

## Microsoft.AI-901.v2026-07-30.q67

|                                                                                                                                                                                   |                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| Exam Code:                                                                                                                                                                        | AI-901                          |
| Exam Name:                                                                                                                                                                        | Microsoft Azure AI Fundamentals |
| Certification Provider:                                                                                                                                                           | Microsoft                       |
| Free Question Number:                                                                                                                                                             | 67                              |
| Version:                                                                                                                                                                          | v2026-07-30                     |
| # of views:                                                                                                                                                                       | 120                             |
| # of Questions views:                                                                                                                                                             | 670                             |
| <a href="https://www.freeqas.com/qa/Microsoft/AI-901/Microsoft.AI-901.v2026-07-30.q67.html">https://www.freeqas.com/qa/Microsoft/AI-901/Microsoft.AI-901.v2026-07-30.q67.html</a> |                                 |

### NEW QUESTION: 1

You have a Microsoft Foundry project that has a generative AI model deployment.

You need to ensure that responses generated by the model minimize costs and remain within a defined length.

Which parameter should you configure?

- A. Temperature
- B. Max Completion Tokens
- C. Top P
- D. Model version settings

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

To minimize cost and keep generated responses within a defined length, configure Max Completion Tokens.

Microsoft's Azure OpenAI / Foundry API reference defines `max_completion_tokens` as an upper bound for the number of tokens that can be generated for a completion. Because generated tokens contribute to usage and response length, limiting completion tokens helps control both output length and cost.

Temperature and Top P control randomness or sampling behavior, not maximum response length. Model version settings do not directly define the generated response length.

### NEW QUESTION: 2

Which Azure OpenAI model can be used to develop code?

- A. microsoft-swinv2-base-patch4-window12-192-22k
- B. Whisper
- C. DALL-E
- D. GPT-4-32k

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

The GPT-4 32k model can be used to develop programming code. It has a large context window, allowing it to process more information and generate more complex code. The model has shown strong performance in coding tasks, even surpassing its predecessor, o3-mini, in some domains.

Reference:

<https://vivekupadhyay1.medium.com/openai-chat-gpt-4-32k-detailed-document-ce3486e4786e>

### NEW QUESTION: 3

#### Hotspot Question

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

NOTE: Each correct selection is worth one point.

Answer Area

| Statements                                                                                                                   | Yes                   | No                    |
|------------------------------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| Generating a response to a user prompt occurs during the inference stage.                                                    | <input type="radio"/> | <input type="radio"/> |
| A generative AI model generates responses by copying stored documents directly from the model's training data.               | <input type="radio"/> | <input type="radio"/> |
| A generative AI model produces output by predicting the next token based on patterns learned from the model's training data. | <input type="radio"/> | <input type="radio"/> |

Microsoft

#### Answer:

Answer Area

| Statements                                                                                                                   | Yes                              | No                               |
|------------------------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| Generating a response to a user prompt occurs during the inference stage.                                                    | <input checked="" type="radio"/> | <input type="radio"/>            |
| A generative AI model generates responses by copying stored documents directly from the model's training data.               | <input type="radio"/>            | <input checked="" type="radio"/> |
| A generative AI model produces output by predicting the next token based on patterns learned from the model's training data. | <input checked="" type="radio"/> | <input type="radio"/>            |

Microsoft

### NEW QUESTION: 4

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

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

Answer: ([SHOW ANSWER](#))

Azure AI Vision is a service within Azure Cognitive Services that enables developers to analyze visual data from images and videos. It includes OCR (Optical Character Recognition) capabilities, allowing you to extract text (both printed and handwritten) from images, PDFs, and scanned documents.

Reference:

<https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/overview-ocr>

### NEW QUESTION: 5

Based on the image provided, here is the transcribed text:

You need to build an AI solution that produces new product images based on written descriptions provided by users.

Which AI workload should you use?

- A. image generation
- B. image analysis
- C. object detection

D. optical character recognition (OCR)

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

The requirement is to produce new product images based on written descriptions. This is an image generation workload, because the AI system is creating entirely new images from natural language prompts.

#### NEW QUESTION: 6

You need to create an AI agent in Microsoft Foundry that follows a specific role and behavior when responding to users.

What should you configure?

A. tokens per minute (TPM)

B. system instructions

C. temperature

D. max completion tokens

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

To create an AI agent that follows a specific role and behavior, you configure system instructions. Microsoft Foundry Agent Service documentation states that agent instructions define goals, constraints, and behavior.

Option A. tokens per minute (TPM) controls throughput quota, not behavior.

Option C. temperature controls response randomness/creativity, not the agent's role.

Option D. max completion tokens controls response length, not the agent's role or behavioral rules.

Therefore, the correct answer is B. system instructions.

#### NEW QUESTION: 7

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

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

What should you do?

A. Export the model

B. Create a new model

C. Add utterances to the None intent

D. Create a prebuilt task entity

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

The None intent is a required intent and can't be deleted or renamed. The intent is meant to categorize utterances that don't belong to any of your other custom intents.

An utterance can be predicted as the None intent if the top scoring intent's score is lower than the None score threshold.

Reference:

<https://learn.microsoft.com/en-us/azure/ai-services/language-service/conversational-language-understanding/concepts/none-intent>

#### NEW QUESTION: 8

Hotspot Question

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

NOTE: Each correct selection is worth one point.

Answer Area

| Statements                                                                                                 | Yes                   | No                    |
|------------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| Voice Live returns only transcribed text.                                                                  | <input type="radio"/> | <input type="radio"/> |
| Voice Live requires you to separately implement speech to text and text to speech services.                | <input type="radio"/> | <input type="radio"/> |
| Voice Live combines speech to text, reasoning, and text to speech into a single conversational experience. | <input type="radio"/> | <input type="radio"/> |

Answer:

Answer Area

| Statements                                                                                                 | Yes                              | No                               |
|------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| Voice Live returns only transcribed text.                                                                  | <input type="radio"/>            | <input checked="" type="radio"/> |
| Voice Live requires you to separately implement speech to text and text to speech services.                | <input type="radio"/>            | <input checked="" type="radio"/> |
| Voice Live combines speech to text, reasoning, and text to speech into a single conversational experience. | <input checked="" type="radio"/> | <input type="radio"/>            |

#### NEW QUESTION: 9

You are developing a web app that processes invoices to calculate expenses.

You need to extract structured fields, including nested values, from the invoices by using a defined schema.

What should you use?

- A. a transcription workflow in Azure Speech in Foundry Tools
- B. an optical character recognition (OCR)-only document processing pipeline
- C. an analyzer in Azure Content Understanding in Foundry Tools
- D. an Azure AI Search service

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

The requirement is to extract structured fields, including nested values, from invoices by using a defined schema. In Azure Content Understanding, an analyzer is the processing unit that defines how content is analyzed, what information is extracted, and how the output is structured, including JSON fields.

Microsoft's Content Understanding document solutions documentation states that Content Understanding uses customizable analyzers to extract essential information, fields, and relationships from documents and forms. Microsoft's quickstart also shows invoice processing with the prebuilt-invoice analyzer to extract structured data from an invoice document.

Why the other options are incorrect:

- A . transcription workflow in Azure Speech is for converting audio to text, not invoice field extraction.
- B . OCR-only document processing can extract text but does not meet the requirement for structured fields and nested values by schema.
- D . Azure AI Search is for indexing and querying content, not defining invoice extraction schemas.

Therefore, the correct answer is C. an analyzer in Azure Content Understanding in Foundry Tools.

#### NEW QUESTION: 10

You are developing an application that extracts fields from PDFs by using Azure Content Understanding in Foundry Tools.

You need to use the Python SDK to submit a PDF for analysis and retrieve the extraction results.

What should you do?

- A. Call analyze() to return the extracted fields synchronously in the same request.

- B.** Submit the PDF to an analyzer and read the results from the request headers.
- C.** Call `begin_analyze()`, and then call `poller.result()` to retrieve the results.
- D.** Use optical character recognition (OCR) to extract text from the PDF and map the fields in the code.

**Answer: C (LEAVE A REPLY)**

To build an application that extracts fields from PDFs using Azure Content Understanding in Foundry Tools, you need to initialize the client library, use `begin_analyze_binary` to submit the file stream, and poll the long-running operation for the result.

Detail from implementation:

Wait for the operation to complete and poll the results

```
result = poller.result()
```

Reference:

<https://lobehub.com/pt-BR/skills/frank-luongt-faos-skills-marketplace-azure-ai-contentunderstanding-py>

### NEW QUESTION: 11

You plan to develop an AI application that will read the license plates of motor vehicles by using Azure AI Foundry. What should you use to develop the application?

- A.** Microsoft Visual Studio Code
- B.** GitHub Actions
- C.** Copilot for Microsoft 365
- D.** Azure AI Studio

**Answer: D (LEAVE A REPLY)**

Azure AI Studio (now Azure AI Foundry) is used to develop an AI application for reading license plates using its tools and models. The platform allows developers to build, customize, and deploy AI agents and applications, including those that process images to extract information like license plate numbers.

Key components of the process

Platform: Azure AI Foundry (formerly Azure AI Studio) is the central platform for building and managing AI applications.

Tools: You can use its integrated development environment (IDE), SDKs, and APIs to develop and customize models.

Models: The platform provides access to a wide range of AI models, including computer vision models capable of optical character recognition (OCR), which can be used to read text from images.

Workflow:

Create an Azure AI Foundry project in Azure.

Use the model catalog to find and deploy a suitable computer vision or OCR model.

Integrate the model into your application using the SDK or API.

For a license plate application, you would feed images to the model, and the model would output the recognized text from the license plate.

Reference:

<https://azure.microsoft.com/en-us/products/ai-foundry>

### NEW QUESTION: 12

You plan to create an AI application that will read the license plates of motor vehicles by using Azure AI Foundry. The solution will be billed via tokens for inputs and outputs to the API. Which deployment option should you use?

- A.** Azure Kubernetes Service (AKS) cluster
- B.** Azure virtual machines
- C.** serverless API
- D.** managed compute

Answer: (SHOW ANSWER)

<https://learn.microsoft.com/en-us/azure/ai-foundry/concepts/deployments-overview>

NEW QUESTION: 13

Hotspot Question

Select the answer that correctly completes the sentence.

Answer Area

Ensuring that human reviewers oversee AI-generated decisions and remain responsible for the final output is an example of the Microsoft responsible AI principle of

|                      |   |
|----------------------|---|
|                      | ▼ |
| accountability       |   |
| fairness             |   |
| privacy and security |   |
| transparency         |   |

Answer:

Answer Area

Ensuring that human reviewers oversee AI-generated decisions and remain responsible for the final output is an example of the Microsoft responsible AI principle of

 Microsoft

|                      |   |
|----------------------|---|
|                      | ▼ |
| accountability       |   |
| fairness             |   |
| privacy and security |   |
| transparency         |   |

NEW QUESTION: 14

What are two purposes of instructions when prompting a generative AI model? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. defines constraints on the model's responses
- B. defines the agent's role and behavior
- C. defines the Azure region where inference occurs
- D. selects which model to use



E. defines the tokens per minute (TPM) allocation for the model

**Answer: A,B (LEAVE A REPLY)**

Microsoft Foundry Agent Service documentation states that instructions define goals, constraints, and behavior for an agent. Therefore, instructions are used to guide how the generative AI model or agent should respond and behave.

Option A is correct because instructions can define constraints the model must follow.

Option B is correct because instructions can define the agent's role and behavior.

Options C, D, and E are incorrect because Azure region, model selection, and TPM allocation are configuration or deployment/resource settings, not purposes of prompt instructions.

#### NEW QUESTION: 15

You need to compare the costs of large language models (LLMs) for a generative AI solution.

What should you use in the Microsoft Foundry portal?

A. Evaluator catalog

B. Model leaderboard

C. Compliance

D. Tools

**Answer: B (LEAVE A REPLY)**

To compare the costs of large language models in Microsoft Foundry portal, use the Model leaderboard.

Microsoft documentation states that the model leaderboard helps compare models across quality, safety, estimated cost, and throughput. It also supports trade-off charts and side-by-side model comparison for features, performance, and estimated cost.

Why the other options are incorrect:

A . Evaluator catalog is for selecting evaluators to measure model or application outputs, not comparing LLM costs.

C . Compliance relates to governance and compliance, not model cost comparison.

D . Tools provides Foundry tools, not benchmarked cost comparison across models.

#### NEW QUESTION: 16

You need to compare the costs of large language models (LLMs) for a generative AI solution.

What should you use in the Microsoft Foundry portal?

A. Compliance

B. Evaluator catalog

C. Tools

D. Model leaderboard

**Answer: D (LEAVE A REPLY)**

To compare the costs, capabilities, and performance of large language models (LLMs) in the Foundry portal, you should use the Model leaderboard.

Why Use the Model Leaderboard

The leaderboard allows you to:

Compare models side-by-side: Evaluate multiple LLMs simultaneously for trade-offs regarding estimated costs, throughput, safety, and quality.

Make informed decisions: Select the most cost-effective model for specific workloads before committing to deployments, reducing the risk of expensive errors or overspending.

Reference:

<https://learn.microsoft.com/en-us/azure/foundry/concepts/manage-costs>

**Valid AI-901 Dumps** shared by PrepPdf.com for Helping Passing AI-901 Exam! PrepPdf.com now offer the **newest AI-901 exam dumps**, the PrepPdf.com AI-901 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-901 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/AI-901-prepaway-exam-dumps.html> (87 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

**NEW QUESTION: 17**

You need to convert written customer notifications into natural-sounding spoken audio that can be played over a phone system.  
Which Azure Speech in Foundry Tools capability should you use?

- A. speaker recognition
- B. speech synthesis
- C. speech recognition
- D. speech translation

**Answer: B (LEAVE A REPLY)**

The requirement is to convert written customer notifications into natural-sounding spoken audio. This is speech synthesis, also known as text to speech. Microsoft's Azure Speech documentation describes text to speech as a capability that converts text into natural-sounding synthesized speech. Therefore, for playing written notifications over a phone system, the correct Azure Speech capability is speech synthesis.

Why the other options are incorrect:

- A . speaker recognition identifies or verifies speakers by voice.
- C . speech recognition converts spoken audio into text.
- D . speech translation translates spoken audio between languages.

**NEW QUESTION: 18**

Hotspot Question

You are developing an application that converts text into spoken audio and saves the synthesized audio to a file by using Azure Speech in Foundry Tools.  
How should you complete the Python code? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

```
import azure.cognitiveservices.speech as speechsdk
...
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
audio_config = speechsdk.audio.
synthesizer = speechsdk.SpeechSynthe.
    speech_config=speech_config,
    audio_config=audio_config
)
```

Microsoft

AudioOutputConfig(filename="output.wav")

AudioOutputConfig(stream)

AudioStreamFormat(wave\_stream\_format=AudioStreamWaveFormat.PCM)

**Answer:**



```
import azure.cognitiveservices.speech as speechsdk
...
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
audio_config = speechsdk.audio:
synthesizer = speechsdk.SpeechSynthesizer(
    speech_config=speech_config,
    audio_config=audio_config
)
...
```

AudioOutputConfig(filename="output.wav")

AudioOutputConfig(stream)

AudioStreamFormat(wave\_stream\_format=AudioStreamWaveFormat.PCM)

NEW QUESTION: 19

You are using the Azure Speech SDK to develop a Python application that supports real-time spoken conversations. Which Azure speech class should you use to configure the connection to the Azure Speech service?

- A. AuditOutputConfig
- B. SpeechSynthesizer
- C. AudioConfig

Answer: (SHOW ANSWER)

NEW QUESTION: 20

Hotspot Question

Select the answer that correctly completes the sentence.

Answer Area

Evaluating model outcomes across demographic groups to reduce bias is an example of the Microsoft responsible AI principle of

accountability

fairness

privacy and security

transparency

Answer:

Answer Area

Evaluating model outcomes across demographic groups to reduce bias is an example of the Microsoft responsible AI principle of

accountability

fairness

privacy and security

transparency

NEW QUESTION: 21

Hotspot Question

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

NOTE: Each correct selection is worth one point.

**Answer Area**

| Statements                                                                                                                   | Yes                   | No                    |
|------------------------------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| Generating a response to a user prompt occurs during the inference stage.                                                    | <input type="radio"/> | <input type="radio"/> |
| A generative AI model generates responses by copying stored documents directly from the model's training data.               | <input type="radio"/> | <input type="radio"/> |
| A generative AI model produces output by predicting the next token based on patterns learned from the model's training data. | <input type="radio"/> | <input type="radio"/> |

**Answer:**

**Answer Area**

| Statements                                                                                                                   | Yes                              | No                               |
|------------------------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| Generating a response to a user prompt occurs during the inference stage.                                                    | <input checked="" type="radio"/> | <input type="radio"/>            |
| A generative AI model generates responses by copying stored documents directly from the model's training data.               | <input type="radio"/>            | <input checked="" type="radio"/> |
| A generative AI model produces output by predicting the next token based on patterns learned from the model's training data. | <input checked="" type="radio"/> | <input type="radio"/>            |

#### NEW QUESTION: 22

What can be used to analyze scanned invoices and extract data, such as billing addresses and the total amount due?

- A. Azure AI Custom Vision
- B. Azure AI Search
- C. Azure OpenAI
- D. Azure AI Document Intelligence

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

The Document Intelligence invoice model uses powerful Optical Character Recognition (OCR) capabilities to analyze and extract key fields and line items from sales invoices, utility bills, and purchase orders.

Reference:

<https://learn.microsoft.com/en-us/azure/ai-services/document-intelligence/prebuilt/invoice>

#### NEW QUESTION: 23

You are developing an application that extracts fields from PDFs by using Azure Content Understanding in Foundry Tools.

You need to use the Python SDK to submit a PDF for analysis and retrieve the extraction results.

What should you do?

- A. Call `begin_analyze()`, and then call `poller.result()` to retrieve the results.
- B. Submit the PDF to an analyzer and read the results from the request headers.
- C. Use optical character recognition (OCR) to extract text from the PDF and map the fields in the code.
- D. Call `analyze()` to return the extracted fields synchronously in the same request.

**Answer: A (LEAVE A REPLY)**

Azure Content Understanding analysis operations are long-running operations in the Python SDK. Microsoft's Python SDK documentation states that analysis operations return a poller, and the SDK provides LROPoller types that handle polling automatically when you call .result().

Therefore, the correct workflow is to submit the PDF by calling begin\_analyze(), receive a poller, and then call:

```
result = poller.result()
```

Option B is incorrect because extraction results are not read from request headers.

Option C is incorrect because the requirement is to use Azure Content Understanding extraction, not build a manual OCR-only mapping pipeline.

Option D is incorrect because the SDK analysis pattern is asynchronous/long-running, not a simple synchronous analyze() call that returns all extracted fields in the same request.

#### NEW QUESTION: 24

You are using the Azure Speech SDK to develop a Python application that supports real-time spoken conversations. Which Azure Speech class should you use to configure the connection to the Azure Speech service?

A. AudioOutputConfig

B. SpeechSynthesizer

C. AudioConfig

D. SpeechConfig

**Answer: (SHOW ANSWER)**

To configure the connection to the Azure Speech service, you should use the `azure.cognitiveservices.speech.SpeechConfig` class.

Implementation Details

Initialization: This class holds authentication and endpoint details, such as your Azure resource key and region.

Usage: You pass this configuration object into your real-time processing classes, like `SpeechRecognizer` (for speech-to-text) or `SpeechSynthesizer` (for text-to-speech).

Reference:

<https://docs.azure.cn/en-us/ai-services/speech-service/how-to-recognize-speech>

#### NEW QUESTION: 25


Select the answer that correctly completes the sentence.

Answer Area



**Answer:**

Answer Area



### NEW QUESTION: 26

You are developing an AI-powered customer support application. Which task is an example of the Microsoft responsible AI principle of inclusiveness?

- A. Encrypt stored customer data and restrict access by using role-based controls.
- B. Evaluate model outputs across demographic groups to reduce bias.
- C. Design the interface to support multiple languages and screen readers.
- D. Provide explanations about how predictions are generated.

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

The inclusiveness principle states that AI systems should empower everyone and engage people.

Designing for multiple languages and screen readers directly reflects this goal.

Screen reader support ensures accessibility for people with visual impairments.

Multi-language interfaces break down barriers for non-native speakers.

Inclusive design proactively addresses physical, cognitive, and linguistic differences.

Reference:

<https://www.netguru.com/blog/fintech-app-accessibility>

### NEW QUESTION: 27

You are explaining generative AI workflows to a colleague. What is the first step of a generative AI workflow?

- A. Deploy a generative AI model.
- B. Fine-tune a model by using feedback.
- C. Train a generative AI model by using data.
- D. Generate outputs based on user prompts.

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

Training a generative AI model using data is a foundational step in the workflow, but it is preceded by data collection and preparation. After the model is trained, subsequent steps involve validation, generation, and refinement.

Steps in a generative AI workflow

1. Data Gathering: Collect a massive, relevant dataset (images, text, audio, etc.) that will be used for training.
2. Data Preparation: Clean and preprocess the raw data. This includes removing errors, standardizing formats, and structuring the data so the model can comprehend it.
- \*-> 3. Model Training: Feed the prepared data into the model. During this stage, the model learns patterns and relationships within the data to understand how to create new content.
4. Model Validation: Test the model on a separate set of data to evaluate its performance and accuracy.
5. Content Generation: Use the trained and validated model to create new content based on prompts.
6. Refinement: Fine-tune the model using human feedback or additional data to improve its output and address specific tasks.

Reference:

<https://www.tricension.com/domainai/proper-training-for-generative-ai/>

### NEW QUESTION: 28

Select the answer that correctly completes the sentence.

Answer Area

The Microsoft responsible AI principle of transparency requires that AI systems

- be explainable to users
- protect sensitive user data
- reduce bias in decisions
- require human oversight



Answer:

Answer Area

The Microsoft responsible AI principle of transparency requires that AI systems

- be explainable to users
- protect sensitive user data
- reduce bias in decisions
- require human oversight

### NEW QUESTION: 29

Hotspot Question

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

NOTE: Each correct selection is worth one point.

Answer Area



Statements

Fairness can be achieved by focusing solely on improving the overall accuracy of an AI model.

Yes

☐

No

☐

Evaluating AI system outputs to identify and reduce bias across demographic groups supports fairness.

☐☐

Ensuring fairness for an AI system means that all users always receive the same output from the system.

☐☐

Answer:

Answer Area



Statements

Fairness can be achieved by focusing solely on improving the overall accuracy of an AI model.

Yes

☐

No

☒

Evaluating AI system outputs to identify and reduce bias across demographic groups supports fairness.

☒☐

Ensuring fairness for an AI system means that all users always receive the same output from the system.

☐☒

### NEW QUESTION: 30

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



NOTE: Each correct selection is worth one point.

Answer Area

| Statements                                                                       | Yes                   | No                    |
|----------------------------------------------------------------------------------|-----------------------|-----------------------|
| An AI generative model is retrained before performing each user request.         | <input type="radio"/> | <input type="radio"/> |
| An AI agent responds by copying and pasting answers stored in a database.        | <input type="radio"/> | <input type="radio"/> |
| An AI agent uses a generative AI model to establish actions based on user input. | <input type="radio"/> | <input type="radio"/> |

Answer:

Answer Area

| Statements                                                                       | Yes                              | No                               |
|----------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| An AI generative model is retrained before performing each user request.         | <input type="radio"/>            | <input checked="" type="radio"/> |
| An AI agent responds by copying and pasting answers stored in a database.        | <input type="radio"/>            | <input checked="" type="radio"/> |
| An AI agent uses a generative AI model to establish actions based on user input. | <input checked="" type="radio"/> | <input type="radio"/>            |

### NEW QUESTION: 31

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

NOTE: Each correct selection is worth one point.

Answer Area

| Statements                                                                                     | Yes                   | No                    |
|------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| The Temperature parameter can be set before deploying a model.                                 | <input type="radio"/> | <input type="radio"/> |
| During inference, the model name is used to route requests to a specific deployment.           | <input type="radio"/> | <input type="radio"/> |
| After a model is deployed, both code and testing tools can be used to interact with the model. | <input type="radio"/> | <input type="radio"/> |

Answer:

Answer Area

| Statements                                                                                     | Yes                              | No                               |
|------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| The Temperature parameter can be set before deploying a model.                                 | <input type="radio"/>            | <input checked="" type="radio"/> |
| During inference, the model name is used to route requests to a specific deployment.           | <input type="radio"/>            | <input checked="" type="radio"/> |
| After a model is deployed, both code and testing tools can be used to interact with the model. | <input checked="" type="radio"/> | <input type="radio"/>            |



**Valid AI-901 Dumps** shared by PrepPdf.com for Helping Passing AI-901 Exam! PrepPdf.com now offer the **newest AI-901 exam dumps**, the PrepPdf.com AI-901 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-901 dumps with Test Engine here:  
<https://www.preppdf.com/Microsoft/AI-901-prepaway-exam-dumps.html> (87 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

**NEW QUESTION: 32**

You need to convert written customer notifications into natural-sounding spoken audio that can be played over a phone system. Which Azure Speech in Foundry Tools capability should you use?

- A. speaker recognition
- B. speech synthesis
- C. speech recognition
- D. speech translation

**Answer: B (LEAVE A REPLY)**

The requirement is to convert written customer notifications into natural-sounding spoken audio.

This is speech synthesis, also known as text to speech.

Microsoft's Azure Speech documentation describes text to speech as a capability that converts text into natural-sounding synthesized speech. Therefore, for playing written notifications over a phone system, the correct Azure Speech capability is speech synthesis.

**NEW QUESTION: 33**

Hotspot Question

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements

In the Foundry playground, you can upload a local image and include text in the same message when prompting a multimodal model.

When using the OpenAI Responses API and a vision-enabled model, images can be provided only as base64-encoded image data.

Prompts that include images require deploying a text-only model because multimodal capabilities are handled by the application layer.

Yes

No

☐

☐

☐

☐

☐

☐

**Answer:**

Answer Area

Statements

In the Foundry playground, you can upload a local image and include text in the same message when prompting a multimodal model.

When using the OpenAI Responses API and a vision-enabled model, images can be provided only as base64-encoded image data.

Prompts that include images require deploying a text-only model because multimodal capabilities are handled by the application layer.

Yes

No

☒

☐

☐

☒

☐

☒

**NEW QUESTION: 34**

You have a Microsoft Foundry project that contains a vision-enabled model deployment.

You are developing an application that sends images to the model.  
You need to ensure that the model can analyze the images.  
In which two formats can you provide the images? Each correct answer presents part of the solution.  
NOTE: Each correct selection is worth one point.

- A. a JSON document that describes the image content
- B. a UTF-8 encoded text description of the image
- C. a publicly accessible URL of the image
- D. a base64 encoded image data string

Answer: C,D (LEAVE A REPLY)

For vision-enabled Azure OpenAI / Microsoft Foundry model requests, image input can be provided by using an image URL or base64-encoded image data. Microsoft's Azure OpenAI REST API reference states that the image content part URL field can contain either a URL of the image or the base64 encoded image data. It also states that the Responses API input\_image.image\_url value can be a fully qualified URL or a base64 encoded image in a data URL.

NEW QUESTION: 35

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

| Answer Area                                                                                                                                          |                       |                       |
|------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| Statements                                                                                                                                           | Yes                   | No                    |
| In Microsoft Foundry Agent Service, setting tool_choice to auto for an agent enables the agent to decide whether to call a tool.                     | <input type="radio"/> | <input type="radio"/> |
| In Microsoft Foundry Agent Service, setting tool_choice to none for an agent means that the model decides whether to call a tool.                    | <input type="radio"/> | <input type="radio"/> |
| In Microsoft Foundry Agent Service, setting tool_choice to required for an agent ensures that the agent must call one or more tools during each run. | <input type="radio"/> | <input type="radio"/> |

Answer:

| Answer Area                                                                                                                                          |                                  |                                  |
|------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| Statements                                                                                                                                           | Yes                              | No                               |
| In Microsoft Foundry Agent Service, setting tool_choice to auto for an agent enables the agent to decide whether to call a tool.                     | <input checked="" type="radio"/> | <input type="radio"/>            |
| In Microsoft Foundry Agent Service, setting tool_choice to none for an agent means that the model decides whether to call a tool.                    | <input type="radio"/>            | <input checked="" type="radio"/> |
| In Microsoft Foundry Agent Service, setting tool_choice to required for an agent ensures that the agent must call one or more tools during each run. | <input checked="" type="radio"/> | <input type="radio"/>            |

NEW QUESTION: 36

You are reviewing best practices for using AI at your company.

Which Microsoft responsible AI principle is each task an example of? To answer, drag the appropriate principles to the correct tasks. 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 match is worth one point.

**Principles**

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

**Answer Area**

Evaluating model outputs to ensure that decisions are NOT biased against specific demographic groups: Principle

Encrypting sensitive customer data and restricting system access to authorized personnel: Principle

Informing users when they are interacting with an AI system and explaining the system's capabilities and limitations: Principle

Testing AI systems under different conditions to reduce unexpected failures: Principle

**Answer:**

**Principles**

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

**Answer Area**

Evaluating model outputs to ensure that decisions are NOT biased against specific demographic groups: Fairness

Encrypting sensitive customer data and restricting system access to authorized personnel: Privacy and security

Informing users when they are interacting with an AI system and explaining the system's capabilities and limitations: Transparency

Testing AI systems under different conditions to reduce unexpected failures: Reliability and safety

#### NEW QUESTION: 37

You are building a knowledge base by using Azure AI Language service's custom answering feature. Which file format can you use to populate the knowledge base?

- A. PPTX
- B. PDF
- C. JPEG
- D. ZIP

**Answer: B (LEAVE A REPLY)**

Azure AI Language service's custom answering can populate its knowledge base using various file formats, including structured text (.txt, .tsv), Microsoft Office documents (.docx, .xlsx, .pptx), and HTML, PDF, and Markdown files.

Reference:

<https://docs.azure.cn/en-us/ai-services/language-service/question-answering/reference/document-format-guidelines>

#### NEW QUESTION: 38

Drag and Drop Question

You are reviewing best practices for using AI at your company.



Which Microsoft responsible AI principle is each task an example of? To answer, drag the appropriate principles to the correct tasks. 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 match is worth one point.

| Principles             | Answer Area                                                                                                           |  |
|------------------------|-----------------------------------------------------------------------------------------------------------------------|--|
| Accountability         | Evaluating model outputs to ensure that decisions are <b>NOT</b> biased against specific demographic groups:          |  |
| Fairness               | Encrypting sensitive customer data and restricting system access to authorized personnel:                             |  |
| Inclusiveness          | Informing users when they are interacting with an AI system and explaining the system's capabilities and limitations: |  |
| Privacy and security   | Testing AI systems under different conditions to reduce unexpected failures:                                          |  |
| Reliability and safety |                                                                                                                       |  |
| Transparency           |                                                                                                                       |  |

Answer:

| Principles             | Answer Area                                                                                                           |                        |
|------------------------|-----------------------------------------------------------------------------------------------------------------------|------------------------|
| Accountability         | Evaluating model outputs to ensure that decisions are <b>NOT</b> biased against specific demographic groups:          | Fairness               |
| Fairness               | Encrypting sensitive customer data and restricting system access to authorized personnel:                             | Privacy and security   |
| Inclusiveness          | Informing users when they are interacting with an AI system and explaining the system's capabilities and limitations: | Transparency           |
| Privacy and security   | Testing AI systems under different conditions to reduce unexpected failures:                                          | Reliability and safety |
| Reliability and safety |                                                                                                                       |                        |
| Transparency           |                                                                                                                       |                        |

#### NEW QUESTION: 39

Select the answer that correctly completes the sentence.



Answer:




#### NEW QUESTION: 40

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

NOTE: Each correct selection is worth one point.



Answer:



#### NEW QUESTION: 41

You are developing an application that converts text into spoken audio and saves the synthesized audio to a file by using Azure Speech in Foundry Tools.

How should you complete the Python code? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

```

import azure.cognitiveservices.speech as speechsdk
...
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
audio_config = speechsdk.audio.
synthesizer = speechsdk.SpeechSynthesizer(
    speech_config=speech_config,
    audio_config=audio_config
)
...

```

- AudioOutputConfig(filename="output.wav")
- AudioOutputConfig(stream)
- AudioStreamFormat(wave\_stream\_format=AudioStreamWaveFormat.PCM)

Answer:

```

import azure.cognitiveservices.speech as speechsdk
...
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
audio_config = speechsdk.audio.
synthesizer = speechsdk.SpeechSynthesizer(
    speech_config=speech_config,
    audio_config=audio_config
)
...

```

- AudioOutputConfig(filename="output.wav")
- AudioOutputConfig(stream)
- AudioStreamFormat(wave\_stream\_format=AudioStreamWaveFormat.PCM)

### NEW QUESTION: 42

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

NOTE: Each correct selection is worth one point.

| Statements                                                                                              | Yes                   | No                    |
|---------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| Fairness can be achieved by focusing solely on improving the overall accuracy of an AI model.           | <input type="radio"/> | <input type="radio"/> |
| Evaluating AI system outputs to identify and reduce bias across demographic groups supports fairness.   | <input type="radio"/> | <input type="radio"/> |
| Ensuring fairness for an AI system means that all users always receive the same output from the system. | <input type="radio"/> | <input type="radio"/> |

Answer:

| Statements                                                                                              | Yes                              | No                               |
|---------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| Fairness can be achieved by focusing solely on improving the overall accuracy of an AI model.           | <input type="radio"/>            | <input checked="" type="radio"/> |
| Evaluating AI system outputs to identify and reduce bias across demographic groups supports fairness.   | <input checked="" type="radio"/> | <input type="radio"/>            |
| Ensuring fairness for an AI system means that all users always receive the same output from the system. | <input type="radio"/>            | <input checked="" type="radio"/> |



**NEW QUESTION: 43**

You have an Azure subscription.

You need to use Azure Content Understanding in Foundry Tools to extract structured data from invoices.

What should you provision?

- A.** A Microsoft Foundry project
- B.** an Azure OpenAI resource
- C.** a Microsoft Foundry resource
- D.** an Azure AI Search service

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

To extract structured data from invoices using Azure Content Understanding in Foundry Tools, you must provision a Microsoft Foundry resource (often referred to as an Azure AI services resource configured within a supported region).

Reference:

<https://learn.microsoft.com/en-us/azure/ai-services/content-understanding/tutorial/create-custom-analyzer>

**NEW QUESTION: 44**

You have a Microsoft Foundry project that contains a vision-enabled model deployment.

You use the Azure OpenAI Responses API to send a prompt to the model.

You need to provide an image for analysis.

Which content item should you include in the request?

- A.** image\_generation
- B.** output\_image
- C.** input\_image
- D.** image\_base64

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

When sending an image for analysis using the Azure OpenAI Responses API within a Microsoft Foundry project, you must include a content item object with the type input\_image in the message payload.

Required Content Object Structure

Unlike the standard Chat Completions API (which uses type: "image\_url"), the Responses API explicitly structures a multimodal image input as a separate block with the following properties:

type: Must be explicitly set to "input\_image".

image\_url: A dictionary or a direct parameter containing the target location of the image.

Reference:

<https://learn.microsoft.com/en-us/azure/foundry/openai/how-to/responses>

**NEW QUESTION: 45**

Hotspot Question

Select the answer that correctly completes the sentence.

Answer Area

An AI workload that produces new content based on user input is an example of

|                        |
|------------------------|
|                        |
| content understanding  |
| generative AI          |
| information extraction |
| text analysis          |



Answer:

Answer Area

An AI workload that produces new content based on user input is an example of

|                        |
|------------------------|
|                        |
| content understanding  |
| generative AI          |
| information extraction |
| text analysis          |



NEW QUESTION: 46

Your company processes customer support emails. You need to implement an AI solution that automatically identifies mentions of people, organizations, and locations in the emails. Which text analysis technique should you use?

- A. Named Entity Recognition (NER)
- B. key phrase extraction
- C. sentiment analysis
- D. summarization

Answer: (SHOW ANSWER)

The adequate text analysis feature is Named Entity Recognition (NER), which is a core prebuilt capability of the Azure AI Language service. Prebuilt Categorization: The prebuilt NER feature automatically parses unstructured text (like emails) to identify and group elements into standard classes, including Person, Organization, and Location. No Training Required: Because it relies on state-of-the-art pretrained transformer models, it can be deployed immediately out-of-the-box without requiring data labeling or custom model training. High Extensibility: If your emails contain highly specialized terms (such as proprietary product IDs or industry-specific roles), you can easily scale to Custom NER within the same architecture. Reference: <https://learn.microsoft.com/en-us/azure/ai-services/language-service/named-entity-recognition/overview>

**Valid AI-901 Dumps** shared by PrepPdf.com for Helping Passing AI-901 Exam! PrepPdf.com now offer the **newest AI-901 exam dumps**, the PrepPdf.com AI-901 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-901 dumps with Test Engine here:

<https://www.preppdf.com/Microsoft/AI-901-prepaway-exam-dumps.html> (87 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 47

You have a Microsoft Foundry project that contains a vision-enabled model deployment.

You are developing an application that sends images to the model.

You need to ensure that the model can analyze the images.

In which two formats can you provide the images? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. a base64-encoded image data string
- B. a JSON document that describes the image content
- C. a UTF-8 encoded text description of the image
- D. a publicly accessible URL of the image

**Answer: A,D (LEAVE A REPLY)**

Images can be sent to these models either by passing a publicly accessible image URL or by sending Base64-encoded image data directly in the message payload.

Reference:

<https://learn.microsoft.com/en-us/azure/foundry-classic/foundry-models/how-to/use-chat-multi-modal>

#### NEW QUESTION: 48

You have a Microsoft Foundry project that contains a vision-enabled model deployment.

You need to develop an application that sends a message containing text and an image URL. The solution must ensure the quickest response time.

Which message structure should you include in the request?

- A. a user message that includes only a text item and sends an image item in a separate request
- B. a system message that includes both a text item and an image item in the content array
- C. a system message that includes only a text item and sends an image item in a separate request
- D. a user message that includes both a text item and an image item in the content array

**Answer: D (LEAVE A REPLY)**

For a vision-enabled model, Microsoft documentation shows that multimodal prompts can include both text and image content in the same user message content array. Azure OpenAI vision models accept multimodal image-and-text input and return natural language responses.

The correct request structure is a user message with a content array containing both items, for example:

```
{
  "role": "user",
  "content": [
    {"type": "input_text", "text": "What is in this image?"},
    {"type": "input_image", "image_url": image_url}
  ]
}
```

This avoids sending separate requests and is therefore the best choice for quickest response time.  
A and C are incorrect because they send the text and image separately.  
B is incorrect because the image/question should be sent as a user message, while system messages are used for instructions and behavior.

**NEW QUESTION: 49**

Drag and Drop Question

You have a Microsoft Foundry project named project1 that contains an Azure OpenAI resource named Resource1.  
To Resource1, you deploy a gpt-4.1-mini model by using a model deployment named my-mini- gpt.  
You need to connect to my-mini-gpt from an application.  
How should you complete the Python code? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.  
NOTE: Each correct selection is worth one point.

Values

gpt-4.1-mini

my-mini-gpt

project1

resource1

Answer Area

```
client = OpenAI(api_key="...",
base_url="https:// [ ] .openai.azure.com/openai/v1/")
response = client.responses.create(
model="", [ ]
...
)
```

Answer:

Values

gpt-4.1-mini

my-mini-gpt

project1

resource1

Answer Area

```
client = OpenAI(api_key="...",
base_url="https:// [ resource1 ] .openai.azure.com/openai/v1/")
response = client.responses.create(
model="", [ my-mini-gpt ]
...
)
```

**NEW QUESTION: 50**

Hotspot Question

Select the answer that correctly completes the sentence.

When content is submitted to Azure Content Understanding in Foundry Tools, the analysis is


synchronous

asynchronous

returned only as unstructured plain text

limited to optical character recognition (OCR)-only processing

Answer:



#### NEW QUESTION: 51

Hotspot Question

Select the answer that correctly completes the sentence.




Answer:



#### NEW QUESTION: 52

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

NOTE: Each correct selection is worth one point.





### Answer:

Answer Area

| Statements                                                                                                   | Yes                              | No                               |
|--------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| Human-in-the-loop practices provide accountability for AI-generated decisions.                               | <input checked="" type="radio"/> | <input type="radio"/>            |
| Deploying an AI system to a production environment eliminates the need for ongoing monitoring.               | <input type="radio"/>            | <input checked="" type="radio"/> |
| Disclosing the team that designed and deployed an AI system provides accountability for the system's output. | <input checked="" type="radio"/> | <input type="radio"/>            |

### NEW QUESTION: 53

What are two common use cases for generative AI solutions? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. classifying email messages as spam or non-spam
- B. generating draft responses for customer service agents
- C. predicting sales revenue based on historical data
- D. creating original artwork from textual descriptions

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

B: Generative AI tools can assist customer service agents by automatically drafting responses to customer inquiries, providing personalized and context-aware solutions. This helps agents improve efficiency, reduce workload, and enhance customer satisfaction.

D: DALL E3, developed by OpenAI, is a cutting-edge generative AI model designed to create images from text descriptions.

Reference:

<https://medium.com/@eastgate/top-7-generative-ai-models-tools-for-text-image-and-video-creation-237a0de231b1>

### NEW QUESTION: 54

Hotspot Question

You are developing an application that converts text into spoken audio and saves the synthesized audio to a file by using Azure Speech in Foundry Tools.

How should you complete the Python code? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

```
import azure.cognitiveservices.speech as speechsdk
```

```
...
```

```
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
```

```
audio_config = speechsdk.audio.
```

AudioOutputConfig(filename="output.wav")

AudioOutputConfig(stream)

AudioStreamFormat(wave\_stream\_format=AudioStreamWaveFormat.PCM)

```
synthesizer = speechsdk.SpeechSynthesizer(
```

```
    speech_config=speech_config,
```

```
    audio_config=audio_config
```

```
)
```

**Answer:**

```
import azure.cognitiveservices.speech as speechsdk
```

```
...
```

```
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
```

```
audio_config = speechsdk.audio.
```

AudioOutputConfig(filename="output.wav")

AudioOutputConfig(stream)

AudioStreamFormat(wave\_stream\_format=AudioStreamWaveFormat.PCM)

```
synthesizer = speechsdk.SpeechSynthesizer(
```

```
    speech_config=speech_config,
```

```
    audio_config=audio_config
```

```
)
```

#### NEW QUESTION: 55

Verifying that machine learning models do NOT show racial or gender bias is an example of which Microsoft responsible AI principle?

A. fairness

B. privacy and security

C. safety

D. reliability

Answer: A (LEAVE A REPLY)

Fairness and inclusiveness in Azure Machine Learning: The fairness assessment component of the Responsible AI dashboard enables data scientists and developers to assess model fairness across sensitive groups defined in terms of gender, ethnicity, age, and other characteristics.

Reference:

<https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai?view=azureml-api-2>

NEW QUESTION: 56

Select the answer that correctly completes the sentence.

Answer Area

To define an agent's role and behaviors, you must configure a \_\_\_\_\_ for the agent.

- deployment slot
- embedding index
- fine tuning job
- system prompt

Answer:

Answer Area

To define an agent's role and behaviors, you must configure a \_\_\_\_\_ for the agent.

- deployment slot
- embedding index
- fine tuning job
- system prompt

NEW QUESTION: 57

Drag and Drop Question

You are developing an application that extracts structured information from different types of content by using Azure Content Understanding in Foundry Tools. You need to extract scanned invoices in the PDF format and voicemail recordings in the WAV format. Which type of analyzer should you use for each content type? To answer, drag the appropriate analyzer types to the correct content types. Each analyzer type 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.

| Analyzer Types    | Answer Area                                |
|-------------------|--------------------------------------------|
| audio analyzer    | Scanned invoices: <input type="text"/>     |
| document analyzer | Voicemail recordings: <input type="text"/> |
| image analyzer    |                                            |
| video analyzer    |                                            |

Answer:

Analyzer types

audio analyzer

document analyzer

image analyzer

video analyzer

Answer Area

Scanned invoices:

document analyzer

Voicemail recordings:

audio analyzer

NEW QUESTION: 58

Hotspot Question

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

NOTE: Each correct selection is worth one point.

Answer Area

| Statements                                                                                                                            | Yes                   | No                    |
|---------------------------------------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| In the Foundry playground, you can upload a local image and include text in the same message when prompting a multimodal model.       | <input type="radio"/> | <input type="radio"/> |
| When using the OpenAI Responses API and a vision-enabled model, images can be provided only as base64-encoded image data.             | <input type="radio"/> | <input type="radio"/> |
| Prompts that include images require deploying a text-only model because multimodal capabilities are handled by the application layer. | <input type="radio"/> | <input type="radio"/> |

Answer:

Answer Area

Statements

Yes

No

In the Foundry playground, you can upload a local image and include text in the same message when prompting a multimodal model.

☒

☐

When using the OpenAI Responses API and a vision-enabled model, images can be provided only as base64-encoded image data.

☐

☒

Prompts that include images require deploying a text-only model because multimodal capabilities are handled by the application layer.

☐

☒

NEW QUESTION: 59

You have a Microsoft Foundry project that contains an agent named Agent1.

You need to ensure that Agent1 always calls an Azure function when the agent responds to user input.

To what should you set tool\_choice for Agent1?

- A. required
- B. auto
- C. none

**Answer: A (LEAVE A REPLY)**

Yes, you should set tool\_choice="required" if you want to force the model to pick a tool.

Reference:

<https://learn.microsoft.com/en-us/azure/foundry/agents/concepts/tool-best-practice>

## NEW QUESTION: 60

Drag and Drop Question

You have a Microsoft Foundry project named project1 that contains an Azure OpenAI resource named Resource1.

To Resource1, you deploy a gpt-4.1-mini model by using a model deployment named my-mini-gpt.

You need to connect to my-mini-gpt from an application.

How should you complete the Python code? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

NOTE: Each correct selection is worth one point.

**Values**

gpt-4.1-mini

my-mini-gpt

project1

resource1

**Answer Area**

```
client = OpenAI(api_key="...",
base_url="https://.openai.azure.com/openai/v1/",)
response = client.responses.create(
    model=" ",
    ...
)
```

**Answer:**

**Values**

gpt-4.1-mini

my-mini-gpt

project1

resource1

**Answer Area**

```
client = OpenAI(api_key="...",
base_url="https://resource1.openai.azure.com/openai/v1/",)
response = client.responses.create(
    model="my-mini-gpt ",
    ...
)
```



### NEW QUESTION: 61

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

NOTE: Each correct selection is worth one point.

| Statements                                                                                                                                           | Yes                   | No                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| In Microsoft Foundry Agent Service, setting tool_choice to auto for an agent enables the agent to decide whether to call a tool.                     | <input type="radio"/> | <input type="radio"/> |
| In Microsoft Foundry Agent Service, setting tool_choice to none for an agent means that the model decides whether to call a tool.                    | <input type="radio"/> | <input type="radio"/> |
| In Microsoft Foundry Agent Service, setting tool_choice to required for an agent ensures that the agent must call one or more tools during each run. | <input type="radio"/> | <input type="radio"/> |

Answer:

| Statements                                                                                                                                           | Yes                              | No                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| In Microsoft Foundry Agent Service, setting tool_choice to auto for an agent enables the agent to decide whether to call a tool.                     | <input checked="" type="radio"/> | <input type="radio"/>            |
| In Microsoft Foundry Agent Service, setting tool_choice to none for an agent means that the model decides whether to call a tool.                    | <input type="radio"/>            | <input checked="" type="radio"/> |
| In Microsoft Foundry Agent Service, setting tool_choice to required for an agent ensures that the agent must call one or more tools during each run. | <input checked="" type="radio"/> | <input type="radio"/>            |

**Valid AI-901 Dumps** shared by PrepPdf.com for Helping Passing AI-901 Exam! PrepPdf.com now offer the **newest AI-901 exam dumps**, the PrepPdf.com AI-901 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-901 dumps with Test Engine here:

<https://www.preppdf.com/Microsoft/AI-901-prepaway-exam-dumps.html> (87 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 62

You are developing an application that analyzes voicemail recordings by using Azure Content Understanding in Foundry Tools.

You need to extract a transcript and structured information from the recordings.

Which type of analyzer should you use?

- A. document analyzer
- B. video analyzer
- C. audio analyzer
- D. image analyzer

Answer: C (**LEAVE A REPLY**)

Voicemail recordings are audio content. Azure Content Understanding analyzers define what type of content to process, including documents, images, audio, or video, and what elements to extract, including transcripts and structured fields.

Microsoft's custom analyzer documentation also shows an audio example based on prebuilt- audio for processing customer support call recordings, which is the same content type as voicemail recordings.

Therefore, to extract a transcript and structured information from voicemail recordings, you should use an audio analyzer.

### NEW QUESTION: 63

You are developing an application that analyzes voicemail recordings by using Azure Content Understanding in Foundry Tools. You need to extract a transcript and structured information from the recordings. Which type of analyzer should you use?

- A. image analyzer
- B. video analyzer
- C. audio analyzer
- D. document analyzer

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

#### NEW QUESTION: 64

In Azure Machine Learning, what are two ensemble methods for combining models in automated machine learning (automated ML)? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. voting
- B. computer vision
- C. stacking
- D. classification
- E. regression

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

In Azure Machine Learning's automated machine learning (AutoML), two prominent ensemble methods for combining models are Voting and Stacking.

Reference:

<https://learn.microsoft.com/en-us/azure/machine-learning/concept-automated-ml> 1

#### NEW QUESTION: 65

Select the answer that correctly completes the sentence.

Answer Area

defines which fields to extract when analyzing content

- A keyword list
- Optical character recognition (OCR)-only processing
- A schema
- A synchronous API call

Microsoft

**Answer:**

Answer Area

defines which fields to extract when analyzing content

- A keyword list
- Optical character recognition (OCR)-only processing
- A schema
- A synchronous API call

Microsoft

#### NEW QUESTION: 66

You need to build an AI solution that generates marketing email drafts based on a short description of a product and its target audience. Which AI workload should you use?

- A. speech recognition
- B. text classification
- C. generative AI
- D. computer vision

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

A Generative AI workload is exactly what should be used for this scenario. Generative AI models (such as LLMs) are specifically designed to interpret prompts and create original, context-aware content, making them the perfect fit for drafting tailored marketing emails based on a product description and target audience.

Reference:

<https://azure.microsoft.com/en-us/products/ai-foundry/tools/content-understanding>

#### NEW QUESTION: 67

You are authoring a Conversational Language Understanding application to support a music festival.

You want users to be able to ask questions about scheduled shows, such as: "Which act is playing on the main stage?" The question "Which act is playing on the main stage?" is an example of which type of element?

- A. an intent
- B. an utterance
- C. a domain
- D. an entity

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

In Conversational Language Understanding (CLU), an utterance is the spoken or typed input provided by a user in a conversation. It represents a discrete element of their communication, which can be a single sentence, a question, a command, or any form of human language.?

It's a natural language phrase such as "book 2 tickets to Seattle next Tuesday". Example utterances are added to train the model and the model predicts on new utterance at runtime.

Reference:

<https://learn.microsoft.com/en-us/azure/ai-services/language-service/conversational-language-understanding/glossary>

**Valid AI-901 Dumps** shared by PrepPdf.com for Helping Passing AI-901 Exam! PrepPdf.com now offer the **newest AI-901 exam dumps**, the PrepPdf.com AI-901 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-901 dumps with Test Engine here:

<https://www.preppdf.com/Microsoft/AI-901-prepaway-exam-dumps.html> (87 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)