

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

HOTSPOT

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The completed sentence is:
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.
Accountability means people and organizations remain responsible for AI systems and their effects. Human review and responsibility for final decisions are examples of accountability.
The other options are incorrect:
fairness focuses on avoiding bias and treating people fairly.
privacy and security focuses on protecting data and restricting access.
transparency focuses on explaining AI use, capabilities, and limitations.

NEW QUESTION 4

HOTSPOT

You are developing a voice application that listens for spoken commands and converts them into text 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.

Answer Area

```
import azure.cognitiveservices.speech as speechsdk
...
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
recognizer = speechsdk.SpeechRecognizer(speech_config=speech_config)
recognizer.recognize_once()
recognizer.speak_text_async("Ready")
recognizer.start_continuous_recognition()
recognizer.start_keyword_recognition()
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The correct Python method is:

```
recognizer.recognize_once()
```

Completed code:

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

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

```
recognizer = speechsdk.SpeechRecognizer(speech_config=speech_config)
```

```
recognizer.recognize_once()
```

Microsoft's Azure Speech documentation for Foundry Tools explains that Speech to text is used for real-time speech recognition and converting spoken audio into text. The Python Speech SDK uses a SpeechRecognizer object for recognition from audio input, such as a microphone.

Why the other options are incorrect:

`recognizer.speak_text_async("Ready")` is incorrect because speaking text is text-to-speech, not speech-to-text recognition.

`recognizer.start_continuous_recognition()` can be used for continuous recognition, but the code shown is asking for the basic method to recognize spoken input and convert it to text from the SpeechRecognizer.

`recognizer.start_keyword_recognition()` is used for keyword/wake-word recognition, not general speech-to-text transcription of spoken commands.

Therefore, the correct answer is:

```
recognizer.recognize_once()
```

NEW QUESTION 5

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

Answer: A

NEW QUESTION 6

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_base64`
- B. `image_generation`
- C. `output_image`
- D. `input_image`

Answer: D

NEW QUESTION 7

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

NEW QUESTION 8

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

NEW QUESTION 9

DRAG DROP

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 Fairness Inclusiveness Privacy and security Reliability and safety Transparency	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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Principles	Answer Area
Accountability Fairness Inclusiveness Privacy and security Reliability and safety Transparency	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 10

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

Answer: B

NEW QUESTION 10

HOTSPOT
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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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 14

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 16

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 21

You are developing an AI-powered customer support application.

Which task is an example of the Microsoft responsible AI principle of inclusiveness?

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

Answer: B

NEW QUESTION 25

HOTSPOT

Select the answer that correctly complete the sentence.

Answer Area

Information extraction solutions that detect and read text in scanned documents and images rely on

computer vision

image generation

sentiment analysis

speech synthesis

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Information extraction solutions that detect and read text in scanned documents and images rely on

computer vision

image generation

sentiment analysis

speech synthesis

NEW QUESTION 29

DRAG DROP

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

- audio analyzer

document analyzer

image analyzer

video analyzer

Answer Area

Scanned invoices:

Voicemail recordings:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Analyzer Types

- audio analyzer

document analyzer

image analyzer

video analyzer

Answer Area

Scanned invoices: document analyzer

Voicemail recordings: audio analyzer

NEW QUESTION 31

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

Explanation:

To compare the costs of large language models in Microsoft Foundry portal, use theModel leaderboard.
Microsoft documentation states that themodel leaderboardhelps compare models acrossquality, 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 catalogis for selecting evaluators to measure model or application outputs, not comparing LLM costs.
* C. Compliancerelates to governance and compliance, not model cost comparison.
* D. Toolspromotes Foundry tools, not benchmarked cost comparison across models.

NEW QUESTION 33

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 37

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: CD

NEW QUESTION 41

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
The Temperature parameter can be set before deploying a model.	☐	☐
During inference, the model name is used to route requests to a specific deployment.	☐	☐
After a model is deployed, both code and testing tools can be used to interact with the model.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The Temperature parameter can be set before deploying a model.	☐	☒
During inference, the model name is used to route requests to a specific deployment.	☐	☒
After a model is deployed, both code and testing tools can be used to interact with the model.	☒	☐

NEW QUESTION 46

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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 FirstTry
- * AI-901 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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