Accelerate Your Certification with Microsoft AI-901 Practice Questions

Last chance to prepare smart! Get your hands on free Microsoft Azure AI Fundamentals (Updated Version) exam PDF questions. Study real AI-901 dumps with verified answers and fast-track your certification success with PrepBolt AI-901 exam pdf questions and answers.

Thank you for Downloading AI-901 exam PDF Demo

https://prepbolt.com/AI-901.html

QUESTIONS & ANSWERS

DEMO VERSION

(LIMITED CONTENT)

Question 1

Question Type: MultipleChoice

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?

Options:

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 the Model leaderboard.

Microsoft documentation states that the model leaderboard helps compare models across quality, 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 catalog is for selecting evaluators to measure model or application outputs, not comparing LLM costs. C . Compliance relates to governance and compliance, not model cost comparison. D . Tools provides Foundry tools, not benchmarked cost comparison across models.

Question 2

Question Type: MultipleChoice

You are developing an AI-powered customer support application.

Which task is an example of the Microsoft responsible AI principle of inclusiveness?

Options:

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

Explanation:

The Microsoft responsible AI principle of inclusiveness means AI systems should be designed to empower and engage everyone, including people with different abilities, languages, and accessibility needs.

Therefore, designing the interface to support multiple languages and screen readers is an example of inclusiveness.

Why the other options are incorrect:

A . Provide explanations about how predictions are generated = Transparency C . Evaluate model outputs across demographic groups to reduce bias = Fairness D . Encrypt stored customer data and restrict access by using role-based controls = Privacy and security

Question 3

Question Type: MultipleChoice

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?

Options:

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.

Question 4

Question Type: MultipleChoice

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?

Options:

A- image generation
B- image analysis
C- object detection
D- optical character recognition (OCR)

Answer:

A

Explanation:

The requirement is to produce new product images based on written descriptions. This is an image generation workload, because the AI system is creating entirely new images from natural language prompts.

Why the other options are incorrect:

B . image analysis is used to examine and interpret existing images.

C . object detection is used to identify and locate objects within an existing image.

D . optical character recognition (OCR) is used to extract text from images or scanned documents.

Since the solution must generate new visual content from user-provided descriptions, the correct answer is:

A . image generation

Question 5

Question Type: DragDrop

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.

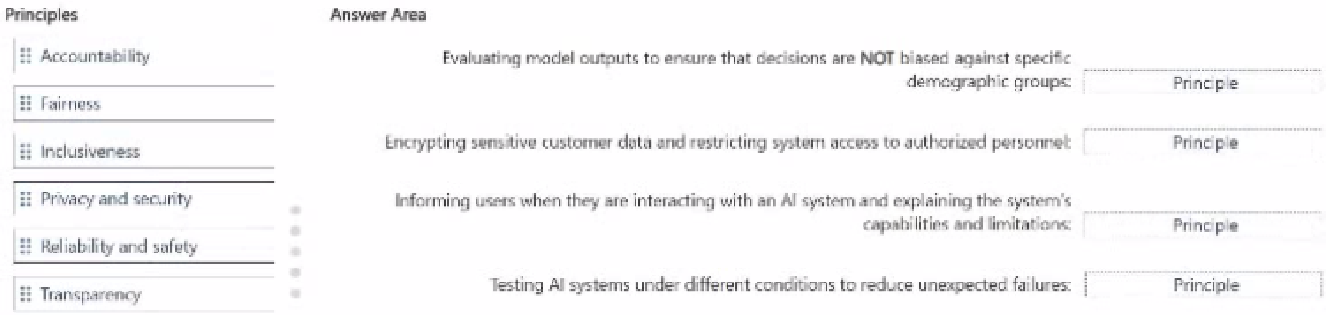

Principles

- Accountability
- Fairness
- Inclusiveness
- Privacy and security
- Reliability and safety
- Transparency

Answer Area

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

Answer:

See the Answer in the Premium Version!

Question 6

Question Type: MultipleChoice

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?

Options:

A- sentiment analysis
B- Named Entity Recognition (NER)
C- summarization
D- key phrase extraction

Answer:

B

Explanation:

The correct text analysis technique is Named Entity Recognition (NER).

Microsoft defines NER as a feature that identifies and categorizes entities in unstructured text, including people, places, and organizations.

Sentiment analysis detects positive, negative, or neutral opinion. Summarization creates shorter versions of text. Key phrase extraction identifies important phrases, but it does not specifically classify mentions as people, organizations, or locations.

Question 7

Question Type: MultipleChoice

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?

Options:

A- speech to text batch transcription
B- speech to text real-time transcription
C- text to speech
D- speech translation

Answer:

A

Explanation:

For thousands of recorded support calls stored as audio files in Azure Storage, the correct capability is speech to text batch transcription.

Microsoft states that batch transcription is designed to transcribe a large amount of audio data in storage, including audio files in Azure Blob Storage, and that files can be processed concurrently to reduce turnaround time.

Real-time transcription is for live audio, not large stored batches. Text to speech converts text into audio. Speech translation translates speech between languages, but the requirement is to generate transcripts.

Question 8

Question Type: MultipleChoice

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?

Options:

A- Temperature
B- Max Completion Tokens
C- Top P
D- Model version settings

Answer:

B

Explanation:

To minimize cost and keep generated responses within a defined length, configure Max Completion Tokens.

Microsoft's Azure OpenAI / Foundry API reference defines max_completion_tokens as an upper bound for the number of tokens that can be generated for a completion. Because generated tokens contribute to usage and response length, limiting completion tokens helps control both output length and cost.

Temperature and Top P control randomness or sampling behavior, not maximum response length. Model version settings do not directly define the generated response length.

Question 9

Question Type: Hotspot

Select the answer that correctly complete the sentence.

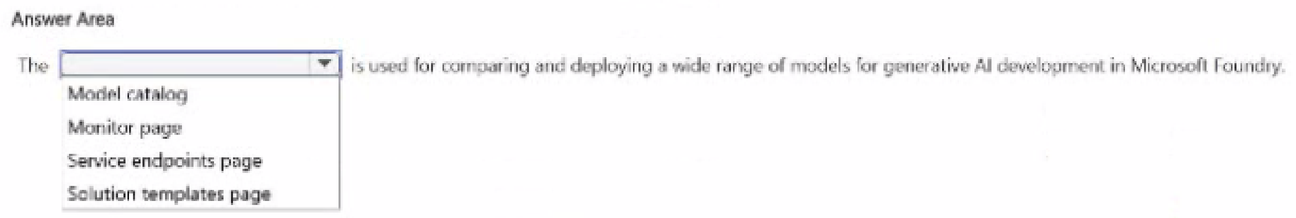

Answer Area

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

Answer:

See the Answer in the Premium Version!

Thank You for trying AI-901 PDF Demo

To try our AI-901 practice exam software visit link below

https://prepbolt.com/AI-901.html

Start Your AI-901 Preparation

Use Coupon “SAVE50” for extra 50% discount on the purchase of Practice Test Software. Test your AI-901 preparation with actual exam questions.