

Microsoft Azure AI Fundamentals

Microsoft AI-900

Version Demo

Total Demo Questions: 63

Total Premium Questions: 633

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Describe Artificial Intelligence workloads and considerations	89
Topic 2, Describe fundamental principles of machine learning on Azure	169
Topic 3, Describe features of computer vision workloads on Azure	136
Topic 4, Describe features of Natural Language Processing (NLP) workloads on Azure	177
Topic 5, Describe features of generative AI workloads on Azure	50
Topic 6, Mix Questions	12
Total	633

QUESTION NO: 1

You have an Internet of Things (IoT) device that monitors engine temperature.

The device generates an alert if the engine temperature deviates from expected norms. Which type of AI workload does the device represent?

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

ANSWER: C

Explanation:

Anomaly detection is the appropriate AI workload because the device evaluates engine-temperature telemetry to identify values or patterns that differ from an established baseline of normal behavior. In this scenario, the important outcome is an alert when a measurement is unexpectedly high, low, or otherwise inconsistent with expected operating conditions. This is a common use of anomaly detection for time-series sensor data from Internet of Things devices.

Establishing expected norms can account for factors such as normal temperature ranges, trends over time, and expected variation during operation. Once a current reading departs sufficiently from that expected pattern, the system can flag it for investigation or trigger an automated response. This supports equipment monitoring and predictive-maintenance workflows, helping identify possible faults before they result in service disruption or damage. Microsoft describes anomaly detection as identifying unusual patterns or observations in data, including monitoring scenarios involving telemetry and time-series data. Learn more in the [Microsoft anomaly detection architecture guidance](#) and the [Azure AI Anomaly Detector overview](#).

QUESTION NO: 2

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

Which service should you use to process the documents?

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

ANSWER: D

Explanation:

QnA Maker is the appropriate service for this scenario because it was designed to build a conversational question-and-answer knowledge base from existing content. It can process sources such as FAQ content, product documentation, manuals, and web pages, organizing the information into question-and-answer pairs that a chatbot can use to respond to natural-language user questions. After the knowledge base is created and published, a website chatbot can submit a user's question to the QnA Maker endpoint and receive the best matching answer from the ingested document content. This directly addresses the requirement to process documents so that responses are grounded in the supplied information rather than being manually authored for every possible question.

QnA Maker is a legacy Azure service in the context of this exam-style question. Microsoft has since replaced it with the custom question answering capability in Azure AI Language, which provides the current approach for creating and maintaining document-based question-answering projects. The underlying use case remains the same: turn curated organizational content into a knowledge base that supports conversational answers. See Microsoft's [custom question answering overview](#) and [QnA Maker overview](#).

QUESTION NO: 3

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 principle of transparency requires that people can understand an AI system's capabilities, limitations, behavior, and the basis for its outcomes. In a loan-approval scenario, an automated recommendation or decision can materially affect a person's access to credit and financial opportunities. The organization therefore needs to be able to communicate meaningful information about the factors that influenced the result, such as the relevant application data and how the system evaluated it. Explainability also supports appropriate human oversight, auditing, and the ability to investigate unexpected or disputed outcomes. It does not require exposing every technical implementation detail; rather, the information provided should be suitable for the stakeholders and the impact of the decision. Microsoft identifies transparency and accountability as principles that help ensure people understand how AI systems operate and can make informed decisions about their use. The explicit requirement that decision factors be explainable maps directly to transparency. See [Microsoft's Responsible AI principles and approach](#) and the [Microsoft Learn module on responsible AI principles](#).

QUESTION NO: 4 - (HOTSPOT)

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Implementing filters to block harmful content as part of a chat solution that uses generative AI is an example of the

	▼	Microsoft responsible AI principle.
accountability		
fairness		
privacy and security		
transparency		

ANSWER:

Answer Area

Implementing filters to block harmful content as part of a chat solution that uses generative AI is an example of the

	▼	Microsoft responsible AI principle.
accountability		
fairness		
privacy and security		
transparency		

Explanation:

Implementing filters to block harmful content in a generative AI chat solution is an example of the **privacy and security** responsible AI principle. Generative AI systems can accept untrusted prompts and can produce content that may be unsafe, abusive, or otherwise harmful if suitable protections are not put in place. Content filters are a practical safeguard that helps

protect users and the application environment by identifying harmful input and output categories and preventing that material from being processed or displayed when appropriate.

Privacy and security is concerned with protecting people, data, and systems throughout an AI solution's lifecycle. For a chat application, this includes applying technical controls that reduce misuse and help the service operate in a protected, trustworthy manner. Filters are one layer of defense used alongside other measures such as access controls, monitoring, data protection, and abuse prevention. Applying these safeguards is particularly important for generative AI because responses are dynamically generated and the solution may be exposed to a wide variety of user requests.

Microsoft provides configurable content filtering for generative AI services to help organizations detect and mitigate potentially harmful content. These controls support responsible deployment by helping teams establish safer boundaries for conversations and reduce exposure to risks in their applications. Microsoft's documentation describes content filtering and related harm-mitigation capabilities for Azure OpenAI Service at [Azure OpenAI content filtering](#). Microsoft also explains its broader responsible AI commitments, including privacy and security, at [Microsoft Responsible AI](#).

QUESTION NO: 5

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

NOTE: Each correct selection is worth one point.

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

ANSWER: A C

Explanation:

Coefficient of determination (R2) and root mean squared error (RMSE) are both standard metrics for evaluating regression models, which predict continuous numeric values. R2 measures the proportion of variation in the observed target values that is explained by the model's predictions. Its value is commonly interpreted relative to a baseline model that predicts the mean target value; a value closer to 1 generally indicates that the model accounts for more of the variation in the data. RMSE measures the typical magnitude of prediction errors by squaring the differences between predicted and actual values, averaging those squared errors, and then taking the square root. Because it is expressed in the same units as the target variable, RMSE is particularly useful for assessing the practical size of a model's errors. Squaring the errors also gives greater weight to larger deviations, helping identify models that make substantial prediction mistakes. For RMSE, lower values represent predictions that are closer to the actual outcomes. Microsoft's Azure Machine Learning documentation lists R2 and RMSE among the metrics reported for regression evaluation. See [Azure Machine Learning regression metrics](#) and [Understand regression metrics](#).

QUESTION NO: 6

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

NOTE: Each correct answer is worth one point.

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

ANSWER: A C

Explanation:

Grounding data can be used to constrain the output of the chatbot. Grounding supplies relevant, trusted information—such as technical documentation, knowledge-base articles, or support procedures—as context for a user's question. A retrieval-augmented generation design retrieves pertinent content and sends it with the prompt, helping the model produce answers that are relevant to the organization's approved data. Azure OpenAI on your data supports this approach by connecting a model to enterprise data, commonly through Azure AI Search. Grounding improves answer relevance and can support citations, but it does not make a generative model infallible.

The chatbot might respond with inaccurate data. GPT-3.5 generates text by predicting likely token sequences; it does not independently verify factual claims or guarantee that every response is correct. Output can be incomplete, misleading, or plausible-sounding but inaccurate, particularly if the supplied grounding content is missing, irrelevant, ambiguous, or misunderstood. A responsible production solution should therefore evaluate responses and apply appropriate safeguards, such as source citations, monitoring, human review where warranted, and clear user guidance. Microsoft documents both grounding with organizational data and the need to evaluate generative-AI output. See [Azure OpenAI on your data](#) and [Evaluate generative AI applications](#).

QUESTION NO: 7

Which term is used to describe uploading your own data to customize an Azure OpenAI model?

- A. completion
- B. grounding
- C. fine -tuning
- D. prompt engineering

ANSWER: C

Explanation:

Fine-tuning is the process of customizing a supported Azure OpenAI base model by training it further with a dataset that you provide. The dataset contains task-specific examples, commonly formatted as inputs and the expected outputs, enabling the model to learn the desired behavior, terminology, style, structure, or response patterns for a particular use case. After the training data is uploaded and a fine-tuning job completes, the resulting customized model can be deployed and called by an application.

Fine-tuning is especially useful when an application needs consistent output for a narrowly scoped task and has enough high-quality examples to demonstrate that task clearly. It can also help reduce the amount of repeated instructional content needed in each request, because the model has learned the patterns from the supplied training examples. Azure OpenAI documentation describes fine-tuning as creating a customized version of a base model using training data, which directly matches the scenario of uploading your own data to customize the model. See [Azure OpenAI fine-tuning concepts](#) and [how to fine-tune Azure OpenAI models](#).

QUESTION NO: 8

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** provide the two parts needed for this FAQ-style chatbot solution. QnA Maker was designed to create a knowledge base of predefined question-and-answer pairs, including content imported from FAQ pages, product documentation, or manually authored responses. It identifies the user's intent by matching their submitted question to the most relevant entry and returns the stored answer. This is well suited to handling repetitive, simple enquiries that would otherwise require a telephone operator.

Azure Bot Service supplies the conversational bot layer that receives messages from users and connects the bot to supported communication channels. It enables the question-answering capability to be delivered as an interactive chat experience, allowing customers to obtain common answers through a self-service interface before escalation to a human agent is necessary.

Although QnA Maker has been retired for new usage and its capabilities are available through custom question answering in Azure AI Language, QnA Maker remains the applicable service name in this question's set of choices and describes the required predefined-answer knowledge-base function. See [Azure AI Language question answering overview](#) and [Azure Bot Service overview](#).

QUESTION NO: 9

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

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

NOTE: Each correct selection is worth one point.

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

ANSWER: C E

Explanation:

speech-to-text is required because verbal commands are received as audio. Before the solution can determine what action to perform, it must recognize the user's spoken words and convert the audio stream into text. Azure AI Speech provides speech recognition for this purpose, enabling an application to process commands such as requests to turn on a light or adjust a thermostat.

language modeling is also required to determine the meaning of the recognized text. A smart-device control application must identify the user's intent, such as switching on, switching off, or changing a setting, and extract relevant details such as the target device, room, or requested value. In Azure AI services, conversational language understanding uses trained language models to classify intents and extract entities from a user's utterance, allowing the application to map the command to an appropriate device action. Together, speech recognition converts the spoken command into usable text, and language modeling provides the understanding needed to act on that command. For more information, see [Azure AI Speech speech-to-text documentation](#) and [Azure conversational language understanding documentation](#).

QUESTION NO: 10

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 a supported scenario because Azure Form Recognizer, now called Azure AI Document Intelligence, analyzes document content and returns structured data. Its prebuilt invoice model is designed for invoices and can extract common fields, including the invoice ID or invoice number, invoice date, vendor and customer details, totals, taxes, due dates, and line items. This enables applications to automate invoice processing without manually reading each document or first training a custom model for standard invoice layouts.

Identify the retailer from a receipt is also supported because the prebuilt receipt model extracts structured information from sales receipts. The returned fields include the merchant name, which identifies the retailer, along with transaction date and time, purchased items, subtotal, tax, and total. These prebuilt document models are intended for extracting text, layout, tables, and meaningful business fields from documents such as invoices and receipts. Microsoft documents the service and its document-analysis capabilities in the [Azure AI Document Intelligence overview](#), and lists merchant name among the fields returned by the [prebuilt receipt model](#).

QUESTION NO: 11

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

ANSWER: A C

Explanation:

Identify the retailer from a receipt is a supported document-processing scenario. Form Recognizer, now called Azure AI Document Intelligence, provides a prebuilt receipt model that analyzes receipt documents and returns structured fields. These include merchant information, which represents the retailer or business associated with the receipt, along with transaction date, purchased items, taxes, totals, and other relevant receipt data. Applications can use this extracted merchant field to automate expense processing and financial workflows.

Extract the invoice number from an invoice is also supported through Azure AI Document Intelligence's prebuilt invoice model. This model recognizes common invoice structures and extracts normalized fields, including the invoice ID, vendor and customer information, invoice dates, line items, subtotal, tax, and total amount. The invoice ID field supplies the invoice-number information needed for indexing, validation, accounts-payable automation, and integration with business systems. Both scenarios involve converting information contained in semi-structured business documents into machine-readable data, which is a core capability of the service. Microsoft documents the available prebuilt models and their extracted fields in the [Azure AI Document Intelligence overview](#) and the [prebuilt invoice model documentation](#).

QUESTION NO: 12

Which two pieces of information are returned by an object detection solution? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. an object label
- B. a bounding box
- C. a translated text value
- D. a predicted numeric value
- E. a speech waveform

ANSWER: A B

Explanation:

An object label is returned to identify the detected object, such as a car, bicycle, or person. **A bounding box** is returned to indicate the object's location in the image by using coordinates and dimensions.

Object detection differs from image classification because it identifies individual objects and their positions. See [Object detection concepts in Azure AI Vision](#).

QUESTION NO: 13

Stating the source of the data used to train a model is an example of which responsible AI principle?

- A. fairness
- B. transparency
- C. reliability and safety
- D. privacy and security

ANSWER: B

Explanation:

transparency is correct. In responsible AI, transparency means providing people with meaningful information about an AI system, including how it was created, what it is intended to do, and important factors that can affect its outputs. Identifying the source of training data is a form of data provenance: it documents where influential data originated and gives stakeholders context for assessing the model's intended use, likely behavior, constraints, and potential limitations.

This information supports appropriate human understanding and oversight. For example, a user or reviewer can use training-data source information to determine whether the model was developed using data relevant to the scenario in which it will be deployed. Clear documentation also makes it easier to communicate material facts about the system to affected people and to support scrutiny throughout the AI lifecycle. Microsoft identifies transparency as one of its responsible AI principles and describes it as ensuring people can understand AI system capabilities, limitations, and behavior. Microsoft Learn also discusses transparency as providing information that helps users understand how AI systems are developed and used. See [Microsoft Responsible AI principles and approach](#) and [Microsoft Learn: Responsible AI principles](#).

QUESTION NO: 14

You use drones to identify where weeds grow between rows of crops to send an instruction for the removal of the weeds. This is an example of which type of computer vision?

- A. scene segmentation
- B. optical character recognition (OCR)
- C. object detection

ANSWER: C

Explanation:

Object detection is the appropriate computer vision capability because the drone imagery must be used to identify weeds as objects and establish their positions among the crop rows. Detection models return both a class label, such as a weed, and location information, commonly represented by bounding-box coordinates. This spatial output enables a downstream system to produce targeted removal instructions—for example, directing precision spraying equipment or autonomous machinery to the areas where weeds were found. In agricultural imaging, this capability supports selective treatment by turning visual observations into actionable locations rather than simply providing a general description of the scene. Azure AI Vision describes object detection as identifying and locating objects in an image, which aligns directly with detecting weeds and determining where they grow. Learn more in the [Azure AI Vision object detection documentation](#) and the [Azure AI Vision overview](#).

QUESTION NO: 15

Which service should you use to extract text, key/value pairs, and table data automatically from scanned documents?

- A. Form Recognizer
- B. Text Analytics
- C. Language Understanding
- D. Custom Vision

ANSWER: A

Explanation:

Form Recognizer is the correct service for automatically extracting text, key/value pairs, and table data from scanned documents. The service has since been renamed Azure AI Document Intelligence, but “Form Recognizer” remains the expected answer when this terminology is used in an exam question. It applies optical character recognition together with document-layout analysis to turn image-based or scanned documents into structured, machine-readable results.

Its models can recognize printed and handwritten text, identify the spatial layout of pages, detect tables and their cells, and extract labeled fields as key/value pairs. This makes it suitable for common business-document workflows involving invoices, receipts, forms, contracts, identity documents, and custom document types. Applications can use the returned structured data to automate validation, data entry, routing, and downstream processing without requiring a person to manually transcribe each document. Microsoft’s documentation specifically describes Azure AI Document Intelligence as extracting text, key-value pairs, tables, and document structure from documents. See the [Azure AI Document Intelligence overview](#) and [document layout analysis documentation](#).

QUESTION NO: 16

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, commonly called text to speech, converts written text into synthesized spoken audio. An automated voice that reads back a credit card number entered into a telephone by using a numeric keypad is a speech synthesis use case because the application can represent the entered digits as text and use a synthetic voice to speak them back to the caller. This capability is frequently used in interactive voice response applications to confirm information and deliver prompts without a human operator.

An AI character in a computer game that speaks audibly to a player is also a complete speech synthesis scenario. The game can provide dialogue as predefined text or dynamically generated text, then convert that text to natural-sounding audio at runtime. Azure AI Speech text to speech offers neural voices designed to render text as lifelike speech, supporting interactive applications, accessibility features, bots, and other experiences that need spoken output. The core requirement in both scenarios is that text or textual data must be rendered as audio for a listener. Microsoft documents text to speech as a Speech service capability that converts text into synthesized speech: [Azure AI Speech text to speech](#). For the broader set of Azure AI Speech capabilities, see the [Azure AI Speech overview](#).

QUESTION NO: 17

You have a dataset that contains information about taxi journeys that occurred during a given period. You need to train a model to predict the fare of a taxi journey.

What should you use as a feature?

- A. the number of taxi journeys in the dataset
- B. the trip distance of individual taxi journeys
- C. the fare of individual taxi journeys
- D. the trip ID of individual taxi journeys

ANSWER: B

Explanation:

The trip distance of individual taxi journeys is an appropriate feature because a feature is an input variable supplied to a machine learning model to help it learn a relationship with the value to be predicted. Here, the prediction target is the taxi fare, which is a numeric value; therefore, this is a regression scenario. Trip distance describes a characteristic of each individual journey and is generally available as part of the journey data. It also has a meaningful relationship with the fare: longer trips commonly result in higher charges, subject to factors such as time, route, and local pricing rules. During training, the model can learn how trip distance, potentially alongside other journey characteristics, relates to historical fare values. When predicting a new journey's fare, its trip distance can then be provided as an input to the trained model. Microsoft describes regression as a machine learning technique for predicting numeric values and distinguishes the input features from the label, which is the value the model is trained to predict. Microsoft's taxi fare regression tutorial likewise uses trip information as model input to predict a fare amount. See [Understand regression in Microsoft Learn](#) and [Microsoft's taxi fare prediction tutorial](#).

QUESTION NO: 18 - (DRAG DROP)

DRAG DROP

You need to scan the news for articles about your customers and alert employees when there is a negative article. Positive articles must be added


to a press book.

Which natural language processing tasks should you use to complete the process? To answer, drag the appropriate tasks to the correct locations.


Each task may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:



ANSWER:



Explanation:

The workflow first needs to identify which news articles are relevant to a customer. **Entity recognition** is the appropriate task for the Filter stage because it detects and categorizes named items in unstructured text, including organization names. By extracting entities from each news article, the process can recognize whether the article mentions one of the organization's customers and pass only those relevant articles forward. Microsoft Azure AI Language provides named entity recognition specifically for locating entities such as people, organizations, locations, events, and products in text. This supports a reliable customer-mention filter without requiring the text to follow a fixed format. See [Named Entity Recognition in Azure AI Language](#).

After an article is confirmed to mention a client, the workflow must determine the tone of its coverage. **Sentiment analysis** is the appropriate task for the Assess stage because it evaluates text and returns sentiment classifications, such as positive, neutral, or negative, with associated confidence scores. That result directly drives the two business actions in the diagram: negative coverage can create an employee alert, while positive coverage can be collected in a press book. Sentiment analysis is designed for evaluating the opinion or emotional tone expressed in written content, making it suitable for automated review of news articles at scale. Microsoft documents this capability in [Sentiment analysis and opinion mining in Azure AI Language](#).

QUESTION NO: 19

You build a QnA Maker bot by using a frequently asked questions (FAQ) page.

You need to add professional greetings and other responses to make the bot more user friendly. What should you do?

- A. Increase the confidence threshold of responses
- B. Enable active learning
- C. Create multi-turn questions
- D. Add chit-chat

ANSWER: D

Explanation:

Add chit-chat is the appropriate action because QnA Maker supports a dedicated chit-chat capability for common conversational interactions, including greetings, farewells, expressions of thanks, and other small-talk messages. An FAQ-based knowledge base supplies answers to the organization's domain-specific questions, but it does not inherently provide the conversational responses that make interactions feel polished and welcoming. Adding a professional chit-chat personality inserts a predefined set of question-and-answer pairs into the knowledge base, which can then be tailored to the desired tone and organizational voice.

This is especially useful when the bot must respond naturally to messages such as "hello," "good morning," or "thank you," even when those messages are unrelated to the FAQ content. The chit-chat content is managed as part of the knowledge base, so its responses can be reviewed and customized before publishing the bot. Microsoft documents that chit-chat adds conversational personality to a QnA Maker knowledge base through predefined personality-based content. See [Add chit-chat to a QnA Maker knowledge base](#) and the [QnA Maker overview](#).

QUESTION NO: 20

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

Azure Machine Learning: This service is used for building, training, and deploying custom machine learning models, not for chatbot Q & A automation.

- C. Translator

Translator: This service performs language translation, which is not required for answering predefined questions unless multilingual support is specifically needed.

Therefore, to implement a chatbot that can answer simple, repetitive user questions and reduce the load on human operators, you combine Azure Bot Service (for interaction) with the Language Service (for question-answering intelligence).

- D. Language Service

ANSWER: A D

Explanation:

Azure Bot Service and Azure AI Language Service are the appropriate services for a telephone-support chatbot that responds to common questions using predefined answers. Azure Bot Service supplies the conversational application

platform. It enables a bot to receive user messages, manage the conversation, and connect the completed bot to supported communication channels. This provides the customer-facing interaction layer that can handle routine requests before an operator needs to become involved.

Azure AI Language Service provides the question answering capability needed for the bot's knowledge component. Question answering is designed to return answers from a curated knowledge base, including manually authored question-and-answer pairs, FAQ content, and supported documents. An organization can define the standard answers for frequent operator inquiries and use natural-language matching to select the most relevant response to a caller's question. The bot can then deliver that response automatically. This combination implements the standard architecture of a conversational interface backed by a managed question-and-answer knowledge base, reducing repetitive work for telephone operators. See the [Azure Bot Service overview](#) and [Azure AI Language question answering overview](#).

QUESTION NO: 21

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

- A. Face
- B. Custom Vision
- C. Computer Vision
- D. Form Recognizer

ANSWER: C

Explanation:

Computer Vision is the appropriate service because Azure AI Vision provides prebuilt image-analysis models that can identify well-known brands and logos in photographs. This capability is designed for scenarios in which an application submits an image to an API and receives analysis results without requiring the organization to collect training images, label data, or train and maintain its own model.

Brand detection is part of Azure AI Vision's image-analysis features. When a recognizable logo or brand mark is present, the service can return the brand name along with details such as its confidence score and the region of the image in which it was detected. This makes it a direct fit for identifying established brands in general digital photographs. In AI-900 terminology, Computer Vision is the commonly used service name for these prebuilt visual-analysis capabilities; Microsoft now groups these capabilities under Azure AI Vision. Microsoft's documentation describes Azure AI Vision as a service for extracting visual information from images and documents, including brand detection. See the [Azure AI Vision overview](#) and the [brand detection documentation](#).

QUESTION NO: 22

You need to predict the sea level in meters for the next 10 years. Which type of machine learning should you use?

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

ANSWER: B

Explanation:

Regression is the appropriate machine learning type because the required result is a numeric measurement: sea level expressed in meters. Regression is a supervised machine learning task that learns from labeled historical data to predict a continuous numerical value. For this scenario, training data could include prior sea-level measurements over time, alongside

relevant variables such as dates, temperatures, ocean conditions, or climate indicators. The trained model can then estimate sea-level values for future dates covering the next 10 years.

Although predicting values at future points in time is often called time-series forecasting, the underlying machine learning task remains regression when the target is a continuous number. The central distinction is the form of the output: meters can take a range of numeric values, including decimal values, rather than belonging to a fixed set of categories. Microsoft describes regression as predicting numeric values, and Azure Machine Learning supports regression tasks for continuous outcomes. See [Microsoft Learn: Regression](#) and [Microsoft Learn: Automated ML regression](#).

QUESTION NO: 23

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:

The correct natural language processing scenarios are "a smart device in the home that responds to questions such as, 'What will the weather be like today?'" and "a website that uses a knowledge base to interactively respond to users' questions." Natural language processing enables applications to work with human language, whether it is spoken or written. A home device responding to a weather question must recognize and interpret a person's request, identify its intent, and formulate a useful response. This can involve speech recognition to convert spoken input to text, language understanding to determine what information is requested, and speech synthesis to deliver a spoken answer.

A knowledge-base website that responds interactively to questions also processes written natural-language input. It identifies the meaning of a user's question and retrieves or generates an appropriate answer from available knowledge. Azure AI Language includes question-answering capabilities designed to create conversational experiences over structured question-and-answer content. More broadly, Azure AI Language provides natural-language features for analyzing and understanding text, while Azure AI Speech supports speech-to-text and text-to-speech interactions. For more information, see the [Azure AI Language documentation](#) and the [Azure AI Speech documentation](#).

QUESTION NO: 24 - (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:

The correct option is **clustering** (an unsupervised machine learning task).

RFM analysis groups customers into segments based on similarities in their recency, purchase frequency, and monetary value. Because the groups are identified from the data rather than assigned predefined labels, this is clustering.

QUESTION NO: 25

A company employs a team of customer service agents to provide telephone and email support to customers. The company develops a webchat bot to provide automated answers to common customer queries.

Which business benefit should the company expect as a result of creating the webchat bot solution?

- A. increased sales
- B. a reduced workload for the customer service agents
- C. improved product reliability

ANSWER: B

Explanation:

A reduced workload for the customer service agents is the expected business benefit. A webchat bot can automatically handle high-volume, routine customer questions, such as requests for opening hours, order status, account information, standard policy details, and basic troubleshooting. Customers receive immediate responses through the chat channel without every common query needing to become a telephone call or email interaction handled by a person.


Reducing these repetitive contacts lets the support team focus its time on cases that genuinely require human expertise, empathy, investigation, or decision-making. This can improve the scalability of customer support because the organization can serve more customers without increasing the agent workload at the same rate. It can also provide assistance outside normal support hours when the bot is made available continuously. Microsoft describes Azure Bot Service as a managed service for developing and deploying conversational bots across channels, and its virtual-agent architecture describes using a bot to answer frequently asked questions and hand off more complex conversations to human agents. See the [Azure Bot Service overview](#) and [Microsoft conversational bot architecture guidance](#).

QUESTION NO: 26 - (HOTSPOT)

Select the answer that correctly completes the sentence.

Answer Area


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



ANSWER:

Answer Area

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



Explanation:

Regression is the correct machine learning task because the requested prediction is a vehicle's miles-per-gallon value. Miles per gallon is a numeric measurement that can take many possible values, including decimal values, rather than belonging to a fixed set of named groups. A model can use vehicle characteristics such as weight, engine power, engine size, fuel-system details, and other measured attributes to learn patterns from historical examples and estimate the expected fuel efficiency of another vehicle.

This is a supervised learning scenario: the training data contains examples of vehicles with known feature values and known miles-per-gallon results. During training, the model learns the relationship between those input features and the numeric

target. When presented with the features of a new vehicle, it produces an estimated numeric fuel-efficiency value. The output can vary across a continuous range, which is the defining characteristic of a regression prediction.

Microsoft describes regression as the machine learning approach used to predict numerical values. Typical regression targets include a house price, a sales amount, a temperature, a demand quantity, or another measured value. Fuel economy follows exactly the same pattern because the desired result is a number representing how far a vehicle can travel per unit of fuel. Therefore, the sentence is correctly completed with **regression**. For Microsoft's overview of machine-learning task types, see [What is machine learning?](#) and the [Azure Machine Learning automated ML documentation](#).

QUESTION NO: 27

What is a use case for classification?

- A. predicting how many cups of coffee a person will drink based on how many hours the person slept the previous night.
- B. analyzing the contents of images and grouping images that have similar colors
- C. predicting whether someone uses a bicycle to travel to work based on the distance from home to work
- D. predicting how many minutes it will take someone to run a race based on past race times

ANSWER: C

Explanation:

Predicting whether someone uses a bicycle to travel to work based on the distance from home to work is a classification use case because the required prediction is a discrete label: the person either uses a bicycle for commuting or does not. Classification is a supervised machine learning technique that learns from examples in which each input has an associated category, then assigns the most appropriate category to new inputs. In this scenario, distance from home to work is an input feature, while bicycle use is the target label. Although the distance itself is numeric, classification can use numeric, categorical, and other feature types to predict a categorical outcome. This is specifically a binary classification problem, as there are two possible classes. Such predictions are useful when a decision, category, or membership result is needed rather than a measured quantity. Microsoft describes classification as a machine learning task for predicting categories, including two-class and multiclass classification. For further details, see [Microsoft Learn: Classification](#) and [Microsoft Learn: Fundamentals of machine learning](#).

QUESTION NO: 28

A smart device that responds to the question. "What is the stock price of Contoso, Ltd.?" is an example of which AI workload?

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

ANSWER: D

Explanation:

natural language processing is the correct AI workload because the smart device needs to interpret a question expressed in human language, identify the user's intent, and recognize the company name as the entity for which stock-price information is requested. Natural language processing enables applications to work with written or spoken language, including understanding text, extracting meaning, answering questions, and supporting conversational interactions. In a smart-assistant scenario, the request might first be converted from speech to text if it was spoken aloud, but processing the meaning of "What is the stock price of Contoso, Ltd.?" remains a natural language processing task. The application can then use the identified intent and entity to retrieve the relevant current price and formulate a response. Microsoft characterizes

Azure AI Language as providing natural-language capabilities for building applications that understand and analyze text, while Azure AI Speech supports speech-enabled experiences that can be used alongside language capabilities in voice-driven applications. See the [Azure AI Language overview](#) and the [Azure AI Speech overview](#).

QUESTION NO: 29

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

- A. AI systems must use only publicly available data.
- B. AI systems must protect the interests of the company.
- C. AI systems must be understandable.
- D. AI systems must keep personal details public.

ANSWER: C

Explanation:

AI systems must be understandable. This statement represents Microsoft's responsible AI principle of transparency. Transparency means that people interacting with, deploying, or affected by an AI system should be able to understand its purpose, capabilities, limitations, and the appropriate interpretation of its outputs. It also involves communicating relevant information about how the system was developed and evaluated, so stakeholders can make informed decisions about whether and how to use it. Understandability supports trust in AI by reducing the "black box" nature of automated outputs and enabling users and organizations to recognize when human review, additional context, or caution is needed. For example, a system should clearly communicate that an output is AI-generated, describe important limitations such as potential inaccuracies, and provide meaningful information about how it is intended to be used. Transparency is one of Microsoft's six responsible AI principles, alongside fairness, reliability and safety, privacy and security, inclusiveness, and accountability. Microsoft Learn describes transparency as ensuring that AI systems and their capabilities are understandable to people, while Microsoft's responsible AI resources explain how this principle is applied across the AI lifecycle. See [Microsoft Learn: Responsible AI principles](#) and [Microsoft AI principles and approach](#).

QUESTION NO: 30

You are building a Conversational Language Understanding model for an e-commerce business. You need to ensure that the model detects when utterances are outside the intended scope of the model.

What should you do?

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

ANSWER: C

Explanation:

Add utterances to the None intent. In Conversational Language Understanding, the None intent is designed to capture utterances that do not belong to any of the application's defined, domain-specific intents. For an e-commerce application, populate this intent with representative examples of requests outside the supported shopping experience, such as unrelated questions or conversational messages that do not concern products, orders, returns, or payments. These examples establish a meaningful boundary around the application's intended capabilities.

During training, CLU uses the None-intent examples to learn patterns associated with out-of-scope input. When a user submits a similar utterance, the model can classify it as None rather than assigning it inaccurately to a supported business intent. A well-designed None intent should contain varied, realistic out-of-domain utterances that reflect messages the

application may receive in production. This supports a safer conversational experience because the application can recognize that it should ask the user to rephrase, redirect them to supported tasks, or provide an appropriate fallback response. Microsoft describes the None intent as the intent used for utterances that do not map to the other intents in a project. See [None intent in Conversational Language Understanding](#) and [Build a CLU model](#).

QUESTION NO: 31

What is a form of unsupervised machine learning?

- A. multiclass classification
- B. clustering
- C. binary classification
- D. regression

ANSWER: B

Explanation:

Clustering is a form of unsupervised machine learning because it identifies meaningful groups within data that has no preassigned labels. Rather than learning a mapping from inputs to a known outcome, a clustering algorithm evaluates characteristics of the data points—such as similarity, distance, or shared patterns—and places related items into the same group. The resulting groups are inferred from the data itself, so they can reveal structure that was not defined in advance.

For example, an organization could use clustering to segment customers according to purchasing activity, product preferences, or engagement patterns. The organization supplies the customer data but does not need to state beforehand which customer belongs to which segment. After clustering, the discovered segments can support tailored marketing, service design, or further analysis.

Microsoft Learn identifies clustering as an unsupervised machine learning task and describes it as grouping observations with similar characteristics. This makes it appropriate when the objective is to explore unlabeled data and discover natural groupings. See [Microsoft Learn: Fundamentals of machine learning](#) and [What is automated machine learning?](#)

QUESTION NO: 32

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. QnA Maker
- B. Azure Bot Service
- C. Form Recognizer
- D. Anomaly Detector
- E. Azure AI Language question answering

ANSWER: B E

Explanation:

Azure Bot Service and Azure AI Language question answering provide the two core capabilities required for this solution. Azure Bot Service supports the conversational application layer: developers use it to build, deploy, manage, and connect a bot to user-facing channels such as a website, Microsoft Teams, or other supported messaging endpoints. It therefore provides the interface through which users submit natural-language questions and receive responses.

Azure AI Language question answering provides the knowledge-base question-and-answer capability. It can create a question-answering project from curated sources such as FAQ content, documents, or manually authored question-and-answer pairs, then identify an appropriate answer for a user's natural-language query. A bot can call this capability during a conversation and return the selected answer to the user. This is the current Azure service for managed knowledge-base question answering; it supersedes the legacy QnA Maker service. Combining the bot's conversational channel integration with Azure AI Language question answering produces a natural-language bot that can answer questions from a knowledge base. See the [Azure Bot Service overview](#) and [Azure AI Language question answering overview](#).

QUESTION NO: 33

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, also called speech-to-text, converts spoken audio into written text. Providing closed captions for recorded or live videos is a direct speech-recognition scenario because the service identifies words spoken in the audio and outputs text that can be displayed as captions. This supports accessibility for people who are deaf or hard of hearing and also enables viewers to follow content in situations where audio cannot be played. Azure AI Speech provides real-time and batch speech-to-text capabilities that can be used to generate captions from live streams or recorded media.

Creating a transcript of a telephone call or meeting is also a speech-recognition scenario. The input is conversational speech, and the required output is a textual record that can be stored, searched, reviewed, or analyzed. Azure AI Speech supports transcription from audio sources and includes capabilities suited to conversation and meeting transcription. In practical implementations, the resulting transcript can support documentation, compliance review, call analytics, and improved content discoverability. Microsoft describes Azure AI Speech speech-to-text as transcribing audio streams to text and documents transcription features for both real-time and recorded audio. See [Speech to text documentation](#) and the [Azure AI Speech overview](#).

QUESTION NO: 34

You are developing a generative AI solution that provides answers to employees about company policies.

Which two actions can help reduce the risk of the solution providing unsupported answers? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Ground the model with approved policy documents.
- B. Instruct the model to state when the supplied information does not contain an answer.
- C. Remove all system messages.
- D. Assume that generated responses are always factually correct.
- E. Use only a larger maximum response length.

ANSWER: A B

Explanation:

Ground the model with approved policy documents provides relevant and authoritative context that the model can use when preparing an answer. This retrieval-augmented generation pattern helps align responses with organizational information.

Instruct the model to state when the supplied information does not contain an answer establishes a clear response behavior for cases where retrieved content is insufficient. This can reduce unsupported responses by directing the solution to acknowledge limitations rather than invent details.

Generative AI output should be evaluated and monitored because grounding and instructions improve response quality but do not guarantee correctness. See [Azure OpenAI on your data](#) and [Prompt engineering techniques](#).

QUESTION NO: 35

Which scenario is an example of a webchat bot?

- A. Determine whether reviews entered on a website for a concert are positive or negative, and then add a thumbs up or thumbs down emoji to the reviews.
- B. Translate into English questions entered by customers at a kiosk so that the appropriate person can call the customers back.
- C. Accept questions through email, and then route the email messages to the correct person based on the content of the message.
- D. From a website interface, answer common questions about scheduled events and ticket purchases for a music festival.

ANSWER: D

Explanation:

From a website interface, answer common questions about scheduled events and ticket purchases for a music festival. is an example of a webchat bot because it describes a conversational application that is presented directly within a website. Visitors can type natural-language questions in a chat interface and receive automated, real-time responses about common topics, such as event times, ticket availability, or purchasing guidance. This lets the festival provide self-service support at any time while giving users an interaction model similar to messaging with a customer-service representative. A bot can use a curated question-and-answer knowledge source to respond consistently to frequently asked questions, and it can be connected to a website through a web chat channel. Microsoft describes Azure Bot Service as a managed service for developing and deploying conversational bots that communicate with users through channels. The Web Chat channel is designed specifically to embed a bot chat experience in a web page. For more information, see the [Azure Bot Service overview](#) and [Connect a bot to Web Chat](#).

QUESTION NO: 36

Which two tasks can be performed by using the Translator service? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Translate product descriptions from English to Japanese.
- B. Detect the language of a customer comment.
- C. Extract tables from a purchase order.
- D. Identify the location of a bicycle in an image.
- E. Convert a telephone call into a transcript.

ANSWER: A B


Explanation:

Translate product descriptions from English to Japanese is supported because Translator converts text between supported languages. An application can submit source text and specify a target language to obtain translated text.

Detect the language of a customer comment is also supported. Translator can identify the language of submitted text, which is useful when an application must determine the source language before translating or routing content. See the [Azure Translator documentation](#).

QUESTION NO: 37

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. Regression predicts a continuous numerical value, such as a house price, sales amount, temperature, or demand level. The chart plots the value predicted by the model against the corresponding true, or observed, value in evaluation data. A diagonal reference line represents perfect predictions: every point on that line has a predicted value equal to its true value. Therefore, points concentrated close to the line indicate that the regression model's estimates generally align well with actual values. The chart can also reveal patterns relevant to regression quality, such as predictions consistently being too high or too low over part of the value range, or errors increasing for larger target values. In Azure Machine Learning, regression runs are evaluated using regression-oriented visualizations and metrics, including predicted-versus-true behavior, mean absolute error, root mean squared error, and the coefficient of determination. This makes the chart a direct visual assessment of how accurately a model predicts a numeric target. See [Understand automated ML results in Azure Machine Learning](#) and [What is automated machine learning?](#).

QUESTION NO: 38

You need to create a model that labels a collection of your personal digital photographs. Which Azure AI service should you use?

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

ANSWER: D

Explanation:

Azure AI Custom Vision is the appropriate service because the requirement is to create a tailored model that applies labels to photographs. Custom Vision supports custom image classification: you provide representative images, assign tags that represent the labels you want, and train a model to recognize those labels in new images. For example, a personal photo collection could be tagged with categories such as family members, pets, holidays, locations, or events, provided the training images are labeled consistently and sufficiently represent each category. After training, the model returns predicted tags and associated confidence scores for images it has not previously seen, enabling automated organization of the collection. This approach is specifically intended for scenarios in which the image categories are defined by the user rather than limited to a fixed set of general-purpose visual descriptions. Microsoft describes Custom Vision as a service for building custom image classification and object-detection models using labeled images. The image classification workflow—create a project, upload and tag images, train, evaluate, and use the resulting model—directly fulfills the need to create a photograph-labeling model. See [Azure AI Custom Vision overview](#) and [Build an image classifier](#).

QUESTION NO: 39

Which two capabilities can be used to build a voice-enabled virtual assistant? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. speech-to-text
- B. text-to-speech
- C. object detection
- D. regression
- E. document classification

ANSWER: A B

Explanation:

Speech-to-text converts a user's spoken request into text that can be processed by the virtual assistant. **Text-to-speech** converts the assistant's textual response into synthesized audio that can be played to the user.

These Azure AI Speech capabilities support a natural voice interaction: the user speaks, the application processes the recognized text, and the solution responds audibly. See the [Azure AI Speech overview](#).

QUESTION NO: 40

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

NOTE: Each correct selection is worth one point.

- A. knowledgeability
- B. decisiveness

- C. inclusiveness
- D. fairness
- E. opinionatedness
- F. reliability and safety

ANSWER: C D F

Explanation:

Microsoft identifies six foundational principles for responsible AI: fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. Therefore, **inclusiveness**, **fairness**, and **reliability and safety** are three valid guiding principles. Inclusiveness means designing AI to empower and engage people with diverse backgrounds, abilities, and circumstances, rather than creating barriers to participation. Fairness requires AI systems to be designed and assessed to minimize unjustified bias and to avoid producing inequitable outcomes for different groups. Reliability and safety means an AI system should perform consistently under expected conditions, be appropriately tested and monitored, and respond safely when unexpected conditions or failures occur. These principles apply throughout an AI system's lifecycle, from defining intended use and collecting data through development, deployment, monitoring, and governance. They help organizations build AI solutions that deserve user and stakeholder trust while reducing the risk of harm. Microsoft's Responsible AI resources describe these principles and their practical application in Azure-based solutions. See [Microsoft's Responsible AI overview](#) and [Responsible AI in Azure](#).

QUESTION NO: 41

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. Azure Form Recognizer, now called Azure AI Document Intelligence, enables organizations to create models tailored to their own business documents. By supplying representative document samples and defining the fields that matter, such as an account number, invoice total, claim identifier, or date, the service learns to extract information from layouts and structures relevant to that organization.

This is particularly valuable when documents use proprietary, industry-specific, or organization-specific formats that are not fully addressed by a prebuilt model. Custom extraction models can handle structured, semi-structured, and unstructured documents. In addition, a composed custom model can combine multiple custom models, allowing Document Intelligence to identify which supported document type was submitted and use the appropriate model for extraction. This provides a practical way to support several related form layouts or document types within a business workflow while maintaining fields and extraction behavior designed for those documents. Microsoft documents custom model capabilities at [Custom models in Azure AI Document Intelligence](#) and describes combining models at [Composed custom models](#).

QUESTION NO: 42

A website must provide spoken versions of its written news articles for users who prefer to listen to the content.

Which capability should you use?

- A. speech to text

- B. text to speech
- C. language detection
- D. optical character recognition (OCR)

ANSWER: B

Explanation:

Text to speech is appropriate because it converts written text into spoken audio by using a synthesized voice. The website can submit article text and play the returned audio to users.

Speech to text performs the reverse task by converting spoken audio into text. See [Azure AI Speech text to speech](#).

QUESTION NO: 43

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. improves response quality by giving GPT-3.5 relevant, current, and authoritative context at inference time. For example, retrieved passages from approved product documentation or an organizational knowledge base can help the model formulate an answer that is specific to the user's request and grounded in supplied facts. This pattern is central to retrieval-augmented generation.

Provide additional examples to prompts. is an effective few-shot prompting technique. Examples demonstrate the intended task, output structure, tone, formatting, and level of detail. GPT-3.5 can use these demonstrations to more reliably produce responses that follow the desired pattern, especially when a task has specialized formatting or classification rules.

Modify system messages. improves quality because system messages establish high-priority behavioral instructions for the model. They can define the assistant's role, audience, response constraints, style, safety requirements, and instructions for using supplied context. Clear and focused system instructions make generated responses more consistent and aligned with the solution's intended behavior. Microsoft recommends combining clear instructions, grounding context, and examples as prompt-engineering practices for improving generative AI outputs. See [Prompt engineering techniques](#) and [System message design](#).

QUESTION NO: 44

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:

Combine multiple datasets. and **Remove records that have missing values.** are actions associated with ingesting and preparing data for a machine learning workload. Data ingestion makes data from relevant sources available for the project, which commonly includes bringing together related data stored in separate files, databases, or other repositories. Combining the sources creates a more complete dataset that can support training and evaluation.

Data preparation then makes the ingested data suitable for machine learning. A core preparation activity is data cleaning, including identifying and handling missing values. Removing incomplete records is one valid approach when those records are not appropriate for the intended analysis or when their absence will not introduce unacceptable bias. Other preparation approaches may impute or otherwise transform missing values, but removal remains a recognized cleaning operation. Preparing data can also involve transformations, formatting, and feature-related processing so the resulting data is consistent and usable by training workflows.

Microsoft's data science guidance describes preparation as the stage in which data is cleaned, transformed, and combined before modeling. Azure Machine Learning supports managed data assets and data preparation workflows for using data reliably in machine learning jobs. See [Prepare data for machine learning](#) and [Data concepts in Azure Machine Learning](#).

QUESTION NO: 45

Which Azure Machine Learning capability should you use to quickly build and deploy a predictive model without extensive coding?

- A. ML pipelines
- B. Copilot
- C. DALL-E
- D. automated machine learning (automated ML)

ANSWER: D

Explanation:

automated machine learning (automated ML) is designed to accelerate the development of predictive machine learning models while minimizing the amount of manual coding and experimentation required. You provide a labeled dataset, specify the prediction task—such as classification, regression, or time-series forecasting—and configure evaluation criteria. Automated ML then runs multiple training iterations, selecting and testing appropriate algorithms, applying supported data preparation and feature-engineering steps, and tuning hyperparameters to identify a high-performing model.

This capability is especially appropriate when a team needs to move quickly from data to a validated predictive model without hand-writing and comparing many individual training scripts. Azure Machine Learning records the experiment runs and their metrics, allowing users to review the generated models and select the best result. The chosen model can then be registered and deployed to an online endpoint or batch endpoint for inference, supporting the complete build-and-deploy scenario. Microsoft describes automated ML as a process that automates the time-consuming, iterative tasks of machine learning model development. See [What is automated machine learning?](#) and [Configure automated ML training](#).

QUESTION NO: 46

You have an Azure subscription that uses Azure OpenAI.

You need to create an original image of a rural scene to use on a website. What should you do?

- A. From Azure AI Foundry, deploy a GPT-3.5 Turbo model and provide instructions to create the image of the rural scene.
- B. From Microsoft Bing, search the term "rural scene" and download the results.
- C. From GitHub Copilot, provide instructions to create the image of the rural scene.
- D. From Azure AI Foundry, deploy a DALL-E model, and provide instructions to create the rural scene.

ANSWER: D

Explanation:

From Azure AI Foundry, deploy a DALL-E model, and provide instructions to create the rural scene. is correct because DALL-E is Azure OpenAI's image-generation model family, designed to produce new images from natural-language prompts. The task specifically requires an original visual asset for a website, so the appropriate workflow is to use a text-to-image model and describe the desired scene in a prompt. For example, the prompt can specify elements such as rolling fields, a farmhouse, a dirt road, lighting, artistic style, and the desired composition. The deployed image model processes those instructions and returns generated image output rather than merely describing a scene in text.

Azure AI Foundry provides an environment for deploying and working with Azure AI models, including Azure OpenAI image-generation capabilities. After deployment, an application or user can submit the scene description through the image-generation API or supported development experience and use the generated result according to the applicable service terms and responsible AI requirements. Microsoft identifies DALL-E as an Azure OpenAI image model and documents the process for generating images from prompts. See [Azure OpenAI image models](#) and [Generate images with Azure OpenAI](#).

QUESTION NO: 47 - (DRAG DROP)

Match the types of machine learning to the appropriate scenarios.

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

NOTE: Each correct selection is worth one point.

Machine Learning Types	Answer Area
Facial detection	Machine Learning Type Separate images of polar bears and brown bears.
Facial recognition	Machine Learning Type Determine the location of a bear in a photo.
Image classification	Machine Learning Type Determine which pixels in an image are part of a bear.
Object detection	
Optical character recognition (OCR)	
Semantic segmentation	

ANSWER:

Machine Learning Types

Facial detection

Facial recognition

Image classification

Object detection

Optical character recognition (OCR)

Semantic segmentation

Answer Area

Image classification

Separate images of polar bears and brown bears.

Object detection

Determine the location of a bear in a photo.

Semantic segmentation

Determine which pixels in an image are part of a bear.

Explanation:

Image classification is the appropriate machine learning task for separating photographs of polar bears from photographs of brown bears. It evaluates the image as a whole and assigns a class label based on its visual content. In this scenario, the useful result is a label identifying the bear category represented in each image, rather than a precise outline or coordinate for the animal.

Object detection is the correct task when the requirement is to determine a bear's location in a photograph. In addition to recognizing that a bear is present, object detection identifies its position in the image. Detection results are commonly represented with a labeled bounding box and a confidence score, which makes this approach useful when image location information is required.

Semantic segmentation is the correct task for identifying the pixels that are part of a bear. This task applies a label to individual pixels, enabling the model to distinguish the bear's actual visual region from surrounding pixels such as sky, snow, trees, or other background. Pixel-level output provides more detailed spatial information than an image-level category or a rectangular detection region. These distinctions reflect standard computer vision task types used with Azure AI solutions. See the [Azure AI Vision overview](#) and Microsoft's [object detection concepts](#) for related Azure AI Vision documentation.

QUESTION NO: 48

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 appropriate services for creating a bot that lets users query a knowledge base in natural language. Azure AI Language includes the question answering capability, which is designed to build a conversational question-and-answer layer over existing content. A question answering project can use knowledge sources to identify the intent and meaning of a user's phrased question, locate the most relevant answer, and return that answer to the calling application. This supplies the natural language processing and knowledge-retrieval portion of the solution.

Azure Bot Service supplies the bot platform and Azure integration needed for the conversational application itself. It enables developers to build, deploy, manage, and connect bots to communication channels such as a website chat interface or Microsoft Teams. The bot receives the user's message, sends the query to the question answering capability in Language Service, and presents the resulting response in the conversation. This combination supports the standard architecture for a knowledge-base bot: Azure Bot Service provides the user-facing conversation and channel connectivity, while Language

Service provides natural-language question answering over the knowledge base. For more information, see [Question answering in Azure AI Language](#) and the [Azure Bot Service overview](#).

QUESTION NO: 49

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:

Azure Machine Learning designer supports visual construction of machine learning workflows on a drag-and-drop canvas. A dataset is a data input that can be placed on the canvas and connected to subsequent workflow steps. It supplies the tabular, image, text, or other data used by a pipeline for preparation, training, scoring, or evaluation. In current Azure Machine Learning terminology, these inputs are commonly managed as data assets, but the dataset term is used in the designer and in many AI-900 exam materials.

A module is a reusable operation that performs a specific machine learning or data-processing task in the visual workflow. Examples include cleaning or transforming data, splitting data, training a model, scoring a model, and evaluating results. By dragging datasets and modules onto the canvas and connecting their input and output ports, users define the sequence and flow of an end-to-end machine learning pipeline without needing to author every processing step in code.

Microsoft describes designer as a drag-and-drop interface for building machine learning pipelines and explains that pipelines are composed of data and components arranged on a canvas. See [Azure Machine Learning designer overview](#) and [Create a pipeline in Azure Machine Learning designer](#).

QUESTION NO: 50

Your company manufactures widgets.

You have 1.000 digital photos of the widgets.

You need to identify the location of the widgets within the photos. What should you use?

- A. Computer Vision Spatial Analysis
- B. Custom Vision object detection
- C. Custom Vision classification
- D. Computer Vision Image Analysis

ANSWER: B

Explanation:

Custom Vision object detection is the appropriate service because the requirement is to determine where widgets appear within each photograph. An object detection model is trained using labeled examples in which each widget is marked with a bounding box. After training, the model evaluates a new image and returns each detected widget's tag, confidence score,

and bounding-box coordinates. Those coordinates identify the object's position and extent in the image, directly satisfying the location requirement.

This approach is particularly suitable for manufactured widgets because the organization can train the model on its own images, visual variants, backgrounds, orientations, and production-specific characteristics. The collection of 1,000 photographs can provide the source material for preparing and labeling a representative training set. Microsoft describes Custom Vision object detection as a capability for locating and identifying objects in images, and its detector workflow uses regions (bounding boxes) to tag the objects before model training. See the [Azure AI Custom Vision overview](#) and [Build an object detector with Custom Vision](#).

QUESTION NO: 51

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 the appropriate AI workload because the insurance claim reports consist of unstructured written language, and the goal is to identify important terms that represent the reports' main subjects. In Azure AI, Natural Language Processing (NLP) enables systems to analyze, interpret, and derive useful information from text. A directly relevant capability is key phrase extraction, available through Azure AI Language. It evaluates text and returns important phrases that describe its central points. These extracted phrases can support automated summaries, document indexing, classification workflows, and downstream review of large collections of claim reports. This scenario does not require a custom model merely to identify salient terminology; Azure AI Language provides prebuilt text-analysis features that can perform the task. Microsoft describes key phrase extraction as identifying the main concepts in unstructured text, making it well suited to reports and other document-based content. For more detail, see the [Azure AI Language key phrase extraction overview](#) and Microsoft Learn's module on [analyzing text with Azure AI Language](#).

QUESTION NO: 52

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, a dataset can be placed on the authoring canvas as the input data for a machine learning workflow. The dataset provides the tabular or other registered data that downstream steps use for preparation,

training, scoring, and evaluation. Connecting the dataset to subsequent steps makes the data flow explicit in the visual pipeline.

A module can also be dragged onto the canvas. Modules are reusable operations that perform a specific task in the workflow, such as transforming data, selecting features, splitting data, training a model, scoring predictions, or evaluating model performance. The designer's visual authoring experience is based on arranging these data inputs and processing or machine-learning modules, then connecting their input and output ports to form a complete pipeline. This supports building and running machine learning workflows without writing the pipeline definition entirely in code.

Microsoft's designer documentation describes the canvas-based approach to constructing pipelines from data assets and components/modules. See [Azure Machine Learning designer overview](#) and [Create a pipeline in Azure Machine Learning designer](#).

QUESTION NO: 53

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:

To consume a published inference pipeline in Azure Machine Learning designer, a client needs both the address of the scoring service and a way to authenticate its request. The REST endpoint supplies the service URI to which the client application sends an HTTP request containing the input data to be scored. This endpoint is the published pipeline's callable interface, enabling applications to obtain predictions programmatically.

The authentication key supplies the credential needed to authorize the request to the deployed web service. A consumer combines the endpoint URI and key in its request code: the URI determines the destination, while the key verifies that the caller is permitted to invoke the scoring service. In Azure Machine Learning designer deployment workflows, the endpoint and authentication key are available from the deployed service's Consume information and are used to construct sample request code. This reflects the core deployment pattern for real-time inference: send properly formatted input to the scoring URI and include the required authorization credential. For details, see [Deploy a model with Azure Machine Learning designer](#) and [Authenticate to Azure Machine Learning online endpoints](#).

QUESTION NO: 54

QUESTION NO: 81

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
- E. Question answering in Azure AI Language

ANSWER: E

Explanation:

The question answering feature in Azure AI Language is the current Azure service for creating an FAQ-based conversational experience. It lets you create a knowledge base, known as a project, from existing question-and-answer content such as FAQ documents, web pages, and other supported files. When a user submits a natural-language question, the service identifies the best matching answer from that curated knowledge base. An application or bot can then present that answer to the user, making it well suited to support, help desk, and informational bots built around frequently asked questions.

Microsoft describes custom question answering as a cloud-based natural-language processing capability that creates a conversational layer over data. Its workflow includes importing content, reviewing and improving question-and-answer pairs, deploying the project, and querying the deployed endpoint from a client application or bot. This directly meets the requirement to create a bot from an FAQ document. QnA Maker was the predecessor to this capability, but Microsoft retired the QnA Maker service on March 31, 2025. For a current Azure solution, use question answering in Azure AI Language. See [What is custom question answering?](#) and [QnA Maker documentation and retirement guidance](#).

QUESTION NO: 55

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

Which Azure AI service should you use?

A. Azure AI Language

Azure AI Language deals with text-based tasks such as language understanding, sentiment analysis, and key phrase extraction — not image processing.

B. Azure AI Computer Vision

Azure AI Computer Vision provides prebuilt image analysis models (e.g., object detection, tag generation, scene description), but it cannot learn custom categories unique to your dataset. It's great for general image recognition but not for specialized labeling tasks.

C. Azure AI Document Intelligence

Azure AI Document Intelligence (Form Recognizer) is used to extract information from structured or semi-structured documents such as forms, invoices, or receipts — not photographs.

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

D. Azure AI Custom Vision

ANSWER: D

Explanation:

Azure AI Custom Vision is the appropriate service because the requirement is to create a model that applies labels to a specific collection of photographs. Custom Vision is designed for custom image-classification solutions: you provide representative images, assign each image one or more meaningful tags, train the model, and then submit new images for classification. For example, personal photographs could be tagged with categories such as family, pet, vacation, or an organization scheme that is unique to the photo collection. The trained classifier learns the visual characteristics associated with those tags and returns predicted labels for images it has not previously seen. This workflow directly matches a requirement to create a tailored labeling model rather than simply consume a general-purpose image-analysis result. Microsoft describes Custom Vision as a service for building custom image classifiers and object detectors from labeled image data. Image classification assigns tags to an entire image, making it a suitable approach when the goal is to categorize photographs. See the [Azure AI Custom Vision overview](#) and Microsoft Learn guidance on [classifying images with Custom Vision](#).

QUESTION NO: 56

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

NOTE: Each correct selection is worth one point.

- A. a telephone answering service that has a pre-recorder message
- B. a chatbot that provides users with the ability to find answers on a website by themselves
- C. telephone voice menus to reduce the load on human resources
- D. a service that creates frequently asked questions (FAQ) documents by crawling public websites

ANSWER: B C

Explanation:

A chatbot that provides users with the ability to find answers on a website by themselves is a conversational AI workload because it enables a person to interact with an application using natural-language messages. The chatbot can interpret a user's question, identify the intent or relevant information, and return an appropriate answer in a conversational experience. This supports self-service interactions without requiring the user to navigate a site manually or contact a person directly.

Telephone voice menus to reduce the load on human resources are also a conversational AI workload when callers interact with an automated voice agent to state needs, select services, obtain answers, or route requests. Conversational AI can use speech recognition to convert spoken input into text, natural language capabilities to understand the request, and speech synthesis to provide a spoken response. These capabilities allow organizations to offer automated, natural-language interactions over voice channels as well as text channels. Microsoft's conversational language understanding service is designed to build applications that understand user utterances and determine their intended meaning. Azure Bot Service supports building and connecting conversational bots across communication channels. See [Conversational language understanding documentation](#) and [Azure Bot Service overview](#).

QUESTION NO: 57

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, also called text to speech, converts written text into spoken audio generated by a synthetic voice. "an automated voice that reads back a credit card number entered into a telephone by using a numeric keypad" is a speech-synthesis scenario because the digits collected by the telephone system can be represented as text and spoken back to the caller for confirmation. This is a typical use of text to speech in an interactive voice response workflow, where the system needs to provide an audible response without a human agent recording every possible message.

"an AI character in a computer game that speaks audibly to a player" is also a direct use of speech synthesis. A game can supply dialogue text dynamically or from prewritten scripts, and a text-to-speech service renders that dialogue as audio in a selected synthetic voice. This supports responsive character interactions and enables spoken content without relying solely on prerecorded sound files. Azure AI Speech provides text-to-speech capabilities that turn text into natural-sounding synthesized speech and supports application scenarios involving voice-enabled experiences. For details, see [Azure AI Speech text to speech](#) and the [Azure AI Speech overview](#).

QUESTION NO: 58

Which two actions support the Microsoft fairness principle for responsible AI? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Evaluate outcomes for different demographic groups.
- B. Use representative data during training and evaluation.
- C. Exclude affected groups from evaluation.
- D. Use protected characteristics without review.
- E. Prioritize the largest group regardless of impact on other groups.

ANSWER: A B

Explanation:

Evaluate outcomes for different demographic groups helps identify whether the solution produces unjustified disparities for groups of people. **Use representative data during training and evaluation** helps ensure that relevant populations are appropriately considered when the solution is developed and assessed.

Excluding affected groups from testing and using protected characteristics without review can increase the risk of unfair outcomes. Microsoft defines fairness as ensuring that systems do not produce unjustified bias or discrimination. See [Responsible AI guidance in Azure](#).

QUESTION NO: 59

Which two statements about training and test datasets are correct? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. The training dataset is used to fit the machine learning solution.
- B. The test dataset is used to evaluate performance on previously unseen data.
- C. The test dataset must contain only input features and no labels.
- D. The training dataset is used only after deployment.
- E. The same rows must be used for training and testing.

ANSWER: A B

Explanation:

The training dataset is used to fit the machine learning solution because it contains the examples from which patterns are learned. **The test dataset is used to evaluate performance on previously unseen data** because its rows are withheld from the fitting process.

Using the same records for both activities can produce overly optimistic results because performance is measured on data already seen during training. See [Train machine learning models in Azure Machine Learning](#).

QUESTION NO: 60

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
- D. ZIP

ANSWER: A

Explanation:

PDF is the correct file format because QnA Maker knowledge bases can be populated from supported document files, including PDF documents. This allows an organization to use existing FAQ-style materials, manuals, help documents, or other semi-structured content as a source rather than manually entering every question-and-answer pair. QnA Maker extracts text from the uploaded document and uses document structure and content to help generate knowledge-base entries. This is particularly useful when content already exists in a maintained document and must be made available through a conversational question-answering experience.

Although QnA Maker has been retired for new implementations, its replacement, Azure AI Language custom question answering, follows the same fundamental approach of importing supported content sources to create a question-answering knowledge base. Microsoft's documentation for the current service identifies files and other sources that can provide content for a project, and PDF is a supported document type. See [Data sources for custom question answering](#) and [Azure AI Language custom question answering overview](#).

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:

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 tasks that required an Enterprise workspace under the legacy Azure Machine Learning workspace editions referenced by this exam question. The Enterprise edition provided Azure Machine Learning studio's visual, low-code experiences. In particular, the automated machine learning interface enabled users to configure an experiment, select the target and training settings, submit runs, and compare generated model results without implementing the complete workflow in code. The designer interface provided a drag-and-drop canvas for composing, configuring, and executing machine learning pipelines visually. These studio GUI capabilities were the distinguishing Enterprise-only features being tested. Azure Machine Learning's service and workspace model have evolved, so current documentation does not frame access around the former Basic and Enterprise editions; nevertheless, the documented capabilities explain the intended legacy answer. See [Automated ML in Azure Machine Learning](#) and [Azure Machine Learning designer](#).

QUESTION NO: 62

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 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 supported ways to build a question-answer knowledge base for a chatbot. A URL can be provided as a content source, allowing the service to ingest suitable FAQ-style or support-page content and extract question-and-answer pairs for review and refinement. This is useful when an organization already maintains customer-facing help content online. Direct authoring is also supported, enabling authors to add an answer and one or more alternate phrasings of a question. This is important for business-specific questions that are not represented in existing documents or webpages. Predefined chat content can be added to provide responses to common conversational inputs, such as greetings and informal exchanges, so that the chatbot can handle interactions outside the core business FAQ. Although QnA Maker has been retired, these authoring patterns are represented in Azure AI Language custom question answering, its successor. Microsoft documents supported content sources, including URLs and manually authored pairs, in [Data sources for custom question answering](#), and documents the use of predefined conversational datasets in [Add chat to a project](#).

QUESTION NO: 63

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 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:

Azure AI Language custom question answering lets authors build a knowledge base for a chatbot from several supported content sources. **Manually enter the questions and answers.** is a supported approach because an author can create editorial question-and-answer pairs directly in a custom question answering project. This is useful for approved business FAQs, policies, and responses that must use precise wording. **Generate the questions and answers from an existing webpage.** is also supported: a project can use a public URL as a source, allowing question answering to extract content from an FAQ-style webpage and make it available as knowledge-base content. **Import chat content from a predefined data source.** is supported through the service's predefined chat datasets. These add common conversational exchanges, such as greetings and casual interactions, to complement organization-specific answers. Together, direct authoring, URL-based content ingestion, and predefined chat provide practical ways to create the answer content needed for a simple business chatbot. Microsoft documents supported custom question answering sources, including URLs and manually authored content, in its [custom question answering project guidance](#). For predefined conversational datasets, see [Add chat to a custom question answering project](#).