

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 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. AudioConfig
- B. SpeechSynthesizer
- C. AuditOutputConfig

Answer: C

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

DRAG DROP

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	Answer Area
gpt-4.1-mini	<pre>client = OpenAI(api_key="...", base_url="https://<project1>.openai.azure.com/openai/v1/",) response = client.responses.create(model="<my-mini-gpt>", ...)</pre>
my-mini-gpt	
project1	
resource1	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

For Azure OpenAI in Microsoft Foundry, the base_url uses the Azure OpenAIresource namein the endpoint format:

https:// < resource-name > .openai.azure.com/openai/v1/

In the question, the Azure OpenAI resource is namedResource1, so the first blank must beresource1. Microsoft documentation for Azure OpenAI v1 endpoints confirms that the endpoint must use the ...openai.azure.com/openai/v1/ path.

For the model parameter, Azure OpenAI requires thedeployment name, not the underlying model name. Microsoft states that Azure OpenAI always requires the deployment name when calling APIs, even when the parameter is named model.

The deployed model isgpt-4.1-mini, but the deployment name ismy-mini-gpt. Therefore, the second blank must be:

model="my-mini-gpt"

So the correct selections are:

base_url blank =resource1

model blank =my-mini-gpt

NEW QUESTION 5

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: AB

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
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 7

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 8

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 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

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

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Principles	Answer Area
Accountability	Evaluating model outputs to ensure that decisions are NOT biased against specific demographic groups: Fairness
Fairness	
Inclusiveness	Encrypting sensitive customer data and restricting system access to authorized personnel: Privacy and security
Privacy and security	Informing users when they are interacting with an AI system and explaining the system's capabilities and limitations: Transparency
Reliability and safety	
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	Answer Area
input_image	<pre>input=[{ "role": "user", "content": [{"type": "image_url", "image_url": input_image, "text": "What is in this image? Provide 3 bullet points."}, {"type": "text", "text": "What is in this image? Provide 3 bullet points."},], },]</pre>
input_text	
input_url	
output_image	
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 13

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 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.

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

- 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.

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 19

Passing Certification Exams Made Easy

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

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 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
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 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.	☒	☐

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. 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 31

HOTSPOT

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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 35

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 38

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 43

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. Toolsprovides Foundry tools, not benchmarked cost comparison across models.

NEW QUESTION 44

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. auto
- B. none
- C. required

Answer: C

NEW QUESTION 45

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 50

.....

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](#)