



Microsoft AI-901

Microsoft Azure AI Fundamentals (Updated Version)

Thank You for Downloading AI-901 Updated Exam Questions

<https://certs4sale.com/microsoft/ai-901-pdf-exam-dumps>

AI-901

Microsoft Azure AI Fundamentals

RESPONSIBLE AI, MACHINE LEARNING PRINCIPLES, AZURE AI FOUNDRY & COMPUTER VISION

Microsoft

VENDOR

5

PRACTICE QUESTIONS

2026

UPDATED EDITION

About This Preparation Material

This document provides key study questions and verified technical explanations for the Microsoft AI-901 examination: *Microsoft Azure AI Fundamentals*.

Review the scenario-based questions, recommended options, and comprehensive explanations on the following pages to evaluate your understanding of responsible AI principles, generative AI prompt engineering, and Azure AI content processing.

Question 1

SINGLE CHOICE

Your organization is developing an AI solution to screen job applications automatically. You want to ensure the model makes fair hiring decisions across all demographic groups. Which of the following actions would best help you address fairness concerns in this AI system?

- A. Train the model on historical hiring data without any modifications to ensure it learns real-world patterns
- B. Audit the training data for demographic bias, test model predictions across different demographic groups, and establish performance metrics that measure fairness outcomes
- C. Remove all personally identifiable information from the model to prevent any discrimination
- D. Use the largest possible dataset available, as larger datasets automatically eliminate bias
- E. Deploy the model to a limited user group first and monitor only the overall accuracy metric

Answer: B**Explanation:**

The correct approach to addressing fairness involves multiple steps: auditing training data for potential bias, testing model performance across demographic groups to identify disparate impacts, and establishing fairness-specific performance metrics—not just overall accuracy. This aligns with responsible AI principles for fairness in AI solutions.

Why other options are incorrect:

- Training on unmodified historical data often perpetuates existing societal biases.
- Removing only PII does not address algorithmic bias in decision-making patterns.
- Large datasets do not automatically eliminate bias; they can amplify it if biased patterns exist.
- Monitoring only overall accuracy misses fairness issues that may disproportionately affect specific groups.

Question 2

SINGLE CHOICE

Your organization is developing an AI solution that analyzes customer support tickets to automatically categorize them by urgency and assign them to the appropriate team. The solution uses historical ticket data for training. You discover that the model performs significantly worse on tickets submitted by customers from certain geographic regions. What principle of responsible AI is most directly addressed by investigating and correcting this issue?

- A. Fairness
- B. Transparency
- C. Accountability
- D. Inclusiveness
- E. Privacy

Answer: A**Explanation:**

The scenario describes a model that performs unequally across different customer groups based on geographic region. This is a classic fairness issue—ensuring that AI systems make equitable decisions across all demographic and geographic groups. Fairness in AI specifically addresses algorithmic bias and unequal performance across different population segments.

Why the other options are incorrect: Transparency refers to explainability and understanding how the model makes decisions; Accountability is about responsibility for outcomes; Inclusiveness addresses designing AI that serves all users including those with disabilities; Privacy concerns data protection and confidentiality.

Question 3

MULTIPLE CHOICE

A healthcare AI system deployed in Azure uses patient medical imaging to detect potential anomalies. The model must operate reliably across different imaging equipment manufacturers, patient populations, and clinical settings. Which considerations are most critical to ensure the solution meets reliability and safety requirements? (Select all that apply.)

- A. Testing the model's performance across diverse imaging equipment, patient demographics, and clinical environments
- B. Encrypting all patient medical data both in transit and at rest
- C. Establishing error thresholds and fallback procedures when the model's confidence is low
- D. Obtaining explicit consent from patients before the system analyzes their images
- E. Providing detailed explanations of how the model identified each anomaly to clinicians

Answer: A, C**Explanation:**

Reliability and safety in AI systems require rigorous testing across diverse real-world conditions (different equipment, demographics, environments) and establishing appropriate confidence thresholds with documented fallback procedures when the model is uncertain. These measures ensure the system performs dependably and safely.

Why the other options are incorrect: Encrypting data is a privacy/security measure, not a reliability consideration; Obtaining consent addresses privacy; Providing explanations to clinicians relates to transparency, not safety. While these are all important responsible AI principles, they don't specifically address reliability and safety.

Question 4

MULTIPLE CHOICE

You are building a generative AI customer service chatbot using Azure AI Foundry. You want the model to respond politely but concisely, prioritize resolving common billing issues, and never suggest refunds that exceed company policy. What are the most effective ways to enforce these behaviors? (Select all that apply.)

- A. Include detailed instructions in the system prompt that define the desired tone, priorities, and policy constraints
- B. Rely exclusively on the model's base training to understand company policies and billing procedures
- C. Use temperature and top-p parameters to control the randomness and diversity of responses
- D. Post-process model outputs to validate that responses comply with policy before sending to customers
- E. Implement a separate legal review step that approves every response before it reaches users

Answer: A, C, D**Explanation:**

Effective prompt engineering for generative AI involves clear system prompts (to define behavior and constraints), careful tuning of model parameters like temperature and top-p (to control randomness), and validation or filtering of outputs to ensure policy compliance. These are practical, scalable approaches used in production systems.

Why the other options are incorrect: Relying solely on base training is insufficient because models don't reliably understand specific company policies without explicit guidance; Requiring legal review of every response is impractical and not a standard implementation pattern for customer service chatbots.

Question 5

SINGLE CHOICE

Your team is implementing an AI solution that extracts structured data from handwritten medical forms using Azure AI Foundry's Content Understanding capabilities. The solution must handle forms in multiple languages, varying handwriting styles, and poor image quality from photocopies. Which approach best addresses these challenges?

- A. Use Content Understanding's multimodal capabilities to process the image and apply post-processing logic to handle language variants and image quality issues
- B. Deploy only optical character recognition (OCR) because handwritten forms require language-specific models
- C. Manually transcribe all forms to ensure perfect accuracy before using the extraction system
- D. Train a custom machine learning model from scratch on your specific form types
- E. Require all submitted forms to be rewritten using typed text instead of handwriting

Answer: A**Explanation:**

Azure AI Foundry's Content Understanding is designed to handle complex, real-world documents including handwritten content, multiple languages, and poor image quality. It uses multimodal AI capabilities that can interpret visual input intelligently, then post-processing logic can address specific edge cases. This is the most practical, scalable solution for this scenario.

Why the other options are incorrect: OCR alone cannot reliably extract from handwritten forms; Manual transcription is not scalable; Training a custom model from scratch requires significant data and effort when a pre-built solution exists; Requiring rewritten forms eliminates the purpose of automating form processing.

End of Practice Questions

You have completed the practice question set for Microsoft AI-901:
Microsoft Azure AI Fundamentals.

Key Exam Domains Covered:

- Responsible AI Principles (Fairness, Reliability & Safety)
- Dataset Bias Auditing & Demographic Performance Testing
- Generative AI Controls (System Prompts, Temperature, Output Filtering)
- Azure AI Foundry Content Understanding & Multimodal Processing

THANK YOU FOR DOWNLOADING AI-901 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our AI-901 Exam Product

Visit Our Site to Purchase the Full Set of Actual AI-901 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

<https://certs4sale.com/microsoft/ai-901-pdf-exam-dumps>