

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

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

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

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

gpt-4.1-mini

my-mini-gpt

project1

resource1

Answer Area

```
client = OpenAI(api_key="...",
base_url="https://project1.openai.azure.com/openai/v1/",)
response = client.responses.create(
model="my-mini-gpt",
...
)
```

- 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 OpenAI resource name in 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 6

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 7

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 8

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 9

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 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
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.	☐	☐

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

Explanation:



NEW QUESTION 18

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 19

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 20

HOTSPOT

Select the answer that correctly completes the sentence.



- 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 24

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 28

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

Voice Live returns only transcribed text.

Yes

☐

No

☐

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 29

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 32

HOTSPOT

You are developing an application that analyze invoices by using Azure Content Understanding in Foundry Tools. You need to ensure that the application retrieves the analysis results after processing completes. 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

```
poller = client.begin_analyze(  
    analyzer_id="invoice",  
    input_url=url  
)  
result = poller.
```

get_results

result

status

wait

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
poller = client.begin_analyze(  
    analyzer_id="invoice",  
    input_url=url  
)  
result = poller.  
    ( )
```

get_results

result

status

wait

NEW QUESTION 34

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 36

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

To define an agent's role and behaviors, you must configure a

deployment slot

embedding index

fine-tuning job

system prompt

for the agent.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To define an agent's role and behaviors, you must configure a

deployment slot

embedding index

fine-tuning job

system prompt

 for the agent.

NEW QUESTION 39
.....

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](#)