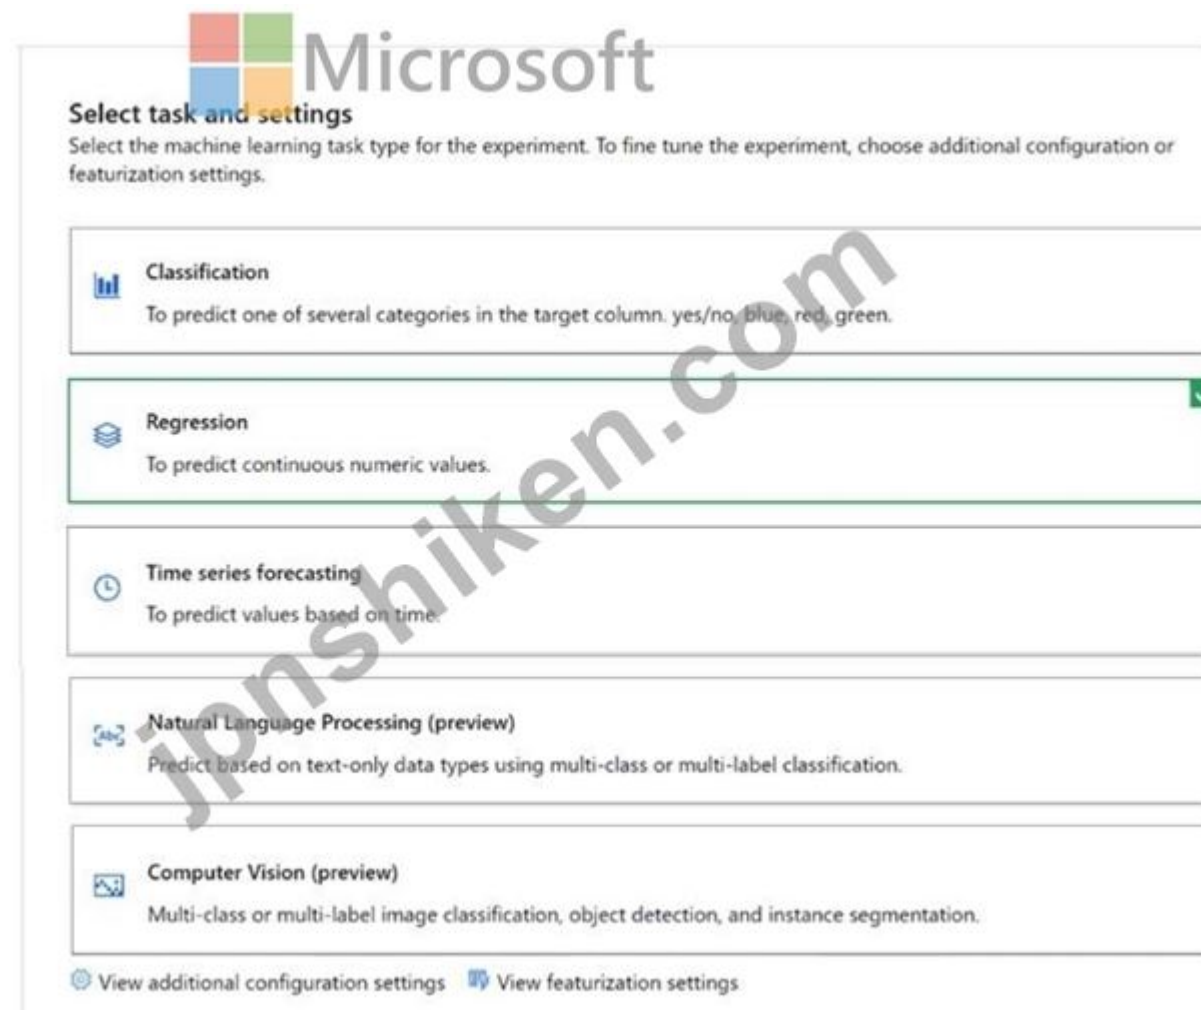
Answer Area 

Select task and settings
Select the machine learning task type for the experiment. To fine tune the experiment, choose additional configuration or featurization settings.

-  **Classification**
To predict one of several categories in the target column: yes/no, blue, red, green.
-  **Regression**
To predict continuous numeric values.
-  **Time series forecasting**
To predict values based on time.
-  **Natural Language Processing (preview)**
Predict based on text-only data types using multi-class or multi-label classification.
-  **Computer Vision (preview)**
Multi-class or multi-label image classification, object detection, and instance segmentation.

[View additional configuration settings.](#) [View featurization settings.](#)

Explanation:



Microsoft

Select task and settings
Select the machine learning task type for the experiment. To fine tune the experiment, choose additional configuration or featurization settings.

- Classification**
To predict one of several categories in the target column. yes/no, blue, red, green.
- Regression** ✓
To predict continuous numeric values.
- Time series forecasting**
To predict values based on time.
- Natural Language Processing (preview)**
Predict based on text-only data types using multi-class or multi-label classification.
- Computer Vision (preview)**
Multi-class or multi-label image classification, object detection, and instance segmentation.

[View additional configuration settings](#) [View featurization settings](#)

According to the Microsoft Azure AI Fundamentals (AI-900) and Azure Machine Learning documentation, regression is the appropriate machine learning task when the goal is to predict continuous numeric values- such as prices, sales amounts, or other measurable quantities.

In this scenario, the dataset contains motor vehicle sales data, and the objective is to predict vehicle sale values. A vehicle's sale value (price) is a continuous numeric variable, meaning it can take on a wide range of possible numbers (for example, \$15,000, \$28,500, \$42,300). Regression models are designed to analyze relationships between input features (like make, model, mileage, age, fuel type) and a continuous output variable (price).

Automated Machine Learning (AutoML) in Azure simplifies this process by automatically testing multiple regression algorithms (e.g., Linear Regression, Random Forest, Gradient Boosted Trees) and hyperparameters to find the best-performing model.

Let's evaluate why the other options are incorrect:

- * Classification - Used for predicting discrete categories (e.g., car type: sedan, SUV, or truck). Sale value prediction is not categorical.
- * Time series forecasting - Used when predicting future values based on time-dependent data (e.g., sales over months or years). The question focuses on predicting price based on vehicle characteristics, not over time.
- * Natural Language Processing (NLP) - Deals with text-based data, not numeric vehicle data.
- * Computer Vision - Applies to image-based tasks (e.g., detecting car types from photos).

Therefore, per Microsoft Learn's Automated ML task selection guidance, when predicting a numeric output like vehicle sale value, the correct machine learning task type is Regression.

Final answer: Regression

質問: 146

次のアプリがあります:

- * アプリ1: 腫瘍の画像セットを使用して、腫瘍が良性か悪性かを特定し、治療法を提案します。

- * アプリ2: カメラの画像を使用して、農場内を移動する個々の家畜を追跡します
- * アプリ3: 看板の写真からブランドを識別します

各アプリは何を使用しますか? 回答するには、回答領域で適切なオプションを選択してください。
注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

App1:

Image classification

Linear regression

Object detection

Optical character recognition (OCR)

App2:

Image classification

Linear regression

Object detection

Optical character recognition (OCR)

App3:


Image classification

Linear regression

Object detection

Optical character recognition (OCR)

正解:



Explanation:



Let's analyze each application in the context of Microsoft Azure AI Fundamentals (AI-900) and computer vision model types.

- * App1 - Uses a set of images of tumors to identify whether the tumors are benign or malignant and suggest a treatment # Image classificationThis application is performing image classification, where each image (of a tumor) is assigned to a single predefined category - benign or malignant. Image classification models learn patterns from labeled training images and predict the correct class for new ones. In this case, the model identifies the type of tumor, a classic binary classification scenario.
- * App2 - Uses images from cameras to track individual livestock as they move around a farm # Object detectionThis scenario describes object detection, which not only identifies what objects (in this case, animals) are in an image but also locates them by drawing bounding boxes. Tracking movement requires detecting the position of each animal frame by frame. Object detection models are well-suited for use cases involving counting, tracking, or monitoring objects in a visual scene.
- * App3 - Identifies brands in photographs of billboards # Optical character recognition (OCR)This app involves reading and interpreting text (brand names, slogans, or logos) from images of billboards. Optical Character Recognition (OCR), part of Azure AI Vision, extracts textual information from images or scanned documents. Once extracted, that text can be analyzed to identify brand names or keywords.

Summary:

- * App1 # Image classification
- * App2 # Object detection
- * App3 # Optical character recognition (OCR)

質問: 147

文を正しく完成させる答えを選択してください。

Answer Area



正解:



Explanation:

Text extraction.

According to the Microsoft Azure AI Fundamentals (AI-900) study guide and Microsoft Learn documentation for Azure AI Vision (formerly Computer Vision), text extraction-also known as Optical Character Recognition (OCR)-is the computer vision capability that detects and extracts printed or handwritten text from images and video frames.

In this scenario, a traffic monitoring system collects vehicle registration numbers (license plates) from CCTV footage. These registration numbers are alphanumeric text that must be read and converted into digital form for processing, storage, or analysis. The Azure AI Vision service's OCR (text extraction) feature performs this function. It analyzes each frame from the video feed, detects text regions (the license plates), and converts the visual text into machine-readable text data.

This process is widely used in Automatic Number Plate Recognition (ANPR) systems that support law enforcement, toll booths, and parking management solutions. The OCR model can handle variations in font, lighting, and angle to accurately extract license plate numbers.

The other options describe different vision capabilities:

- * Image classification assigns an image to a general category (e.g., "car," "truck," or "bike"), not text extraction.
- * Object detection identifies and locates objects in images using bounding boxes (e.g., detecting the car itself), but not the text written on the car.
- * Spatial analysis tracks people or objects in a defined physical space (e.g., counting individuals entering a building), not reading text.

Therefore, for a traffic monitoring system that identifies vehicle registration numbers from CCTV footage, the most accurate Azure AI Vision capability is Text extraction (OCR).

質問: 148

機械学習のタイプを適切なシナリオに一致させます。

答えるには、適切な機械学習タイプを左側の列から右側のシナリオにドラッグします。

各機械学習タイプは、1回使用することも、複数回使用することも、まったく使用しないこともできます。

注：正しい選択はそれぞれ1ポイントの価値があります。

Learning Types

Classification

Clustering

Regression

Answer Area

Learning Type

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

Learning Type

Segment customers into different groups to support a marketing department.

Learning Type

Predict whether a student will complete a university course.

正解:

Learning Types

Classification

Clustering

Regression

Answer Area

Regression

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

Clustering

Segment customers into different groups to support a marketing department.

Classification

Predict whether a student will complete a university course.

Explanation:

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Describe features of common AI workloads", there are three primary supervised and unsupervised machine learning types: Regression, Classification, and Clustering. Each type of learning addresses a different kind of problem depending on the data and desired prediction output.

* Regression - Regression models are used to predict numeric, continuous values. The study guide specifies that "regression predicts a number." In the scenario "Predict how many minutes late a flight will arrive based on the amount of snowfall," the output (minutes late) is a continuous numeric value.

Therefore, this is a regression problem. Regression algorithms like linear regression or decision tree regression estimate relationships between variables and predict measurable quantities.

* Clustering - Clustering falls under unsupervised learning, where the model identifies natural groupings or patterns in unlabeled data. The official AI-900 training material states that "clustering is used to find groups or segments of data that share similar characteristics." The scenario "Segment customers into different groups to support a marketing department" fits this description because the goal is to group customers based on behavior or demographics without predefined labels. Thus, it is a clustering problem.

* Classification - Classification is a supervised learning method used to predict discrete categories or labels. The AI-900 content defines classification as "predicting which category an item belongs to." The scenario "Predict whether a student will complete a university course" requires a yes/no (binary) outcome, which is a classic classification problem. Examples include logistic regression, decision trees, or neural networks trained for categorical prediction.

In summary:

- * Regression # Predicts continuous numeric outcomes.
- * Clustering # Groups data by similarities without predefined labels.
- * Classification # Predicts discrete or categorical outcomes.

Hence, the correct and verified mappings based on the official AI-900 study material are:

- * Regression # Flight delay prediction
- * Clustering # Customer segmentation
- * Classification # Course completion prediction

質問: 149

あなたは、Microsoft Teams で表示できる会話型 AI ソリューションを構築する予定です。Microsoft Cortana と Amazon Alexa どのサービスを使用する必要がありますか？

- A. Azure ボット サービス
- B. Azure Cognitive Search
- C. 言語サービス
- D. スピーチ

正解: ([正解を表示します](#))

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and Microsoft Learn module "Describe features of conversational AI workloads on Azure," the Azure Bot Service is the dedicated Azure service for building, connecting, deploying, and managing conversational AI experiences across multiple channels - such as Microsoft Teams, Cortana, and Amazon Alexa.

The Azure Bot Service integrates with the Bot Framework SDK to design intelligent chatbots that can communicate with users in natural language. It also connects seamlessly with other Azure Cognitive Services, such as Language Service (LUIS) for intent understanding and Speech Service for voice input/output.

The question specifies that the conversational AI must be accessible through multiple platforms, including Microsoft Teams, Cortana, and Alexa. Azure Bot Service supports this multi-channel communication model out of the box, allowing developers to configure a single bot that interacts through many endpoints simultaneously.

Other options:

- * B. Azure Cognitive Search: Used for information retrieval and knowledge mining, not conversational AI.
- * C. Language Service: Provides natural language understanding, key phrase extraction, sentiment analysis, etc., but doesn't handle multi-channel communication.
- * D. Speech: Provides speech-to-text and text-to-speech conversion but is not a chatbot platform.

Therefore, the best solution for building and deploying a multi-channel conversational AI system is Azure Bot Service, as clearly defined in Microsoft's AI-900 learning content.

質問: 150

大量の非構造化データからデータ間の関係を抽出することは、どのタイプの AI ワークロードの例ですか？

- A. コンピュータ ビジョン
- B. ナレッジマイニング
- C. 自然言語処理 (NLP)
- D. 異常検出

正解: B ([コメントを发表する](#))

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.

質問: 151

デジタル写真の自動キャプションを生成するために使用できるコンピューター ビジョン機能はどれですか？

- A. テキストを認識します。
- B. 画像の説明をします。
- C. 関心のある領域を特定します。
- D. オブジェクトを検出します。

正解: ([正解を表示します](#))

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.

有効的な**AI-900J**問題集はJPNTTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AI-900J**試験問題集を提供します。JPNTTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **336**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」

質問: 152

GPT-3.5 を使用する生成 AI ソリューションによって返される応答の品質を向上させる 3 つのアクションはどれですか？それぞれの正解は完全な解決策を示します。

注: 各正解は 1 ポイントの価値があります。

- A. プロンプトにグラウンディング データを追加します。
- B. プロンプトに追加の例を提供します。
- C. トークン化を変更します。
- D. トレーニング データをプロンプトに追加します。
- E. システムメッセージを変更します。

正解: ([正解を表示します](#))

To improve the quality and relevance of responses generated by a generative AI solution using GPT-3.5, the following three actions are emphasized in the Microsoft Learn: Azure OpenAI Service best practices and AI-

900/AI-102 training materials:

- * (A) Add grounding data to prompts: Grounding ensures that the model's output is based on factual, domain-specific information. By adding context or external knowledge sources, responses become more accurate and aligned with the organization's data rather than relying on the model's general training corpus.
- * (B) Provide additional examples to prompts: Also known as few-shot prompting, this method demonstrates desired response patterns by including examples in the prompt. This significantly improves output quality, consistency, and adherence to desired formats.

* (E) Modify system messages: In Azure OpenAI chat completions, system messages define the model's behavior, style, and tone. Adjusting system messages allows fine-tuning of the model's response quality, ensuring it follows context or persona guidelines.

The remaining options are incorrect:

* (C) Modify tokenization is a low-level text-processing technique not used to improve model response quality directly.

* (D) Add training data to prompts is not possible at runtime since GPT-3.5 models are pre-trained; only prompt engineering can influence output behavior.

Therefore, based on Azure OpenAI and AI-900 guidance, the three best ways to enhance generative

質問: 153

次の各ステートメントについて、ステートメントがtrueの場合は、[はい]を選択します。それ以外の場合は、[いいえ]を選択します。

注：正しい選択はそれぞれ1ポイントの価値があります。

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.	☒	☐

正解:

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.

有効的な**AI-900J**問題集はJPNTTest.com提供され、**AI-900J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AI-900J**試験問題集を提供します。JPNTTest.com AI-900J試験問題集はもう更新されました。ここで**AI-900J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AI-900J-mondaishu> **336**問、**30%ディスカウント**、特別な割引コード：**JPNshiken**」