

Isaca AAIA

ISACA Advanced in AI Audit (AAIA)

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/aaia>

Latest Version: 7.1

Question: 1

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

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

When auditing the transparency of an AI system, which of the following would be the MOST effective way to understand the model's decision-making process?

- A. Evaluating the diversity of the training data set
- B. Analyzing the complexity of the algorithms used
- C. Assessing the computational cost of the model
- D. Reviewing the explainability of AI outputs

Answer: D

Explanation:

Transparency in AI systems is a key requirement to ensure trust, accountability, and ethical compliance. According to the ISACA AAIA™ Study Guide under the "AI Governance and Risk Management" section, understanding the decision-making process of an AI system falls under the principle of explainability. Explainability refers to the degree to which an observer can understand the internal mechanics of an AI system and the rationale behind its outputs.

"Reviewing the explainability of AI outputs allows auditors and stakeholders to determine whether model decisions are interpretable and justifiable. High transparency means stakeholders can trace how and why a decision was made."

While algorithm complexity and computational cost are technical considerations, they do not directly facilitate the audit of decision-making transparency. Similarly, training data diversity is essential for bias reduction but does not explain how decisions are derived. Therefore, option D is the most aligned with auditing transparency.

Question: 3

Which of the following is the BEST way to support the development and design of high-risk AI systems?

- A. Regularly back up the AI system's data to a secure, offsite location.
- B. Conduct regular training sessions for users on data privacy.
- C. Ensure the availability of trustworthy data sets.
- D. Implement multi-factor authentication (MFA) for all users accessing the AI system.

Answer: C

Explanation:

The AAIA™ Study Guide emphasizes that the foundation of any high-performing and ethical AI system lies in the quality and integrity of its data. For high-risk AI systems, such as those used in healthcare, finance, or criminal justice, it is essential to base models on trustworthy data. This ensures reliable predictions, reduces bias, and mitigates risk.

"Trustworthy datasets are characterized by accuracy, completeness, consistency, and ethical sourcing. In high-risk AI applications, ensuring data quality at every stage is crucial to system reliability and compliance."

While backups, user training, and MFA are important for security and operational resilience, they do not address the core challenge of ensuring model accuracy and fairness at the development and design phase. Therefore, option C is the most effective practice.

Question: 4

In order to streamline operations, a bank has deployed an AI application to automatically detect and prevent further fraud on accounts. However, customers have voiced concerns that their usual transactions are being rejected. Which of the following is the MOST likely cause of the false positives?

- A. Consent is not properly managed.
- B. Data versioning controls were not developed.
- C. Compute scale training was not performed.
- D. The hyperparameters are not optimized.

Answer: D

Explanation:

False positives in fraud detection AI systems often stem from poorly optimized hyperparameters. Hyperparameters control aspects of the model's learning process such as the learning rate, decision thresholds, and complexity penalties. When these parameters are not tuned correctly, the model can become overly sensitive and flag normal behavior as suspicious, leading to customer complaints.

"Hyperparameter tuning is essential to balance sensitivity and specificity in AI models. Improper tuning can result in a high rate of false positives or negatives, particularly in systems like fraud detection that require nuanced pattern recognition."

Options A and B relate to data governance but do not directly cause false positives in predictions. Option C (compute scale training) may affect model efficiency, not accuracy. Thus, D is the most appropriate answer.

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: "AI Operations and Performance," Subsection: "Model Tuning and Optimization"

Question: 5

An IS auditor is evaluating an organization's incident management program to ensure it is sufficiently prepared to manage AI-related incidents. Which of the following is MOST important for the auditor to validate?

- A. The program mandates retraining AI systems after incidents are investigated.
- B. The program uses past AI-related incidents and resolutions to categorize current incidents.
- C. The program includes processes to respond to AI model drift and data integrity attacks.

D. The program prioritizes incidents based on alignment with industry leading practices.

Answer: C

Explanation:

AI-related incidents often differ significantly from traditional IT incidents due to their dependence on data, model behavior, and algorithm performance. According to the AAIA™ Study Guide, incident management programs must include capabilities specifically tailored to AI, such as detecting and mitigating model drift and safeguarding against data poisoning or integrity attacks.

“AI incident response frameworks must account for issues unique to machine learning, including model drift, adversarial inputs, and data integrity breaches. An effective program incorporates detection, response, and recovery mechanisms for these AI-specific threats.”

While options A and B contribute to

improving incident response over time, and option D suggests best-practice alignment, only option C directly addresses active response capabilities for high-risk, real-time AI vulnerabilities.

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: “AI Governance and Risk Management,” Subsection: “Incident and Risk Management in AI Contexts”

Thank You for Trying Our Product
Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>