Free Questions for AAIA

Shared by Brennan on 06-01-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

An IS auditor is evaluating an organization's data governance controls for its AI system. Which of the following represents the GREATEST risk in this context?

Options:

A- Inconsistent data management practices

B- Lack of procedures for automated data backup

C- Limited frequency of AI system performance and data accuracy reviews

D- Inadequate controls over data accuracy and privacy compliance

Answer:

D

Explanation:

In AI systems, data accuracy and privacy are foundational to both performance and regulatory compliance. Inadequate controls over data accuracy and privacy compliance (D) pose the greatest risk, as they can lead to incorrect decisions, legal violations, regulatory penalties, and significant reputational damage. AAIA emphasizes that data governance programs must ensure data is accurate, secure, lawfully processed, and appropriately protected throughout the AI life cycle.

Option A (inconsistent practices) is concerning but is often a symptom of underlying governance weaknesses; its impact is most severe when it affects accuracy and privacy. Option B focuses on backup procedures, which are important for availability and resilience, but not as central to AI decision quality and legal risk. Option C (limited performance and accuracy reviews) is serious but again narrower than outright inadequacy of key controls over accuracy and privacy.

ISACA, AAIA Exam Content Outline -- Domain 1: Privacy and Data Governance Programs (accuracy, privacy, and governance controls).

ISACA data governance and privacy guidance for AI systems.

Question 2

Question Type: MultipleChoice

An IS auditor reviews an AI tool using K-means to cluster customers. One cluster shows very high spending but low product diversity. What should the auditor recommend?

Options:

A- Document the algorithm failed because high spending customers did not exhibit high product diversity.

B- Treat the cluster as a potentially valid segment of loyal customers with limited product interest.

C- Increase the number of clusters to better capture variations in spending behavior.

D- Replace K-means clustering with a supervised learning model for more accurate analysis.

Answer:

B

Explanation:

K-means clustering is an unsupervised learning technique that groups data based on similarity. Discovering a cluster with high spending but low product diversity is a plausible and meaningful business insight: it may represent loyal customers who repeatedly purchase a narrow range of products. The auditor should therefore recommend treating this cluster as a potentially valid segment (B), subject to further business analysis and controls where appropriate.

Option A is incorrect because this pattern does not imply algorithm failure. Option C (adding more clusters) might overcomplicate the segmentation without evidence that the current clustering is deficient. Option D misunderstands the purpose of clustering; a supervised model would require labeled outcomes and is not necessarily ''more accurate'' for exploratory segmentation. AAIA's content on AI in audit processes stresses that auditors must interpret AI-driven insights critically, not assume anomalies equal errors.

ISACA, AAIA Exam Content Outline -- Domain 3: AI in Audit Processes (AI tools, use of clustering and analytics in audit).

ISACA analytics guidance on interpreting unsupervised learning outcomes in an audit context.

Question 3

Question Type: MultipleChoice

Which of the following should be an IS auditor's GREATEST concern when reviewing an anomaly

detection process implemented for a high-risk AI system?

Options:

A- Failure to identify anomalies that can bias training data

B- Lack of regular quality reviews for training data

C- Infrequent updates to anomaly detection algorithms

D- Inadequate staff training on the use of the system

Answer:

A

Explanation:

In a high-risk AI system, anomaly detection often serves as a frontline control to flag irregularities in input data and model behavior. The GREATEST concern for an IS auditor is when the process fails to identify anomalies that can bias training data (A), because undetected anomalies can fundamentally distort model learning and outputs. This can lead to systemic bias, incorrect decisions, safety risks, and regulatory breaches.

Option B (lack of regular quality reviews) is serious but is partially addressed if anomaly detection is effective. Option C (infrequent updates) may degrade detection performance over time but is less critical than outright failure to detect harmful anomalies. Option D (staff training) is important for operational effectiveness but still secondary to the technical failure to catch bias-inducing anomalies. AAIA stresses that data integrity and monitoring controls are paramount in high-risk contexts.

ISACA, AAIA Exam Content Outline -- Domain 2: AI Operations (Supervision of AI Solutions, data monitoring, and anomaly detection).

ISACA AI risk materials focusing on high-risk AI oversight and data integrity.

Question 4

Question Type: MultipleChoice

Which of the following is the MOST important reason to establish AI governance structures that extend beyond regulatory compliance?

Options:

A- To align with global AI data privacy standards

B- To mitigate reputational risk associated with public scrutiny of AI systems

C- To ensure ethical integrity throughout the AI life cycle

D- To establish guardrails limiting AI system functionality to approved use cases

Answer:

C

Explanation:

While regulatory compliance is essential, AAIA underlines that ethical integrity must guide AI design, deployment, and monitoring. Regulations often lag behind technological capabilities; thus, relying solely on compliance leaves gaps in areas such as fairness, transparency, human dignity, and societal impact. The MOST important reason to extend governance structures beyond compliance is to ensure ethical integrity throughout the AI life cycle (C) --- from data collection and model design to deployment, monitoring, and retirement.

Option A focuses on privacy only, a subset of broader ethical considerations. Option B is valid but secondary; reputational protection is often a consequence of doing the right thing ethically. Option D (guardrails) is part of governance, but the overarching rationale for those guardrails is to uphold ethical principles. Therefore, comprehensive ethical stewardship is the key driver.

ISACA, AAIA Exam Content Outline -- Domain 5: Ethical and Legal Considerations in AI (ethical AI principles, beyond compliance).

ISACA AI ethics and governance guidance emphasizing proactive, values-driven AI oversight.

Question 5

Question Type: MultipleChoice

Which of the following BEST demonstrates effective coordination to ensure comprehensive oversight of an AI system deployed across multiple jurisdictions?

Options:

A- Focusing oversight activities on technical anomaly detection metrics

B- Establishing joint oversight plans and communication channels between agencies

C- Centralizing responsibility under a single supervisory authority

D- Relying on automated processes for anomaly detection and documentation

Answer:

B

Explanation:

When AI systems span multiple jurisdictions, they are subject to different regulatory regimes, cultural expectations, and risk tolerances. AAIA emphasizes that oversight must be coordinated to avoid gaps or overlaps. Establishing joint oversight plans and communication channels between agencies (B) ensures that relevant authorities share information, align on expectations, and collectively monitor AI risks, enabling coherent and comprehensive oversight.

Option A focuses too narrowly on technical metrics without ensuring cross-jurisdiction coordination. Option C may not be feasible or lawful, as jurisdictional sovereignty often prevents centralizing authority. Option D emphasizes automation but does not address governance and coordination. Thus, the best demonstration of effective oversight for cross-jurisdiction AI deployments is formal joint oversight and structured communication.

ISACA, AAIA Exam Content Outline -- Governance of AI (roles, responsibilities, coordination among stakeholders).

ISACA materials addressing multi-jurisdictional AI risk, regulatory alignment, and oversight structures.

Question 6

Question Type: MultipleChoice

Which of the following is the PRIMARY objective of performing adversarial testing on AI models?

Options:

A- Validating AI incident response plans
B- Determining key risk indicators (KRIs)
C- Fostering security awareness
D- Identifying control gaps

Answer:

D

Explanation:

Adversarial testing involves simulating real-world attacks or malicious inputs against AI models (e.g., adversarial examples, poisoning, evasion) to identify how the system behaves under intentional misuse or hostile conditions. The primary objective is to discover weaknesses and control gaps (D) in the model and its surrounding processes---such as inadequate input validation, insufficient monitoring, or missing safeguards against adversarial inputs.

While results from adversarial testing may inform incident response planning (A), KRI definition (B), or security awareness (C), those are secondary benefits. AAIA's coverage of AI threats and vulnerabilities emphasizes adversarial testing as a control validation and gap-identification mechanism, directly addressing AI-specific risk exposure.

ISACA, AAIA Exam Content Outline -- Domain 1 and Domain 2: Threats and Vulnerabilities Specific to AI; Testing Techniques for AI Solutions.

ISACA guidance on adversarial testing and AI security posture assessment.

Question 7

Question Type: MultipleChoice

An organization is using a large language model (LLM) to assist in evaluating loan applications, but the training data used is known to be incomplete. Which of the following is the GREATEST associated risk?

Options:

A- Unfair loan decisions
B- Delays in loan approval
C- Reduced customer satisfaction
D- Increased manual processing of applications

Answer:

A

Explanation:

Incomplete training data often leads to underrepresentation of certain applicant types, products, or scenarios. In credit and lending, this typically translates into systematic bias: some groups are evaluated on richer historical patterns, while others are evaluated on sparse or unrepresentative information. The greatest associated risk is therefore unfair loan decisions (A), which can manifest as unjustified rejections, inappropriate pricing, or inconsistent risk assessments.

While delays (B), reduced satisfaction (C), or increased manual work (D) may occur, they are secondary operational issues. AAIA highlights that for financial services, the central risks include fairness, discrimination, regulatory compliance, and reputational impact. Incomplete data directly undermines fairness and can violate lending regulations and internal risk appetite.

ISACA, AAIA Exam Content Outline -- Domain 1: AI Governance and Risk (risk categories, including fairness and discriminatory outcomes).

ISACA AI ethics content on data completeness and representativeness in decisioning systems.

Question 8

Question Type: MultipleChoice

An IS auditor is testing an AI model used for determining insurance premiums and eligibility. Which of the following is the MOST effective testing method to identify bias in algorithm outputs?

Options:

A- Regression testing
B- Cross-cluster analysis
C- Disparate impact analysis
D- Predictive analytics

Answer:

C

Explanation:

Disparate impact analysis is a specific technique used to detect whether model decisions disproportionately disadvantage certain protected or sensitive groups (e.g., by gender, age, ethnicity, or other attributes). For an AI model determining premiums and eligibility, fairness and non-discrimination are critical regulatory and ethical requirements, and AAIA content highlights fairness and bias evaluation as core elements of AI governance and risk management.

Regression testing (A) checks that changes do not introduce defects in previously functioning components, not fairness. Cross-cluster analysis (B) may reveal patterns but is not inherently a fairness test. Predictive analytics (D) is a broad term for forward-looking analysis, not a method specifically designed to detect bias. Therefore, disparate impact analysis is the most appropriate and targeted method to identify bias in the insurance AI model's outputs.

ISACA, AAIA Exam Content Outline -- Domain 1: AI Governance and Risk (bias, fairness, and transparency considerations).

ISACA AI ethics and fairness guidance used in AAIA training materials.

Question 9

Question Type: MultipleChoice

An IS auditor reviewed an AI-enabled software for processing a bank's financial information and discovered errors in the training data. Which of the following would BEST mitigate this risk?

Options:

A- Functional testing of the application
B- Data quality testing
C- User interface testing
D- Model validation on benchmark data

Answer:

B

Explanation:

When errors are discovered in training data, the primary risk is that the model has learned from incorrect or low-quality inputs, which can systematically distort outputs (e.g., misstatements of financial information, incorrect classifications, or erroneous flags). Data quality testing directly targets this issue by assessing completeness, accuracy, consistency, validity, and integrity of the data used to train, validate, and test AI models.

Functional testing (A) focuses on whether the application behaves according to requirements, not specifically on the underlying data quality. User interface testing (C) concerns usability and presentation, not training data. Model validation on benchmark data (D) helps assess performance but may not reveal or correct rootcause issues in the underlying training data if those benchmarks are also flawed or unrepresentative. AAIA content emphasizes that data

quality controls are foundational to AI reliability and must be treated as a core operational control.

ISACA, Advanced in AI Audit (AAIA) Exam Content Outline -- Domain 2: Data Management Specific to AI (data quality, data balancing, data security).

ISACA AI guidance on data quality controls as the basis of trustworthy AI outcomes.

Question 10

Question Type: MultipleChoice

Which of the following is the PRIMARY advantage of using K-fold cross validation when evaluating the performance of a machine learning (ML) model?

Options:

A- It facilitates performing regressions on smaller data sets.

B- It helps minimize computational costs when evaluating complex models.

C- It enables the reduction of model bias by setting the K variable to higher values.

D- It uses multiple training and testing cycles to minimize overfitting.

Answer:

D

Explanation:

The primary advantage of K-fold cross validation is that it uses multiple train/test splits, cycling through all folds so that each observation is used both for training and testing at different points. This process provides a more reliable estimate of model performance and reduces the risk of overfitting to a single split (option D). It is an established best practice in model evaluation and aligns with AAIA's emphasis on testing techniques for AI solutions and data analytics.

Option A is not specific to regressions; cross validation can be used for classification and other models as well. Option B can actually increase computational cost since multiple models are trained. Option C misunderstands bias--variance trade-offs; increasing K doesn't simply ''reduce model bias.'' The key advantage remains the use of repeated, varied splits to better assess generalization and guard against overfitting.

ISACA, AAIA Exam Content Outline -- Domain 2: AI Operations (Testing Techniques for AI

Solutions; AI-specific testing).

ISACA data analytics content used in AAIA prep covering cross validation as a standard evaluation method.

To Get Premium Files for AAIA Visit

https://www.p2pexams.com/products/aaia

For More Free Questions Visit

https://www.p2pexams.com/isaca/pdf/aaia