

Microsoft

Exam Questions AI-901

Microsoft Azure AI Fundamentals (Updated Version)



NEW QUESTION 1

HOTSPOT

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)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

AudioOutputConfig(filename="output.wav")
The question specifically states the application must save the synthesized audio to a file. In the Azure Speech SDK for Python, speechsdk.audio.AudioOutputConfig(filename="output.wav") directs the synthesizer to write the generated speech output directly to a WAV file on disk — which is exactly the requirement.
Why the other options are wrong:
AudioOutputConfig(stream)— This routes audio output to an in-memory audio stream object, not a file. It is used when you want to process or play the audio programmatically without saving it to disk.
AudioStreamFormat(wave_stream_format=AudioStreamWaveFormat.PCM)— This class defines the format of an audio stream (e.g., PCM encoding, sample rate). It is used when configuring custom audio streams, not when specifying a file output destination. It is not a valid argument for speechsdk.audio. in this context.
The correct and complete line is:
audio_config = speechsdk.audio.AudioOutputConfig(filename="output.wav")

NEW QUESTION 2

HOTSPOT

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

Answer Area

Statements	Yes	No
Human-in-the-loop practices provide accountability for AI-generated decisions.	☐	☐
Deploying an AI system to a production environment eliminates the need for ongoing monitoring.	☐	☐
Disclosing the team that designed and deployed an AI system provides accountability for the system's output.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Human-in-the-loop practices provide accountability for AI-generated decisions.	☒	☐
Deploying an AI system to a production environment eliminates the need for ongoing monitoring.	☐	☒
Disclosing the team that designed and deployed an AI system provides accountability for the system's output.	☒	☐

NEW QUESTION 3

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: B

Explanation:

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 4

HOTSPOT

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

Answer Area

Statements	Yes	No
An AI generative model is retrained before performing each user request.	☐	☐
An AI agent responds by copying and pasting answers stored in a database.	☐	☐
An AI agent uses a generative AI model to establish actions based on user input.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
An AI generative model is retrained before performing each user request.	☐	☒
An AI agent responds by copying and pasting answers stored in a database.	☐	☒
An AI agent uses a generative AI model to establish actions based on user input.	☒	☐

NEW QUESTION 5

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

After deploying a vision-enabled GPT model in Microsoft Foundry,
you can configure an application to send requests to the

endpoint of the model
evaluation pipeline of model
Foundry playground
training dataset of the model

.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

After deploying a vision-enabled GPT model in Microsoft Foundry,
you can configure an application to send requests to the

endpoint of the model
evaluation pipeline of model
Foundry playground
training dataset of the model

.

NEW QUESTION 6

HOTSPOT

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

Answer Area

Statements	Yes	No
Evaluators in Microsoft Foundry replace the need for configuring token limits.	☐	☐
Evaluators in Microsoft Foundry can assess the quality and safety of responses generated by a generative AI model.	☐	☐
Evaluators in Microsoft Foundry can retrain a deployed generative AI model automatically when quality issues are detected.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area		
Statements	Yes	No
Evaluators in Microsoft Foundry replace the need for configuring token limits.	☐	☒
Evaluators in Microsoft Foundry can assess the quality and safety of responses generated by a generative AI model.	☒	☐
Evaluators in Microsoft Foundry can retrain a deployed generative AI model automatically when quality issues are detected.	☐	☒

NEW QUESTION 7

DRAG DROP

You have a Microsoft Foundry project that contains a model deployment.

You are developing an application that sends an image and a user question to the model. You need to send both text and image content in the same request so the model can return an answer.

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. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

input_image

input_text

input_url

output_image

output_text

Answer Area

```
input=[
    {
        "role": "user",
        "content": [
            {"type": "text", "text": "What is in this image? Provide 3 bullet points."},
            {"type": "image_url", "image_url": image_url}
        ]
    }
],
print(response.output_text)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Values

input_image

input_text

input_url

output_image

output_text

Answer Area

```
input=[
    {
        "role": "user",
        "content": [
            {"type": "text", "text": "What is in this image? Provide 3 bullet points."},
            {"type": "image_url", "image_url": image_url}
        ]
    }
],
print(response.output_text)
```

NEW QUESTION 8

HOTSPOT

Select the answer that correctly complete the sentence.

Answer Area


The  is used for comparing and deploying a wide range of models for generative AI development in Microsoft Foundry.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

The  is used for comparing and deploying a wide range of models for generative AI development in Microsoft Foundry.

NEW QUESTION 9

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

Explanation:

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 10

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

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.	☐	☐
In Microsoft Foundry Agent Service, setting tool_choice to none for an agent means that the model decides whether to call a tool.	☐	☐
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.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 10

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

When using the OpenAI Responses API and a vision-enabled model, you can include an image in a request by providing the image as

a base64-encoded image data

a CSV file attachment

an MP4 video stream

a shared access signature (SAS) token

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

When using the OpenAI Responses API and a vision-enabled model, you can include an image in a request by providing the image as

a base64-encoded image data

a CSV file attachment

an MP4 video stream

a shared access signature (SAS) token

NEW QUESTION 14

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

NEW QUESTION 16

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Voice Live returns only transcribed text.	☐	☐
Voice Live requires you to separately implement speech to text and text to speech services.	☐	☐
Voice Live combines speech to text, reasoning, and text to speech into a single conversational experience.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Voice Live returns only transcribed text.	☐	☒
Voice Live requires you to separately implement speech to text and text to speech services.	☐	☒
Voice Live combines speech to text, reasoning, and text to speech into a single conversational experience.	☒	☐

NEW QUESTION 21

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

A keyword list

Optical character recognition (OCR)-only processing

A schema

A synchronous API call

defines which fields to extract when analyzing content

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A keyword list

Optical character recognition (OCR)-only processing

A schema

A synchronous API call

defines which fields to extract when analyzing content

NEW QUESTION 24

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

NEW QUESTION 28

Your company has thousands of recorded customer support calls in multiple languages stored as audio files in Azure Storage. You need to generate text transcripts of all the recordings. Which Azure Speech in Foundry Tools capability should you use?

- A. speech to text batch transcription
- B. speech to text real-time transcription
- C. text to speech
- D. speech translation

Answer: A

NEW QUESTION 33

HOTSPOT

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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 36

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AI-901 Practice Exam Features:

- * AI-901 Questions and Answers Updated Frequently
- * AI-901 Practice Questions Verified by Expert Senior Certified Staff
- * AI-901 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * AI-901 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AI-901 Practice Test Here](#)