

Microsoft

Exam Questions AI-901

Microsoft Azure AI Fundamentals (Updated Version)



NEW QUESTION 1

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

Explanation:

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 2

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 3

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 4

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Fairness can be achieved by focusing solely on improving the overall accuracy of an AI model.	☐	☒
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 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
System prompts can be used to authorize users.	☐	☐
A system prompt is used to reduce tokens per minute (TPM).	☐	☐
A system prompt guides the behavior of a generative AI model.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
A system prompt, also called a system message, is used to guide how a generative AI model behaves, including its tone, role, response style, and output format. Microsoft states that system messages help steer an Azure OpenAI chat model toward the desired behavior, tone, and output format.
Statement 1: System prompts can be used to authorize users. = No
Authorization is handled by application identity, access control, authentication, and security configuration. A system prompt can guide model behavior, but it is not an authorization mechanism.
Statement 2: A system prompt is used to reduce tokens per minute (TPM). = No
TPM is a quota or throughput setting for model deployments. A system prompt does not reduce or configure TPM.
Statement 3: A system prompt guides the behavior of a generative AI model. = Yes
This is correct. Microsoft documentation says system messages help steer the model's behavior, tone, and output format.

NEW QUESTION 7
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 8

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 9

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 10

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 10
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": "image_url", "url": "input_image"},
            {"type": "text", "text": "What is in this image? Provide 3 bullet points."},
        ],
    },
]
```

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", "image_url": image_url}
    ]
  }
],
print(response.output_text)
```

NEW QUESTION 14

HOTSPOT

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

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.

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.

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.

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

Yes

No

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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.

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.

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.

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

Yes

No

NEW QUESTION 18

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 the new Microsoft Foundry portal, you must fine-tune a model before you can deploy the model.

In the new Microsoft Foundry portal, you must fine-tune a model before you can deploy the model.

In the new Microsoft Foundry portal, you can test a model from the model catalog only after you deploy the model.

In the new Microsoft Foundry portal, you can test a model from the model catalog only after you deploy the model.

In the new Microsoft Foundry portal, you can deploy a model from the model catalog only after retraining the model.

In the new Microsoft Foundry portal, you can deploy a model from the model catalog only after retraining the model.

Yes

No

Passing Certification Exams Made Easy

visit - <https://www.surepassexam.com>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Statement 1:In the new Microsoft Foundry portal, you must fine-tune a model before you can deploy the model.=No
Fine-tuning is optional. Microsoft's Foundry model deployment documentation describes deploying Foundry Models directly from the model catalog for inference. It does not require fine-tuning first.
Statement 2:In the new Microsoft Foundry portal, you can test a model from the model catalog only after you deploy the model.=No
Microsoft documentation states thatsome Foundry Tools are available to try via the model catalog without a project, and Foundry playgrounds are used for prototyping and validation before production. Therefore, the statement using ??only after you deploy?? is too restrictive.
Statement 3:In the new Microsoft Foundry portal, you can deploy a model from the model catalog only after retraining the model.=No
Retraining/fine-tuning is not required before deployment. Microsoft states that after you deploy a Foundry Model, you can interact with it in the Foundry Playground and use it from code, and the deployment workflow starts by selecting a model from the model catalog and choosingDeploy.

NEW QUESTION 22

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 25

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 27

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 28

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 32

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 36

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 41

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 45

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

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

NEW QUESTION 50

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 52

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 56

.....

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](#)