

Microsoft

Exam Questions AI-901

Microsoft Azure AI Fundamentals (Updated Version)



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

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 .

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
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 4
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 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
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

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.

Yes

No

☐

☒

☐

☐

☒

☒

☐

NEW QUESTION 6

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

After deploying a vision-enabled GPT model in Microsoft Foundry, you can configure an application to send requests to the .

endpoint of the model

evaluation pipeline of model

Foundry playground

training dataset of the model

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

After deploying a vision-enabled GPT model in Microsoft Foundry, you can configure an application to send requests to the .

endpoint of the model

evaluation pipeline of model

Foundry playground

training dataset of the model

NEW QUESTION 7

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 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
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 9

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 Fairness Inclusiveness Privacy and security Reliability and safety Transparency	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	Answer Area
Accountability	Evaluating model outputs to ensure that decisions are NOT biased against specific demographic groups
Fairness	
Inclusiveness	Encrypting sensitive customer data and restricting system access to authorized personnel
Privacy and security	
Reliability and safety	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

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	
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 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	
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

In Microsoft Foundry Agent Service, setting `tool_choice` to `auto` for an agent enables the agent to decide whether to call a tool.

Yes
☒

No
☐

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 12

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 15

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 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	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 22

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

Vocicemail recordings

audio analyzer

NEW QUESTION 25

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 28

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 31

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 33

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 34
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 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
Generating a response to a user prompt occurs during the inference stage.	☐	☐
A generative AI model generates responses by copying stored documents directly from the model's training data.	☐	☐
A generative AI model produces output by predicting the next token based on patterns learned from the model's training data.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Generating a response to a user prompt occurs during the inference stage.	☒	☐
A generative AI model generates responses by copying stored documents directly from the model's training data.	☐	☒
A generative AI model produces output by predicting the next token based on patterns learned from the model's training data.	☒	☐

NEW QUESTION 42

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 43

.....

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](#)