"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

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

Answer: D

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

Question: 3

The PRIMARY purpose of maintaining an audit trail in AI systems is to:

A. Facilitate transparency and traceability of decisions.
B. Analyze model accuracy and fairness.
C. Measure computational efficiency.
D. Ensure compliance with regulatory standards for AI.

Answer: A

Question: 4

A digital bank utilizes an AI system to generate credit scores. Which of the following would BEST mitigate the risk of sudden and unexplained changes in a borrower's credit score?

A. Ensuring the system is periodically reviewed and calibrated by human experts to maintain stability in predictions

B. Using only data from the last six months to one year to avoid outdated information affecting the credit score

C. Allowing the AI to operate fully autonomously to prevent processing delays

D. Obtaining and validating the credit scores from third-party agencies to cross-check AI-generated results

Answer: A

Question: 5

Which of the following AI system characteristics would BEST help an IS auditor evaluate the system's algorithm?

A. The AI system algorithm uses training data to inform decision output.

B. The AI system provides multiple options for model training.

C. The AI system provides transparent justification of decisions.

D. The AI system uses archived transaction data to provide decisions.

Answer: C

Question: 6

An IS auditor is auditing an organization’s data governance framework. The primary objective is to provide assurance that data management practices are standardized to support a trustworthy AI system. Which of the following should be the auditor's MOST important consideration?

A. Retention of stored data

B. Portability of data

C. Data practices for training models

D. Accountability for data management

Answer: D

Question: 7

Which of the following is an IS auditor's MOST important course of action when determining whether source data should be entered into approved generative AI tools to assist with an audit?

A. Validate that the tool is leveraging the latest model.

B. Validate that the tool provides a privacy notice.

C. Determine whether any AI model hallucinations have occurred.

D. Determine whether the information is reliable.

Answer: D

Question: 8

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

Answer: D

Question: 9

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

Question: 10

A healthcare organization uses patient data to train an AI model for early disease detection. Which of the following practices provides the BEST assurance that personal data is secure and its integrity is maintained?

A. Encrypting stored data to reduce exposure and log access
B. Updating the AI model with new data and tracking changes
C. Implementing strict data access controls and conducting security tests
D. Anonymizing patient data and performing regular quality checks

Answer: D

Question: 11

Which of the following should be done FIRST when an attacker exfiltrates sensitive information from an AI model?

A. Implement rate limiting and query restrictions to reduce exploitation attempts.
B. Isolate impacted systems until the attack vector is identified.
C. Rebuild the AI model using a more secure architecture.
D. Inform regulators and affected stakeholders of a potential data breach.

Answer: B

Question: 12

Which of the following is the MOST important purpose of conducting a risk assessment for AI models within an organization?

A. Categorizing data used by the AI model
B. Defining mitigation strategies for AI deployment
C. Monitoring AI model performance on an ongoing basis
D. Determining whether AI model outputs align with established use cases

Answer: B

Question: 13

An organization is adopting AI for its procurement and inventory teams, raising concern from stakeholders that they will lose their jobs due to AI. Which of the following is the BEST way for the IS auditor to assess whether the potential negative impacts were minimized?

A. Review human-centered design practices to determine how they were considered.
B. Review the AI roadmap for short-term and long-term milestones.
C. Review how the project management team collected feedback in engagement activities.
D. Review the current state assessment of how AI may impact the organization.

Answer: A

Question: 14

An IS auditor is looking to expedite reporting for an audit with complex issues. Which of the following would be the MOST effective way for the auditor to use generative AI?

A. Developing action items discussed in closing meetings for management action plans
B. Developing a draft of an executive summary based on detailed findings and audit scope
C. Revising audit conclusions with precise verbiage to describe the audit observations
D. Revising audit background and scope information based on new information from management

Answer: B

Question: 15

Which of the following is the PRIMARY purpose of an AI acceptable use policy?

A. Establishing guidance on the ethical use of AI
B. Outlining AI usage monitoring procedures
C. Educating employees on where to find and how to use AI tools
D. Explaining the distinction between different types of AI

Answer: A

Question: 16

While evaluating a complex machine learning (ML) model used for regulatory compliance in a financial institution, which of the following should the IS auditor do to BEST ensure transparency?

A. Document sources and data processes.
B. Create dashboards to show outputs.
C. Provide periodic model audit reports.
D. Use tools that explain model decisions.

Answer: D

Question: 17

Which of the following is the GREATEST challenge facing IS auditors evaluating the explainability of generative AI models?

A. Differences of opinion regarding model types
B. Difficulties in preventing the input of biased data
C. Performance issues due to excessive computation
D. Algorithms changing as AI continues to learn

Answer: D

Question: 18

For a sales promotion, an AI system sorts customer attributes into several categories by analyzing transaction history. Verifying which of the following would BEST validate the effectiveness of this process?

A. Stress tests are regularly conducted to maintain consistent AI performance.
B. The applied methodology adequately reflects business objectives.
C. Sensitive attributes are converted to other data types prior to input.
D. Sampling of AI output is conducted to identify unusual decisions.

Answer: B

Question: 19

A bank uses a video-based know your customer (KYC) verification process. Cybercriminals exploit this process by using deepfake technology to impersonate bank customers. Which of the following countermeasures is the BEST way for the bank to mitigate this risk?

A. Requesting additional identity and address documents for verification
B. Leveraging AI-based liveness detection during video verification
C. Encrypting all customer data and communication
D. Discontinuing the use of the video-based verification process

Answer: B

Question: 20

Which of the following key performance indicators (KPIs) are MOST important when evaluating whether an AI model meets business objectives?

A. Cost of resources required for AI model training
B. Number of users interacting with the AI model
C. Frequency of AI model retraining
D. AI model accuracy in predicting actual outcomes

Answer: D

Question: 21

When auditing a research agency's use of generative AI models for analyzing scientific data, which of the following is MOST critical to evaluate in order to prevent hallucinatory results and ensure the accuracy of outputs?

A. The effectiveness of data anonymization processes that help preserve data quality
B. The algorithms for generative AI models designed to detect and correct data bias before processing
C. The frequency of data audits verifying the integrity and accuracy of inputs
D. The measures in place to ensure the appropriateness and relevance of input data for generative AI models

Answer: D

Question: 22

When converting data categories before training an AI model, which of the following scenarios represents the GREATEST risk?

A. One-hot encoding the data attribute car colors for the options red, blue, green, black, white
B. Creating dummy variables for the data attribute dog breed for the options labrador, terrier, beagle
C. One-hot encoding the data attribute customer rewards category for the options economy, business, first class
D. Creating dummy variables for the data attribute product flavor for the options vanilla, chocolate, strawberry, banana

Answer: C

Question: 23

A retail organization uses an AI model to forecast inventory based on customer purchasing trends and updates the model quarterly. The model recently failed to recognize a surge in demand during a popular shopping season. Which of the following issues does this situation BEST demonstrate?

A. Limited data set diversity impacting model training
B. Data drift impacting system forecasting
C. Overfitting issues due to a small training data set
D. Lack of outlier checks in data affecting forecast accuracy

Answer: B

Question: 24

Which of the following metrics are the BEST indication of a mature and effective approach to an organization's data governance program for its AI systems?

A. Number of AI projects completed within the last fiscal year
B. Percentage of AI models with documented data lineage
C. Frequency of data quality audits on the organization's data sets
D. Total budget allocated to AI initiatives across all departments

Answer: B

Question: 25

Which of the following BEST detects model drift or unexpected changes in AI model outputs?

A. Standardization of AI configurations
B. Anomaly monitoring
C. AI model documentation reviews
D. AI model retraining

Answer: B

Question: 26

To confirm the fairness of AI model decisions, the BEST way to collect reliable evidence during an AI audit is by:

A. Analyzing system metadata.
B. Testing the model with a curated sample data set.
C. Interviewing developers.
D. Observing the system's interactions with end users.

Answer: B

Question: 27

An organization's system development process has been enhanced with AI. Which of the following features presents the GREATEST risk?

A. The AI allocates resources for new system development projects.
B. Non-technical users are validating AI results.
C. The AI personalizes applications for the user.
D. All codes are generated by AI without human oversight.

Answer: D

Question: 28

Which of the following is the GREATEST risk associated with using AI in audit planning?

A. Increased planning costs
B. Scope creep
C. Incomplete data
D. Limited knowledge

Answer: C

Question: 29

An IS auditor notes the combined number of records utilized within the training, validation, and testing data sets exceeds the total number of records in the original data set. Which of the following is MOST important for the auditor to determine?

A. Whether the training, validation, and testing data sets were created in the correct order
B. Whether data leakage occurred from utilizing overlapping records in the data sets
C. Whether a sufficient number of records were utilized in the training data set
D. Whether the validation data set utilized the same number of records as the training data sets

Answer: B

Question: 30

During an audit of an investment organization's AI-powered software, an IS auditor identifies a potential security risk. What is the GREATEST risk associated with staff exfiltrating organizational data to a generative AI tool?

A. Data contamination due to biased AI model outputs
B. Unauthorized data disclosure
C. Potential business disruptions
D. Excessive reliance on AI-generated insights

Answer: B

Question: 31

Which of the following is the MOST important reason to perform regular ethical reviews of AI systems?

A. To improve the accuracy and performance of the systems
B. To align AI system development with organizational values and principles
C. To ensure the systems align with the preservation of individual rights
D. To identify and mitigate potential data drift within models

Answer: C

Question: 32

Which of the following is the MOST important task when gathering data during the AI system development process?

A. Stratifying the data B. Isolating the system C. Cleaning the data
D. Training the system

Answer: C

Question: 33

Which of the following controls would MOST effectively mitigate worst-case service disruption scenarios affecting an AI-based application system?

A. Performing periodic tabletop exercises
B. Implementing a kill chain process in the event of disruption
C. Updating key risk indicators (KRIs) regularly
D. Including a range of AI disruption scenarios in the disaster recovery plan (DRP)

Answer: D

Question: 34

An IS auditor reviewing documentation for an AI model notes that the modeler utilized a K-means clustering algorithm, which clusters data into categories for correlations and analysis. Which of the following is the MOST important risk for the auditor to consider?

A. K-means clustering is not a common data clustering method due to its complexity and difficulty categorizing data correctly.
B. K-means clustering requires the modeler to supervise the learning analysis, which can introduce bias.
C. K-means clustering algorithms are significantly sensitive to outliers and dependent on the similarity of units of measure.
D. K-means clustering determines the number of clusters for the modeler without supervision.

Answer: C

Question: 35

An organization uses an AI-powered tool to detect and respond to cybersecurity threats in real time.
An IS auditor finds that the tool produces excessive false positives, increasing the workload of the security team. Which of the following techniques should the auditor recommend to BEST evaluate

the tool's effectiveness in managing this issue?

A. Use a log analysis tool to examine the types and frequency of alerts generated.
B. Implement a benchmarking tool to compare the system's alerting capability with industry standards.
C. Conduct penetration testing to assess the system's ability to detect genuine threats.
D. Deploy a machine learning (ML) validation tool to increase the model's accuracy and performance.

Answer: D

Question: 36

Which of the following is the MOST important consideration when auditing the data used for training an AI model?

A. Timeliness
B. Predictability
C. Representativeness
D. Understandability

Answer: C

Question: 37

An IS auditor is testing an AI-based fraud detection system that flags suspicious transactions and finds that the system has a high false positive rate. Which of the following testing methods should be prioritized to BEST optimize the detection rate?

A. Regression testing
B. Cross-validation testing
C. Substantive testing
D. Benford's Law analysis

Answer: B

Question: 38

The BEST way to prevent sensitive information disclosure by large language model (LLM) chatbots is through:

A. Manual monitoring
B. Access controls
C. Data sanitization
D. Data masking

Answer: D

Question: 39

A car manufacturer uses an AI model to predict maintenance needs for its vehicles. Which of the following techniques can an IS auditor apply to MOST effectively verify the AI model's decisions to stakeholders?

A. Using neural network visualization to show how the AI model processes data through its layers B. Using K-means algorithms to group vehicles based on mileage or engine temperature for maintenance patterns
C. Utilizing support vector machines (SVM) to classify vehicles based on maintenance urgency
D. Using local interpretable model-agnostic explanation (LIME) to analyze how specific features contribute to predictions

Answer: D

Question: 40

Which of the following do supervised AI learning models PRIMARILY use to train algorithms?

A. Unlabeled data sets
B. Clustered data sets
C. Labeled data sets
D. Randomized data sets

Answer: C

Question: 41

From a data appropriateness and bias perspective, which of the following should be of GREATEST concern when reviewing an AI model used in a credit scoring system?

A. The model incorporates the applicant's loan history to assess spending habits.
B. The model utilizes historical credit data to predict future credit behavior.
C. The model considers the applicant's income level as a key factor in the credit decision.
D. The model uses postal codes as a primary factor in determining creditworthiness.

Answer: D

Question: 42

The PRIMARY objective of auditing AI systems is to:

A. Identify biases and decision transparency.
B. Maximize system efficiency and throughput.
C. Optimize user experience and interface satisfaction.
D. Minimize algorithm latency and information storage impacts.

Answer: A

Question: 43

When auditing a machine learning (ML) solution, false positives can BEST be assessed by examining the level of:

A. Precision
B. Completeness
C. Accuracy
D. Recall

Answer: A

Question: 44

An IS auditor is performing an inventory audit for a manufacturing organization. Which of the following would BEST enable the auditor to identify types of products without assistance from organizational staff?

A. Natural language processing
B. Speech modeling
C. Robotic process automation (RPA)
D. Computer vision

Answer: D

Question: 45

An IS auditor notes that an AI model achieved significantly better results on training data than on test dat a. Which of the following problems with the model has the IS auditor identified?

A. Underfitting
B. Overfitting
C. Generalization
D. Bias

Answer: B

Question: 46

An AI social media platform uses an algorithm to increase user engagement that could unintentionally promote divisive content. Which of the following is the BEST course of action to mitigate this risk?

A. Introduce controls allowing individuals to customize content preferences.
B. Suspend the algorithm until concerns are addressed.
C. Obtain users' consent for the content they wish to view.
D. Regularly audit and adjust algorithms to reduce biases.

Answer: D

Question: 47

Which of the following is the MOST important risk for an IS auditor to consider when reviewing the adoption

of an AI system?

A. Costs associated with AI system maintenance
B. Immaturity of AI systems in the industry
C. Bias in AI system decision making
D. Resistance to the use of AI technology

Answer: C

Question: 48

An organization uses an AI image generation platform to create promotional materials. An IS auditor identifies that the platform includes copyrighted images in its training dat
a. Which of the following is the auditor's BEST recommendation to address this issue?

A. Implement a manual review process to ensure no copyrighted images are used in generated outputs.
B. Use a platform that certifies the provenance and licensing of its training data.
C. Label all AI-generated images to disclaim the possibility of third-party content.
D. Suspend the use of the platform until the training data is sanitized.

Answer: B

Question: 49

An IS auditor uses an internally developed generative AI tool to prepare a status update for audit stakeholders. Which of the following is the auditor’s MOST appropriate course of action?

A. Compare results with a publicly available generative AI tool to ensure outputs are similar.
B. Assess whether the information provided is complete and accurate.
C. Regenerate the results to ensure similar outputs are provided.
D. Share and review the results with management.

Answer: B

Question: 50

Which of the following is the PRIMARY benefit of implementing a robust data governance framework specific to AI solutions in an organization?

A. It focuses on enhancing the accuracy and reliability of AI model predictions.
B. It accelerates AI implementation timelines by fully automating data preparation processes.
C. It fosters adherence to industry regulations while minimizing the risk of data breaches and privacy

violations.

D. It reduces the need for human oversight, ensuring seamless and autonomous data governance.

Answer: C

Question: 51

Which of the following BEST ensures that an AI system complies with user data ownership rights under privacy regulations?

A. Applying data clustering techniques to anonymize data sets

B. Enforcing strict data retention policies to limit storage duration

C. Implementing a transparent data consent management process

D. Regularly conducting AI system performance testing for accuracy

Answer: C

Question: 52

When using off-the-shelf AI models, which of the following is the MOST appropriate way for organizations to approach vendor management?

A. Ensure a minimum of three quotes have been obtained for market research and comparison.

B. Establish responsibility and clear terms for model updates and support.

C. Only use models from vendors with globally recognized accreditation.

D. Use the vendor only if the contract has been reviewed by the information security department.

Answer: B

Question: 53

Which of the following is the PRIMARY reason IS auditors must be aware that generative AI may return different investment recommendations from the same set of data?

A. Limitations can arise in the quantification of risk profiles.

B. Neural node access varies each time the process is executed.

C. Computational logic is based on probabilities.

D. Servers are reconfigured periodically.

Answer: C

Question: 54

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.

Answer: C

Question: 55

The PRIMARY objective of machine learning (ML) in data processing is to:

A. Analyze data sets to identify visual patterns and trends.
B. Enhance the explainability of AI model outputs.
C. Perform actions that would typically require human intelligence.
D. Draw statistical inferences for creating artificial human intelligence.

Answer: C

Question: 56

In the context of an AI implementation, which of the following actions is MOST critical for an organization's change management program?

A. Ensuring the organization has a dedicated AI governance committee
B. Reviewing documentation for AI system changes, updates, and patches
C. Conducting a comprehensive risk assessment specific to AI-related changes
D. Verifying that all employees have completed mandatory AI ethics training

Answer: C

Question: 57

Which of the following testing techniques would BEST validate whether an organization's data governance program effectively ensures data quality and integrity for AI model training and deployment?

A. Performing a business impact analysis (BIA) to assess the consequences of AI model failure B. Reviewing the organization's AI software development life cycle documentation

C. Conducting a penetration test to identify vulnerabilities in the model

D. Assessing data lineage to verify the traceability of data sources

Answer: D

Question: 58

Which of the following controls MOST effectively helps to ensure an AI model is resilient against external threats?

A. AI data set anonymization

B. Monitoring of AI model developers

C. Monitoring of AI access logs

D. AI model configuration testing

Answer: D

Question: 59

A generative AI system has a validation control in place to reject inappropriate questions by checking them against built-in ethical standards. Which of the following enables malicious actors to circumvent this control through prompt engineering?

A. Submitting the same questions in a foreign language translated by another AI-based system

B. Presenting theoretical situations to justify the reason for asking the questions

C. Asking the same questions later when the algorithm has changed after further learning

D. Randomly placing keywords unrelated to the main topic

Answer: B

Question: 60

Which of the following is the MOST effective way an IS auditor could use generative AI to plan an audit of a new database storing transactional data?

A. Identifying separation of duties conflicts for database data changes
B. Developing architecture diagrams
C. Identifying technology-specific risk and considerations
D. Summarizing meeting transcripts from interviews with database administrators (DBAs)

Answer: C

Question: 61

Which of the following is MOST important for an IS auditor to review during an AI system audit in order to determine compliance with intellectual property and data rights?

A. Data performance metrics
B. Data usage agreements
C. Use of open-source intellectual property
D. Model runtime efficiency logs

Answer: B

Question: 62

Which use case for an AI model to be used by a food delivery service would pose ethical risk to the organization?

A. Correlating time, cost, delivery distance, and customer satisfaction metrics to issue coupons to customers receiving substandard service
B. Basing driver retention and termination decisions on the number of delivered orders per total hours worked as compared to an industry benchmark
C. Comparing total food preparation and delivery time to an industry benchmark to set key performance and risk indicators for individual restaurants
D. Using customer service metrics for service speed and food quality to predict customer retention and forecast revenue

Answer: B

Question: 63

An IS auditor is interviewing management about implemented controls around machine learning (ML) models deployed in the production environment. Which of the following schedules for reviewing the performance of a deployed model would be of GREATEST concern to the auditor?

A. After changes to hardware and software platforms
B. After functionality changes
C. One time prior to migrating to production
D. On an annual recurring basis

Answer: C

Question: 64

Which of the following correctly summarizes the conclusions of the model card excerpt provided?

Model Card – Electrical Grid Predictive Maintenance Model
Model Information:
Description: AI model designed to predict maintenance needs for electrical grid components, reduce unplanned downtime, and improve grid reliability.
Inputs: Real-time sensor data, historical maintenance records, and operational logs.
Outputs: Maintenance needs predictions for 60 & 90 days.Evaluation:
Approach: Cross-validation and validation of accuracy, precision, and recall.
Results: Accuracy 72%; Precision 60%; Recall 95%; F1 76%

A. The AI model correctly predicts maintenance needs 95% of the time.
B. The electrical grid uptime is expected to be 72% of the time.
C. Grid failure is predicted to occur after 90 days.
D. F1 indicates that the model identifies true maintenance needs 76% of the time.

Answer: D

Question: 65

Which of the following is MOST important to have in place when initially populating data into a data frame for an AI model?

A. The box charts, histograms, scatterplots, and Venn diagrams that identify correlations and outliers
B. The code for separating data into training and testing data sets
C. An analysis of exploratory data that checks for incorrect data types, null values, and duplicate entries

D. An approved risk assessment for including, excluding, or subsequently dropping data attributes from the model

Answer: C

Question: 66

Which of the following controls helps mitigate the risk of competitors poisoning data utilized by a machine learning (ML) model performing sentiment analysis of product reviews?

A. Peer reviewing code that acquires product reviews from social media posts

B. Hiring a marketing firm to text links to customers requesting product reviews for monetary compensation

C. Augmenting the unbalanced product review data set with the use of oversampling by the model developer

D. Requiring customers to authenticate access to their accounts prior to writing product reviews

Answer: D

Question: 67

An IS auditor is auditing an AI system that predicts inventory needs. The system recently failed to predict a stock outage for a key product. Which of the following audit tests would BEST validate the system's accuracy?

A. Unit testing of the forecasting algorithm

B. Load testing during peak sales periods

C. Sensitivity analysis on input variables

D. Historical testing with past sales data

Answer: D

Question: 68

An IS auditor is auditing a financial system in which a generative AI tool is used to identify trends in batches of 4,000 rows, while the generative AI tool has a limit of 3,000 tokens. Which of the following is the GREATEST concern?

A. The AI will process only a portion of the data set.

B. The AI will prioritize high-value entries.
C. The AI will reject the data set and not analyze the data.
D. The AI output will be biased toward the first 3,000 tokens.

Answer: D

Question: 69

Which of the following is MOST important to review in order to gain assurance that an AI model is performing without biases?

A. AI training data
B. AI development environment
C. AI model adaptability
D. AI model temperature

Answer: A

Question: 70

Which of the following is MOST important for an IS auditor to consider when identifying AI risk in a know your customer (KYC) application within a banking organization?

A. Intellectual property leakage and invalidation
B. Benchmarking against peer organizations
C. Incident response plan
D. Business disruption and financial impact

Answer: D

Question: 71

Which of the following is the PRIMARY objective of AI governance?

A. Implementing compliance and ethics controls for AI initiatives
B. Defining clear roles and responsibilities for AI development, use, and oversight
C. Ensuring controls over AI are designed well and operate effectively
D. Promoting a positive return on investment (ROI) from AI projects

Answer: B

Question: 72

Which of the following strategies used by modelers to enhance data accuracy has the GREATEST risk of bias and information loss?

A. Filling blank attributes in records with the mean, median, or mode within a grouping
B. Identifying and deleting duplicate entries in the data set
C. Separating multiple data attributes within one field into individual attribute columns
D. Placing numerical data into bins or buckets for a manageable quantity of correlations and result analyses

Answer: A

Question: 73

During a pre-implementation risk assessment, an AI model is determined to present a significant risk of bias and potential harm in excess of the organization's risk tolerance. Which of the following is the MOST appropriate response?

A. Postpone deployment until the risk can be safely managed.
B. Enhance the data that the model is trained on.
C. Obtain board approval for an exception.
D. Revisit the risk tolerance to ensure it is appropriate.

Answer: A

Question: 74

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

Question: 75

Which of the following will provide the BEST evidence to support the alignment of an AI model with an organization's business objectives?

A. AI model vulnerability assessment
B. AI change management requests
C. AI model inventory
D. AI acceptable use policy

Answer: C

Question: 76

The PRIMARY purpose of utilizing neural networks in AI is to:

A. Improve the user interface.
B. Increase computational power.
C. Mimic human decision making.
D. Minimize maintenance costs.

Answer: C

Question: 77

When an IS auditor is reviewing results from an AI system, which of the following would cause the GREATEST risk?

A. Inability to identify where an AI system is housed
B. System output not being checked for inconsistencies
C. Cascading failures of AI system outputs
D. Difficulty of documenting AI algorithm processes

Answer: B

Question: 78

When auditing an AI system, which of the following steps ensures that AI model behavior is aligned with

organizational objectives?

A. Algorithm debugging
B. Data transformation
C. Model training
D. Problem framing

Answer: D

Question: 79

An organization is using information gathered from customer accounts to train its AI chatbot. Which of the following is the GREATEST risk associated with this practice?

A. Disclosure of personal information
B. AI bias
C. Transparency
D. AI model hallucinations

Answer: A

Question: 80

When reviewing contracts or other lengthy documentation in the planning phase, which of the following tools would BEST extract relevant information?

A. Robotic process automation (RPA)
B. Autoregressive sequencing model
C. Predictive analytics
D. Natural language processing

Answer: D

Question: 81

The GREATEST benefit of using AI auditing techniques over traditional methods is that AI auditing techniques can:

A. eliminate the need for human intervention.

B. ensure full compliance with regulations.
C. identify complex data patterns.
D. significantly reduce data bias.

Answer: C

Question: 82

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

Question: 83

An organization deploys an AI recruitment platform to screen job applicants. The IS auditor identifies that the platform's decisions may be influenced by model bias. Which of the following risk mitigation strategies is BEST for the auditor to recommend?

A. Implement a process to periodically test the AI system for biases and adjust parameters as needed.
B. Suspend the use of the AI system until the training data can be verified for fairness and compliance.
C. Retrain the AI model using an external data set certified for inclusivity and fairness.
D. Require manual reviews of all AI-generated recruitment decisions before hiring is finalized.

Answer: A

Question: 84

Which of the following is the MOST important course of action for an organization prior to allowing end users to utilize an AI tool?

A. Develop an AI policy with guidelines on appropriate use.
B. Determine the impact to the disaster recovery plan (DRP).
C. Implement baseline performance metrics.

D. Ensure a cybersecurity insurance clause is in place to include the use of AI.

Answer: A

Question: 85

Which of the following is MOST important to consider when deciding whether to implement an AI solution?

A. The cost of AI implementation
B. The speed of AI implementation
C. The space required for AI hardware
D. The ethical implications of AI

Answer: D

Question: 86

A healthcare organization uses data clustering to group patients by medical history for personalized treatment recommendations. Which of the following is the GREATEST privacy risk associated with this practice?

A. The clustering requires more data, increasing the risk of a privacy breach.
B. Clustering increases the complexity of the model, making data harder to anonymize.
C. Irrelevant features in the data may result in inaccurate or biased treatments.
D. Clusters can reveal sensitive personal information depending on how the information is presented.

Answer: D

Question: 87

An organization is evaluating change management practices for AI-based decision support models. Which of the following BEST demonstrates effective AI-focused change management?

A. Engaging an independent expert to review the model's accuracy and precision on a quarterly basis
B. Assigning a single data science team member to adjust the model in order to establish accountability
C. Documenting model updates and retraining sessions to ensure traceability
D. Deploying two separate copies of the model after each adjustment to compare results

Answer: C

Question: 88

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

Answer: A

Question: 89

When utilizing a machine learning (ML) model to predict whether a wind turbine electricity generator will fail, which model evaluation metric should be the PRIMARY focus?

A. Precision

B. Specificity

C. Accuracy

D. Recall

Answer: D

Question: 90

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