

DUMPSBOSS.

Microsoft Azure AI Fundamentals

Microsoft AI-900

Version Demo

Total Demo Questions: 62

Total Premium Questions: 624

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co

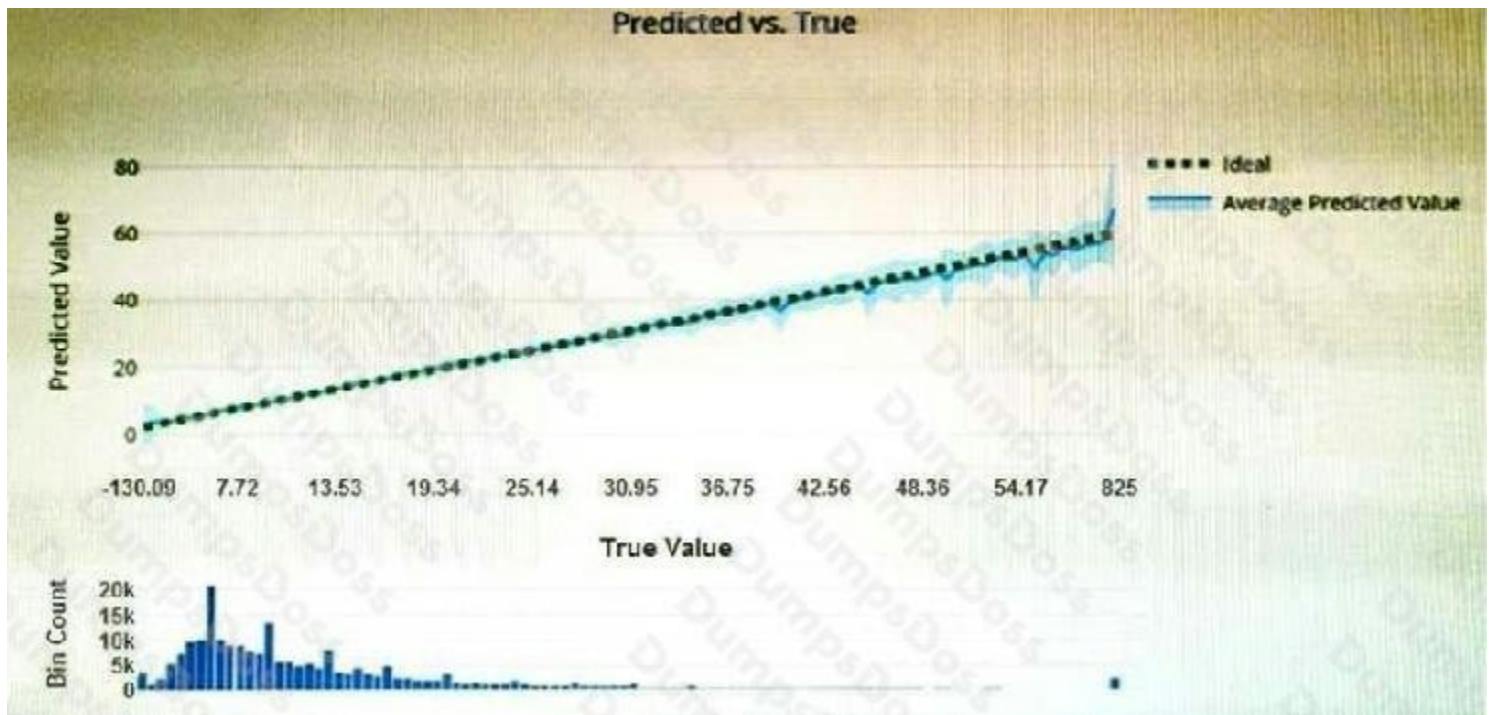
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Describe Artificial Intelligence workloads and considerations	105
Topic 2, Describe fundamental principles of machine learning on Azure	166
Topic 3, Describe features of computer vision workloads on Azure	135
Topic 4, Describe features of Natural Language Processing (NLP) workloads on Azure	163
Topic 5, Describe features of generative AI workloads on Azure	51
Topic 6, Mix Questions	4
Total	624

QUESTION NO: 1

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



Which type of model is the chart used to evaluate?

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

ANSWER: B

Explanation:

A Predicted vs. True chart is used to evaluate a regression model. In regression, the model predicts a continuous numeric value, such as a price, quantity, temperature, score, or demand forecast. The chart compares each value predicted by the model with the actual known value from the test data. When predictions are accurate, the plotted points tend to fall close to the diagonal reference line where the predicted value equals the true value. The closer the points are to that line, the better the model's predictions generally are.

This visualization is commonly used in Azure Machine Learning model evaluation for regression tasks because it gives a quick visual indication of prediction error, bias, and overall fit. Microsoft's automated machine learning documentation describes regression model evaluation using predicted versus true values, along with related metrics such as mean absolute error, root mean squared error, and coefficient of determination. Therefore, the model type evaluated by a Predicted vs. True chart is regression. See [Microsoft Learn: Understand automated ML results](#) and [Microsoft Learn: What is automated machine learning?](#)

QUESTION NO: 2

What is an example of the Microsoft responsible AI principle of transparency?

- A. helping users understand the decisions made by an AI system
- B. ensuring that developers are accountable for the solutions they create
- C. ensuring that opportunities are allocated equally to all applicants

D. ensuring that the privileged data of users is stored in a secure manner

ANSWER: A

Explanation:

“helping users understand the decisions made by an AI system” is the correct example of the Microsoft responsible AI principle of transparency. In Microsoft’s responsible AI framework, transparency means that people should be able to understand how an AI system works, what it is intended to do, what its limitations are, and how its outputs or decisions are produced. This does not always mean exposing every internal technical detail of a model, but it does mean providing enough meaningful information so users, developers, and stakeholders can interpret the system’s behavior and use it appropriately. For example, an AI-powered decision system should provide explanations, confidence information, or other supporting context so that users can understand why a recommendation or prediction was made. Microsoft describes transparency as helping stakeholders comprehend AI system capabilities and limitations, which directly aligns with helping users understand AI-driven decisions. See Microsoft’s responsible AI principles in the [Microsoft Cloud Adoption Framework guidance for responsible AI](#) and the [Microsoft Responsible AI](#) overview.

QUESTION NO: 3

You need to generate images based on user prompts. Which Azure OpenAI model should you use?

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

ANSWER: B

Explanation:

DALL-E is the correct Azure OpenAI model family for generating images from natural-language user prompts. In Azure OpenAI, DALL-E models are designed specifically for image generation scenarios where a user describes the desired visual output in text, and the model creates one or more images that match that prompt. This aligns directly with the requirement to “generate images based on user prompts,” because the task is not simply understanding, summarizing, or classifying text; it is producing visual content from textual input.

Microsoft documents DALL-E as the Azure OpenAI image generation model capability, including support for creating images from text prompts through the Azure OpenAI service. For AI-900, the key concept is recognizing the appropriate model type for the workload: DALL-E maps to image generation. You can review Microsoft’s Azure OpenAI model documentation at [Azure OpenAI Service models](#) and the image generation guidance at [Generate images with Azure OpenAI](#).

QUESTION NO: 4

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

NOTE: Each correct selection is worth one point.

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

ANSWER: A C E

Explanation:

Generate the questions and answers from an existing webpage, Manually enter the questions and answers, and Import chat content from a predefined data source are correct. QnA Maker knowledge bases, and the newer Azure AI Language custom question answering feature that replaced it, support creating question-answer pairs from structured or semi-structured content sources such as FAQ pages and support webpages. This lets the service extract likely question-and-answer pairs automatically so they can be reviewed and edited before publishing. They also support direct authoring, where you manually add custom questions and answers in the portal for business-specific scenarios that may not already exist in documentation. In addition, predefined chat content can be added to provide conversational responses for common social or casual interactions, helping the bot respond more naturally when users ask greetings or small-talk questions. Microsoft documents these supported content sources and authoring approaches for question answering, including URL-based content, manual question-answer editing, and chat datasets. References: [Data sources for question answering](#) and [Add chat to a project](#).

QUESTION NO: 5

In which two scenarios can you use the Azure AI Document Intelligence service (formerly Form Recognizer)? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

ANSWER: A B

Explanation:

Azure AI Document Intelligence is designed to analyze documents and extract structured information from them, including key-value pairs, tables, selection marks, and fields from common business documents. "Extract the invoice number from an invoice." is correct because the service includes a prebuilt invoice model that can identify and extract invoice-specific fields such as invoice ID or number, customer and vendor details, invoice date, due date, subtotal, tax, and total. This is one of the core use cases for Document Intelligence: turning semi-structured business documents into usable data.

"Identify the retailer from a receipt." is also correct because Azure AI Document Intelligence provides a prebuilt receipt model that extracts common receipt fields, including merchant or retailer name, transaction date and time, items, taxes, and totals. This allows applications to process scanned or photographed receipts and automatically capture important business information without building a custom model from scratch. Microsoft documents these prebuilt models as part of the Document Intelligence service capabilities. See [Azure AI Document Intelligence overview](#) and [receipt model documentation](#).

QUESTION NO: 6

You are developing a chatbot solution in Azure.

Which service should you use to determine a user's intent?

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

ANSWER: D

Explanation:

Language Understanding (LUIS)

is the correct choice because it is the Azure service designed to interpret natural language utterances and map them to intents. In a chatbot scenario, an intent represents what the user wants to accomplish, such as booking an appointment, checking an order status, or resetting a password. LUIS supports intent classification and entity extraction, allowing a bot to understand both the user's goal and the important details needed to fulfill the request. This capability is exactly what is needed when a chatbot must determine the meaning behind a user's message rather than simply process fixed commands.

Microsoft has since moved these capabilities into Azure AI Language through conversational language understanding, but the concept remains the same: use a language understanding capability to classify user input into intents and extract relevant entities. For the wording and options in this question, Language Understanding (LUIS)

is the service that matches intent detection for a chatbot. See Microsoft's documentation on [Language Understanding \(LUIS\)](#) and the current [conversational language understanding](#) capability.

QUESTION NO: 7

You plan to use Azure Machine Learning Studio and automated machine learning (automated ML) to build and train a model. What should you create first?

- A. a Jupyter notebook
- B. a Machine Learning workspace
- C. a registered dataset
- D. a Machine Learning designer pipeline

ANSWER: B

Explanation:

A Machine Learning workspace is correct because it is the foundational Azure resource for using Azure Machine Learning Studio and automated machine learning. The workspace acts as the central place where Azure Machine Learning stores and manages the assets needed for the machine learning lifecycle, including experiments, jobs, models, environments, data assets, and compute resources. When you open Azure Machine Learning Studio, you work within a selected workspace, and automated ML jobs are created and run in that workspace context. Before you can configure an automated ML experiment, select or create compute, connect data, train models, or register outputs, the workspace must already exist. Microsoft describes the Azure Machine Learning workspace as the top-level resource for Azure Machine Learning, providing a centralized place to work with all artifacts created when using the service. For more information, see [Microsoft Learn: What is an Azure Machine Learning workspace?](#) and [Microsoft Learn: Set up AutoML to train machine learning models.](#)

QUESTION NO: 8

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

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

ANSWER: C

Explanation:

Content filtering is the correct implementation for identifying hateful responses returned by a generative AI solution. In Azure AI services, content filtering evaluates model inputs and outputs for harmful content categories, including hate and unfairness, sexual content, violence, and self-harm. For generative AI scenarios, this is especially important because the system must assess not only what users submit, but also what the model generates before that response is shown to an end user. Azure OpenAI includes a content filtering system that classifies text by severity levels and can be configured to block or handle content that exceeds defined safety thresholds. This makes it the appropriate mechanism for detecting hateful generated responses in real time.

Microsoft also provides Azure AI Content Safety, which is designed to detect harmful user-generated and AI-generated content across text and images. It can be used to build safety controls around generative AI applications and identify hate-related content consistently. For more details, see Microsoft's documentation on [Azure OpenAI content filtering](#) and the [Azure AI Content Safety overview](#).

QUESTION NO: 9

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

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

NOTE: Each correct selection is worth one point.

- A. Azure 80l Service
- B. Azure Machine Learning
- C. Translator
- D. Language Service

ANSWER: A D

Explanation:

Azure 80l Service and Language Service are the correct services for this scenario. In Azure, the bot-hosting and conversational interface component is Azure Bot Service, which provides the framework, channels, and integration points needed to build and deploy a chatbot that can interact with users through a chat experience. It is commonly used to create automated assistants that reduce the need for human agents to answer routine questions.

Language Service is also required because it provides the natural language capabilities needed to interpret user questions and return predefined answers. In particular, Azure AI Language includes question answering features that let you create a knowledge base of questions and answers, then return the most appropriate response when a user asks a related question. This is well suited to FAQ-style or simple support scenarios where users ask predictable questions and the chatbot responds with curated answers.

Together, Azure Bot Service provides the conversational bot application layer, while Language Service provides the language understanding and question-answering capability. Microsoft describes Azure Bot Service as a platform for building conversational AI experiences, and Azure AI Language as a service for text understanding and question answering. References: [Azure Bot Service overview](#) and [Question answering in Azure AI Language](#).

QUESTION NO: 10

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

- A. Presence penalty
- B. Temperature
- C. Stop sequence

D. Max response

ANSWER: D

Explanation:

Max response is the correct parameter to configure when you want a chat solution using an Azure OpenAI GPT-3.5 model to produce more verbose responses. In Azure OpenAI Studio and related configuration experiences, this setting controls the maximum number of tokens the model can generate in its reply. Tokens are the units of text the model processes and produces, so increasing the maximum response length gives the model more capacity to generate longer, more detailed answers. If this limit is set too low, the response can be shortened or cut off even when the prompt asks for a detailed explanation.

This aligns with the Azure OpenAI completions and chat completions behavior, where the request body includes a maximum token setting that limits generated output length. For chat-based solutions, tuning this value is a direct way to allow more extensive responses, assuming the prompt also asks for sufficient detail. Microsoft documents these request parameters in the Azure OpenAI REST API reference and explains token-based generation behavior in its Azure OpenAI concepts documentation. See [Azure OpenAI Service REST API reference](#) and [Tokens in Azure OpenAI Service](#).

QUESTION NO: 11

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

NOTE; Each correct selection is worth one point.

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

ANSWER: B D

Explanation:

Natural language processing workloads involve interpreting, analyzing, generating, or responding to human language in spoken or written form. A smart device in the home that responds to questions such as, "What will the weather be like today?" is a valid NLP scenario because the system must process a user's natural-language question, determine the user's intent, and provide an appropriate response. In Azure, this type of experience commonly combines speech capabilities, such as speech-to-text and text-to-speech, with language understanding or conversational AI.

A website that uses a knowledge base to interactively respond to users' questions is also a valid NLP workload because it uses user-entered natural language to retrieve and present relevant answers. This maps closely to Azure AI Language features such as question answering, where a knowledge base can be used to support conversational question-and-answer experiences. Microsoft describes Azure AI Language as supporting capabilities for understanding and analyzing text, while Azure AI Speech supports speech-enabled interactions. See [Azure AI Language documentation](#) and [Azure AI Speech documentation](#).

QUESTION NO: 12

You are building a knowledge base by using QnA Maker. Which file format can you use to populate the knowledge base?

- A. PDF
- B. PPTX
- C. XML

ANSWER: A

Explanation:

PDF is correct because QnA Maker supports using documents as sources for a knowledge base, and PDF is one of the supported file formats for importing content. When you create or populate a QnA Maker knowledge base, you can add content from sources such as FAQ URLs or supported files. For PDF files, QnA Maker can extract structured question-and-answer content or derive question-answer pairs from appropriately formatted document content, making it a valid way to seed a knowledge base. This aligns with Microsoft's guidance that QnA Maker knowledge bases can be built from semi-structured content such as product manuals, FAQ documents, and other supported file types. Although QnA Maker has been succeeded by Azure AI Language custom question answering for newer implementations, the AI-900 concept remains the same: supported document formats can be used to populate a question answering knowledge base. Microsoft documents supported QnA Maker data sources and file types in its guidance on [data sources and content](#), and also describes the newer custom question answering capability in [Azure AI Language question answering](#).

QUESTION NO: 13

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

A. providing examples

Providing examples (few-shot prompting) shows the model sample inputs and outputs to guide formatting or style, not factual context.

B. fine-tuning

Fine-tuning involves retraining the model with labeled data to permanently adjust its behavior — it's not a prompt-based technique.

C. grounding data

D. system messages

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.

ANSWER: C

Explanation:

Grounding data is the correct answer because it describes supplying a generative AI model with relevant contextual information so its response is based on specific facts rather than only on its pretrained general knowledge. In Azure OpenAI and Microsoft generative AI guidance, grounding is commonly used to improve response quality, accuracy, and relevance by including source material such as product documentation, support articles, policy text, database results, or other retrieved content in the prompt or conversation context. This helps the model generate answers that are aligned to the user's actual scenario and reduces the likelihood of unsupported or generic responses. Microsoft's prompt engineering guidance explains that prompts can include grounding context to guide the model toward more accurate completions, especially when the required information is outside the model's built-in training data. This approach is also central to retrieval augmented generation patterns, where external data is retrieved and then provided to the model as grounding context before it generates an answer. For more detail, see Microsoft Learn on [prompt engineering with Azure OpenAI](#) and [Azure OpenAI on your data](#).

QUESTION NO: 14

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

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

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

Which type of model is this?

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

ANSWER: A

Explanation:

Regression is correct because the model is being used to predict the value of ColumnE, and the question states that the prediction is based on other numeric columns. In Microsoft's machine learning terminology, regression is a supervised learning approach used when the label being predicted is a numeric value. The model learns patterns and relationships between input features, such as the other numeric columns in the dataset, and a numeric target value, such as ColumnE. Common examples of regression include predicting a price, revenue amount, temperature, delivery time, or any other measurable quantity.

This matches the AI-900 concept that machine learning tasks are selected based on the kind of output you need. When the output is a number, the appropriate model type is regression. Microsoft Learn describes regression as a technique for predicting numeric labels, and Azure Machine Learning also categorizes regression as the task type for estimating continuous numeric values from input data. For more detail, see [Microsoft Learn: Types of machine learning](#) and [Microsoft Learn: Regression in Azure Machine Learning](#).

QUESTION NO: 15

Your company is exploring the use of voice recognition technologies in its smart home devices. The company wants to identify any barriers that might unintentionally leave out specific user groups.

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

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

ANSWER: C

Explanation:

Inclusiveness is the correct Microsoft responsible AI guiding principle for this scenario. The key phrase is that the company wants to identify barriers that might unintentionally leave out specific user groups. In Microsoft's responsible AI guidance, inclusiveness means designing AI systems so they empower and engage everyone, including people with different abilities, backgrounds, languages, accents, ages, or usage contexts. For voice recognition in smart home devices, this could involve checking whether the system works well for users with speech impairments, regional accents, children's voices, older adults, multilingual households, or people who need non-voice alternatives to control the device.

This principle encourages teams to proactively consider who may be excluded by the design, data, or interaction model of an AI system, and then improve the solution so more people can use it effectively. In a voice-enabled smart home product, applying inclusiveness might include testing with diverse speech samples, supporting varied dialects and languages, and providing accessible fallback controls. Microsoft describes inclusiveness as one of its core responsible AI principles, focused on ensuring AI technologies address a broad range of human needs and experiences. See Microsoft's guidance on [responsible AI inclusiveness](#) and the Microsoft Learn module on [responsible AI principles](#).

QUESTION NO: 16

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

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

NOTE: Each correct selection is worth one point.

- A. Use a graphical user interface (GUI) to run automated machine learning experiments.
- B. Create a compute instance to use as a workstation.
- C. Use a graphical user interface (GUI) to define and run machine learning experiments from Azure Machine Learning designer.
- D. Create a dataset from a comma-separated value (CSV) file.

ANSWER: A C

Explanation:

Use a graphical user interface (GUI) to run automated machine learning experiments and use a graphical user interface (GUI) to define and run machine learning experiments from Azure Machine Learning designer are the correct selections for this question. In the context of the older Azure Machine Learning workspace editions used by this exam item, the Enterprise workspace tier enabled the studio-based, low-code/no-code experiences. Automated ML in the GUI lets users configure, run, and compare machine learning training jobs without writing full training code. Azure Machine Learning designer provides a visual drag-and-drop canvas for building and running machine learning pipelines. These capabilities were associated with the Enterprise workspace experience when Basic and Enterprise workspace editions were still distinguished.

Microsoft has since changed Azure Machine Learning, and the current workspace documentation no longer presents Basic and Enterprise workspace editions in the same way. However, for this legacy exam-style question, the expected enterprise-only tasks are the studio GUI experiences for automated ML and designer. Current references for the relevant Azure Machine Learning capabilities are available in the [Automated ML documentation](#) and the [Azure Machine Learning designer documentation](#).

QUESTION NO: 17

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

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

ANSWER: C

Explanation:

Regression is correct because predicting the animal population of an area means estimating a numeric value, such as the number of animals expected in a region at a future point in time. In machine learning, regression is used when the label or target value is continuous or quantitative. An Azure Machine Learning regression model can learn relationships between historical population counts and relevant features such as season, weather patterns, food availability, migration trends, land use, and predator levels, then use those learned patterns to predict a future population count.

This aligns with Microsoft's description of regression as a machine learning task for predicting numeric values. Azure Machine Learning supports regression scenarios through approaches such as automated machine learning, designer pipelines, and custom training jobs, where the goal is to train a model that outputs a number rather than a category or group. For AI-900, the key concept is recognizing that a prediction involving a specific count, amount, price, quantity, or measurement is typically a regression problem. See Microsoft's guidance on [regression in Azure Machine Learning](#) and the Microsoft Learn module on [machine learning fundamentals](#).

QUESTION NO: 18

Which Azure Cognitive Services service can be used to identify documents that contain sensitive information?

- A. Custom Vision
- B. Conversational Language Understanding
- C. Form Recognizer
- D. Azure AI Language (PII detection)

ANSWER: D

Explanation:

Azure AI Language (PII detection) is the correct choice because its Personally Identifiable Information detection capability is designed to find, categorize, and optionally redact sensitive information in text. In a document-processing scenario, the document text can be analyzed with this feature to identify content such as names, addresses, phone numbers, email addresses, government-issued identifiers, financial information, and other personal or sensitive data types. This directly matches the requirement to identify documents that contain sensitive information, because the service can detect the presence of sensitive entities and return their categories, confidence scores, and locations in the text.

Microsoft describes PII detection as part of Azure AI Language, a cloud-based natural language processing service that can recognize sensitive information and support redaction workflows. This capability is commonly used for privacy, compliance, and data governance scenarios where organizations need to locate or protect personal data in unstructured text. For more details, see the Microsoft Learn overview for [PII detection in Azure AI Language](#) and the broader [Azure AI Language overview](#).

QUESTION NO: 19

You plan to develop a bot that will enable users to query a knowledge base by using natural language processing.

Which two services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Language Service
- B. Azure Bot Service
- C. Form Recognizer
- D. Anomaly Detector

ANSWER: A B

Explanation:

Language Service and Azure Bot Service are the correct services to include for a solution where users can query a knowledge base through natural language. Azure Bot Service provides the framework and managed Azure integration for building, publishing, and connecting a conversational bot to channels such as Microsoft Teams, web chat, and other client applications. It is the component that enables the user-facing conversational experience and handles bot communication.

Language Service, specifically the question answering capability in Azure AI Language, provides the natural language processing layer for a knowledge base. It can create a question-answering project from structured or semi-structured content, match a user's natural language question to the most relevant answer, and return that response to the bot. Together, Azure Bot Service and Language Service form the standard pattern for a knowledge-based conversational application: the bot manages the conversation, while Azure AI Language interprets the user's question and retrieves the best answer from the knowledge base.

Microsoft documents Azure Bot Service as a platform for developing and connecting intelligent bots, and Azure AI Language question answering as a cloud-based NLP capability for building a conversational question-and-answer layer over data. See [Azure Bot Service overview](#) and [Question answering in Azure AI Language](#).

QUESTION NO: 20

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

NOTE: Each correct selection is worth one point.

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

ANSWER: A D

Explanation:

In Azure Machine Learning designer, **dataset** and **module** are correct because the designer canvas is used to visually build machine learning pipelines by placing data assets and processing steps onto the workspace. A dataset provides the input data that the pipeline will use for training, transformation, scoring, or evaluation. A module represents an operation in the workflow, such as data preparation, model training, scoring, or model evaluation. In the designer experience, you construct a pipeline by dragging these building blocks from the asset library or palette onto the canvas, then connecting them to define how data flows through the machine learning process.

This drag-and-drop approach is central to Azure Machine Learning designer because it lets users create end-to-end machine learning workflows without writing code for every step. Microsoft describes designer pipelines as visual workflows made from datasets and components/modules that are arranged and connected on the canvas. For more detail, see Microsoft's overview of the Azure Machine Learning designer at [Microsoft Learn: What is Azure Machine Learning designer?](#) and the designer pipeline guidance at [Microsoft Learn: Create pipelines in Azure Machine Learning designer](#).

QUESTION NO: 21

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

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

Which service should you use to process the documents?

- A. Azure Bot Service

Azure Bot Service provides the chatbot interface and conversation logic but doesn't extract knowledge from documents.

- B. Language Understanding

Language Understanding (LUIS) identifies intents and entities in natural language input, but it's not used to read document content.

C. Text Analytics

Text Analytics is used for key phrase extraction and sentiment analysis, not Q & A creation.

Therefore, the correct service for processing FAQ-style and document-based content into a question-answering bot is QnA Maker.

D. QnA Maker

ANSWER: D

Explanation:

QnA Maker is the correct answer because it was designed to create a question-and-answer knowledge base from existing content such as FAQ pages, product manuals, and supported document formats including Microsoft Word documents. In this scenario, the chatbot needs to answer user questions by using information that already exists in a troubleshooting guide and an FAQ webpage. QnA Maker can ingest those sources, extract question-and-answer pairs, and publish the resulting knowledge base so that a bot can use it to respond to users conversationally.

In Microsoft's current Azure AI portfolio, QnA Maker has been replaced by the question answering capability in Azure AI Language, but for an AI-900 question that includes "QnA Maker" as an available service, that is the intended service for processing FAQ-style and document-based content into answers for a chatbot. The bot interface can then be built separately and connected to the knowledge base. Microsoft documentation describes this scenario under question answering for Azure AI Language and the retired QnA Maker service: [Question answering overview](#) and [QnA Maker overview](#).

QUESTION NO: 22

You are building an AI-based app.

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

Which two principles should you follow? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Implement an Agile software development methodology
- B. Implement a process of AI model validation as part of the software review process
- C. Establish a risk governance committee that includes members of the legal team, members of the risk management team, and a privacy officer
- D. Prevent the disclosure of the use of AI-based algorithms for automated decision making

ANSWER: B C

Explanation:

Implement a process of AI model validation as part of the software review process is correct because responsible AI requires AI systems to be evaluated for quality, reliability, safety, and fairness before and during use. Model validation helps confirm that the solution performs as intended, works consistently across relevant data populations, and is monitored for issues such as unintended bias or performance drift. This aligns with Microsoft's responsible AI principles, especially fairness, reliability and safety, and accountability.

Establish a risk governance committee that includes members of the legal team, members of the risk management team, and a privacy officer is also correct because responsible AI is not only a technical practice; it also requires organizational governance. A cross-functional governance group helps define policies, review risks, ensure privacy and legal requirements are considered, and assign accountability for decisions made throughout the AI lifecycle. Microsoft guidance emphasizes that responsible AI should include governance processes, risk management, transparency, and human oversight so that AI systems are designed and operated in a trustworthy way.

References: [Microsoft Responsible AI overview](#) and [Trusted AI in the Cloud Adoption Framework](#).

QUESTION NO: 23

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

NOTE: Each correct selection is worth one point.

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

ANSWER: A D

Explanation:

Speech synthesis is the text-to-speech capability of Azure AI Speech: it takes text as input and produces spoken audio output using a synthetic voice. Therefore, “an automated voice that reads back a credit card number entered into a telephone by using a numeric keypad” is correct because the digits entered by the caller can be treated as text or structured data, then converted into speech so the system can read them back for confirmation. This is a common interactive voice response scenario for text-to-speech.

“an AI character in a computer game that speaks audibly to a player” is also correct because game dialogue can be generated from written lines and rendered as audible speech. Azure text-to-speech supports natural-sounding voices and can be used to create spoken responses for applications, assistants, characters, and other interactive experiences. Microsoft describes text-to-speech as a Speech service feature that converts text into lifelike synthesized speech, which directly matches both scenarios. See Microsoft’s documentation for [Azure AI Speech text to speech](#) and the [Azure AI Speech overview](#).

QUESTION NO: 24

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

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

NOTE: Each correct selection is worth one point.

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

ANSWER: B C

Explanation:

QnA Maker and Azure Bot Service are the correct services for building a chatbot that answers simple questions with predefined answers. QnA Maker, now represented in Azure AI Language as custom question answering, is designed to create a question-and-answer knowledge base from sources such as FAQ documents, support pages, or manually entered question/answer pairs. When a user asks a question, the service matches the user’s wording to the most relevant stored answer, which is exactly the pattern needed for repeatable, predefined responses.

Azure Bot Service provides the conversational bot framework and hosting capabilities needed to expose that knowledge base to users through channels such as web chat, Microsoft Teams, or other supported messaging interfaces. In this solution, Azure Bot Service handles the bot experience and communication flow, while QnA Maker provides the FAQ-style answer retrieval. Together, they reduce operator workload by allowing users to self-serve answers to common questions before needing human assistance. Microsoft describes Azure Bot Service as a platform for building and connecting

intelligent bots, and custom question answering as a capability for conversational question answering over your own content. References: [Microsoft Learn: Question answering](#) and [Microsoft Learn: Azure Bot Service overview](#).

QUESTION NO: 25

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

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

ANSWER: A

Explanation:

QnA Maker is correct because it was specifically designed to create a conversational question-and-answer experience from existing FAQ-style content. With QnA Maker, you could import sources such as FAQ documents, web pages, or manually defined question-and-answer pairs into a knowledge base, then publish that knowledge base so a bot or client application could return the best matching answer to a user's natural-language question. This matches the scenario exactly: the source content is already a frequently asked questions document, and the desired outcome is a bot that can answer those questions conversationally.

In current Azure services, Microsoft has moved these capabilities into the custom question answering feature of Azure AI Language, which provides the same general pattern of building a knowledge base from question-and-answer content and exposing it for conversational applications. However, for AI-900 questions that use the older service names in the answer choices, QnA Maker remains the expected answer for creating an FAQ-based bot. Microsoft's documentation describes this capability in the [Azure AI Language question answering overview](#) and also notes the QnA Maker service lifecycle. See [Question answering in Azure AI Language](#) and [QnA Maker overview](#).

QUESTION NO: 26

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

NOTE: Each correct selection is worth one point.

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

ANSWER: A D

Explanation:

Extract the invoice number from an invoice is correct because Azure Form Recognizer, now Azure AI Document Intelligence, is designed to extract structured data from documents such as forms, invoices, receipts, and business records. Its prebuilt invoice model can identify common invoice fields including invoice ID or invoice number, vendor details, customer information, dates, totals, and line items without requiring you to train a custom model. This makes invoice field extraction a core use case for the service.

Identify the retailer from a receipt.

is also correct because the service includes a prebuilt receipt model that analyzes sales receipts and extracts key fields such as merchant name, transaction date, transaction time, subtotal, tax, total, and purchased items. In this context, the retailer

corresponds to the merchant or business identified on the receipt, which is exactly the type of document understanding task Azure AI Document Intelligence supports. Microsoft describes the service as using machine learning to extract text, key-value pairs, tables, and fields from documents, with specialized prebuilt models for common document types. For more details, see [Azure AI Document Intelligence overview](#) and the [prebuilt receipt model documentation](#).

QUESTION NO: 27

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

NOTE: Each correct selection is worth one point.

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

ANSWER: A B

Explanation:

Conversational AI workloads are solutions that enable natural interaction between people and software through language, typically by using text chat or voice. A smart device in the home that responds to questions such as “What will the weather be like today?” is a conversational AI scenario because the user asks a natural-language question and the device interprets the request, retrieves or generates an answer, and responds conversationally. This is the kind of voice-assistant interaction commonly associated with conversational AI.

A website that uses a knowledge base to interactively respond to users’ questions is also a conversational AI workload. In Azure, this type of solution is commonly implemented with services such as Azure AI Language question answering and Azure Bot Service, where users ask questions in natural language and receive relevant answers from a configured knowledge source. Microsoft describes bots as apps that users interact with conversationally, and question answering as a capability for creating a conversational layer over data such as FAQs or product documentation. See [Azure Bot Service overview](#) and [Question answering overview](#).

QUESTION NO: 28 - (HOTSPOT)

HOTSPOT

You have the following apps:

- App1: Uses a set of images of tumors to identify whether the tumors are benign or malignant and suggest a treatment
- App2: Uses images from cameras to track individual livestock as they move around a farm
- App3: Identifies brands in photographs of billboards

What does each app use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

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)

ANSWER:

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)

Explanation:

The correct selections are based on the type of computer vision task each app performs. App1 uses medical images of tumors and determines whether each tumor belongs to a predefined category, benign or malignant. That is an image classification workload because the model is trained to assign an image to a label or class. In Azure AI and machine learning scenarios, image classification is commonly used when the goal is to categorize an entire image, such as classifying a medical image, identifying a product type, or sorting photos into known groups. Microsoft describes this pattern in its Custom Vision classification guidance: [Build an image classifier with Custom Vision](#).

App2 needs to track individual livestock as they move around a farm. To do that, the system must detect animals in images and determine where each one is located, usually by using bounding boxes. That matches object detection, which is used when an app must find and locate objects within an image rather than simply categorize the whole image. Microsoft's Custom Vision object detection documentation describes this kind of workload here: [Build an object detector with Custom Vision](#).

App3 identifies brands in photographs of billboards. In this scenario, the brand information is typically present as visible text in the image, so the appropriate capability is optical character recognition. OCR extracts printed or handwritten text from images and makes it available for analysis, searching, or downstream processing. Azure AI Vision includes OCR capabilities for reading text from images, as described here: [OCR in Azure AI Vision](#).

QUESTION NO: 29

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

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

A. transparency

B. inclusiveness

Inclusiveness focuses on making AI accessible to all people, regardless of ability or background.

C. fairness

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

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.

ANSWER: A

Explanation:

transparency is correct because Microsoft's responsible AI guidance states that AI systems should be understandable, and people should be able to know how and why an AI system behaves the way it does. In a loan approval scenario, the AI system's output can have a significant effect on an individual's financial opportunities, so the decision-making process should not be a "black box." If the system uses factors such as income, credit history, debt level, or repayment behavior, those factors should be explainable to the people who build, operate, audit, or are affected by the system.

Microsoft describes transparency as a principle that helps stakeholders understand the capabilities, limitations, and behavior of AI systems. This includes providing information that supports traceability, communication, and explainability. For AI-900, the key wording in the question is that "the factors used to make the decision should be explainable," which directly maps to transparency in Microsoft's responsible AI principles. You can review Microsoft's responsible AI principles at [Microsoft Responsible AI](#) and the Azure AI Fundamentals responsible AI content at [Microsoft Learn](#).

QUESTION NO: 30 - (DRAG DROP)

DRAG DROP

Match the machine learning tasks to the appropriate scenarios.

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

at all.

NOTE: Each correct selection is worth one point.

Select and Place:

Learning Types	Answer Area
Feature engineering	Task Examining the values of a confusion matrix
Feature selection	Task Splitting a date into month, day, and year fields
Model deployment	Task Picking temperature and pressure to train a weather model
Model evaluation	
Model training	

ANSWER:

Learning Types	Answer Area
Feature engineering	Model evaluation Examining the values of a confusion matrix
Feature selection	Feature engineering Splitting a date into month, day, and year fields
Model deployment	Feature selection Picking temperature and pressure to train a weather model
Model evaluation	
Model training	

Explanation:

The correct matches are based on the standard machine learning workflow used in Azure Machine Learning and general ML practice. A confusion matrix is used after a classification model has been trained to compare predicted labels against actual labels. Reviewing those counts helps determine model quality, such as how many predictions were correct or incorrect, so that activity is part of model evaluation. Microsoft documentation discusses evaluation metrics and artifacts such as confusion matrices when assessing model performance in Azure Machine Learning: [Understand automated machine learning results](#).

Splitting a date into month, day, and year fields is feature engineering because it transforms raw data into more informative inputs for a machine learning algorithm. A single date value may not be directly useful to a model, but derived fields such as month or day can reveal seasonal, weekly, or calendar-based patterns. Azure Machine Learning refers to these kinds of transformations as featurization, which prepares data so algorithms can learn from it more effectively. Microsoft describes featurization in automated ML here: [Automated ML featurization](#).

Picking temperature and pressure to train a weather model is feature selection because it involves choosing which input variables are relevant for training. In a weather prediction scenario, temperature and pressure are meaningful predictors, so selecting them as inputs helps define the set of features the model will use. This is different from creating new fields; it is about identifying useful existing variables for the model. Microsoft's Azure Machine Learning guidance on data and feature handling supports this workflow: [Configure automated ML features](#).

QUESTION NO: 31

You have a knowledge base of frequently asked questions (FAQ).

You create a bot that uses the knowledge base to respond to customer requests.

You need to identify what the bot can perform without adding additional skills.

What should you identify?

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

ANSWER: C

Explanation:

Answer questions from multiple users simultaneously is correct because a bot connected to an FAQ knowledge base is designed to receive natural language questions, find the best matching question-and-answer pair in the knowledge base, and return the relevant response to the user. In Azure, this scenario is commonly implemented with Azure AI Language question answering, often exposed through Azure Bot Service. The knowledge base provides the answer retrieval capability, while the bot channel provides the conversational interface for users.

This capability does not require adding extra bot skills because handling conversational requests and sending knowledge-base answers is the core behavior of an FAQ bot. Azure Bot Service is also a cloud-hosted service, so it can support concurrent conversations from different users through the bot runtime and connected channels. In other words, multiple customers can ask FAQ-style questions at the same time, and the bot can route each request to the question answering knowledge base and return matching answers independently. Microsoft describes question answering as a feature for building a conversational question-and-answer layer over existing content, and Azure Bot Service as a platform for building bots that interact with users across channels. See [question answering in Azure AI Language](#) and [Azure Bot Service overview](#).

QUESTION NO: 32

Your company wants to build a recycling machine for bottles. The recycling machine must

automatically identify bottles of the correct shape and reject all other items. Which type of AI workload should the company use?

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

ANSWER: C

Explanation:

Computer vision is the correct AI workload for this scenario because the recycling machine must visually inspect physical items and determine whether each item has the correct bottle shape. Computer vision workloads are designed to process images or video from cameras, detect objects, classify visual content, and make decisions based on visual characteristics such as shape, size, edges, and contours. In this case, the machine would likely use a camera to capture an image of each item, then apply an image classification or object detection model to decide whether the item matches the accepted bottle type.

Microsoft describes computer vision as an area of AI that deals with interpreting visual inputs, including analyzing images and detecting objects. Azure AI Vision and Custom Vision are commonly used for these kinds of tasks, especially when an organization needs to identify or classify objects from images. For example, a custom model could be trained with images of

acceptable bottles and non-bottle items so the recycling machine can automatically accept or reject items based on visual recognition. See Microsoft's overview of [Azure AI Vision](#) and the AI workload guidance for [computer vision](#).

QUESTION NO: 33

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

NOTE; Each correct selection is worth one point.

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

ANSWER: C D

Explanation:

Python and R are the supported scripting languages for adding custom code directly in Azure Machine Learning designer pipelines. Azure Machine Learning designer is a visual pipeline authoring experience, but it also includes modules that let you extend the pipeline with your own scripts when the built-in components do not cover a required data preparation, transformation, feature engineering, training, or evaluation step. Microsoft documents the **Execute Python Script** module for running Python code in designer, including support for input and output datasets within the pipeline. This makes Python a complete solution for custom logic in a designer workflow.

R is also supported through the **Execute R Script** module, which enables users to run R code as part of a designer pipeline. This is useful for teams that already use R for statistical analysis, data manipulation, or modeling and want to reuse that work inside Azure Machine Learning designer. Together, Python and R are the two custom-code language choices associated with designer scripting modules. For more information, see Microsoft's documentation for [Execute Python Script](#) and [Execute R Script](#).

QUESTION NO: 34

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

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

NOTE: Each correct selection is worth one point.

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

ANSWER: C D

Explanation:

When an inference pipeline is published from Azure Machine Learning designer, it is exposed as a callable web service so that applications or scripts can send input data and receive predictions. To consume that published pipeline, you need the REST endpoint because it provides the HTTPS URL to which scoring requests are submitted. This endpoint is the network address of the deployed inference service and is used by client code when calling the pipeline programmatically.

You also need the authentication key because Azure Machine Learning endpoints are typically secured, and the key is included with the request to authorize access to the service. In practice, a client sends an HTTP request to the REST endpoint and includes the authentication key in the request headers, along with the input data in the expected format. These two values together allow a consumer to locate the deployed inference pipeline and access it securely. Microsoft's documentation for consuming Azure Machine Learning endpoints describes using the scoring URI and authentication credentials for endpoint invocation. See [Deploy and score a machine learning model by using an online endpoint](#) and [What is Azure Machine Learning designer?](#).

QUESTION NO: 35

You have insurance claim reports that are stored as text.

You need to extract key terms from the reports to generate summaries. Which type of AI workload should you use?

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

ANSWER: C

Explanation:

Natural language processing is correct because the source content is unstructured text and the required task is to identify important words, phrases, or concepts that can be used to summarize the reports. In Microsoft Azure AI, this type of capability is provided through language services, especially features such as key phrase extraction. Key phrase extraction analyzes written text and returns the main talking points, entities, or significant terms in the content, which is a typical NLP scenario for documents such as insurance claims, support tickets, customer feedback, and case notes.

For AI-900, this maps directly to the natural language processing workload category: using AI to understand, interpret, and process human language in written or spoken form. Azure AI Language includes prebuilt NLP capabilities that can help extract key phrases without building a custom model from scratch. Microsoft describes key phrase extraction as a feature that evaluates unstructured text and returns a list of key phrases, making it well suited for generating report summaries or indexing document content. See Microsoft's overview of [key phrase extraction](#) and the Azure AI Fundamentals guidance on [analyzing text with Azure AI Language](#).

QUESTION NO: 36

You need to create a customer support solution to help customers access information. The solution must support email, phone, and live chat channels. Which type of AI solution should you use?

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

ANSWER: D

Explanation:

A chatbot is the correct type of AI solution for this scenario because it is designed to provide conversational customer support across multiple communication channels. In Microsoft Azure, conversational AI solutions can be built with Azure Bot Service and connected to channels such as web chat, Microsoft Teams, email, and telephony integrations. A chatbot can understand user questions, retrieve relevant answers, guide users through support tasks, and escalate to a human agent

when needed. This directly matches the requirement to help customers access information through email, phone, and live chat.

Microsoft describes Azure Bot Service as a platform for developing, publishing, and managing intelligent bots that can interact naturally with users. For customer support scenarios, bots are commonly combined with language capabilities such as question answering so they can respond to frequently asked questions from a knowledge base. This makes a chatbot the best fit when the goal is an interactive support solution rather than simply analyzing data or images. See Microsoft's documentation for [Azure Bot Service](#) and [AI workloads in Azure](#).

QUESTION NO: 37

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

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

ANSWER: A

Explanation:

Sentiment analysis is correct because it is the natural language processing workload used to evaluate text and determine the attitude, opinion, or emotional tone expressed by the writer. In a social media scenario, posts often contain subjective language that can indicate whether the author feels positive, negative, neutral, or mixed about a topic, brand, event, or experience. Azure AI Language provides sentiment analysis and opinion mining capabilities that analyze written text and return sentiment labels with confidence scores, making it well suited for monitoring customer feedback, public reaction, and social media conversations at scale.

This aligns directly with the app described in the question: the app is not merely extracting words or converting audio, but interpreting the tone conveyed in written posts. Microsoft's Azure AI Language documentation describes sentiment analysis as a feature that detects positive, negative, neutral, or mixed sentiment in text, and AI-900 learning materials classify it as a common NLP workload. For more detail, see Microsoft's documentation on [sentiment analysis and opinion mining](#) and the AI-900 module covering [text analysis with Azure AI Language](#).

QUESTION NO: 38

Which three actions improve the quality of responses returned by a generative AI solution that uses GPT-3.5? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. Add grounding data to prompts.
- B. Provide additional examples to prompts.
- C. Modify tokenization.
- D. Add training data to prompts.
- E. Modify system messages.

ANSWER: A B E

Explanation:

Add grounding data to prompts is correct because GPT-3.5 generates responses from patterns learned during training and from the context supplied at inference time. Supplying relevant grounding content, such as product documentation, policy text, or retrieved knowledge base passages, helps the model answer based on authoritative information and improves factuality. Provide additional examples to prompts is also correct because examples are a core prompt-engineering technique: they show the model the expected style, structure, level of detail, and response pattern, which can make outputs more consistent and better aligned with the task. Modify system messages is correct because system messages define high-priority instructions for the model, such as its role, constraints, tone, safety boundaries, and expected behavior. In Azure OpenAI, Microsoft guidance specifically emphasizes using clear instructions, contextual grounding data, examples, and system messages to improve generated output quality. These techniques are part of prompt engineering and are appropriate ways to improve responses without changing the underlying GPT-3.5 model. For more, see Microsoft's guidance on [prompt engineering with Azure OpenAI](#) and [system message design](#).

QUESTION NO: 39

You need to generate cartoons for use in a brochure. Each cartoon will be based on a text description.

Which Azure OpenAI model should you use?

- A. Codex
- B. DALL-E
- C. GPT-3.5
- D. GPT-4

ANSWER: B

Explanation:

DALL-E is correct because it is the Azure OpenAI model family designed for image generation from natural-language prompts. In this scenario, the requirement is to create cartoons based on text descriptions, which is a text-to-image task. DALL-E models can take a prompt that describes the desired subject, style, composition, and visual details, then generate an image that matches the prompt. For example, a prompt can specify that the image should be in a "cartoon style" and include brochure-friendly details such as colors, characters, background, or mood.

Microsoft's Azure OpenAI documentation identifies DALL-E as the model used for image generation, including creating original images from text prompts. This aligns directly with the need to generate brochure cartoons rather than produce text, code, or analysis. Azure OpenAI provides access to these image-generation capabilities through deployed DALL-E models, allowing applications to programmatically submit prompts and receive generated images. For more information, see Microsoft's documentation on [Azure OpenAI image generation models](#) and the [DALL-E quickstart](#).

QUESTION NO: 40

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

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

ANSWER: B

Explanation:

A custom model can be trained to recognize a variety of form types is correct because Azure Form Recognizer, now Azure AI Document Intelligence, lets you build models that are tailored to your own business documents rather than being limited to Microsoft's prebuilt document types. With custom extraction models, you provide sample documents and label the fields you

want extracted, so the model learns the structure, field names, and values that are specific to your forms. This is especially useful when an organization has proprietary forms, industry-specific documents, or templates that are not covered by prebuilt models.

Custom models can also be used in scenarios where there are multiple document layouts or form variations. Azure AI Document Intelligence supports custom model capabilities such as custom extraction and composed custom models, which help route different document types to the appropriate trained model. This makes custom modeling a flexible approach for recognizing and extracting data from different kinds of forms used by a business. For more detail, see Microsoft's documentation on [custom models in Azure AI Document Intelligence](#) and [composed custom models](#).

QUESTION NO: 41

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

NOTE: Each correct selection is worth one point.

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

ANSWER: A D

Explanation:

Extract the invoice number from an invoice is correct because Azure Form Recognizer, now Azure AI Document Intelligence, is designed to analyze documents and extract structured fields, key-value pairs, tables, and text. Its prebuilt invoice model can identify common invoice fields such as invoice ID or invoice number, customer information, vendor details, due dates, totals, taxes, and line items without requiring a custom model. This makes invoice field extraction one of the standard use cases for the service.

Identify the retailer from a receipt is also correct because the prebuilt receipt model is specifically intended to process sales receipts and return structured data such as merchant name, transaction date and time, subtotal, tax, total, and purchased items. In receipt analysis, the retailer or merchant is a core extracted field, so this scenario aligns directly with the service's capabilities. Microsoft describes Azure AI Document Intelligence as a document processing service that uses machine learning to extract text, layout, and meaningful fields from forms and documents. See the official overview at [Microsoft Learn](#) and the receipt model documentation at [Microsoft Learn receipt model](#).

QUESTION NO: 42

You have insurance claim reports that are stored as text.

You need to extract key terms from the reports to generate summaries.

Which type of AI workload should you use?

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

ANSWER: A

Explanation:

Natural language processing is correct because the source content is written text and the required task is to identify meaningful words and phrases that can be used to summarize the reports. In Microsoft Azure AI terminology, natural language processing covers AI workloads that interpret, analyze, and generate human language. Key phrase extraction is a standard NLP capability: it identifies the main concepts in unstructured text, such as important entities, topics, or terms in claim descriptions, notes, and reports. These extracted phrases can then be used to create indexes, highlight important content, support search, or help generate concise summaries.

For AI-900, this maps directly to the language AI category. Azure AI Language includes features for analyzing text, including key phrase extraction, entity recognition, sentiment analysis, language detection, and summarization-related language understanding scenarios. Since the reports are already stored as text and the goal is to extract key terms, the appropriate workload is natural language processing. Microsoft describes Azure AI Language as a service for building applications that understand and analyze text. See [Azure AI Language documentation](#) and [Microsoft Learn: AI fundamentals](#).

QUESTION NO: 43

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

NOTE: Each correct selection is worth one point.

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

ANSWER: C E

Explanation:

During the data ingestion and data preparation stage of an Azure Machine Learning process, the goal is to collect raw data and transform it into a clean, consistent dataset that can be used for training. **Combine multiple datasets.** is correct because machine learning projects often require data from multiple sources, such as files, databases, or cloud storage. Bringing those sources together into a single usable dataset is a common ingestion and preparation task.

Remove records that have missing values.

is also correct because data preparation includes cleaning the data before model training. Missing values can reduce model quality or cause training errors, so handling them is an important preparation step. Depending on the scenario, missing values might be removed, replaced, or otherwise transformed, but removing records with missing values is a valid data-cleaning action.

Microsoft describes data preparation as including activities such as cleaning, transforming, and combining data so that it is ready for analysis and modeling. Azure Machine Learning also supports working with data assets and datasets as part of preparing data for machine learning workflows. References: [Prepare data for machine learning](#) and [Data concepts in Azure Machine Learning](#).

QUESTION NO: 44 - (HOTSPOT)

HOTSPOT

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

NOTE: Each correct selection is worth one point.

HOT AREA

Answer Area

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

ANSWER:

Answer Area

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

Explanation:

The correct selections are based on what natural language processing is designed to do: analyze, understand, and extract meaning from human language in text or speech. Monitoring online service reviews for profanities involves inspecting written customer content and detecting inappropriate language. That kind of text moderation relies on language analysis, so it fits naturally under natural language processing. Microsoft Azure AI services include language-related capabilities for analyzing text, extracting meaning, and evaluating sentiment or other linguistic patterns. See Microsoft's overview of Azure AI Language at [Azure AI Language documentation](#).

Identifying brand logos in an image is not a language-understanding task. It requires analyzing visual features such as shapes, colors, patterns, and objects inside an image. In Azure AI, this type of workload is associated with computer vision, not natural language processing. Azure AI Vision can analyze images and detect visual content, which is a different AI domain from processing written or spoken language. Microsoft describes image analysis capabilities in the [Azure AI Vision image analysis documentation](#).

Monitoring public news sites for negative mentions of a product is also a natural language processing scenario because it requires reading text and determining whether the language expresses a negative opinion or sentiment about the product. This is commonly called sentiment analysis or opinion mining. Azure AI Language supports sentiment analysis by evaluating text and identifying whether the expressed sentiment is positive, neutral, or negative. Microsoft documents this capability at [Sentiment analysis and opinion mining](#).

QUESTION NO: 45

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

Which service should you use?

- A. Text Analytics
- B. Azure Bot Service

C. Translator

D. Form Recognizer

ANSWER: B

Explanation:

Azure Bot Service is correct because it is designed to build, deploy, and manage conversational AI bots that can communicate with users across multiple channels. In Azure, a bot can be connected to channels such as Microsoft Teams, Web Chat, Direct Line, and Email, allowing the same conversational experience to be exposed in several places without building a separate bot for each channel. This matches the requirement to support email, Microsoft Teams, and webchat from a single conversational AI solution.

Azure Bot Service works with the Bot Framework SDK and related Azure resources to handle bot hosting, channel connectivity, authentication scenarios, and integration with other AI services when needed. For an AI-900 scenario, the key point is that Azure Bot Service is the Azure service specifically associated with conversational AI and multi-channel bot communication. Microsoft describes Azure Bot Service as a managed service for creating conversational interfaces and connecting them to popular channels. You can review Microsoft's overview of the service in the [Azure Bot Service documentation](#) and see supported channel connectivity in the [Bot Service channels documentation](#).

QUESTION NO: 46 - (SIMULATION)

SIMULATION

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

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

ANSWER: See the explanation for the answer

Explanation:

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

QUESTION NO: 47

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

A. Translator Text

B. Text Analytics

C. Speech

D. Language Understanding (LUIS)

ANSWER: D

Explanation:

Language Understanding (LUIS) is correct because it was designed to interpret natural-language user utterances by identifying the user's intent and extracting relevant entities. For an input such as "Call me back later?", the important task is not simply processing text, but understanding what the user wants the application to do. LUIS supports this by mapping phrases to intents, such as a callback request, and by identifying useful details that an application or bot can use to decide

the next action. This is the core pattern used in conversational AI solutions where users express requests in everyday language rather than selecting fixed commands.

In Microsoft's Azure AI ecosystem, this type of capability is now represented by conversational language understanding in Azure AI Language, which is the successor path for building intent-and-entity models. The underlying concept remains the same for this question: use a language understanding service when the goal is to determine the meaning of a user's input in a conversational application. See Microsoft's documentation for [Language Understanding \(LUIS\)](#) and [Conversational Language Understanding](#).

QUESTION NO: 48

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

NOTE: Each correct selection is worth one point.

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

ANSWER: B D

Explanation:

Speech recognition is used when an AI system analyzes spoken audio and converts it into text. In Azure AI Speech, this capability is provided through speech-to-text, which can transcribe audio from microphones, audio files, phone calls, meetings, or video streams. Therefore, providing closed captions for recorded or live videos is a correct scenario because the spoken words in the video are recognized and turned into readable captions for accessibility, search, and content consumption. Creating a transcript of a telephone call or meeting is also a correct scenario because the system listens to the conversation and produces written text from the speech. These are standard speech-to-text workloads supported by Azure AI Speech, including real-time transcription and batch transcription scenarios. Microsoft describes speech-to-text as a capability that converts audio to text and supports use cases such as transcription and captioning. For more detail, see [Microsoft Learn: Speech to text](#) and [Microsoft Learn: Captioning with speech to text](#).

QUESTION NO: 49

You are building an AI system.

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

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

ANSWER: C

Explanation:

Provide documentation to help developers debug code is the correct answer because Microsoft's responsible AI transparency principle is about making AI systems understandable, explainable, and usable by the people who build, operate, and are affected by them. In practice, transparency requires clear communication about how an AI system works, what it is intended to do, what its limitations are, and how its outputs should be interpreted. Documentation that helps developers debug code supports this principle by giving technical teams the information they need to inspect system behavior, diagnose unexpected outputs, understand dependencies, and communicate how the service functions.

For AI-900, transparency is commonly associated with explainability, disclosure, and clear guidance for developers and users. Microsoft describes transparency as helping people understand the capabilities and limitations of AI systems so they can make informed decisions about their use. Good documentation is one of the practical mechanisms for achieving that goal, especially during development, testing, monitoring, and troubleshooting. You can review Microsoft's responsible AI principles in the [Microsoft Learn responsible AI principles module](#) and Microsoft's broader guidance on [Responsible AI](#).

QUESTION NO: 50

Which two tools can you use to call the Azure OpenAI service? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. Azure Command-Line Interface (CLI)
- B. Azure REST API
- C. Azure SDK for Python
- D. Azure SDK for JavaScript

ANSWER: B C D

Explanation:

Azure REST API is correct because Azure OpenAI exposes HTTPS endpoints that applications can call directly by sending requests to a deployed model with the appropriate endpoint, deployment name, API version, authentication, and JSON request body. This is a complete solution for invoking operations such as chat completions or embeddings from any environment that can make HTTP requests. Azure SDK for Python is also correct because Microsoft provides supported client libraries and examples for using Python to authenticate and submit requests to Azure OpenAI deployments without manually constructing every HTTP request. Azure SDK for JavaScript is correct for the same reason: JavaScript applications can use Microsoft-supported client libraries to call Azure OpenAI deployments from Node.js or other supported JavaScript environments. In practice, REST is the language-agnostic approach, while the SDKs provide idiomatic developer experiences for specific programming languages. Microsoft documents REST operations in the [Azure OpenAI REST API reference](#) and provides language-specific SDK guidance in the [Azure OpenAI SDK documentation](#).

QUESTION NO: 51

You plan to build a conversational AI solution that can be surfaced in Microsoft Teams, Microsoft Cortana, and Amazon Alexa.

Which service should you use?

- A. Azure Bot Service
- B. Azure Cognitive Search
- C. Language service
- D. Speech

ANSWER: A

Explanation:

Azure Bot Service is the correct service for building a conversational AI solution that can be exposed through multiple user channels such as Microsoft Teams and other conversational endpoints. In Azure, a bot is the core application that manages dialog flow, receives user messages, sends responses, and can integrate with AI capabilities such as language understanding, question answering, and speech where needed. Azure Bot Service works with the Microsoft Bot Framework and provides the channel connectivity layer that lets the same bot application be surfaced through supported channels without rebuilding the conversational logic for each client experience.

This matches the requirement because the scenario is not only about understanding language, searching content, or converting speech; it is specifically about creating a conversational application and making it available through different platforms. Microsoft documentation describes Azure Bot Service as a managed service for developing, connecting, deploying, and managing intelligent bots, and the Bot Service channels documentation describes how bots can be connected to channels such as Microsoft Teams. See [Azure Bot Service overview](#) and [Bot Service channels](#).

QUESTION NO: 52

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

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

NOTE: Each correct selection is worth one point.

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

ANSWER: B C

Explanation:

QnA Maker and Azure Bot Service are the correct services for building a chatbot that answers simple questions with predefined answers. QnA Maker was designed to create a knowledge base from FAQ content, product manuals, or other structured question-and-answer sources, then return the most appropriate predefined answer when a user asks a natural-language question. In current Azure services, this capability is represented by the custom question answering feature in Azure AI Language, but QnA Maker remains the expected answer for many AI-900 exam-style questions that use the older service name. You can learn more in Microsoft's overview of [question answering in Azure AI Language](#).

Azure Bot Service is also required because it provides the framework and managed service for creating, deploying, and connecting the chatbot to user-facing channels. It handles the bot application layer, conversation flow, and channel integration, while the QnA knowledge base supplies the predefined answers. Together, QnA Maker and Azure Bot Service form a standard FAQ chatbot solution that can reduce the number of routine questions handled by telephone operators. Microsoft describes this bot hosting and channel integration capability in the [Azure Bot Service overview](#).

QUESTION NO: 53

You need to provide content for a business chatbot that will help answer simple user queries.

What are three ways to create question and answer text by using Azure AI Language Service 's question answering? Each correct answer presents a complete solution.

NOTE: Each correct and ask questions by selection is worth one point.

- A. Connect the bot to the Cortana channel using Cortana.
- B. Import chit-chat content from a predefined data source.
- C. Manually enter the questions and answers.
- D. Use Azure Machine Learning Automated ML to train a model based on a file that contains question and answer pairs.
- E. Generate the questions and answers from an existing webpage.

ANSWER: B C E

Explanation:

Import chit-chat content from a predefined data source, Manually enter the questions and answers, and Generate the questions and answers from an existing webpage are correct. Azure AI Language question answering is designed to create a knowledge base of question-and-answer pairs that can be used by an application or bot to respond to user questions. Content can be authored directly by manually adding question-and-answer pairs, which is useful when you need precise control over custom business FAQs, support responses, or internal knowledge. The service can also ingest supported source content, such as an existing webpage, and extract relevant question-and-answer pairs to help build the knowledge base more quickly from existing FAQ or help content. In addition, predefined chit-chat content can be added to give the bot more natural conversational responses for common social interactions such as greetings, thanks, or casual small talk. These are supported content creation methods in the custom question answering capability of Azure AI Language. For more information, see Microsoft Learn on [question answering in Azure AI Language](#) and [adding chit-chat to a project](#).

QUESTION NO: 54

What are two tasks that can be performed by using the Computer Vision service? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Train a custom image classification model.
- B. Detect faces in an image.
- C. Recognize handwritten text.
- D. Translate the text in an image between languages.

ANSWER: B C**Explanation:**

Detect faces in an image is correct because Azure AI Vision, formerly known as Computer Vision, includes image analysis capabilities that can identify visual features in images, including detecting faces and returning information such as face location in the image. This aligns with the AI-900 expectation that computer vision services can analyze images and identify people-related features.

Recognize handwritten text is also correct because the Computer Vision service includes optical character recognition capabilities through the Read OCR feature. The Read feature is designed to extract printed and handwritten text from images and documents, such as photos of notes, forms, receipts, or scanned pages. In Microsoft's Azure AI services, this is a core Computer Vision capability and is commonly tested as OCR or text extraction from images.

Microsoft documents these capabilities in the Azure AI Vision overview and OCR guidance. See [Azure AI Vision overview](#) and [Optical character recognition overview](#).

QUESTION NO: 55

Which two resources can you use to analyze code and generate explanations of code function and code comments? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. the Azure OpenAI DALL-E model
- B. the Azure OpenAI Whisper model
- C. the Azure OpenAI GPT-4 model
- D. the GitHub Copilot service

ANSWER: C D

Explanation:

The Azure OpenAI GPT-4 model is correct because GPT-4 is a generative AI model that can understand natural language and code, making it suitable for analyzing source code, explaining what functions or blocks of code do, and generating meaningful code comments. In Azure OpenAI, GPT models are designed for conversational and text-generation scenarios, and code explanation is a common use case because the model can interpret programming syntax and produce human-readable descriptions. Microsoft's Azure OpenAI documentation describes GPT models as capable of understanding and generating natural language and code-related responses. See [Azure OpenAI Service models](#).

The GitHub Copilot service is also correct because it is specifically built as an AI pair programming assistant. GitHub Copilot can help developers write code, generate comments, explain code, and provide coding suggestions directly in development environments such as Visual Studio Code and Visual Studio. Its purpose is closely aligned with analyzing code and helping users understand or document code functionality. Microsoft and GitHub describe Copilot as an AI coding assistant that can provide code suggestions and help with development tasks. See [What is GitHub Copilot?](#)

QUESTION NO: 56

You are building an AI-based app.

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

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

NOTE: Each correct selection is worth one point.

- A. Implement an Agile software development methodology
- B. Implement a process of AI model validation as part of the software review process
- C. Establish a risk governance committee that includes members of the legal team, members of the risk management team, and a privacy officer
- D. Prevent the disclosure of the use of AI-based algorithms for automated decision making

ANSWER: B C**Explanation:**

Implement a process of AI model validation as part of the software review process is correct because responsible AI requires systems to be reliable and safe. Validation helps confirm that the model performs as intended, behaves consistently across expected scenarios, and is reviewed before or during release so that risks can be identified and addressed. This aligns with Microsoft's responsible AI principle of reliability and safety, which emphasizes that AI systems should operate dependably and be evaluated appropriately.

Establish a risk governance committee that includes members of the legal team, members of the risk management team, and a privacy officer is also correct because responsible AI includes accountability. Accountability means an organization should have people, processes, and governance structures in place to oversee AI systems, manage risks, and ensure responsible use throughout the AI lifecycle. Including legal, privacy, and risk stakeholders supports informed oversight for compliance, ethical review, and risk management. Microsoft describes responsible AI through principles including fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. See Microsoft's overview of [responsible AI principles](#) and the [Microsoft Responsible AI](#) guidance.

QUESTION NO: 57

What are two tasks that can be performed by using computer vision? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Predict stock prices.
- B. Detect brands in an image.

- C. Detect the color scheme in an image
- D. Translate text between languages.
- E. Extract key phrases.

ANSWER: B C

Explanation:

Computer vision focuses on analyzing and interpreting visual content such as images and videos. **Detect brands in an image.** is correct because Azure AI Vision includes image analysis capabilities that can identify visual features and recognize brands or logos in supported image scenarios. This is a typical computer vision use case: the service examines pixels and visual patterns to determine what objects, entities, or brand-related elements appear in the image.

Detect the color scheme in an image is also correct because image analysis can extract visual attributes such as dominant foreground and background colors, accent colors, and whether an image is black and white. These capabilities help applications understand the overall visual composition of an image, which is a core computer vision task. In Azure, these features are part of Azure AI Vision image analysis, which provides prebuilt models for extracting information from images without requiring a custom model for common visual analysis tasks. Microsoft describes Azure AI Vision as a service for processing images and returning information about visual features, including objects, brands, and color-related metadata. For more details, see [Azure AI Vision overview](#) and [Image analysis concepts](#).

QUESTION NO: 58

You need to predict the income range of a given customer by using the following dataset.

First Name	Last Name	Age	Education Level	Income Range
Orlando	Gee	45	University	25,000-50,000
Keith	Harris	36	High school	25,000-50,000
Donna	Carreras	52	University	50,000-75,000
Janet	Gates	21	University	75,000-100,000
Lucy	Harrington	68	High school	50,000-75,000

Which two fields should you use as features? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Education Level
- B. Last Name
- C. Age
- D. Income Range
- E. First Name

ANSWER: A C

Explanation:

Education Level and Age are the correct fields to use as features because they are input variables that can help a machine learning model learn patterns associated with a customer's income range. In supervised machine learning, the value you want to predict is the label, while the columns used to make that prediction are features. Here, the prediction target is Income Range, so the model should learn from other customer attributes that may have a meaningful relationship with income. Age can capture life-stage and career-progression patterns, while Education Level can capture qualification-related differences that often correlate with earning potential. These are the kinds of structured, relevant attributes that are typically selected as predictors when training a classification model for a categorical outcome such as an income band. Microsoft describes

features as input data used by a model and labels as the known output values used during training; this distinction is central to preparing data for Azure Machine Learning and automated ML scenarios. See Microsoft's overview of machine learning basics at [Machine learning concepts](#) and automated ML training data guidance at [What is automated ML?](#)

QUESTION NO: 59 - (DRAG DROP)

Match the Azure Cognitive Services to the appropriate AI workloads.

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

NOTE: Each correct match is worth one point.

Services	Answer Area
Custom Vision ▼	<input type="text"/> Identify objects in an image.
Face ▼	<input type="text"/> Automatically import data from an invoice to a database.
Form Recognizer ▼	<input type="text"/> Identify people in an image.

ANSWER:

Services

Custom Vision ▼

Face ▼

Form Recognizer ▼

Answer Area

Custom Vision ▼

Identify objects in an image.

Form Recognizer ▼

Automatically import data from an invoice to a database.

Face ▼

Identify people in an image.

Explanation:

The correct matches follow the intended purpose of each Azure AI service. Azure AI Custom Vision is designed for building custom image classification and object detection models. When an organization needs to recognize specific objects in images, such as products, tools, defects, or other business-specific items, Custom Vision lets users train a model with labeled images and then use that model to detect or classify those objects. This aligns directly with the workload to identify objects in an image. Microsoft describes these capabilities in the [Custom Vision documentation](#).

Azure AI Document Intelligence, formerly known as Form Recognizer, is the Azure service used to extract structured information from documents. It can process invoices and return fields such as vendor name, invoice number, dates, totals, and line items. That extracted data can then be stored in an application or database, which makes it the right service for automatically importing data from an invoice to a database. Microsoft documents these document extraction capabilities in the [Azure AI Document Intelligence overview](#).

The Azure AI Face service is built for detecting, analyzing, verifying, and identifying human faces in images. For a workload that needs to identify people in an image, the Face service provides the relevant facial recognition and identification functionality. Microsoft covers these capabilities in the [Face and identity recognition documentation](#). Therefore, the matching is Custom Vision for object identification, Form Recognizer for invoice data extraction, and Face for identifying people in images.

QUESTION NO: 60

Which machine learning technique can be used for anomaly detection?

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

ANSWER: C

Explanation:

The correct answer is “A machine learning technique that analyzes data over time and identifies unusual changes.” Anomaly detection is used to identify values, events, or patterns that differ significantly from expected behavior. In many business and operational scenarios, the data being monitored is time-series data, such as sales totals, application metrics, sensor readings, financial transactions, or network activity. The technique learns or evaluates what normal behavior looks like and then flags unexpected spikes, drops, outliers, or shifts that may require investigation.

In Azure AI-900 terms, anomaly detection is one of the common AI workloads and is commonly associated with detecting unusual patterns in data. Microsoft describes anomaly detection as a capability for finding abnormalities in time-series data, including situations such as monitoring system performance, fraud indicators, or equipment telemetry. Azure AI services include anomaly detection capabilities designed to help identify abnormal changes automatically from historical and current data. For more details, see Microsoft Learn’s overview of [Anomaly Detector](#) and the AI-900 module on [fundamental AI concepts](#).

QUESTION NO: 61

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

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

NOTE: Each correct selection is worth one point.

- A. Use a graphical user interface (GUI) to run automated machine learning experiments.
- B. Create a compute instance to use as a workstation.
- C. Use a graphical user interface (GUI) to define and run machine learning experiments from Azure Machine Learning designer.
- D. Create a dataset from a comma-separated value (CSV) file.

ANSWER: A C

Explanation:

In the Azure Machine Learning workspace edition model referenced by this AI-900 objective, the tasks that require an enterprise workspace are “Use a graphical user interface (GUI) to run automated machine learning experiments” and “Use a graphical user interface (GUI) to define and run machine learning experiments from Azure Machine Learning designer.” These features are part of the richer visual authoring experience in Azure Machine Learning studio. Automated machine learning provides a guided interface for selecting data, choosing the prediction task, training candidate models, and comparing results without manually writing all model-training code. Azure Machine Learning designer provides a visual, drag-and-drop environment for building machine learning pipelines by connecting datasets, transformations, training modules, and evaluation steps. Both capabilities are associated with the enterprise workspace experience in the context of this exam topic because they enable low-code/no-code model development through the studio interface. Microsoft documents automated machine learning as a way to automate model selection and tuning, and documents designer as a visual canvas for creating machine learning pipelines. See [Automated machine learning](#) and [Azure Machine Learning designer](#).

QUESTION NO: 62

Which two actions can you perform by using the Azure OpenAI DALL-E model? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. Create images.
- B. Use optical character recognition (OCR).
- C. Detect objects in images.

D. Modify images.

E. Generate captions for images.

ANSWER: A D

Explanation:

Create images. and Modify images. are correct because DALL-E models in Azure OpenAI are designed for image generation scenarios. With DALL-E, a user can submit a natural language prompt and the model creates a new image that matches the prompt, such as a product concept, illustration, scene, or design idea. Microsoft describes DALL-E as a model family that generates images from text prompts, which directly supports creating images as a complete solution.

Modify images. is also a valid DALL-E capability in supported image workflows, particularly with DALL-E 2-style image editing and variation features. These capabilities let users provide an existing image and use prompts or related inputs to create altered versions, variations, or edits based on the source image. In Azure OpenAI documentation, DALL-E is positioned as the image-focused generative AI model family, and the image generation guidance covers using the service to create images from prompts and work with image outputs programmatically. For more detail, see Microsoft's Azure OpenAI model documentation and image generation guidance: [Azure OpenAI Service models](#) and [Generate images with DALL-E](#).