

Exam Questions AI-901

Microsoft Azure AI Fundamentals (Updated Version)

<https://www.2passeasy.com/dumps/AI-901/>



NEW QUESTION 1

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

NEW QUESTION 2

HOTSPOT

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 3

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 4

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 5

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 6

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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 7

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

Answer: B

NEW QUESTION 8

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 9

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 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

In Microsoft Foundry Agent Service, setting tool_choice to auto for an agent enables the agent to decide whether to call a tool.

Yes

☐

No

☐

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

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.

Yes

No

☒☐☐☒☒☐

NEW QUESTION 10

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 14

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 19

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

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.

Yes

No

☐☐☐☐☐☐

- 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 20

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 23

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
Generating a response to a user prompt occurs during the inference stage.	☐	☐
A generative AI model generates responses by copying stored documents directly from the model's training data.	☐	☐
A generative AI model produces output by predicting the next token based on patterns learned from the model's training data.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Generating a response to a user prompt occurs during the inference stage.	☒	☐
A generative AI model generates responses by copying stored documents directly from the model's training data.	☐	☒
A generative AI model produces output by predicting the next token based on patterns learned from the model's training data.	☒	☐

NEW QUESTION 26

You have a Microsoft Foundry project that contains a generative AI model deployment. You test the model by using the Foundry playground. You need to develop an application that sends requests to the deployed model. Which information must the application include to call the model?

- A. The model training dataset
- B. The Foundry project display name
- C. The exported playground session history
- D. The model endpoint and authentication credentials

Answer: D

Explanation:

To call a deployed Azure OpenAI model from an application, the app must know the service endpoint and authenticate its request. Microsoft documentation states that Azure OpenAI supports API key authentication or Microsoft Entra ID authentication, and API key authentication requires including the API key in the request. Microsoft quickstart guidance also states that to successfully make a call against Azure OpenAI, you need an endpoint and a key. The application does not need the model training dataset, the Foundry project display name, or exported playground session history to call the deployed model.

NEW QUESTION 28

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AI-901 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AI-901 Product From:

<https://www.2passeasy.com/dumps/AI-901/>

Money Back Guarantee

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