


AI-901 Dumps

Microsoft Azure AI Fundamentals (Updated Version)

<https://www.certleader.com/AI-901-dumps.html>



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

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

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

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
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
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 are developing an application that analyzes voicemail recordings by using Azure Content Understanding in Foundry Tools. You need to extract a transcript and structured information from the recordings. Which type of analyzer should you use?

- A. document analyzer
- B. video analyzer
- C. audio analyzer
- D. image analyzer

Answer: C

NEW QUESTION 8

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 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
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 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": "text", "text": "What is in this image? Provide 3 bullet points."},
      {"type": "image_url", "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_url", "image_url": image_url}
    ]
  }
],
)
print(response.output_text)
```

NEW QUESTION 10

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 11

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	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.	☐	☒

NEW QUESTION 16

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 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 deploy a model from the model catalog only after retraining the model.	☐	☐

- 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 20

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 24

HOTSPOT

Select the answer that correctly complete the sentence.

Answer Area

The is used for comparing and deploying a wide range of models for generative AI development in Microsoft Foundry.

Model catalog

Monitor page

Service endpoints page

Solution templates page

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

The is used for comparing and deploying a wide range of models for generative AI development in Microsoft Foundry.

Model catalog

Monitor page

Service endpoints page

Solution templates page

NEW QUESTION 26

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 27

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 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

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 38

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 43

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

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 48

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 53

You are developing an application that processes voicemail recordings by using Azure Content Understanding in Foundry Tools.
Which feature does Azure Content Understanding use to convert audio to text?

- A. Voice Live
B. key phrase extraction
C. transcription
D. optical character recognition (OCR)

Answer: C

Explanation:

Azure Content Understanding usestranscriptionto convert audio content, such as voicemail recordings, into text. Microsoft??s Azure Content Understanding audio documentation states that transcription converts conversational audio into searchable and analyzable text-based transcripts.
OptionA. Voice Liveis not the feature used by Content Understanding to convert voicemail recordings into text.
OptionB. key phrase extractionidentifies important phrases after text is available; it is not the audio-to-text conversion step.
OptionD. optical character recognition (OCR)is for extracting text from images or documents, not audio.
Therefore, the correct answer isC. transcription.

NEW QUESTION 56

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AI-901 Exam with Our Prep Materials Via below:

<https://www.certleader.com/AI-901-dumps.html>