

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

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 4

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 5

HOTSPOT

Select the answer that correctly completes the sentence.

When content is submitted to Azure Content Understanding in Foundry Tools, the analysis is



- 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 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
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.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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.	☒	☐

NEW QUESTION 7

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 8

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": "text", "text": "What is in this image? Provide 3 bullet points."},
      {"type": "image", "image_url": image_url}
    ]
  }
],
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 9

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 10

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

Answer Area

Statements	Yes	No
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.	☐	☐

- 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 10

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 11

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 12

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 15

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 19

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 21

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 22

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 24

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 27

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 28

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

defines which fields to extract when analyzing content

A keyword list

Optical character recognition (OCR)-only processing

A schema

A synchronous API call

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

defines which fields to extract when analyzing content

A keyword list

Optical character recognition (OCR)-only processing

A schema

A synchronous API call

NEW QUESTION 33

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

for the agent.

deployment slot

embedding index

fine-tuning job

system prompt

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

for the agent.

deployment slot

embedding index

fine-tuning job

system prompt

NEW QUESTION 36

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 38

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 40

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 45

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

An AI workload that produces new content based on user input is an example of

content understanding

generative AI

information extraction

text analysis

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

An AI workload that produces new content based on user input is an example of

content understanding

generative AI

information extraction

text analysis

NEW QUESTION 46

.....

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](#)