

Fast2Test

Pass Your Next Certification Exam Fast!

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

365 days free updates. First attempt guaranteed success.



Instant Download

After Payment, our system will send you the products you purchase in mailbox in a minute after payment. If not received within 2 hours, please contact us.

365 Days Free Updates

Free update is available within 365 days after your purchase. After 365 days, you will get 50% discounts for updating.



Money Back Guarantee

Full refund if you fail the corresponding exam in 60 days after purchasing. And Free get any another product.

Security & Privacy

We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.

We're not the only ones **happy** about Fast2test Practice Materials ...

62316+ customers in 100+ countries use Fast2test Self Test Engine. Meet our customers.

<https://de.fast2test.com>

Anbieter der Studienmaterialien zur IT-Zertifizierung! Sicher, einfach und schnell. 100%-Pass-Garantie!

Exam : **AI-900J**

Title : Microsoft Azure AI
Fundamentals (AI-
900日本語版)

Vendor : Microsoft

Version : DEMO

QUESTION NO: 1

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.

Answer:**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.

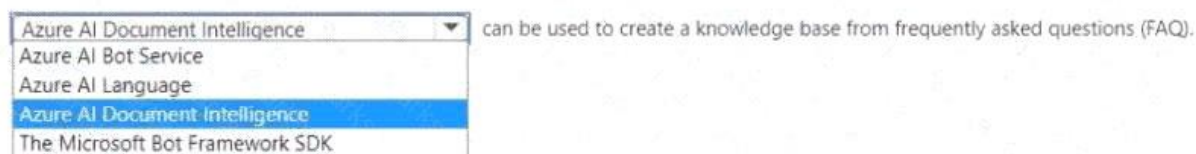
Hence, the correct and verified mappings based on the official AI-900 study material are:

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

QUESTION NO: 2

文を正しく完成させる答えを選択してください。

Answer Area



Answer:

Answer Area



Explanation:

Answer Area



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.

QUESTION NO: 3

文を正しく完成させる答えを選択してください。

Answer Area

The layer uses system inputs and additional context to

metaprompt and grounding

metaprompt and grounding

model

safety system

user experience

mitigate the harmful output of a generative AI model.

Answer:

Answer Area

The layer uses system inputs and additional context to

metaprompt and grounding

metaprompt and grounding

model

safety system

user experience

mitigate the harmful output of a generative AI model.

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.

QUESTION NO: 4

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

Answer: C

Explanation:

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.

QUESTION NO: 5

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

Answer: C

Explanation:

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

QUESTION NO: 6

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

Answer:

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

Explanation:

Answer Area

The Azure portal	Create a Machine Learning workspace
Machine Learning designer	Use a drag-and-drop interface used to train and deploy models
Automated machine learning (automated ML)	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.

QUESTION NO: 7

Form Recognizer でカスタム モデルを使用する利点は何ですか？

- A. オンプレミスにデプロイできるのはカスタム モデルのみです。
- B. カスタム モデルは、さまざまなフォーム タイプを認識するようにトレーニングできます。
- C. カスタム モデルは、事前構築されたモデルよりも安価です。
- D. カスタム モデルは常に高い精度を提供します。

Answer: B

Explanation:

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

As per the AI-900 study guide, a key advantage of a custom model is its flexibility. It can be trained with a set of labeled examples (e.g., invoices, purchase orders, receipts) that match the company's format. Once trained, the model learns where to locate and extract fields such as invoice numbers, dates, or totals-regardless of layout differences between form types. Option B is correct because a custom model can be trained to recognize a variety of form types, making it adaptable for diverse business processes.

Options A, C, and D are incorrect:

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

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

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

QUESTION NO: 8

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。

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

	▼
classification.	
clustering.	
regression.	

Answer:

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

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.

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

Answer:

Answer Area

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

Explanation:

Answer Area

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

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 # YES The designer lets users save their current work as a pipeline draft, enabling them to pause and return later. Microsoft Learn explicitly states that you can "save and publish pipeline drafts before running or deploying them." This functionality ensures workflow continuity, collaboration, and version management-making this statement also True.

* Include Custom JavaScript Functions # NO The Azure Machine Learning Designer allows the integration of Python scripts through the "Execute Python Script" module for custom logic, but it does not support JavaScript. Custom code in the designer environment is limited to Python, as the platform is built for data science and machine learning tasks typically handled in Python-based environments.

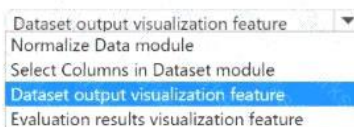
Therefore, this statement is False.

QUESTION NO: 10

文を正しく完成させる答えを選択してください。

Answer Area

In Azure Machine Learning designer, you use the




to explore the distribution of values in potential feature columns.

Answer:

Answer Area

In Azure Machine Learning designer, you use the



to explore the distribution of values in potential feature columns.

Explanation:

Answer Area

In Azure Machine Learning designer, you use the



to explore the distribution of values in potential feature columns.

In Azure Machine Learning Designer, the Dataset output visualization feature is specifically used to explore and understand the distribution of values in potential feature columns before model training begins. This capability is critical for data exploration and preprocessing, two essential stages of the machine learning pipeline described in the Microsoft Azure AI Fundamentals (AI-900) and Azure Machine Learning learning paths.

When a dataset is imported into Azure Machine Learning Designer, users can right-click on the dataset output port and select "Visualize". This launches the dataset visualization pane, which provides detailed statistical summaries for each column, including:

- * Data type (numeric, categorical, string, Boolean)
- * Minimum, maximum, mean, and standard deviation values for numeric columns

* Frequency counts and distinct values for categorical columns

* Missing value counts

This visual inspection helps determine which columns should be used as features, which might need normalization or encoding, and which contain missing or irrelevant data. It is a vital step in ensuring the dataset is clean and ready for model training.

Let's examine why other options are incorrect:

* Normalize Data module is used to scale numeric data, not to visualize distributions.

* Select Columns in Dataset module is used to include or exclude columns, not to analyze them.

* Evaluation results visualization feature is used after model training to interpret performance metrics like accuracy or recall, not data distributions.

Therefore, based on official Microsoft documentation and AI-900 study materials, to explore the distribution of values in potential feature columns, you use the Dataset output visualization feature in Azure Machine Learning Designer.

QUESTION NO: 11

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 bot that responds to queries by internal users is an example of a conversational AI workload.	<input type="radio"/>	<input type="radio"/>
An application that displays images relating to an entered search term is an example of a conversational AI workload.	<input type="radio"/>	<input type="radio"/>
A web form used to submit a request to reset a password is an example of a conversational AI workload.	<input type="radio"/>	<input type="radio"/>

Answer:

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

Explanation:

Statements

Yes

No

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

Yes

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

No

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

No

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials, conversational AI workloads are those that enable interaction between humans and AI systems through natural language conversation, either by text or speech. These workloads are typically implemented using Azure Bot Service, Azure Cognitive Services for Language, and Azure OpenAI Service. The key characteristic of a conversational AI workload is the presence of dialogue-the AI interprets user intent and provides a meaningful, contextual response in a conversation-like manner.

* "A bot that responds to queries by internal users is an example of a conversational AI workload." # YES This fits the definition perfectly. A chatbot that helps employees (internal users) by answering questions about policies, IT issues, or HR procedures is a typical example of conversational AI. It uses natural language understanding to interpret questions and provide automated responses. Microsoft Learn explicitly identifies chatbots as conversational AI solutions designed for both internal and external interactions.

* "An application that displays images relating to an entered search term is an example of a conversational AI workload." # NO This is not conversational AI because there is no dialogue or language understanding involved. It is an example of information retrieval or computer vision if it uses image recognition, but not conversation.

* "A web form used to submit a request to reset a password is an example of a conversational AI workload." # NO A password reset form is a simple UI-driven process that doesn't require AI or conversational logic. It performs a fixed function based on user input but does not understand or respond to natural language.

Therefore, based on the AI-900 study guide, only the first statement is an example of a conversational AI workload, while the second and third statements are not.

QUESTION NO: 12

AIのワークロードを適切なタスクにマッチングさせる。

回答するには、左側の列から適切なAIワークロードを右側のタスクにドラッグしてください。各ワークロードは、1回、複数回、またはまったく使用しない場合があります。注：正解ごとに1ポイントが加算されます。

AI workloads	Answer Area
<ul style="list-style-type: none"> Azure AI Document Intelligence Computer vision Generative AI Natural language processing 	<ul style="list-style-type: none"> 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.

Answer:

AI workloads	Answer Area
<ul style="list-style-type: none"> Azure AI Document Intelligence Computer vision Generative AI Natural language processing 	<ul style="list-style-type: none"> 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.

AI workloads	Answer Area
<ul style="list-style-type: none"> Azure AI Document Intelligence Computer vision Generative AI Natural language processing 	<ul style="list-style-type: none"> 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.

QUESTION NO: 13

デジタル写真内の有名なブランドを識別する事前構築されたソリューションを実装する必要があります。どのアズールアルセン/Tシャツを使うべきですか？

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

Answer: C

Explanation:

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Azure," the Computer Vision service can analyze images to detect objects, landmarks, celebrities, and brands.

The brand detection capability in the Computer Vision Image Analysis API uses pre-trained models to identify well-known brand logos within images. When an image is analyzed, the service returns brand names, confidence scores, and bounding box coordinates where the logos appear.

Let's examine the other options:

- * A. Face: Detects and analyzes human faces, not brand logos.
- * B. Custom Vision: Used for training custom models to recognize unique objects (e.g., company-specific products), not pre-built brand detection.
- * D. Form Recognizer: Extracts text and data from structured or semi-structured documents like invoices and receipts.

Thus, since the question specifies identifying well-known brands using a pre-built AI model, the correct Azure service is Computer Vision.

QUESTION NO: 14

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。

Answer: A D

Explanation:

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.

QUESTION NO: 15

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

Answer: D

Explanation:

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.

QUESTION NO: 16

同様の購買習慣を持つ人々のグループを特定するには、どのタイプの機械学習を使用する必要がありますか？

- A.分類
- B.回帰
- C.クラスタリング

Answer: C

Explanation:

According to the Microsoft Azure AI Fundamentals (AI-900) official study guide and the Microsoft Learn module "Describe features of common AI workloads", clustering is a type of unsupervised machine learning used to group data points that share similar characteristics. In unsupervised learning, the data provided to the model does not have predefined labels or outcomes. Instead, the algorithm identifies inherent patterns or groupings within the dataset based on similarities in input features.

In this scenario, the task is to identify groups of people who have similar purchasing habits. There is no predefined label such as "buyer type" or "purchase category." The goal is to discover hidden patterns-such as grouping customers by spending behavior, preferred products, or frequency of purchases. This is precisely what clustering algorithms are designed to do.

Clustering is commonly used in:

- * Customer segmentation for marketing analytics.
- * Market basket analysis to find associations in purchasing patterns.
- * Recommender systems to identify similar user profiles.

* Anomaly detection when outliers deviate from natural clusters.

Typical algorithms for clustering include K-means, Hierarchical clustering, and DBSCAN. These models analyze multidimensional data to form clusters that maximize intra-group similarity and minimize inter-group similarity.

By contrast:

* Classification (A) is a supervised learning method that predicts a categorical label (e.g., whether a customer will churn or not). It requires labeled training data.

* Regression (B) is used to predict continuous numeric values (e.g., sales revenue, temperature).

Since the question focuses on discovering groups of similar customers without prior labels, the correct type of machine learning is Clustering.

QUESTION NO: 17

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

Answer:

Answer Area

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

Explanation:

Answer Area

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

According to the Microsoft Azure AI Fundamentals (AI-900) official study materials and the Microsoft Learn module: "Describe features of common AI workloads", conversational AI solutions like chatbots can be created using various methods-not only through custom code. Azure provides both no-code/low-code and developer-focused approaches. For instance, users can design chatbots using Power Virtual Agents, which requires no programming knowledge, or they can use Azure Bot Service with the Bot Framework SDK for fully customized scenarios. Hence, the statement "Chatbots can only be built by using custom code" is False (No) because Azure supports multiple levels of technical involvement for building bots.

The second statement is True (Yes) because the Azure Bot Service is designed specifically to host, manage, and connect conversational bots to users across different channels. Microsoft Learn explicitly explains that the service provides integrated hosting, connection management, and telemetry for bots built using the Bot Framework or Power Virtual Agents. It acts as the foundation for deploying, scaling, and managing chatbot workloads in Azure. The third statement is also True (Yes) because Azure Bot Service supports integration with Microsoft Teams, among many other channels such as Skype, Facebook Messenger, Slack, and web chat. Microsoft documentation states that Azure-hosted bots can communicate directly with Teams users through the Teams channel, enabling intelligent virtual assistants within the Teams environment.

QUESTION NO: 18

Azure Cognitive Services を適切な AI ワークロードに一致させます。

回答するには、適切なサービスを左側の列から右側のワークロードにドラッグします。各サービスは、1 回または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: それぞれの正しい一致は 1 ポイントの価値があります。

Answer:

Explanation:

The correct matches are Custom Vision, Form Recognizer, and Face - each corresponding to a distinct capability under Azure Cognitive Services as described in the Microsoft Azure AI

Fundamentals (AI-900) official study guide and Microsoft Learn modules on Computer Vision workloads.

* Custom Vision # Identify objects in an image
The Custom Vision service is part of the Azure Cognitive Services suite that enables developers to train custom image classification and object detection models.

Unlike the prebuilt Computer Vision API, Custom Vision allows users to upload their own labeled images and teach the model to recognize specific objects relevant to their business context. The AI-900 syllabus explains that Custom Vision is ideal for tasks such as identifying products on a shelf, categorizing images, or detecting defects in manufacturing.

* Form Recognizer # Automatically import data from an invoice to a database
Form Recognizer is a document processing AI service that extracts structured data from forms, receipts, and invoices. It uses optical character recognition (OCR) combined with layout and key-value pair detection to automatically capture information such as invoice numbers, amounts, and vendor names. The AI-900 study materials highlight this service under the Document Intelligence category, emphasizing its ability to streamline data entry and business automation workflows by importing extracted data directly into databases or applications.

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

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

* Custom Vision # Object recognition

* Form Recognizer # Data extraction from forms

* Face # People identification

Final verified configuration:

* Custom Vision # Identify objects in an image

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

* Face # Identify people in an image

QUESTION NO: 19

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

Answer: B

Explanation:

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.

QUESTION NO: 20

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

Answer: B

Explanation:

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.

QUESTION NO: 21

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

Answer: C

Explanation:

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.

QUESTION NO: 22

自然言語処理ワークロードの例として挙げられる 2 つのシナリオはどれですか？
それぞれの正解は完全な解決策を示します。

注記; 正しく選択するたびに 1 ポイントの価値があります。

- A. 自動でヘッドランプを自動車に挿入する組立ライン機械
- B. 「今日の天気はどうか?」などの質問に応答する家庭内のスマート デバイス。
- C. 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。
- D. ナレッジベースを使用してユーザーの質問に対話的に回答する Web サイト

Answer: B D

Explanation:

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

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

* Option B - A smart device in the home that responds to questions such as "What will the weather be like today?" This is an example of an NLP workload because the device must process spoken language (speech-to-text), interpret the user's intent (language understanding), and generate a relevant spoken response (text-to-speech). This workflow involves several Azure Cognitive Services, such as Speech Service for recognizing and synthesizing speech, and Language Understanding (LUIS) for interpreting intent. This aligns with conversational AI and NLP tasks in the AI-900 syllabus.

* Option D - A website that uses a knowledge base to interactively respond to users' questions. This is also an NLP workload because the system interprets text input from users and retrieves appropriate answers from a knowledge base. Microsoft's QnA Maker (now part of the Azure AI Language service) and Azure Bot Service enable such behavior. The model uses NLP to understand the user's question, find the most relevant response, and generate an appropriate reply - key characteristics of natural language processing.

Incorrect options:

* A (assembly line machinery) represents automation or robotics, not NLP.

* C (monitoring temperature to activate a fan) is an example of an IoT (Internet of Things) or

rule-based system, not related to language processing.

QUESTION NO: 23

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

Answer:

Answer Area

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

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.

QUESTION NO: 24


文を正しく完成させる答えを選択してください。

Answer Area



Answer:

Answer Area



Explanation:

Answer Area



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.

QUESTION NO: 25

あなたは Azure で Chabot ソリューションを開発しています。

ユーザーの意図を判断するにはどのサービスを使用する必要がありますか？

- A. 翻訳者
- B. Azure Cognitive Search
- C. スピーチ
- D. 言語

Answer: D

Explanation:

In Azure, the Language service unifies several natural language capabilities, including LUIS, QnA Maker, and Text Analytics, into one comprehensive service. To determine a user's intent in a chatbot, you use the Conversational Language Understanding (CLU) feature of the Language service, which is the evolution of LUIS.

CLU helps chatbots and applications comprehend natural language input by identifying the intent (the purpose of the user's statement) and extracting entities (important details). For example, when a user types "Book a meeting for tomorrow," the model recognizes the intent (BookMeeting) and the entity (tomorrow).

The other options do not determine intent:

- * Translator (A) is used for language translation.
- * Azure Cognitive Search (B) retrieves documents based on search queries.
- * Speech (C) converts audio to text but doesn't analyze meaning.

Thus, to determine a user's intent in a chatbot scenario, the correct service is D. Language.

Reference:Microsoft Learn - "Build conversational language understanding models with Azure Language service."

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

Answer:

Answer Area

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

Explanation:

Answer Area

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

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 # Yes
The 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 # Yes
The 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 # No
While 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.

QUESTION NO: 27

あなたの会社はウィジェットを製造しています。

ウィジェットのデジタル写真が 1,000 枚あります。

写真内のウィジェットの位置を特定する必要があります。

何を使うべきですか？

A. コンピュータビジョンの空間分析

B. Custom Vision オブジェクト検出

C. カスタム ビジョンの分類

D. コンピュータ ビジョン画像解析

Answer: B

Explanation:

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure," object detection is a computer vision technique used to locate and identify objects within an image. It not only determines what objects are present but also where they appear in the image by returning bounding box coordinates around each detected item.

In this scenario, the goal is to identify the location of widgets within digital photos. This requires both recognition (knowing that the object is a widget) and localization (determining its position). The Custom Vision service in Azure allows you to train a model specifically for your own images, making it ideal for recognizing company-specific products such as widgets. By selecting the Object Detection domain in Custom Vision, you can label regions of interest in your training images. The model then learns to detect and locate those objects in new photos.

Let's examine the other options:

* A. Computer Vision Spatial Analysis: Used for people tracking, movement detection, and occupancy analytics in video streams - not for locating products in still images.

* C. Custom Vision classification: This model categorizes an image as a whole (e.g., "contains a widget" or "does not contain a widget") but does not locate objects within the image.

* D. Computer Vision Image Analysis: Provides general image tagging, description, and OCR capabilities but does not pinpoint object locations.

QUESTION NO: 28

自然言語処理を使用して、Microsoft ニュース記事のテキストを処理します。

次の展示に示す出力を受け取ります。

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



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

どのタイプの自然言語処理が実行されましたか？

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

Answer: A

Explanation:

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

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

QUESTION NO: 29

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。

▼
Accuracy
Confidence
Root Mean Square Error
Sentiment

is the calculated probability of a correct image classification.

Answer:

	▼
Accuracy	
Confidence	
Root Mean Square Error	
Sentiment	

is the calculated probability of a correct image classification.

Explanation:

Confidence.

According to the Microsoft Azure AI Fundamentals (AI-900) Official Study Guide and the Microsoft Learn module "Explore computer vision in Microsoft Azure," the confidence score represents the calculated probability that a model's prediction is correct. In image classification, when an AI model analyzes an image and assigns it to a specific category, it also produces a confidence value—a numerical probability (usually between 0 and 1) indicating how certain the model is about its prediction.

For example, if an image classification model identifies an image as a "cat" with a confidence of 0.92, it means the model is 92% certain that the image depicts a cat. The confidence value helps developers and users understand the model's certainty level about its classification output.

Microsoft Learn emphasizes that in Azure Cognitive Services—such as the Custom Vision Service—each prediction result includes both the predicted label (class) and a confidence score. These confidence scores are essential for evaluating model performance and determining thresholds for automated decisions (e.g., accepting predictions only above a 0.8 probability).

Let's evaluate the other options:

- * Accuracy: This is an overall performance metric measuring the percentage of correct predictions across the dataset, not a probability for a single prediction.
- * Root Mean Square Error (RMSE): This is a metric for regression models, not classification tasks. It measures average error magnitude between predicted and actual values.
- * Sentiment: This is a type of prediction (positive, negative, neutral) in text analysis, not a probability metric.

Therefore, based on Microsoft's AI-900 training materials and Azure Cognitive Services documentation, the calculated probability of a correct image classification is called Confidence, which expresses how sure the model is about its prediction for a specific input.

QUESTION NO: 30

コンピュータビジョンを使用して実行できる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. 株価を予測します。
- B. 画像内のブランドを検出します。
- C. 画像の配色を検出する
- D. 言語間でテキストを翻訳します。
- E. キーフレーズを抽出します。

Answer: B C

Explanation:

According to the Microsoft Azure AI Fundamentals study guide and Microsoft Learn module "Identify features of computer vision workloads", computer vision is an AI workload that allows systems to interpret and understand visual information from the world, such as images and videos.

Computer vision tasks typically include:

- * Object detection and image classification (e.g., detecting brands, logos, or items in images)
- * Image analysis (e.g., identifying colors, patterns, or visual features)
- * Face detection and recognition
- * Optical Character Recognition (OCR) for reading text in images

Therefore, both detecting brands and detecting color schemes in an image are clear examples of computer vision tasks because they involve analyzing visual content.

In contrast:

- * A. Predict stock prices # Regression task, not vision-based.
- * D. Translate text between languages # Natural language processing (NLP).
- * E. Extract key phrases # NLP as well.

Thus, the correct computer vision tasks are B and C.

Reference: Microsoft Learn - Identify features and uses of computer vision in Azure AI services

QUESTION NO: 31

Azure OpenAI を使用する Azure サブスクリプションがあります。

ウェブサイトで使用するための田園風景のオリジナル画像を作成する必要があります。何をすべきでしょうか？

A. Azure AI Foundry から GPT-3.5 Turbo

モデルをデプロイし、田園風景の画像を作成するための指示を提供します。

B. Microsoft Bing で「rural scene」という用語を検索し、結果をダウンロードします。

C. GitHub Copilot から、田園風景の画像を作成するための手順を説明します。

D. Azure AI Foundry から DALL-E

モデルを展開し、田園風景を作成するための手順を提供します。

Answer: D

Explanation:

The Azure OpenAI DALL-E model is specifically designed for generating original images from natural language prompts. If you want to create an image of a "rural scene," you can use Azure AI Foundry (formerly Azure OpenAI Studio) to deploy the DALL-E model and provide descriptive instructions such as "create an image of a peaceful rural village with trees and a sunset."

- * A. GPT-3.5 Turbo # Handles text generation, not image creation.
- * B. Bing search # Finds existing images, not generate original ones.
- * C. GitHub Copilot # Assists with writing code, not generating images.

QUESTION NO: 32

サングラスをかけている人の写真をTwitterに投稿するチャリティーイベントを運営しています。

次の要件を満たす写真のみをリツイートするようにする必要があります。

1つまたは複数の面を含めます。

サングラスをかけている人が少なくとも1人含まれている。
画像を分析するために何を使用する必要がありますか？

- A.Faceサービスでの検証操作
- B.Faceサービスでの検出操作
- C.コンピュータービジョンサービスでの画像の説明操作
- D.コンピュータービジョンサービスでの画像の分析操作

Answer: B

Explanation:

The scenario requires two checks on each photo: (1) there is at least one face, and (2) at least one detected face is wearing sunglasses. The Azure AI Face service - Detect operation is purpose-built for this combination. It detects faces and returns per-face attributes, including glasses type, so you can enforce both rules in a single pass. From the official guidance, the Detect API "detects human faces in an image and returns the rectangle coordinates of their locations" and exposes face attributes such as glasses. A concise attribute extract states: "Glasses: NoGlasses, ReadingGlasses, Sunglasses, Swimming Goggles." With this, you can count faces (requirement 1) and then verify that at least one face's glasses attribute equals sunglasses (requirement 2).

By contrast, other options don't align as precisely:

* A. Verify (Face service) compares whether two detected faces belong to the same person. It does not provide content attributes like sunglasses; it requires face inputs for identity/one-to-one scenarios, which doesn't meet your content-filter goal.

* C. Describe Image (Computer Vision) returns a natural-language caption of the whole image. While a caption might mention "a person wearing sunglasses," it's not guaranteed, is not face-scoped, and offers less deterministic filtering than a structured attribute on a detected face.

* D. Analyze Image (Computer Vision) can return tags such as "person" or sometimes "sunglasses," but those tags are image-level and not bound to specific faces. You need to ensure that a detected face (not just any region) is wearing sunglasses. Face-scoped attributes from Face Detect are more reliable for this logic.

Therefore, the most accurate and exam-aligned choice is B. the Detect operation in the Face service, because it allows you to programmatically confirm face presence and per-face sunglasses in a precise, rule-driven workflow.

QUESTION NO: 33

生成 AI

ソリューションの応答品質を向上させるためにコンテキスト情報を提供することは、どのプロンプト エンジニアリング手法の例ですか。

- A. 例を挙げる
- B. 微調整
- C. グラウンディングデータ
- D. システムメッセージ

Answer: C

Explanation:

In Microsoft Azure OpenAI Service and the AI-900/AI-102 study materials, grounding data is

the correct term used to describe the process of providing contextual or external information to improve the accuracy, relevance, and quality of responses generated by a generative AI model such as GPT-3.5 or GPT-4.

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

This technique differs from other prompt engineering concepts:

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

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