PRACTICE TEST SOFTWARE

Isaca

AAIA Exam

ISACA Advanced in AI Audit (AAIA)

DEMO Version

Full Version Features:

• 90 Days Free Updates

• 30 Days Money Back Guarantee

• Instant Download Once Purchased

• 24 Hours Live Chat Support

Full version is available at link below with affordable price.

https://www.practicetestsoftware.com/isaca/aaia

Question 1. (Single Select)

A healthcare organization uses an AI model to analyze patient data and provide diagnostic recommendations. Which of the following MOST effectively detects data drift related to the model's predictions?

A: Comparing incoming patient data distributions with the training data set
B: Applying overrides to allow healthcare professionals to correct the AI model’s recommendations
C: Conducting periodic model retraining to ensure alignment with updated patient data
D: Using adversarial testing to simulate scenarios that stress test the model’s predictions

Correct Answer: A

Explanation:
Detecting data drift is critical in maintaining the reliability and accuracy of AI models, especially in dynamic environments like healthcare where patient populations and data characteristics can change over time. According to the ISACA Advanced in AI Audit™ (AAIA™) Study Guide, data drift refers to changes in the input data's statistical properties compared to the data on which the model was originally trained. If not detected, data drift can degrade model performance and lead to erroneous predictions.
The most effective approach to detect data drift is to continuously compare the statistical distributions of incoming (production) data with those of the training data set. This allows organizations to identify deviations in data patterns, which can be early indicators that the AI model’s predictions may no longer be valid or optimal.
As stated in the AAIA™ Study Guide under "AI Model Monitoring and Maintenance":
“Monitoring input data for distributional changes compared to the model’s training data is an essential step in identifying data drift. Statistical tests and visualizations can help auditors and AI operators detect when the underlying data characteristics have shifted, prompting further investigation or retraining needs.”
While options such as retraining the model (option C) or adversarial testing (option D) are valuable for ongoing performance and robustness, they do not inherently detect data drift—they respond to or stress-test existing issues. Applying overrides (option B) is a human-in-the-loop safeguard, not a method for drift detection.
ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: "AI Model Monitoring and Maintenance," Subsection: "Detection and Management of Data Drift"

Question 2. (Single Select)

Which of the following is an IS auditor MOST likely to use in order to ensure an AI model has the ability to make correct predictions?

A: Adversarial testing
B: Group analysis
C: Latency testing
D: Confusion matrix

Correct Answer: D

Explanation:
The confusion matrix is a key performance evaluation tool in machine learning and AI auditing. According to the AAIA™ Study Guide, a confusion matrix presents detailed information about actual versus predicted classifications, allowing auditors to assess accuracy, precision, recall, and F1 scores.
“A confusion matrix reveals not just how often predictions are correct, but also the types of errors being made—false positives and false negatives—thereby providing a clear view of the model’s predictive reliability.”
Adversarial testing evaluates robustness, group analysis identifies bias across subgroups, and latency testing examines performance speed—not predictive accuracy. Thus, D is the most relevant for ensuring correct predictions.

Question 3. (Single Select)

An organization shares an AI model with external partners. One partner reports that sensitive data has been inadvertently exposed through the model’s outputs. Which of the following is the IS auditor's BEST recommendation?

A: Limit the model's outputs to anonymized results while investigating further.
B: Audit the data pipelines of all partners to identify the source of the leak.
C: Disable the shared model and notify partners of the potential breach.
D: Retrain the model immediately and implement privacy-preserving techniques.

Correct Answer: C

Explanation:

In the case of a potential data exposure through AI model outputs, the first and most responsible action from an auditing and risk standpoint is to halt further risk propagation. According to the AAIA™ Study Guide, immediate containment is vital, especially when regulatory and reputational risks are high.
“Upon detection of a data breach risk, AI models should be immediately disabled from public or partner use, and all relevant parties should be notified as part of a responsible disclosure and containment strategy.”
While options A and D are longer-term remediation steps and B is investigative, none of them provide the urgent containment that is best practice in such a breach context.

Question 4. (Single Select)

A retail organization uses an AI model to analyze customers' purchase history in order to offer personalized discounts. Which of the following practices represents the MOST ethical use of customer data?

A: Utilizing customer purchase data only after obtaining explicit consent and allowing customers to opt out
B: Retaining and analyzing all available customer data to ensure unbiased recommendations
C: Providing the public with access to review and audit the data set of collected customer information
D: Sharing customer purchase data with third-party vendors to improve advertising and communication

Correct Answer: A

Explanation:
The ethical use of customer data is rooted in respecting privacy, maintaining informed consent, and enabling data subjects to exercise control over their personal information. The AAIA™ Study Guide clearly outlines that obtaining explicit consent and providing opt-out capabilities align with principles of data protection and ethical AI.
“Ethical AI implementation includes transparency in data collection, clear consent mechanisms, and the right of users to opt out or control their personal data usage. Retail and consumer applications must ensure that personalized services do not override these data subject rights.”
Options B and D violate principles of minimal data use and consent, while C may create unnecessary privacy exposure. Therefore, A is the correct and most ethical practice.

Question 5. (Single Select)

Which of the following is MOST important to consider when auditing an organization's AI procedures?

A: Frequency of AI system updates to enhance security

B: Employee training on recognized AI best practices

C: Backup and recovery in the event of an AI data breach

D: AI data validation and filtration to prevent data poisoning

Correct Answer: D

Explanation:

The integrity of data fed into AI systems is a critical concern. The AAIA™ Study Guide emphasizes that validation and filtration processes are essential to mitigate the risk of data poisoning—an attack that can manipulate model behavior by injecting malicious inputs.

“Data poisoning represents a major vulnerability in AI pipelines. Effective controls include robust validation, filtration, and monitoring of training data sources. These preventive practices are essential to ensure model reliability and security.”

While options A, B, and C are important operational and training measures, only D addresses a technical risk that can directly compromise model outputs and trustworthiness.