

IAPP.AIGP.v2026-10-12.q96

Exam Code:	AIGP
Exam Name:	IAPP Certified Artificial Intelligence Governance Professional
Certification Provider:	IAPP
Free Question Number:	96
Version:	v2026-10-12
# of views:	102
# of Questions views:	960
https://www.freecram.net/torrent/IAPP.AIGP.v2026-10-12.q96.html	

NEW QUESTION: 1

Which of the following is an obligation of an importer of high-risk AI systems under the EU AI Act?

- A. Provide technical documentation.
- B. Affix the CE marking.
- C. Verify the Declaration of Conformity.
- D. Conduct a data protection impact assessment.

Answer: (SHOW ANSWER)

Importers of high-risk AI systems into the EU have specific responsibilities under the EU AI Act. They are not the parties responsible for affixing the CE marking or providing technical documentation-but they must verify that these have been done by the provider.

From the AI Governance in Practice Report 2025:

"Importers must verify that the appropriate conformity assessment has been carried out, the technical documentation is available, and the CE marking has been affixed." (p. 34-35) Thus:

- * A. Provide technical documentation- done by the provider.
- * B. Affix the CE marking-provider's responsibility.
- * C. Verify the Declaration of Conformity-importer obligation.
- * D. Conduct a DPIA- relevant under data protection laws, not required under the EU AI Act for importers.

NEW QUESTION: 2

Scenario:

An organization is evaluating different AI models for integration into its internal workflows. Before moving forward with a particular AI solution from a third-party vendor, the governance team needs to assess the ethical and operational implications of the model.

The most important policy to assess the operations of an AI model is to follow the:

- A. Acceptable use policy of the model provider
- B. Privacy policy of the model provider

- C. Security policy of the model provider
- D. Code of conduct policy of the model provider

Answer: ([SHOW ANSWER](#))

The correct answer is A. The Acceptable Use Policy (AUP) sets the primary terms and conditions under which the AI model can be used, including limitations, prohibited uses, and operational constraints.

From the AIGP ILT Guide:

"The AUP is the most critical document to assess what you can and cannot do with the AI system. It governs use cases, outlines forbidden uses (e.g., disinformation), and supports responsible AI practices." Also confirmed in the AI Governance in Practice Report 2025 (Third-Party AI Assurance section):

"Reviewing and adhering to the acceptable use policy ensures that the model is being used in a manner that aligns with both provider expectations and ethical AI practices."

NEW QUESTION: 3

All of the following types of testing can help evaluate the performance of a responsible AI system EXCEPT?

- A. Risk probability/severity.
- B. Adversarial robustness.
- C. Statistical sampling.
- D. Decision analysis.

Answer: ([SHOW ANSWER](#))

Risk probability/severity testing is not typically used to evaluate the performance of an AI system. While important for risk management, it does not directly assess an AI system's operational performance.

Adversarial robustness, statistical sampling, and decision analysis are all methods that can help evaluate the performance of a responsible AI system by testing its resilience, accuracy, and decision-making processes under various conditions. Reference: AIGP Body of Knowledge on AI Performance Evaluation and Testing.

NEW QUESTION: 4

The benefit of having a clear process for handling AI-related incidents is that it reduces?

- A. Response times.
- B. Model failures.
- C. Compliance issues.
- D. Documentation requirements.

Answer: ([SHOW ANSWER](#))

The correct answer is A because having a well-defined incident management process enables organizations to respond quickly and effectively when AI-related issues arise. AI governance frameworks emphasize incident management plans as a key component of operational governance, ensuring that risks such as system failures, harmful outputs, or security breaches

are promptly identified, escalated, and resolved. A structured process reduces delays by clearly defining roles, responsibilities, and response procedures, thereby minimizing potential harm and operational disruption. While incident processes may also indirectly support compliance and reduce the impact of failures, their primary benefit is improving responsiveness and coordination. Efficient response times are critical in maintaining trust, ensuring safety, and limiting negative consequences in real-world AI deployments.

NEW QUESTION: 5

What is the best reason for a company to adopt a policy that prohibits the use of generative AI?

- A.** To avoid using technology that cannot be monetized.
- B.** To avoid needing to identify and hire qualified resources.
- C.** To avoid the time necessary to train employees on acceptable use.
- D.** To avoid accidental disclosure of its confidential and proprietary information.

Answer: ([SHOW ANSWER](#))

The correct answer is D because one of the most significant risks associated with generative AI is the potential leakage of sensitive, confidential, or proprietary information. AI governance frameworks emphasize data protection, confidentiality, and secure data handling as critical organizational responsibilities. When employees input sensitive data into generative AI tools, especially external or third-party systems, that information may be stored, reused, or exposed unintentionally. This creates legal, regulatory, and reputational risks. As a result, organizations may restrict or prohibit generative AI use to prevent unauthorized data disclosure. The other options relate to operational or business considerations, which are not primary governance drivers. Protecting sensitive information aligns with core AI governance principles such as privacy, security, and risk mitigation, making it the strongest justification for such a policy.

NEW QUESTION: 6

If it is possible to provide a rationale for a specific output of an AI system, that system can best be described as?

- A.** Accountable.
- B.** Transparent.
- C.** Explainable.
- D.** Reliable.

Answer: ([SHOW ANSWER](#))

If it is possible to provide a rationale for a specific output of an AI system, that system can best be described as explainable. Explainability in AI refers to the ability to interpret and understand the decision-making process of the AI system. This involves being able to articulate the factors and logic that led to a particular output or decision. Explainability is critical for building trust, enabling users to understand and validate the AI system's actions, and ensuring compliance with ethical and regulatory standards. It also facilitates debugging and improving the system by providing insights into its behavior.

NEW QUESTION: 7

When should an ethical impact assessment for AI be performed?

- A. When the system is ready for production and deployment.
- B. When the number of AI system users has reached a predefined threshold.
- C. After the AI system has encountered ethical issues in real-world applications.
- D. At the initial stages of AI system development and continually throughout its lifecycle.

Answer: ([SHOW ANSWER](#))

Ethical impact assessment should begin early and remain iterative throughout the AI lifecycle, making D correct. Waiting until production, a user threshold, or an actual ethical incident would make the assessment reactive rather than preventive. UNESCO describes ethical impact assessment as an "ethics by design" mechanism covering research, design, development, deployment, use, maintenance, monitoring, and eventual termination, while emphasizing its particular value during early research and design. OECD-listed assessment guidance similarly states that the first assessment should occur during design, with reassessments at development, deployment, and other significant lifecycle stages. NIST's AI RMF Playbook also recommends impact assessment at key stages because risks, operating conditions, stakeholders, data, and system capabilities can change over time. Continuous reassessment therefore supports proactive identification and mitigation of emerging harms.

NEW QUESTION: 8

What is the best method to apply to an AI model so that it is statistically unlikely that a specific piece of training data can be identified from a model's output?

- A. Homomorphic encryption.
- B. Data sharding.
- C. Differential privacy.
- D. Data compartmentalization.

Answer: ([SHOW ANSWER](#))

The correct answer is C because differential privacy is specifically designed to ensure that the inclusion or exclusion of any single data point does not significantly affect the output of a model. This provides strong mathematical guarantees that individual training data cannot be reverse-engineered or identified from model outputs. Differential privacy works by introducing controlled noise into the data or model outputs, thereby protecting individual-level information while preserving overall statistical patterns. In AI governance, it is widely recognized as a key privacy-enhancing technology for mitigating risks such as data leakage, membership inference attacks, and re-identification. Other options, such as homomorphic encryption, focus on secure computation, while data sharding and compartmentalization address data management but do not provide formal privacy guarantees at the model output level.

NEW QUESTION: 9

Which of the following is NOT a common type of machine learning?

- A. Deep learning.

- B. Cognitive learning.
- C. Unsupervised learning.
- D. Reinforcement learning.

Answer: ([SHOW ANSWER](#))

The common types of machine learning include supervised learning, unsupervised learning, reinforcement learning, and deep learning. Cognitive learning is not a type of machine learning; rather, it is a term often associated with the broader field of cognitive science and psychology. Reference: AIGP BODY OF KNOWLEDGE and standard AI/ML literature.

NEW QUESTION: 10

According to the GDPR's transparency principle, when an AI system processes personal data in automated decision-making, controllers are required to provide data subjects specific information on?

- A. The existence of automated decision-making and meaningful information on its logic and consequences.
- B. The personal data used during processing, including inferences drawn by the AI system about the data.
- C. The data protection impact assessments carried out on the AI system and legal bases for processing.
- D. The contact details of the data protection officer and the data protection national authority.

Answer: A ([LEAVE A REPLY](#))

The GDPR's transparency principle requires that when personal data is processed for automated decision-making, including profiling, data subjects must be informed about the existence of such automated decision-making. Additionally, they must be provided with meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for them. This requirement ensures that data subjects are fully aware of how their personal data is being used and the potential impacts, thereby promoting transparency and trust in the processing activities.

NEW QUESTION: 11

To maintain fairness in a deployed system, it is most important to?

- A. Protect against loss of personal data in the model.
- B. Monitor for data drift that may affect performance and accuracy.
- C. Detect anomalies outside established metrics that require new training data.
- D. Optimize computational resources and data to ensure efficiency and scalability.

Answer: ([SHOW ANSWER](#))

To maintain fairness in a deployed system, it is crucial to monitor for data drift that may affect performance and accuracy. Data drift occurs when the statistical properties of the input data change over time, which can lead to a decline in model performance. Continuous monitoring and updating of the model with new data ensure that it remains fair and accurate, adapting to any changes in the data distribution. Reference: AIGP Body of Knowledge on Post-Deployment Monitoring and Model Maintenance.

NEW QUESTION: 12

You are the chief privacy officer of a medical research company that would like to collect and use sensitive data about cancer patients, such as their names, addresses, race and ethnic origin, medical histories, insurance claims, pharmaceutical prescriptions, eating and drinking habits and physical activity.

The company will use this sensitive data to build an AI algorithm that will spot common attributes that will help predict if seemingly healthy people are more likely to get cancer. However, the company is unable to obtain consent from enough patients to sufficiently collect the minimum data to train its model.

Which of the following solutions would most efficiently balance privacy concerns with the lack of available data during the testing phase?

- A. Deploy the current model and recalibrate it over time with more data.
- B. Extend the model to multi-modal ingestion with text and images.
- C. Utilize synthetic data to offset the lack of patient data.
- D. Refocus the algorithm to patients without cancer.

Answer: ([SHOW ANSWER](#))

Utilizing synthetic data to offset the lack of patient data is an efficient solution that balances privacy concerns with the need for sufficient data to train the model. Synthetic data can be generated to simulate real patient data while avoiding the privacy issues associated with using actual patient data. This approach allows for the development and testing of the AI algorithm without compromising patient privacy, and it can be refined with real data as it becomes available. Reference: AIGP Body of Knowledge on Data Privacy and AI Model Training.

NEW QUESTION: 13

Your organization is searching for a new way to help accurately forecast sales predictions by various types of customers.

Which of the following is the best type of model to choose if your organization wants to customize the model and avoid lock-in?

- A. A free large language model.
- B. A classic machine learning model.
- C. A proprietary generative AI model.
- D. A subscription-based, multimodal model.

Answer: ([SHOW ANSWER](#))

For customizable, interpretable models that allow organizations to retain control and avoid vendor lock-in, classic ML models (e.g., regression, decision trees, random forests) are optimal.

From the AI Governance in Practice Report 2025:

"Organizations seeking transparency, customizability, and control often prefer classic ML models due to their flexibility and ease of governance." (p. 33)

* A and C may have limited transparency and are often tied to specific providers.

* D involves ongoing costs and limited model control.

NEW QUESTION: 14

A company plans on procuring a tool from an AI provider for its employees to use for certain business purposes.

Which contractual provision would best protect the company's intellectual property in the tool, including training and testing data?

- A. The provider will give privacy notice to individuals before using their personal data to train or test the tool.
- B. The provider will defend and indemnify the company against infringement claims.
- C. The provider will obtain and maintain insurance to cover potential claims.
- D. The provider will warrant that the tool will work as intended.

Answer: ([SHOW ANSWER](#))

To protect the company's intellectual property, the most pertinent contractual provision is ensuring that the AI provider will defend and indemnify the company against infringement claims. This clause means the provider will take responsibility for any intellectual property disputes that arise, thereby safeguarding the company from potential legal and financial repercussions related to the use of the tool. Other options, while beneficial, do not directly address the protection of intellectual property. This concept is detailed in the contractual best practices section of the IAPP AIGP Body of Knowledge.

NEW QUESTION: 15

Random forest algorithms are in what type of machine learning model?

- A. Symbolic.
- B. Generative.
- C. Discriminative.
- D. Natural language processing.

Answer: ([SHOW ANSWER](#))

Random forest algorithms are classified as discriminative models. Discriminative models are used to classify data by learning the boundaries between classes, which is the core functionality of random forest algorithms.

They are used for classification and regression tasks by aggregating the results of multiple decision trees to make accurate predictions.

Reference: The AIGP Body of Knowledge explains that discriminative models, including random forest algorithms, are designed to distinguish between different classes in the data, making them effective for various predictive modeling tasks.

NEW QUESTION: 16

CASE STUDY

Please use the following answer the next question:

A mid-size US healthcare network has decided to develop an AI solution to detect a type of cancer that is most likely arise in adults. Specifically, the healthcare network intends to create a

recognition algorithm that will perform an initial review of all imaging and then route records a radiologist for secondary review pursuant Agreed-upon criteria (e.g., a confidence score below a threshold).

To date, the healthcare network has taken the following steps: defined its AI ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an AI governance committee; assembled a broad, crossfunctional team with clear roles and responsibilities; and created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.

The healthcare network intends to retain a cloud provider to host the solution and a consulting firm to help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.

The most significant risk from combining the healthcare network's existing data with the clinical research partner data is?

- A. Privacy risk.
- B. Security risk.
- C. Operational risk.
- D. Reputational risk.

Answer: (SHOW ANSWER)

The most significant risk from combining the healthcare network's existing data with the clinical research partner data is privacy risk. Combining data sets, especially in healthcare, often involves handling sensitive information that could lead to privacy breaches if not managed properly. De-identified data can still pose re-identification risks when combined with other data sets. Ensuring privacy involves implementing robust data protection measures, maintaining compliance with privacy regulations such as HIPAA, and conducting thorough privacy impact assessments.

Reference: AIGP Body of Knowledge on Data Privacy and Security.

Valid AIGP Dumps shared by EduDump.com for Helping Passing AIGP Exam! EduDump.com now offer the **newest AIGP exam dumps**, the EduDump.com AIGP exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AIGP dumps with Test Engine here: <https://www.edudump.com/exams/IAPP/AIGP/premium/> (219 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 17

What is the most important factor when deciding whether or not to select a proprietary AI model?

- A. What business purpose it will serve.
- B. How frequently it will be updated.
- C. Whether its training data is disclosed.
- D. Whether its system card identifies risks.

Answer: (SHOW ANSWER)

The primary consideration in selecting any AI system, especially a proprietary model, is its fit for business purpose. Whether it serves the intended goals is foundational before evaluating technical or governance features.

From the AI Governance in Practice Report 2025:

"AI governance starts with defining the corporate strategy for AI... and aligning systems with business purpose and operational context." (p. 11)

* B, C, D are relevant for evaluation, but only after confirming business applicability.

NEW QUESTION: 18

A company deploys an AI model for fraud detection in online transactions. During its operation, the model begins to exhibit high rates of false positives, flagging legitimate transactions as fraudulent.

Which is the best step the company should take to address this development?

- A. Dedicate more resources to monitor the model.
- B. Maintain records of all false positives.
- C. Deactivate the model until an assessment is made.
- D. Conduct training for customer service teams to handle flagged transactions.

Answer: ([SHOW ANSWER](#))

When an AI system causes significant false positives, especially in sensitive contexts like fraud detection, the priority is to halt harmful activity and perform a full assessment. Continued use without understanding the fault may cause further customer harm and legal exposure.

From the AI Governance in Practice Report 2025:

"Incident management plans should enable identification, escalation, and system rollback to prevent continued harm from malfunctioning AI systems." (p. 12, 35)

NEW QUESTION: 19

A leading software development company wants to integrate AI-powered chatbots into their customer service platform. After researching various AI models in the market which have been developed by third-party developers, they're considering two options:

Option A - an open-source language model trained on a vast corpus of text data and capable of being trained to respond to natural language inputs.

Option B - a proprietary, generative AI model pre-trained on large data sets, which uses transformer-based architectures to generate human-like responses based on multimodal user input.

Option A would be the best choice for the company because?

- A. It is less expensive to run
- B. It may be better suited for applications requiring customization.
- C. It can handle voice commands and is more suitable for phone-based customer support.
- D. It is built for large-scale, complex dialogues and would be more effective in handling high-volume customer inquiries.

Answer: ([SHOW ANSWER](#))

Open-source models offer more customization flexibility, allowing organizations to fine-tune or adapt the model to fit their own workflows, branding, or compliance needs- making it preferable when deep control is needed.

From the AI Governance in Practice Report 2025:

"Open-source AI allows organizations to review, adapt, and control model behavior in line with organizational needs and policies." (p. 39)

NEW QUESTION: 20

All of the following are potential benefits of using private over public LLMs EXCEPT?

- A. Reduction in time taken for data validation and verification.
- B. Confirmation of security and confidentiality.
- C. Reduction in possibility of hallucinated information.
- D. Application for specific use cases within the enterprise.

Answer: A (LEAVE A REPLY)

Private LLMs offer advantages like customizability, reduced hallucination, confidentiality, and alignment with enterprise-specific tasks, but they do not inherently reduce the time or effort needed for data validation or verification- which remains an essential step regardless of model privacy.

From the AI risk and quality sections:

"Ensuring the quality of the data... is highly contextual and must be validated regardless of the model's deployment environment." (p. 17)

- * B, C, Dare legitimate benefits of private LLMs.
- * A is incorrect - validation still requires time and resources.

NEW QUESTION: 21

Retrieval-Augmented Generation (RAG) is defined as?

- A. Combining LLMs with private knowledge bases to improve their outputs.
- B. Reducing computational processing requirements of the LLMs.
- C. Applying advanced filtering techniques to the LLMs.
- D. Fine tuning LLMs to minimize biased outputs.

Answer: (SHOW ANSWER)

Retrieval-Augmented Generation (RAG) enhances Large Language Models (LLMs) by integrating external, up-to-date, or proprietary information into the generation pipeline-allowing the model to fetch relevant facts from a trusted knowledge source at query time.

Though RAG is not defined directly in the IAPP documents, it is a widely recognized technique in AI governance for ensuring more accurate and contextually grounded outputs, especially in unregulated or high-stakes environments where hallucinations are a concern.

- * B, C, and D describe optimization or bias mitigation-not the core function of RAG.

NEW QUESTION: 22

CASE STUDY

A premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent. The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.

It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.

The organization has a large procurement team that is responsible for the contracting of technology solutions.

One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company deploy technology solutions into the organization's operations in a responsible, cost-effective manner.

The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

All of the following are potential negative consequences created by using the AI tool to help make hiring decisions EXCEPT?

- A. Automation bias
- B. Candidate quality
- C. Privacy violations
- D. Disparate impacts

Answer: (SHOW ANSWER)

The correct answer is B. "Candidate quality" is not a negative consequence of using AI—rather, it is the intended benefit of using such tools (e.g., more efficient filtering of strong candidates).

From the AIGP ILT Guide:

"Automation bias, disparate impact, and privacy risks are well-documented concerns in AI-assisted hiring.

These risks may arise when AI models replicate biases present in training data or obscure the decision logic." AI Governance in Practice Report 2025 (Bias and Fairness Section) also warns:

"Improper AI use in hiring can lead to disparate impact, where neutral criteria disproportionately disadvantage protected groups." Candidate quality is a goal, not a risk, making B the correct answer for what is not a negative outcome.

NEW QUESTION: 23

Which of the following may be permissible uses of an AI system under the EU AI Act EXCEPT?

- A. Using biometrics in abduction cases
- B. Detecting emotions in a telemedicine session

- C. Improving the response time for emergency services
- D. Detecting emotions in a workplace for employee morale

Answer: ([SHOW ANSWER](#))

The correct answer is D. Emotion recognition in the workplace is flagged as unacceptable or highly restricted under the EU AI Act due to its intrusive nature and potential for misuse.

From the AIGP ILT Guide - EU AI Act Training Module:

"AI systems that monitor individuals' emotions in the workplace or educational settings are listed among prohibited or strictly limited practices under Article 5." AI Governance in Practice Report 2025 supports this interpretation:

"Emotion recognition systems, especially in sensitive contexts such as employment or education, raise significant concerns under EU fundamental rights law and are likely to be restricted." Other uses listed—such as emergency response or emotion detection in healthcare—may fall under lawful and beneficial uses, especially when justified by public interest.

NEW QUESTION: 24

What is a key opportunity for companies deploying proprietary Agentic AI models?

- A. Increased immunity from intellectual property claims.
- B. Transferring risk to the cloud vendor.
- C. Greater ability to customize guardrails.
- D. Reduced need for internal AI governance.

Answer: ([SHOW ANSWER](#))

C is the strongest answer because proprietary deployment can give an organization greater control over system architecture, permissions, agent behavior, data access, escalation rules, and technical guardrails. This control is especially valuable for agentic AI, where systems may autonomously invoke tools, access organizational data, and execute multistep workflows. IAPP guidance emphasizes that agentic systems require purpose-specific guardrails, including controls limiting the data and actions available to an agent. Proprietary control can therefore enable safeguards to be tailored closely to organizational risk tolerances and business requirements. It does not provide immunity from intellectual-property claims, eliminating A. Organizations cannot simply transfer accountability or all risk to a cloud provider, eliminating B. Greater agent autonomy actually increases the importance of internal governance, monitoring, authorization, and human oversight rather than reducing it, eliminating D.

NEW QUESTION: 25

Which of the following deployments of generative AI best respects intellectual property rights?

- A. The system produces content that is modified to closely resemble copyrighted work.
- B. The system categorizes and applies filters to content based on licensing terms.
- C. The system provides attribution to creators of publicly available information.
- D. The system produces content that includes trademarks and copyrights.

Answer: ([SHOW ANSWER](#))

Respecting intellectual property rights means adhering to licensing terms and ensuring that generated content complies with these terms. A system that categorizes and applies filters based on licensing terms ensures that content is used legally and ethically, respecting the rights of content creators. While providing attribution is important, categorization and application of filters based on licensing terms are more directly tied to compliance with intellectual property laws. This principle is elaborated in the IAPP AIGP Body of Knowledge sections on intellectual property and compliance.

NEW QUESTION: 26

Pursuant to the White House Executive Order of November 2023, who is responsible for creating guidelines to conduct red-teaming tests of AI systems?

- A.** National Institute of Standards and Technology (NIST).
- B.** National Science and Technology Council (NSTC).
- C.** Office of Science and Technology Policy (OSTP).
- D.** Department of Homeland Security (DHS).

Answer: ([SHOW ANSWER](#))

The White House Executive Order of November 2023 designates the National Institute of Standards and Technology (NIST) as the responsible body for creating guidelines to conduct red-teaming tests of AI systems. NIST is tasked with developing and providing standards and frameworks to ensure the security, reliability, and ethical deployment of AI systems, including conducting rigorous red-teaming exercises to identify vulnerabilities and assess risks in AI systems.

Reference: AIGP BODY OF KNOWLEDGE, sections on AI governance and regulatory frameworks, and the White House Executive Order of November 2023.

NEW QUESTION: 27

According to November 2023 White House Executive Order, which of the following best describes the guidance given to governmental agencies on the use of generative AI as a workplace tool?

- A.** Limit access to specific uses of generative AI.
- B.** Impose a general ban on the use of generative AI.
- C.** Limit access of generative AI to engineers and developers.
- D.** Impose a ban on the use of generative AI in agencies that protect national security.

Answer: ([SHOW ANSWER](#))

The November 2023 White House Executive Order provides guidance that governmental agencies should limit access to specific uses of generative AI. This means that generative AI tools should be used in a controlled manner, where their applications are restricted to well-defined, approved use cases that ensure the security, privacy, and ethical considerations are adequately addressed. This approach allows for the benefits of generative AI to be harnessed while mitigating potential risks and abuses.

Reference: AIGP BODY OF KNOWLEDGE, sections on AI governance and risk management, and the White House Executive Order of November 2023.

NEW QUESTION: 28

CASE STUDY

Please use the following answer the next question:

A mid-size US healthcare network has decided to develop an AI solution to detect a type of cancer that is most likely arise in adults. Specifically, the healthcare network intends to create a recognition algorithm that will perform an initial review of all imaging and then route records a radiologist for secondary review pursuant agreed-upon criteria (e.g., a confidence score below a threshold).

To date, the healthcare network has taken the following steps: defined its AI ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an AI governance committee; assembled a broad, crossfunctional team with clear roles and responsibilities; and created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.

The healthcare network intends to retain a cloud provider to host the solution and a consulting firm to help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.

In the design phase, which of the following steps is most important in gathering the data from the clinical research partner?

- A. Perform a privacy impact assessment.
- B. Combine only anonymized data.
- C. Segregate the data sets.
- D. Review the terms of use.

Answer: (SHOW ANSWER)

Reviewing the terms of use is essential when gathering data from a clinical research partner. This step ensures that the healthcare network complies with all legal and contractual obligations related to data usage. It addresses data ownership, usage limitations, consent requirements, and privacy obligations, which are critical to maintaining ethical standards and avoiding legal repercussions. This review helps ensure that the data is used in a manner consistent with the agreements made and the regulatory environment, which is fundamental for lawful and ethical AI development. Reference: AIGP Body of Knowledge on Legal and Regulatory Considerations.

NEW QUESTION: 29

You asked a generative AI tool to recommend new restaurants to explore in Boston, Massachusetts that have a specialty Italian dish made in a traditional fashion without spinach and wine. The generative AI tool recommended five restaurants for you to visit.

After looking up the restaurants, you discovered one restaurant did not exist and two others did not have the dish.

This information provided by the generative AI tool is an example of what is commonly called?

- A. Prompt injection.
- B. Model collapse.

C. Hallucination.

D. Overfitting.

Answer: (SHOW ANSWER)

In the context of AI, particularly generative models, "hallucination" refers to the generation of outputs that are not based on the training data and are factually incorrect or non-existent. The scenario described involves the generative AI tool providing incorrect and non-existent information about restaurants, which fits the definition of hallucination. Reference: AIGP BODY OF KNOWLEDGE and various AI literature discussing the limitations and challenges of generative AI models.

NEW QUESTION: 30

A company is creating a mobile app to enable individuals to upload images and videos, and analyze this data using ML to provide lifestyle improvement recommendations. The signup form has the following data fields:

- 1.First name
- 2.Last name
- 3.Mobile number
- 4.Email ID
- 5.New password
- 6.Date of birth
- 7.Gender

In addition, the app obtains a device's IP address and location information while in use.

What GDPR privacy principles does this violate?

- A. Purpose Limitation and Data Minimization.
- B. Accountability and Lawfulness.
- C. Transparency and Accuracy.
- D. Integrity and Confidentiality.

Answer: (SHOW ANSWER)

The GDPR privacy principles that this scenario violates are Purpose Limitation and Data Minimization.

Purpose Limitation requires that personal data be collected for specified, explicit, and legitimate purposes and not further processed in a manner that is incompatible with those purposes. Data Minimization mandates that personal data collected should be adequate, relevant, and limited to what is necessary in relation to the purposes for which they are processed. In this case, collecting extensive personal information (e.g., IP address, location, gender) and potentially using it beyond the necessary scope for the app's functionality could violate these principles by collecting more data than needed and possibly using it for purposes not originally intended.

NEW QUESTION: 31

CASE STUDY

Please use the following answer the next question:

XYZ Corp., a premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent.

The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.

It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.

The organization has a large procurement team that is responsible for the contracting of technology solutions.

One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company are responsible for integrating and deploying technology solutions into the organization's operations in a responsible, cost-effective manner.

The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

If XYZ does not deploy and use the AI hiring tool responsibly in the United States, its liability would likely increase under all of the following laws EXCEPT?

- A. Anti-discrimination laws.
- B. Product liability laws.
- C. Accessibility laws.
- D. Privacy laws.

Answer: B ([LEAVE A REPLY](#))

In the United States, the use of AI hiring tools must comply with anti-discrimination laws, accessibility laws, and privacy laws to avoid increasing liability. Anti-discrimination laws (A) ensure that hiring practices do not unlawfully discriminate against protected classes. Accessibility laws (C) require that hiring tools are accessible to all applicants, including those with disabilities. Privacy laws (D) govern the handling of personal data during the hiring process. Product liability laws (B), however, typically apply to the safety and reliability of physical products and would not generally increase liability specifically related to the responsible use of AI hiring tools in the employment context.

Valid AIGP Dumps shared by EduDump.com for Helping Passing AIGP Exam! EduDump.com now offer the **newest AIGP exam dumps**, the EduDump.com AIGP exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AIGP

NEW QUESTION: 32

CASE STUDY

Please use the following answer the next question:

ABC Corp, is a leading insurance provider offering a range of coverage options to individuals. ABC has decided to utilize artificial intelligence to streamline and improve its customer acquisition and underwriting process, including the accuracy and efficiency of pricing policies.

ABC has engaged a cloud provider to utilize and fine-tune its pre-trained, general purpose large language model ("LLM"). In particular, ABC intends to use its historical customer data-including applications, policies, and claims-and proprietary pricing and risk strategies to provide an initial qualification assessment of potential customers, which would then be routed .. human underwriter for final review.

ABC and the cloud provider have completed training and testing the LLM, performed a readiness assessment, and made the decision to deploy the LLM into production. ABC has designated an internal compliance team to monitor the model during the first month, specifically to evaluate the accuracy, fairness, and reliability of its output. After the first month in production, ABC realizes that the LLM declines a higher percentage of women's loan applications due primarily to women historically receiving lower salaries than men.

During the first month when ABC monitors the model for bias, it is most important to?

- A. Continue disparity testing.
- B. Analyze the quality of the training and testing data.
- C. Compare the results to human decisions prior to deployment.
- D. Seek approval from management for any changes to the model.

Answer: ([SHOW ANSWER](#))

During the first month of monitoring the model for bias, it is most important to continue disparity testing.

Disparity testing involves regularly evaluating the model's decisions to identify and address any biases, ensuring that the model operates fairly across different demographic groups.

Reference: Regular disparity testing is highlighted in the AIGP Body of Knowledge as a critical practice for maintaining the fairness and reliability of AI models. By continuously monitoring for and addressing disparities, organizations can ensure their AI systems remain compliant with ethical and legal standards, and mitigate any unintended biases that may arise in production.

NEW QUESTION: 33

Decreasing the complexity of a machine learning model reduces variance and?

- A. Increases bias.
- B. Increases accuracy.
- C. Decreases tolerance.
- D. Decreases interpretability.

Answer: A ([LEAVE A REPLY](#))

The correct answer is A because of the fundamental bias-variance tradeoff in machine learning. When model complexity is reduced, the model becomes simpler and less flexible, which decreases variance because it is less sensitive to fluctuations in the training data. However, this simplification comes at the cost of increased bias, meaning the model may oversimplify relationships and fail to capture underlying patterns accurately.

This tradeoff is a core concept in AI fundamentals and directly impacts model performance, reliability, and governance decisions. From an AI governance perspective, understanding this balance is critical when evaluating model risk, as high bias can lead to systematic errors and fairness issues, while high variance can result in instability and unpredictability in outputs. Proper model tuning aims to balance both for optimal and responsible performance.

NEW QUESTION: 34

In the context of AI governance, what is the primary concern raised by the opacity of many AI systems?

- A.** It ensures that AI systems will make decisions without any human involvement, increasing the risk of harmful outcomes.
- B.** It prevents regulators from gaining legal access to the proprietary algorithms of private companies, limiting enforcement actions.
- C.** It complicates efforts to trace and justify how specific outputs are generated, which can undermine accountability and informed oversight.
- D.** It indicates that the system is inherently unreliable, meaning its predictions or recommendations cannot be trusted in practice.

Answer: ([SHOW ANSWER](#))

Opacity creates a governance problem because stakeholders may be unable to understand, trace, or adequately explain how an AI system arrived at a particular output. This weakens accountability, effective human oversight, auditing, contestability, and risk management. NIST's AI Risk Management Framework connects trustworthy AI with transparency, accountability, explainability, and interpretability, including documenting how systems operate and how outputs are interpreted. Option A incorrectly equates opacity with autonomous operation; an opaque system can still involve humans. Option B concerns regulatory access and intellectual-property issues rather than the fundamental governance problem created by opacity. Option D is also too absolute: lack of interpretability does not automatically make a model unreliable. Therefore, the principal concern is difficulty tracing and justifying outputs, making C the best answer.

NEW QUESTION: 35

What is the most significant risk of deploying an AI model that can create realistic images and videos?

- A.** Copyright infringement.
- B.** Security breaches.
- C.** Downstream harms.

D. Output cannot be protected.

Answer: ([SHOW ANSWER](#))

The greatest risk from AI systems generating realistic synthetic media is downstream harm, such as deepfakes, misinformation, reputational damage, and erosion of trust.

From the AI Governance in Practice Report 2025:

"With generative AI, downstream harms such as deception, reputational damage, misinformation, and manipulation can emerge even if original use was lawful." (p. 55-56)

NEW QUESTION: 36

In the machine learning context, feature engineering is the process of?

- A. Converting raw data into clean data.
- B. Creating learning schema for a model apply.
- C. Developing guidelines to train and test a model.
- D. Extracting attributes and variables from raw data.

Answer: ([SHOW ANSWER](#))

In the machine learning context, feature engineering is the process of extracting attributes and variables from raw data to make it suitable for training an AI model. This step is crucial as it transforms raw data into meaningful features that can improve the model's accuracy and performance. Feature engineering involves selecting, modifying, and creating new features that help the model learn more effectively. Reference: AIGP Body of Knowledge on AI Model Development and Feature Engineering.

NEW QUESTION: 37

What is the 1956 Dartmouth summer research project on AI best known as?

- A. A meeting focused on the impacts of the launch of the first mass-produced computer.
- B. A research project on the impacts of technology on society.
- C. A research project to create a test for machine intelligence.
- D. A meeting focused on the founding of the AI field.

Answer: ([SHOW ANSWER](#))

The 1956 Dartmouth summer research project on AI is best known as a meeting focused on the founding of the AI field. This conference is historically significant because it marked the formal beginning of artificial intelligence as an academic discipline. The term "artificial intelligence" was coined during this event, and it laid the foundation for future research and development in AI. Reference: The AIGP Body of Knowledge highlights the importance of the Dartmouth Conference as a pivotal moment in the history of AI, which established AI as a distinct field of study and research.

NEW QUESTION: 38

Which of the following compliance related controls within an organization is most easily adapted to identify AI risks?

- A. Privacy training.

- B. Penetration testing.
- C. Transfer risk assessments.
- D. Privacy impact assessments.

Answer: D ([LEAVE A REPLY](#))

The correct answer is D because Privacy Impact Assessments are already structured processes designed to identify risks related to data use, processing, and potential harm to individuals. These assessments can be readily adapted to evaluate AI-specific risks, such as bias, automated decision-making impacts, and data protection concerns. AI governance frameworks emphasize leveraging existing compliance mechanisms to efficiently integrate AI risk management without duplicating processes. Privacy impact assessments align closely with AI risk evaluation because they examine how personal data is collected, used, and protected throughout the system lifecycle. Other options, such as penetration testing or training, focus on narrower objectives like security or awareness and are not comprehensive tools for identifying broader AI risks. Adapting PIAs supports a risk-based, scalable, and governance-aligned approach to managing AI systems.

NEW QUESTION: 39

CASE STUDY

Please use the following answer the next question:

A local police department in the United States procured an AI system to monitor and analyze social media feeds, online marketplaces and other sources of public information to detect evidence of illegal activities (e.g., sale of drugs or stolen goods). The AI system works by surveilling the public sites in order to identify individuals that are likely to have committed a crime. It cross-references the individuals against data maintained by law enforcement and then assigns a percentage score of the likelihood of criminal activity based on certain factors like previous criminal history, location, time, race and gender.

The police department retained a third-party consultant assist in the procurement process, specifically to evaluate two finalists. Each of the vendors provided information about their system's accuracy rates, the diversity of their training data and how their system works. The consultant determined that the first vendor's system has a higher accuracy rate and based on this information, recommended this vendor to the police department.

The police department chose the first vendor and implemented its AI system. As part of the implementation, the department and consultant created a usage policy for the system, which includes training police officers on how the system works and how to incorporate it into their investigation process.

The police department has now been using the AI system for a year. An internal review has found that every time the system scored a likelihood of criminal activity at or above 90%, the police investigation subsequently confirmed that the individual had, in fact, committed a crime. Based on these results, the police department wants to forego investigations for cases where the AI system gives a score of at least 90% and proceed directly with an arrest.

When notifying an accused perpetrator, what additional information should a police officer provide about the use of the AI system?

- A. Information about the accuracy of the AI system.
- B. Information about how the accused can oppose the charges.
- C. Information about the composition of the training data of the system.
- D. Information about how the individual was identified by the AI system.

Answer: ([SHOW ANSWER](#))

When notifying an accused perpetrator, the police officer should provide information about how the individual was identified by the AI system. This transparency is crucial for maintaining trust and ensuring that the accused understands the basis of the charges against them. Information about the accuracy, how to oppose the charges, and the composition of the training data, while potentially relevant, do not directly address the immediate need for the accused to understand the specific process that led to their identification. Reference:

AIGP Body of Knowledge on AI Transparency and Explainability.

NEW QUESTION: 40

A French medical research center wishes to develop an AI-based system that will predict the risk of serious diseases based on patients' genetic data. To do so, it contracts with a technology company and provides it with patients' data previously obtained by the center during research.

To guarantee compliance when processing special categories of personal data, the medical research center must ensure that?

- A. The AI-based system is designed for the purposes of preventive medicine.
- B. The patients' health and genetic data is anonymized.
- C. The patients have given explicit consent to using the data.
- D. The technology company is located in the EU and is not cloud-based.

Answer: ([SHOW ANSWER](#))

Genetic and health information are special categories of personal data under Article 9 of the GDPR, for which processing is generally prohibited unless a valid Article 9 exception applies. Among the available answers, explicit consent provides the clearest applicable condition: Article 9(2)(a) permits processing where the data subject has explicitly consented to processing for specified purposes. Importantly, explicit consent is not the only possible legal condition under the GDPR; scientific research and healthcare processing may sometimes rely on other Article 9 conditions where applicable law and safeguards permit. However, A merely describing a preventive-medicine purpose does not itself establish the necessary legal condition. D is irrelevant because EU establishment and non-cloud processing do not establish lawfulness. B would materially change the analysis only if the data were genuinely anonymized so that they were no longer personal data.

NEW QUESTION: 41

Which of the following AI uses is best described as human-centric?

- A. Pattern recognition algorithms are used to improve the accuracy of weather predictions, which benefits many industries and everyday life.

- B.** Autonomous robots are used to move products within a warehouse, allowing human workers to reduce physical strain and alleviate monotony.
- C.** Machine learning is used for demand forecasting and inventory management, ensuring that consumers can find products they want when they want them.
- D.** Virtual assistants are used adapt educational content and teaching methods to individuals, offering personalized recommendations based on ability and needs.

Answer: ([SHOW ANSWER](#))

Human-centric AI focuses on improving the human experience by addressing individual needs and enhancing human capabilities. Option D exemplifies this by using virtual assistants to tailor educational content to each student's unique abilities and needs, thereby supporting personalized learning and improving educational outcomes. This use case directly benefits individuals by providing customized assistance and adapting to their learning pace and style, aligning with the principles of human-centric AI.

Reference: AIGP BODY OF KNOWLEDGE, sections on trustworthy AI and human-centric AI principles.

NEW QUESTION: 42

What is most likely the first action that a developer takes to map, plan and scope an AI project?

- A.** Define the business case and perform a cost benefit analysis answering the question of why AI
- B.** Use a test, evaluation, verification, validation TEVV process
- C.** Perform an algorithmic impact assessment leveraging PIAs
- D.** Determine feasibility and optionality of redress

Answer: ([SHOW ANSWER](#))

The correct answer is A because the first step in any AI project lifecycle is to clearly define the business objective and justify the use of AI. AI governance frameworks emphasize that planning begins with understanding the purpose, value, and necessity of the system before moving into design or risk assessment stages. Establishing a business case ensures alignment with organizational goals, identifies expected benefits, and evaluates whether AI is the appropriate solution. This step corresponds to the planning phase of the AI lifecycle, where objectives and intended outcomes are documented. In contrast, TEVV processes occur during development and testing, while impact assessments and redress considerations arise later as part of risk management and governance. Starting with a clear "why AI" foundation ensures responsible, efficient, and goal-oriented system development.

NEW QUESTION: 43

An AI start-up is developing a system for automated loan approvals. The team wants to minimize risks of bias and regulatory non-compliance. They have already identified potential stakeholders, including regulators and consumer groups.

What is the most appropriate sequence of next steps?

- A.** Conduct harm analysis,
Benchmark system performance,

Proceed to deployment if the system performs similarly.

B. Launch a small-scale pilot,

Gather user feedback,

Afterward analyze harms.

C. Perform a probability/severity analysis,

Apply a risk mitigation hierarchy to address the most severe risks,

Conduct pre-deployment pilot testing.

D. Enable explainability tools to reassure users,

Launch pilot,

Map risks during deployment.

Answer: (SHOW ANSWER)

The correct answer is C because it follows a structured, risk-based AI governance approach aligned with best practices and regulatory expectations. AI governance frameworks emphasize identifying and assessing risks through probability and severity analysis before deployment. This allows organizations to prioritize the most critical risks, such as bias and discrimination in loan approvals. Applying a risk mitigation hierarchy ensures that controls are implemented proactively rather than reactively. Only after risks are assessed and mitigated should pilot testing occur to validate system performance in a controlled environment. Other options either delay risk assessment until after deployment or focus on limited aspects like benchmarking or explainability without addressing core risk management principles. A proactive, risk-first approach is essential for compliance, fairness, and responsible AI deployment.

NEW QUESTION: 44

A company developing and deploying its own AI model would perform all of the following steps to monitor and evaluate the model's performance EXCEPT?

A. Publicly disclosing data with forecasts of secondary and downstream harms to stakeholders.

B. Setting up automated tools to regularly track the model's accuracy, precision and recall rates in real-time.

C. Implementing a formal incident response plan to address incidents that may occur during system operation.

D. Establishing a regular schedule for human evaluation of the model's performance, including qualitative assessments.

Answer: (SHOW ANSWER)

While transparency is encouraged, publicly disclosing forecasts of secondary harms is not a required or standard practice for internal performance evaluation. Risk assessments and reporting typically remain internal or shared with regulators.

From the AI Governance in Practice Report 2025:

"Organizations must assess secondary risks... but disclosure is subject to context, regulatory requirements, and risk management discretion." (p. 30)

NEW QUESTION: 45

A shipping service based in the US is looking to expand its operations into the EU. It utilizes an in-house developed multimodal AI model that analyzes all personal data collected from shipping senders and recipients, and optimizes shipping routes and schedules based on this data. As they expand into the EU, all of the following descriptions should be included in the technical documentation for their AI model EXCEPT?

- A. A general description of the AI system.
- B. A description of the prioritization of the risks of deployment of the AI system.
- C. A description of the appropriateness of the performance metrics for the specific AI system.
- D. A detailed description of the elements of the AI system and of the process for its development.

Answer: ([SHOW ANSWER](#))

The EU AI Act outlines what must be included in technical documentation for high-risk systems. These requirements are designed to support conformity assessment, transparency, and traceability.

From the AI Governance in Practice Report 2025:

"It mandates drawing up technical documentation... must include a general description of the AI system, the intended purpose, and a detailed description of the elements and development process." (p. 34)

"Documentation... includes training, testing, evaluation procedures, and appropriateness of performance metrics." (p. 34-35) The risk management system is addressed separately through a risk management plan, not within the technical documentation itself.

Thus:

- * A, C, and D are explicitly required in the technical documentation.
- * B, while important, is part of the risk management process, not a required section of technical documentation.

NEW QUESTION: 46

Scenario:

A U.S.-based AI governance professional is evaluating resources from the National Institute of Standards and Technology (NIST) to guide the organization's AI risk assessment strategy. They are particularly interested in programs focused on assessing AI-specific impacts.

The main purpose of NIST's Assessing Risks and Impacts of AI (ARIA) program is to:

- A. Provide a suite of resources to manage risks
- B. Pilot new standards for AI red-teaming
- C. Promote interoperability across AI systems
- D. Offer a regulatory sandbox for risk reporting

Answer: ([SHOW ANSWER](#))

The correct answer is A. The ARIA program by NIST is explicitly designed to support stakeholders in understanding and managing the risks and impacts of AI systems.

From the AIGP ILT Guide - U.S. Risk Frameworks Module:

"NIST's ARIA program develops and pilots assessment tools for AI risks and impacts, aimed at improving organizational capacity for responsible AI use." Also cited in the AI Governance in Practice Report 2025 (Frameworks Section):

"ARIA supports and aligns with the AI Risk Management Framework by helping organizations assess AI harms, safety concerns, and societal implications." ARIA is not a red-teaming or sandbox program—it's an assessment and governance resource.

Valid AIGP Dumps shared by EduDump.com for Helping Passing AIGP Exam! EduDump.com now offer the **newest AIGP exam dumps**, the EduDump.com AIGP exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AIGP dumps with Test Engine here: <https://www.edudump.com/exams/IAPP/AIGP/premium/> (219 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 47

All of the following are unique characteristics of AI that require a comprehensive approach to governance EXCEPT?

- A. Autonomy
- B. Automation
- C. Adaptability
- D. Speed and scale
- E. Superintelligence

Answer: (SHOW ANSWER)

The correct answer is E - Superintelligence.

While the other options (Autonomy, Automation, Adaptability, Speed and Scale) are commonly cited as real-world characteristics that affect governance today, superintelligence remains a theoretical concept.

From the AIGP ILT Guide and AI Governance in Practice Report 2025:

"Core AI characteristics—such as automation, adaptability, speed, and autonomy—require active governance due to their impact on decision-making, legal liability, and risk."

However, superintelligence is speculative and not a current feature of AI systems under governance frameworks like the EU AI Act or NIST RMF.

Thus, it's not a current characteristic requiring governance in real-world enterprise settings.

NEW QUESTION: 48

Which of the following use cases would be best served by a non-AI solution?

- A. A non-profit wants to develop a social media presence.
- B. A business analyst wants to forecast future cost overruns and underruns.
- C. A customer service agency wants to automate answers to common questions.

Answer: (SHOW ANSWER)

Developing a social media presence for a non-profit is best served by non-AI solutions. This task primarily involves content creation, community engagement, and strategic planning, which are effectively managed by human expertise and traditional marketing tools. AI is more suitable for tasks requiring automation, large-scale data analysis, and personalized recommendations, such as e-commerce personalization, forecasting cost overruns, or automating customer service responses. Reference: AIGP Body of Knowledge on AI Use Cases and Applications.

NEW QUESTION: 49

Which of the following considerations is the most important in mitigating the potential of bias in training and testing data?

- A. Using privacy-enhancing tools to protect the data.
- B. Validating that the data is representative.
- C. Receiving explicit consent to use the data.
- D. Assessing the sufficiency of the third-party data.

Answer: (SHOW ANSWER)

The correct answer is B because ensuring that training and testing data is representative is the most critical factor in mitigating bias in AI systems. AI governance frameworks emphasize that biased or unrepresentative datasets can lead to discriminatory outcomes, particularly when certain demographic groups are underrepresented or overrepresented. Representative data helps ensure that the model performs fairly and accurately across different populations. While privacy-enhancing tools and consent address legal and ethical data use, they do not directly prevent bias in model outcomes. Similarly, assessing the sufficiency of third-party data focuses on quantity rather than fairness or distribution. Effective bias mitigation begins with evaluating whether the dataset reflects the diversity and characteristics of the real-world population the AI system will impact.

NEW QUESTION: 50

A company that deploys AI but is not currently a provider or developer intends to develop and market its own AI system.

Which obligation would then be likely to apply?

- A. Implementing a risk management framework.
- B. Conducting an impact assessment including a post-deployment monitoring plan.
- C. Developing documentation on the system, the potential risks and the safeguards applied.
- D. Developing a reporting plan for any observed algorithmic discrimination or harms to individuals' rights and freedoms.

Answer: (SHOW ANSWER)

Once a company moves from being a deployer to also acting as a provider or developer, it assumes new obligations under regulations like the EU AI Act. One of the core requirements for providers is to produce and maintain technical documentation, including descriptions of the model, associated risks, and mitigation strategies.

From the AI Governance in Practice Report 2025:

"Providers of high-risk AI systems must draw up technical documentation demonstrating the system's conformity with the requirements... including potential risks and safeguards applied." (p. 34)

"This documentation must be available before placing the system on the market." (p. 35)

NEW QUESTION: 51

Scenario:

A public sector agency is reviewing proposed AI use cases for improving services. It wants to prioritize implementations that deliver value but minimize unintended negative consequences. When evaluating which AI use cases to implement, an organization should consider all of the following EXCEPT:

- A. Related TEVV (test, evaluate, verify, validate) and system metrics
- B. The users and their expectations
- C. Equitable access to the AI tool
- D. Potential positive and negative impacts of the system

Answer: (SHOW ANSWER)

The correct answer is A. While TEVV is important in later lifecycle phases, it is not the primary consideration when evaluating and prioritizing use cases.

From the AIGP Body of Knowledge - Use Case Assessment Module:

"Use case evaluation focuses on value, impact, fairness, and accessibility-technical testing considerations come later." ILT Guide confirms:

"Organizations should first assess whether the AI system provides equitable outcomes and aligns with stakeholder expectations. TEVV is part of implementation, not initial prioritization." Thus, A is not a top-level consideration during use case selection.

NEW QUESTION: 52

Scenario:

A distributor operating in the EU is responsible for selling imported high-risk AI systems to businesses. The distributor wants to ensure they fulfill all applicable obligations under the EU AI Act.

All of the following are obligations of a distributor of high-risk AI systems under the EU AI Act EXCEPT?

- A. Corrective actions
- B. Verification of CE marking
- C. Registration in EU Database
- D. Communication with national authorities

Answer: (SHOW ANSWER)

The correct answer is C. Registration in the EU database is an obligation of providers of high-risk AI systems- not distributors.

From the AIGP ILT Guide - Roles & Obligations Module:

"Distributors must verify CE marking, ensure instructions for use are provided, inform authorities of risks, and take corrective action when necessary. However, registration duties in the EU database lie with the provider." Also from the AI Governance in Practice Report 2025:

"The AI Act differentiates responsibilities for developers, providers, importers, and distributors. Only providers of high-risk systems are obligated to register their systems in the EU AI Database." Distributors focus on verification and communication, not formal registration.

NEW QUESTION: 53

CASE STUDY

Please use the following answer the next question:

ABC Corp, is a leading insurance provider offering a range of coverage options to individuals. ABC has decided to utilize artificial intelligence to streamline and improve its customer acquisition and underwriting process, including the accuracy and efficiency of pricing policies.

ABC has engaged a cloud provider to utilize and fine-tune its pre-trained, general purpose large language model ("LLM"). In particular, ABC intends to use its historical customer data-including applications, policies, and claims-and proprietary pricing and risk strategies to provide an initial qualification assessment of potential customers, which would then be routed to a human underwriter for final review.

ABC and the cloud provider have completed training and testing the LLM, performed a readiness assessment, and made the decision to deploy the LLM into production. ABC has designated an internal compliance team to monitor the model during the first month, specifically to evaluate the accuracy, fairness, and reliability of its output. After the first month in production, ABC realizes that the LLM declines a higher percentage of women's loan applications due primarily to women historically receiving lower salaries than men.

The best approach to enable a customer who wants information on the AI model's parameters for underwriting purposes is to provide?

- A.** A transparency notice.
- B.** An opt-out mechanism.
- C.** Detailed terms of service.
- D.** Customer service support.

Answer: (SHOW ANSWER)

The best approach to enable a customer who wants information on the AI model's parameters for underwriting purposes is to provide a transparency notice. This notice should explain the nature of the AI system, how it uses customer data, and the decision-making process it follows.

Providing a transparency notice is crucial for maintaining trust and compliance with regulatory requirements regarding the transparency and accountability of AI systems.

Reference: According to the AIGP Body of Knowledge, transparency in AI systems is essential to ensure that stakeholders, including customers, understand how their data is being used and how decisions are made. This aligns with ethical principles of AI governance, ensuring that customers are informed and can make knowledgeable decisions regarding their interactions with AI systems.

NEW QUESTION: 54

Retraining an LLM can be necessary for all of the following reasons EXCEPT?

- A. To minimize degradation in prediction accuracy due to changes in data.
- B. Adjust the model's hyper parameters to a specific use case.
- C. Account for new interpretations of the same data.
- D. To ensure interpretability of the model's predictions.

Answer: ([SHOW ANSWER](#))

Retraining an LLM (Large Language Model) is primarily done to improve or maintain its performance as data changes over time, to fine-tune it for specific use cases, and to incorporate new data interpretations to enhance accuracy and relevance. However, ensuring interpretability of the model's predictions is not typically a reason for retraining. Interpretability relates to how easily the outputs of the model can be understood and explained, which is generally addressed through different techniques or methods rather than through the retraining process itself. References to this can be found in the IAPP AIGP Body of Knowledge discussing model retraining and interpretability as separate concepts.

NEW QUESTION: 55

Which of the following most encourages accountability over AI systems?

- A. Determining the business objective and success criteria for the AI project.
- B. Performing due diligence on third-party AI training and testing data.
- C. Defining the roles and responsibilities of AI stakeholders.
- D. Understanding AI legal and regulatory requirements.

Answer: ([SHOW ANSWER](#))

Defining the roles and responsibilities of AI stakeholders is crucial for encouraging accountability over AI systems. Clear delineation of who is responsible for different aspects of the AI lifecycle ensures that there is a person or team accountable for monitoring, maintaining, and addressing issues that arise. This accountability framework helps in ensuring that ethical standards and regulatory requirements are met, and it facilitates transparency and traceability in AI operations. By assigning specific roles, organizations can better manage and mitigate risks associated with AI deployment and use.

NEW QUESTION: 56**CASE STUDY**

Please use the following answer the next question:

A local police department in the United States procured an AI system to monitor and analyze social media feeds, online marketplaces and other sources of public information to detect evidence of illegal activities (e.g., sale of drugs or stolen goods). The AI system works by surveilling the public sites in order to identify individuals that are likely to have committed a crime. It cross-references the individuals against data maintained by law enforcement and then assigns a percentage score of the likelihood of criminal activity based on certain factors like previous criminal history, location, time, race and gender.

The police department retained a third-party consultant assist in the procurement process, specifically to evaluate two finalists. Each of the vendors provided information about their system's accuracy rates, the diversity of their training data and how their system works. The consultant determined that the first vendor's system has a higher accuracy rate and based on this information, recommended this vendor to the police department.

The police department chose the first vendor and implemented its AI system. As part of the implementation, the department and consultant created a usage policy for the system, which includes training police officers on how the system works and how to incorporate it into their investigation process.

The police department has now been using the AI system for a year. An internal review has found that every time the system scored a likelihood of criminal activity at or above 90%, the police investigation subsequently confirmed that the individual had, in fact, committed a crime. Based on these results, the police department wants to forego investigations for cases where the AI system gives a score of at least 90% and proceed directly with an arrest.

What is the best reason the police department should continue to perform investigations even if the AI system scores an individual's likelihood of criminal activity at or above 90%?

- A. Because the department did not perform an impact assessment for this intended use.
- B. Because AI systems that affect fundamental civil rights should not be fully automated.
- C. Because investigations may identify additional individuals involved in the crime.
- D. Because investigations may uncover information relevant to sentencing.

Answer: (SHOW ANSWER)

The best reason for the police department to continue performing investigations even if the AI system scores an individual's likelihood of criminal activity at or above 90% is that AI systems affecting fundamental civil rights should not be fully automated. Human oversight is essential to ensure that decisions impacting civil liberties are made with due consideration of context and mitigating factors that an AI might not fully appreciate. This approach ensures fairness, accountability, and adherence to legal standards. Reference: AIGP Body of Knowledge on AI Ethics and Human Oversight.

NEW QUESTION: 57

The best method to ensure a comprehensive identification of risks for a new AI model is?

- A. An environmental scan.
- B. Red teaming.
- C. Integration testing.
- D. An impact assessment.

Answer: (SHOW ANSWER)

The most comprehensive way to identify a full range of risks - legal, ethical, operational, and societal - for a new AI model is through a formal impact assessment, such as a Data Protection Impact Assessment (DPIA) or Algorithmic Impact Assessment.

From the AI Governance in Practice Report 2025:

"Risk-based approaches are often distilled into organizational risk management efforts, which put impact assessments at the heart of deciding whether harm can be reduced." (p. 29)

"DPIAs... help organizations identify, analyze and minimize data-related risks and demonstrate accountability." (p. 30)

- * A. Environmental scans too general.
- * B. Red teaming is useful for adversarial risk but not broad.
- * C. Integration testing focuses on technical/system compatibility, not overall risk.

NEW QUESTION: 58

MULTI-SELECT

Please select 3 of the 5 options below. No partial credit will be given.

The use of paid generative AI public tools is appealing because?

- A. They are convenient to adopt.
- B. They have additional privacy and security controls.
- C. They have frequent enhancements of new features.
- D. They provide transparency in models and in decision-making.
- E. They eliminate concerns about the data used to generate outputs.

Answer: ([SHOW ANSWER](#))

Paid public generative AI services can be attractive because organizations can adopt them rapidly without developing and maintaining foundational models themselves. Commercial offerings may also provide organizational privacy, security, access-management, and administrative controls beyond those available through basic public access, while centrally managed services can receive frequent provider-driven feature enhancements. These characteristics support A, B, and C. However, payment does not guarantee transparency into proprietary model architecture, training data, weights, or decision processes, so D is incorrect. Nor does a paid service eliminate concerns regarding training-data provenance, intellectual property, personal information, confidential prompts, data retention, or generated outputs; consequently, E is incorrect. AI governance still requires vendor due diligence, contractual safeguards, acceptable-use controls, security assessment, and continuing monitoring when organizations depend on third-party generative AI services.

NEW QUESTION: 59

Your management consulting firm is planning to use an AI system to support its employees.

Which category of operator applies to the firm in this context?

- A. Authorized representative.
- B. Distributor.
- C. Provider.
- D. Deployer.

Answer: ([SHOW ANSWER](#))

Under the EU AI Act, a deployer is a natural or legal person, public authority, agency, or other body that uses an AI system under its authority, except where the system is used in a personal,

non-professional activity. The consulting firm is using an existing AI system to support its employees and is therefore acting as a deployer. It is not a provider because the scenario does not state that the firm develops the AI system, has it developed, or places it on the market under its own name or trademark. It is not a distributor because it is not making the system available in the supply chain. An authorized representative acts on behalf of a provider established outside the EU. Consequently, D correctly describes the firm's role.

NEW QUESTION: 60

Scenario:

An enterprise is evaluating multiple third-party generative AI tools to integrate into its platform. As part of its AI governance policy, it is assessing the most effective methods to reduce risks related to bias, data misuse, and liability when using third-party solutions.

All of the following are commonly adopted processes and policies in reducing potential risks introduced by third-party AI tools or applications EXCEPT:

- A. Including clauses in the procurement agreement for buyers of generative AI tools to put certain liabilities on the tool supplier
- B. Allowing publicly available information and personally identifiable information (PII) to be incorporated into the prompt
- C. Requiring an independent third-party bias audit for third-party generative AI tools
- D. Requiring new use cases of the generative AI tools or applications to be reviewed and approved by the generative AI governance body

Answer: (SHOW ANSWER)

The correct answer is B. Allowing PII to be freely entered into prompts without safeguards is considered a major privacy and security risk and is not a responsible governance practice.

From the AIGP ILT Guide - Generative AI & Third-Party Risk Management:

"Use of personal or sensitive information in AI prompts can result in unintended exposure, regulatory breaches, and downstream liability." The AI Governance in Practice

Report 2025 highlights:

"PII should be minimized or protected by design. Prompt engineering should prevent entry of personally identifiable data unless legally and technically safeguarded." A, C, and D are established best practices under responsible AI procurement and use.

NEW QUESTION: 61

Scenario:

A company is using different types of AI systems to enhance consumer engagement. These include chatbots, recommendation engines, and automated content generation tools.

Which of the following situations would be least likely to raise concerns under existing consumer protection laws?

- A. An AI algorithm being used in a credit decision-making process by a financial institution
- B. An AI customer service system claiming that it is as accurate as a human support agent
- C. An AI tool using scraped digital content to generate news summaries on a publishing website

D. An online platform offering recommendations to its users by displaying user-specific content and targeted advertisements

Answer: (SHOW ANSWER)

The correct answer is D. Personalized content and advertisements, as long as properly disclosed and non-deceptive, are not generally a consumer protection issue under current legal regimes.

From the AI Governance in Practice Report 2025 (Consumer Protection Section):

"Standard practices like targeted advertising and recommendations are widely accepted provided they comply with transparency and consent requirements." Meanwhile, credit decision-making and misleading AI performance claims (Answers A and B) have already led to regulatory enforcement.

The AIGP ILT Guide highlights:

"Deceptive claims, biased financial decisions, and unauthorized data use may violate consumer protection and privacy laws. Advertising personalization is routine but must be disclosed appropriately."

Valid AIGP Dumps shared by EduDump.com for Helping Passing AIGP Exam! EduDump.com now offer the **newest AIGP exam dumps**, the EduDump.com AIGP exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AIGP dumps with Test Engine here: <https://www.edudump.com/exams/IAPP/AIGP/premium/> (219 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 62

What is the best method to proactively train an LLM so that there is mathematical proof that no specific piece of training data has more than a negligible effect on the model or its output?

- A.** Clustering.
- B.** Transfer learning.
- C.** Differential privacy.
- D.** Data compartmentalization.

Answer: C (LEAVE A REPLY)

Differential privacy is a technique used to ensure that the inclusion or exclusion of a single data point does not significantly affect the outcome of any analysis, providing a way to mathematically prove that no specific piece of training data has more than a negligible effect on the model or its output. This is achieved by introducing randomness into the data or the algorithms processing the data. In the context of training large language models (LLMs), differential privacy helps in protecting individual data points while still enabling the model to learn effectively. By adding noise to the training process, differential privacy provides strong guarantees about the privacy of the training data.

Reference: AIGP BODY OF KNOWLEDGE, pages related to data privacy and security in model training.

NEW QUESTION: 63

Which of the following best defines an "AI model"?

- A. A system that applies defined rules to execute tasks.
- B. A system of controls that is used to govern an AI algorithm.
- C. A corpus of data which an AI algorithm analyzes to make predictions.
- D. A program that has been trained on a set of data to find patterns within the data.

Answer: (SHOW ANSWER)

An AI model is best defined as a program that has been trained on a set of data to find patterns within that data. This definition captures the essence of machine learning, where the model learns from the data to make predictions or decisions. Reference: AIGP BODY OF KNOWLEDGE, which provides a detailed explanation of AI models and their training processes.

NEW QUESTION: 64

Scenario:

A European AI technology company was found to be non-compliant with certain provisions of the EU AI Act.

The regulator is considering penalties under the enforcement provisions of the regulation.

According to the EU AI Act, which of the following non-compliance examples could lead to fines of up to

€15 million or 3% of annual worldwide turnover(whichever is higher)?

- A. In case of AI Act prohibitions
- B. In case of breach of a provider ' s obligations for high-risk AI systems
- C. In case of the supply of misleading information to notified bodies in reply to a request
- D. In case of a breach of AI Act prohibition by the Union institutions, bodies, offices and agencies

Answer: (SHOW ANSWER)

The correct answer is B. The EU AI Act assigns a tiered penalty system based on the severity of the violation. A breach of obligations related to high-risk AI systems falls into the mid-tier category, triggering fines of €15 million or 3% of annual global turnover.

From the AIGP ILT Guide - EU AI Act Module:

"Providers of high-risk AI systems must comply with strict documentation, testing, monitoring, and registration obligations. Breaches of these result in significant fines of up to €15 million or 3% of turnover." AI Governance in Practice Report 2025 supports this:

"Non-compliance with obligations under Title III (high-risk systems) leads to financial penalties under Article

71(3) of the EU AI Act."

Note: The highest penalty (€35 million or 7%) applies to prohibited AI uses, not to obligations for high-risk systems.

NEW QUESTION: 65

All of the following issues are unique for proprietary AI model deployments EXCEPT?

- A. The acquisition of training data.
- B. The cost of AI chips.
- C. The potential for bias.
- D. The necessity of performing conformity assessments.

Answer: ([SHOW ANSWER](#))

Bias is a common risk across both proprietary and open-source models, and not unique to proprietary deployments. All AI systems - regardless of origin - require evaluation for fairness, accuracy, and representativeness.

From the AI Governance in Practice Report 2025:

"Bias, discrimination and fairness challenges are present in both open and closed models, regardless of how the model is sourced." (p. 41)

NEW QUESTION: 66

Scenario:

An organization wants to leverage its existing compliance structures to identify AI-specific risks as part of an ongoing data governance audit.

Which of the following compliance-related controls within an organization is most easily adapted to identify AI risks?

- A. Privacy training
- B. Penetration testing
- C. Transfer risk assessments
- D. Privacy impact assessments

Answer: ([SHOW ANSWER](#))

The correct answer is D - Privacy impact assessments (PIAs). These are directly adaptable for identifying risks in AI systems, particularly around data usage, bias, and individual impacts.

From the AIGP ILT Guide - Risk Management Module:

"PIAs and DPIAs are existing tools used in privacy compliance that can be extended to evaluate the risks of AI, including fairness, explainability, and legality." AI Governance in Practice Report 2025 further explains:

"Organizations can adapt privacy impact assessments to evaluate the ethical, legal, and technical risks posed by AI systems. They provide a structured and recognized method." PIAs are preferable over general security practices (like pen testing) which do not address algorithmic bias or legal compliance directly.

NEW QUESTION: 67

According to the EU AI Act, providers of what kind of machine learning systems will be required to register with an EU oversight agency before placing their systems in the EU market?

- A. AI systems that are harmful based on a legal risk-utility calculation.
- B. AI systems that are "strong" general intelligence.
- C. AI systems trained on sensitive personal data.
- D. AI systems that are high-risk.

Answer: (SHOW ANSWER)

According to the EU AI Act, providers of high-risk AI systems are required to register with an EU oversight agency before these systems can be placed on the market. This requirement is part of the Act's framework to ensure that high-risk AI systems comply with stringent safety, transparency, and accountability standards.

High-risk systems are those that pose significant risks to health, safety, or fundamental rights. Registration with oversight agencies helps facilitate ongoing monitoring and enforcement of compliance with the Act's provisions. Systems categorized under other criteria, such as those trained on sensitive personal data or exhibiting "strong" general intelligence, also fall under scrutiny but are primarily covered under different regulatory requirements or classifications.

NEW QUESTION: 68

What is the technique to remove the effects of improperly used data from an ML system?

- A. Data cleansing.
- B. Model inversion.
- C. Data de-duplication.
- D. Model disgorgement.

Answer: (SHOW ANSWER)

The correct answer is D, model disgorgement. This technique refers to removing or eliminating the influence of improperly obtained, biased, or unlawfully used data from a trained machine learning model. It is increasingly discussed in AI governance and regulatory enforcement, particularly where models have been trained on data collected without proper consent or in violation of legal requirements. Instead of merely cleaning datasets, model disgorgement may require retraining or discarding the model entirely to ensure that problematic data no longer influences outputs. This aligns with accountability and compliance principles in AI governance, where organizations must ensure lawful data use throughout the AI lifecycle. Other options like data cleansing and de-duplication address data quality but do not fully remove learned patterns already embedded in trained models.

NEW QUESTION: 69

During the first month when the company monitors the model for bias, it is most important to?

- A. Continue disparity testing.
- B. Provide regular awareness training.
- C. Analyze the quality of the training and testing data.
- D. Document the results of final decisions made by the human underwriter.

Answer: (SHOW ANSWER)

The initial deployment phase of an AI model is critical for post-deployment monitoring. When tracking for bias, the most important task is to continue disparity testing to determine whether outputs differ across protected groups.

From the AI Governance in Practice Report 2025:

"Performance monitoring protocols... should include mechanisms to assess and measure disparities in outcomes across different demographic groups." (p. 12)

"Bias may not be evident during pre-deployment testing but can emerge in real-world use." (p. 41)

* B. Awareness training is helpful, but not a technical bias mitigation activity.

* C. Analyzing training data is a pre-deployment task.

* D. Documenting human decisions may support auditability but doesn't detect bias in AI outputs.

NEW QUESTION: 70

What is the primary purpose of conducting ethical red-teaming on an AI system?

A. To improve the model's accuracy.

B. To simulate model risk scenarios.

C. To identify security vulnerabilities.

D. To ensure compliance with applicable law.

Answer: (SHOW ANSWER)

The primary purpose of conducting ethical red-teaming on an AI system is to simulate model risk scenarios.

Ethical red-teaming involves rigorously testing the AI system to identify potential weaknesses, biases, and vulnerabilities by simulating real-world attack or failure scenarios. This helps in proactively addressing issues that could compromise the system's reliability, fairness, and security. Reference: AIGP Body of Knowledge on AI Risk Management and Ethical AI Practices.

NEW QUESTION: 71

CASE STUDY

Please use the following answer the next question:

Good Values Corporation (GVC) is a U.S. educational services provider that employs teachers to create and deliver enrichment courses for high school students. GVC has learned that many of its teacher employees are using generative AI to create the enrichment courses, and that many of the students are using generative AI to complete their assignments.

In particular, GVC has learned that the teachers they employ used open source large language models ("LLM") to develop an online tool that customizes study questions for individual students. GVC has also discovered that an art teacher has expressly incorporated the use of generative AI into the curriculum to enable students to use prompts to create digital art.

GVC has started to investigate these practices and develop a process to monitor any use of generative AI, including by teachers and students, going forward.

Which of the following risks should be of the highest concern to individual teachers using generative AI to ensure students learn the course material?

A. Financial cost.

B. Model accuracy.

C. Technical complexity.

D. Copyright infringement.

Answer: (SHOW ANSWER)

The highest concern for individual teachers using generative AI to ensure students learn the course material is model accuracy. Ensuring that the AI-generated content is accurate and relevant to the curriculum is crucial for effective learning. If the AI model produces inaccurate or irrelevant content, it can mislead students and hinder their understanding of the subject matter. Reference: According to the AIGP Body of Knowledge, one of the core risks posed by AI systems is the accuracy of the data and models used. Ensuring the accuracy of AI-generated content is essential for maintaining the integrity of the educational material and achieving the desired learning outcomes.

NEW QUESTION: 72

What is the best reason for a company adopt a policy that prohibits the use of generative AI?

- A. Avoid using technology that cannot be monetized.
- B. Avoid needing to identify and hire qualified resources.
- C. Avoid the time necessary to train employees on acceptable use.
- D. Avoid accidental disclosure to its confidential and proprietary information.

Answer: [\(SHOW ANSWER\)](#)

The primary concern for a company adopting a policy prohibiting the use of generative AI is the risk of accidental disclosure of confidential and proprietary information. Generative AI tools can inadvertently leak sensitive data during the creation process or through data sharing. This risk outweighs the other reasons listed, as protecting sensitive information is critical to maintaining the company's competitive edge and legal compliance. This rationale is discussed in the sections on risk management and data privacy in the IAPP AIGP Body of Knowledge.

NEW QUESTION: 73

A US hospital plans to develop an AI that will review available patient data in order to propose an initial diagnosis to licensed physicians. The hospital will implement a policy that requires physicians to consider the AI proposal, but conduct their own physical examinations prior to making a final diagnosis.

An important ethical concern with this plan is?

- A. Whether patients will receive an economic benefit from the use of AI.
- B. Whether the AI was trained on a representative dataset.
- C. Whether physicians understand how the AI works.
- D. Whether the AI will have an error rate comparable to human physicians.

Answer: [B \(LEAVE A REPLY\)](#)

The core ethical concern when deploying diagnostic AI in a healthcare setting is ensuring fairness and accuracy across diverse patient populations. If the AI is trained on a dataset that is not representative of the population it will serve, it risks reinforcing health disparities and leading to misdiagnoses.

From the AI Governance in Practice Report 2025:

"Training datasets lacking in diversity can produce outputs that systematically underperform for certain groups... this can lead to inaccurate or biased outcomes in healthcare settings." (p. 41)

"Bias, discrimination and fairness challenge... inadequate or nonrepresentative training data can result in AI systems that propagate historical disparities." (p. 42) While physician oversight may reduce risk, biased data can still shape clinical decision-making.

* A- Economic benefit is not central to ethical risk here.

* C- Important but less critical than data representativeness.

* D- Error rate matters but is addressed via validation; it's not the core ethical issue.

NEW QUESTION: 74

Which risk management framework/guide/standard focuses on value-based engineering methodology?

A. ISO/IEC Guide 51 (Safety).

B. ISO 31000 Guidelines (Risk Management).

C. IEEE 7000-2021 Standard Model Process for Addressing Ethical Concerns during System Design.

D. Council of Europe Human Rights, Democracy, and the Rule of Law Assurance Framework (HUDERIA) for AI Systems.

Answer: (SHOW ANSWER)

The IEEE 7000-2021 Standard focuses on a value-based engineering methodology for addressing ethical concerns during system design. This standard guides engineers and organizations in integrating ethical considerations into the design and development processes of AI systems, ensuring that these technologies are developed responsibly and align with human values. Reference: AIGP Study Material, section on risk management frameworks and standards.

NEW QUESTION: 75

MULTI-SELECT

Please select 3 of the 5 options below. No partial credit will be given.

In cooperation with an Italian University, a US technology company called Nudge 2078, with no offices or employees in other countries, is developing an AI system which subtly encourages "good" behaviors, such as saving money and exercising. The two companies plan to conduct some experiments on volunteers. However, the study protocol is misleading on the purpose of the study, the experiments that will be conducted and its relevance to the application.

Which regulations are most likely applicable?

A. The EU AI Act.

B. The Digital Services Act.

C. The General Data Protection Act.

D. The Federal Trade Commission Act.

E. Payment Card Industry PCI Standards.

Answer: (SHOW ANSWER)

The correct answers are A, C, and D because multiple regulatory regimes apply based on jurisdiction, data use, and deceptive practices. The EU AI Act applies because the system is being developed and tested in cooperation with an EU-based university and involves behavioral

influence, which may fall under prohibited or high-risk AI practices, especially where manipulation or deception is present. The General Data Protection Regulation applies due to the processing of personal data of EU-based participants, particularly given the misleading study protocol, which violates principles such as transparency and informed consent. The Federal Trade Commission Act applies to the US company because deceptive or unfair practices, such as misleading participants about the purpose of experiments, fall under FTC enforcement. The Digital Services Act and PCI standards are not relevant to this specific use case.

NEW QUESTION: 76

All of the following analyses are part of a deactivating risk strategy EXCEPT?

- A. Understanding regulatory requirements related to data retention.
- B. Evaluating business continuity and financial risks.
- C. Understanding algorithmic methodology in AI model.
- D. Evaluating up and downstream system dependencies.

Answer: ([SHOW ANSWER](#))

The correct answer is C because understanding algorithmic methodology is primarily a design, development, and model evaluation activity rather than part of a deactivation or sunset risk strategy. A deactivating risk strategy focuses on safely retiring or discontinuing an AI system, ensuring that operational, legal, and business risks are properly managed. This includes evaluating regulatory requirements such as data retention and deletion obligations, assessing business continuity and financial implications, and understanding dependencies with upstream and downstream systems to avoid disruption. These steps ensure a controlled and compliant system shutdown. In contrast, analyzing algorithmic methodology is relevant during system development, validation, and explainability efforts, not during the decommissioning phase of the AI lifecycle.

Valid AIGP Dumps shared by EduDump.com for Helping Passing AIGP Exam! EduDump.com now offer the **newest AIGP exam dumps**, the EduDump.com AIGP exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AIGP dumps with Test Engine here: <https://www.edudump.com/exams/IAPP/AIGP/premium/> (219 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 77

What is the primary purpose of an AI impact assessment?

- A. To determine whether a conformity assessment is needed
- B. To escalate the findings to the appropriate owner(s)
- C. To identify and measure the benefits of an AI system
- D. To anticipate and manage the potential risks and harms of an AI system

Answer: D ([LEAVE A REPLY](#))

The correct answer is D. AI Impact Assessments are primarily used to identify and manage risks and harms associated with AI systems.

From the AIGP Body of Knowledge:

"The goal of an AI impact assessment is to ensure that risks are identified, evaluated, and mitigated prior to or during development and deployment." As further confirmed in the AI Governance in Practice Report 2025 (Part III):

"Risk-based tools like DPIAs and Algorithmic Impact Assessments help identify potential risks to individuals and society, enabling organizations to implement mitigation plans and safeguards." While benefits may be noted in such assessments, the core objective is to manage risks and promote responsible AI.

NEW QUESTION: 78

Which of the following is a foundational characteristic of effective AI governance?

- A. Engagement of a cross-functional team
- B. Reliance on tested vendor management processes
- C. Thorough reviews of a company's public filings with experts
- D. Uniform policies and procedures across developer, deployer and user roles

Answer: (SHOW ANSWER)

The correct answer is Engagement of a cross-functional team. Effective AI governance requires collaboration among various organizational functions including legal, compliance, IT, ethics, and data science.

From the AIGP Body of Knowledge:

"AI governance cannot be siloed—it requires input and oversight from across departments... A cross-functional team ensures that ethical, technical, legal, and operational risks are all appropriately managed." Also confirmed in the ILT Participant Guide:

"Cross-functional teams allow organizations to bring in different perspectives... Legal, compliance, and technical experts must work together to ensure responsible AI outcomes."

NEW QUESTION: 79

MULTI-SELECT

Please select 3 of the 5 options below. No partial credit will be given.

From a governance perspective, which of the following correctly describe the responsibilities of AI developers or deployers?

- A. AI developers are responsible for reducing biases in the system, such as discrimination against certain demographics.
- B. AI developers remain fully liable for the AI system's compliance with regulations even after deployment by others.
- C. AI developers must provide clear documentation explaining how the system functions for its operators.
- D. AI deployers must conduct independent bias assessments prior to release of a system.
- E. AI deployers must handle ongoing monitoring and auditing of the AI system.

Answer: (SHOW ANSWER)

The correct answers are A, C, and E because they reflect the shared but distinct responsibilities between AI developers and deployers in governance frameworks. Developers are responsible for designing systems that minimize bias and ensuring fairness during development. They are also required to provide clear documentation, including instructions for use, system capabilities, and limitations, to support transparency and proper operation. Deployers, on the other hand, are responsible for overseeing the system once it is in use, including continuous monitoring, auditing, and ensuring it performs as intended in real-world conditions.

Option B is incorrect because liability is typically shared and depends on roles and context, not solely on developers. Option D is incorrect because independent bias assessments are not always strictly required of deployers before release.

NEW QUESTION: 80

Which stakeholder is responsible for lawful collection of data for the training of the foundational AI model?

- A. The marketing agency.
- B. The tech company.
- C. The data aggregator.
- D. The marketing agency's client.

Answer: (SHOW ANSWER)

Data aggregators are third parties that collect and license data from various sources. They are responsible for ensuring the lawful collection and proper usage rights of the data they distribute - especially when such data is used to train foundational AI models.

From the AI Governance in Practice Report 2025:

"As organizations have neither proximity to how third-party data was first collected nor direct control over the data governance practices of third parties, an organization can benefit from carrying out its own legal due diligence and third-party risk management." (p. 19)

"Legal due diligence may include verification of the personal data's lawful collection by the data broker..." (p. 19) This confirms that data aggregators bear the legal and ethical burden to verify that data has been lawfully collected and is appropriately licensed for use, including in AI training.

A . The marketing agency and D. its client may use data, but they rely on upstream providers for its lawful origin.

B . The tech company may train the model but depends on lawful sourcing by data aggregators.

NEW QUESTION: 81

CASE STUDY

A company is considering the procurement of an AI system designed to enhance the security of IT infrastructure. The AI system analyzes how users type on their laptops, including typing speed, rhythm and pressure, to create a unique user profile. This data is then used to authenticate users and ensure that only authorized personnel can access sensitive resources.

When prioritizing the updates to its policies, rules and procedures to include the new AI system for user authentication, the organization should:

- A. Update third-party data sharing policies
- B. Update security controls for sensitive data
- C. Ensure that any personal data used is only processed for a specific and lawful purpose
- D. Reduce the complexity of the policy to make it easier for non-technical employees to understand

Answer: C ([LEAVE A REPLY](#))

The correct answer is C. This action ties directly into principles of data minimization, purpose limitation, and lawfulness of processing, which are central to privacy and AI governance.

From the AIGP Body of Knowledge, Section on Privacy Considerations:

"Personal data must only be processed for specified and lawful purposes. Organizations must consider whether they have a legal basis for processing such data under data protection laws like the GDPR or CCPA." Additionally, AI Governance in Practice Report 2025 emphasizes:

"One of the most significant challenges when designing and developing AI systems is ensuring the data used is appropriate for the intended purpose... Managing unnecessary data, especially data that may contain sensitive attributes, can increase risk."

NEW QUESTION: 82

A company has trained an ML model primarily using synthetic data, and now intends to use live personal data to test the model.

Which of the following is NOT a best practice apply during the testing?

- A. The test data should be representative of the expected operational data.
- B. Testing should minimize human involvement to the extent practicable.
- C. The test data should be anonymized to the extent practicable.
- D. Testing should be performed specific to the intended uses.

Answer: ([SHOW ANSWER](#)**)**

Minimizing human involvement to the extent practicable is not a best practice during the testing of an ML model. Human oversight is crucial during testing to ensure that the model performs correctly and ethically, and to interpret any anomalies or issues that arise. Best practices include using representative test data, anonymizing data to the extent practicable, and performing testing specific to the intended uses of the model.

Reference: AIGP Body of Knowledge on AI Model Testing and Human Oversight.

NEW QUESTION: 83

An app developer is creating a to-do app that compiles users' to-do lists, conducts basic online research about the tasks, prioritizes tasks, and allows users to add descriptions and notes.

However, during the app's download process, it requests the user's name, address, phone number, political and union affiliations.

Which privacy principle relevant to the responsible development of AI systems is most likely violated by this data collection practice?

- A. Data minimization and privacy by design.
- B. Purpose limitation and explicit consent.
- C. Transparency and explainability.
- D. Data portability and subject access rights.

Answer: A ([LEAVE A REPLY](#))

Data minimization requires organizations to collect and process only personal information that is adequate, relevant, and necessary for the specified purpose. A basic to-do application may reasonably require limited account or contact information, but political opinions and trade-union affiliations have no apparent necessity for compiling, researching, or prioritizing tasks. Those attributes are also treated as particularly sensitive under European data-protection law. Privacy by design requires privacy safeguards and data-protection principles to be incorporated into the product from the design stage and reflected in default settings.

Therefore, unnecessary collection at installation directly implicates both minimization and privacy by design.

Purpose limitation may also become relevant, but explicit consent does not by itself cure unjustified overcollection. Transparency, portability, and access rights address different privacy obligations and do not best characterize the defect described here.

NEW QUESTION: 84

CASE STUDY

A global marketing agency is adapting a large language model ("LLM") to generate content for an upcoming marketing campaign for a client's new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.

The marketing agency is accessing the LLM through an application programming interface ("API") developed by a third-party technology company. They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers. Both the marketing agency and the technology company have taken reasonable steps to address AI governance.

The marketing company has:

- * Entered into a contract with the technology company with suitable representations and warranties.
- * Completed an impact assessment on the LLM for this intended use.
- * Built technical guidance on how to measure and mitigate bias in the LLM.
- * Enabled technical aspects of transparency, explainability, robustness and privacy.
- * Followed applicable regulatory requirements.
- * Created specific legal statements and disclosures regarding the use of the AI on its client's advertising.

The technology company has:

- * Provided guidance and resources to developers to address environmental concerns.
- * Build technical guidance on how to measure and mitigate bias in the LLM.
- * Provided tools and resources to measure bias specific to the LLM.

- * Enabled technical aspects of transparency, explainability, robustness and privacy.
- * Mapped and mitigated potential societal harms and large-scale impacts.
- * Followed applicable regulatory requirements and industry standards.
- * Created specific legal statements and disclosures regarding the LLM. including with respect to IP and rights to data.

The agency has taken governance actions such as:

- * Conducting an impact assessment
- * Providing legal disclosures
- * Enabling bias mitigation and explainability
- * Complying with regulatory requirements

Which of the following should be included in the marketing company's disclosures about the use of the LLM EXCEPT?

- A. Intended purpose
- B. Proprietary methods
- C. Compliance with law
- D. Acknowledgement of limitations

Answer: (SHOW ANSWER)

The correct answer is B - Proprietary methods. While transparency is important, organizations are not obligated to disclose proprietary algorithms, methods, or trade secrets in public disclosures.

From the AIGP Body of Knowledge - Transparency & Disclosures:

"AI system users should disclose the purpose, capabilities, limitations, and applicable legal context-but not sensitive IP." AI Governance in Practice Report 2025 (Transparency Section) states:

"Disclosure requirements balance public understanding with the need to protect proprietary business interests.

Proprietary training methods are not expected to be disclosed."

Thus, while it's best practice to disclose the intended purpose, legal compliance, and system limitations, internal proprietary techniques are usually excluded.

NEW QUESTION: 85

The framework set forth in the White House Blueprint for an AI Bill of Rights addresses all of the following EXCEPT?

- A. Human alternatives, consideration and fallback.
- B. High-risk mitigation standards.
- C. Safe and effective systems.
- D. Data privacy.

Answer: B (LEAVE A REPLY)

The White House Blueprint for an AI Bill of Rights focuses on protecting civil rights, privacy, and ensuring AI systems are safe and effective. It includes principles like data privacy (D), human alternatives (A), and safe and effective systems (C). However, it does not specifically address high-risk mitigation standards as a distinct category (B).

NEW QUESTION: 86

Which of the following is an example of a high-risk application under the EU AI Act?

- A. A resume scanning tool that ranks applicants.
- B. An AI-enabled inventory management tool.
- C. A government-run social scoring tool.
- D. A customer service chatbot tool.

Answer: ([SHOW ANSWER](#))

The EU AI Act categorizes certain applications of AI as high-risk due to their potential impact on fundamental rights and safety. High-risk applications include those used in critical areas such as employment, education, and essential public services. A government-run social scoring tool, which assesses individuals based on their social behavior or perceived trustworthiness, falls under this category because of its profound implications for privacy, fairness, and individual rights. This contrasts with other AI applications like resume scanning tools or customer service chatbots, which are generally not classified as high-risk under the EU AI Act.

NEW QUESTION: 87

Which of the following disclosures is NOT required for an EU organization that developed and deployed a high-risk AI system?

- A. The human oversight measures employed.
- B. How an individual may contest a decision.
- C. The location(s) where data is stored.
- D. The fact that an AI system is being used.

Answer: ([SHOW ANSWER](#))

Under the EU AI Act, organizations that develop and deploy high-risk AI systems are required to provide several key disclosures to ensure transparency and accountability. These include the human oversight measures employed, how individuals can contest decisions made by the AI system, and informing individuals that an AI system is being used. However, there is no specific requirement to disclose the exact locations where data is stored. The focus of the Act is on the transparency of the AI system's operation and its impact on individuals, rather than on the technical details of data storage locations.

NEW QUESTION: 88

What is the most important reason to document the results of AI testing?

- A. To support post-deployment maintenance.
- B. To identify areas for red-teaming focus.
- C. To create a verifiable audit trail.
- D. To limit the need for future testing cycles.

Answer: ([SHOW ANSWER](#))

Testing results need to be documented thoroughly to ensure traceability, accountability, and compliance. This is central to enabling audits, investigations, or regulatory inquiries into the system's development and performance.

From the AI Governance in Practice Report 2025:

"Documentation and recordkeeping are essential components... to demonstrate AI system compliance, trace system behavior, and support audits and conformity assessments." (p. 34-35)

"Maintaining audit trails across development and deployment enables transparency and accountability." (p.

12)

* A and B are benefits, but not the primary governance justification.

* D- Limiting future testing is not a recommended goal.

NEW QUESTION: 89

Scenario:

A mid-sized tech firm is building its AI governance program and is exploring ISO/IEC standards that could support consistency in terminology and risk assessment processes across teams. ISO/IEC 22989 and ISO/IEC 42001 can be valuable resources for AI Governance professionals in all of the following ways EXCEPT:

- A. Establishing terminology and describing concepts so that governance team members can communicate with diverse parties and stakeholders from around the world
- B. Being applicable to organizations of any size and industry seeking to use AI responsibly and effectively in their design processes, information systems and controls
- C. Addressing specific issues related to managing procurement processes with third parties that provide or develop AI systems for their organization
- D. Recommending key activities to assess and manage risk: test, evaluate, verify and validate (TEVV)

Answer: C (LEAVE A REPLY)

The correct answer is C. ISO/IEC 22989 and 42001 focus on terminology, risk, and management systems, but do not specifically address procurement-related concerns with third-party vendors.

From the AIGP Body of Knowledge - Standards Section:

"ISO/IEC 22989 defines terminology and foundational concepts. ISO/IEC 42001 provides a management system standard for AI. They are not procurement-focused documents." Also confirmed in the AI Governance in Practice Report 2025:

"These standards help establish common language and risk governance procedures.

Procurement governance typically falls under separate frameworks or sector-specific guidance."

Thus, procurement governance (Option C) is not a central use case for these standards.

NEW QUESTION: 90

MULTI-SELECT

Please select 3 of the 5 options below. No partial credit will be given.

What are the roles and responsibilities of deployers of a proprietary model?

- A. Ethical testing.
- B. Ethical design.
- C. Technical performance.
- D. System documentation.
- E. Regulatory compliance.

Answer: ([SHOW ANSWER](#))

Deployers of proprietary models are not responsible for design, but they are accountable for how the system performs in their context of use, including ensuring ethical behavior, performance, and legal compliance.

From the AI Governance in Practice Report 2025:

"Deployers of AI systems must take reasonable steps to ensure that systems are used ethically, perform safely, and align with applicable laws and standards." (p. 11-12)

"Operational governance... includes performance monitoring protocols, incident management plans, and regulatory oversight." (p. 12) Thus:

- * #A. Ethical testing- Required to mitigate misuse and unintended harms.
- * #B. Ethical design- Belongs to developers/providers, not deployers.
- * #C. Technical performance- Deployers must ensure that AI performs as expected.
- * #D. System documentation- This is the provider's obligation.
- * #E. Regulatory compliance- Deployers must ensure system use complies with applicable laws.

NEW QUESTION: 91

What is the most important reason for documenting risks when developing an AI system?

- A. To provide transparency to stakeholders.
- B. To align with industry standards.
- C. To promote knowledge sharing.
- D. To mitigate potential liability.

Answer: ([SHOW ANSWER](#))

The most critical reason for documenting AI-related risks is to reduce exposure to legal, regulatory, and reputational liabilities. Clear documentation demonstrates that risks were identified, assessed, and addressed, which is essential for accountability and defensibility in the face of audits, litigation, or enforcement actions.

From the AI Governance in Practice Report 2025:

"An effective AI governance model is about collective responsibility... which should encompass oversight mechanisms such as privacy, accountability, compliance." (p. 13)

"Accountability... is based on the idea that there should be a person or entity that is ultimately responsible for any harm resulting from the use of the data, algorithm and AI system's underlying processes." (p. 28) While transparency, alignment with standards, and knowledge sharing are all secondary benefits, risk documentation's primary role is liability mitigation.

Valid AIGP Dumps shared by EduDump.com for Helping Passing AIGP Exam! EduDump.com now offer the **newest AIGP exam dumps**, the EduDump.com AIGP exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AIGP dumps with Test Engine here: <https://www.edudump.com/exams/IAPP/AIGP/premium/> (219 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 92

Scenario:

An organization is developing a powerful general-purpose AI (GPAI) model that has systemic impact. The compliance team is assessing what legal obligations apply under the EU AI Act. Under the EU AI Act, which of the following compliance actions applies only to General Purpose AI models with systemic risk?

- A. Publishing a detailed summary of the data used to train the model
- B. Maintaining up-to-date technical documentation, including testing details
- C. Implementing an intellectual property policy to comply with EU copyright laws
- D. Making information available to downstream providers who integrate the model into their AI systems

Answer: (SHOW ANSWER)

The correct answer is A. Only GPAI models with systemic risk must publish a detailed summary of training data to meet transparency and accountability standards under the EU AI Act.

From the AI Governance in Practice Report 2025 (EU AI Act Section):

"For GPAI systems with systemic risk, providers must publish sufficiently detailed summaries of the content used to train the model." Also, the AIGP ILT Guide confirms:

"The obligation to disclose summaries of training data applies only to systemic-risk GPAI models, not all general-purpose models or high-risk systems." This unique requirement is part of the Act's effort to increase transparency and auditability for powerful foundational models.

NEW QUESTION: 93

According to the Singapore Model AI Governance Framework, all of the following are recommended measures to promote the responsible use of AI EXCEPT?

- A. Determining the level of human involvement in algorithmic decision-making.
- B. Adapting the existing governance structure algorithmic decision-making.
- C. Employing human-over-the-loop protocols for high-risk systems.
- D. Establishing communications and collaboration among stakeholders.

Answer: (SHOW ANSWER)

The Singapore Model AI Governance Framework recommends several measures to promote the responsible use of AI, such as determining the level of human involvement in decision-making, adapting governance structures, and establishing communications and collaboration among stakeholders. However, employing human-over-the-loop protocols is not specifically mentioned in this framework. The focus is more on integrating human oversight appropriately within the

decision-making process rather than exclusively employing such protocols. Reference: AIGP Body of Knowledge, section on AI governance frameworks.

NEW QUESTION: 94

CASE STUDY

Please use the following answer the next question:

A local police department in the United States procured an AI system to monitor and analyze social media feeds, online marketplaces and other sources of public information to detect evidence of illegal activities (e.g., sale of drugs or stolen goods). The AI system works by surveilling the public sites in order to identify individuals that are likely to have committed a crime. It cross-references the individuals against data maintained by law enforcement and then assigns a percentage score of the likelihood of criminal activity based on certain factors like previous criminal history, location, time, race and gender.

The police department retained a third-party consultant assist in the procurement process, specifically to evaluate two finalists. Each of the vendors provided information about their system's accuracy rates, the diversity of their training data and how their system works. The consultant determined that the first vendor's system has a higher accuracy rate and based on this information, recommended this vendor to the police department.

The police department chose the first vendor and implemented its AI system. As part of the implementation, the department and consultant created a usage policy for the system, which includes training police officers on how the system works and how to incorporate it into their investigation process.

The police department has now been using the AI system for a year. An internal review has found that every time the system scored a likelihood of criminal activity at or above 90%, the police investigation subsequently confirmed that the individual had, in fact, committed a crime. Based on these results, the police department wants to forego investigations for cases where the AI system gives a score of at least 90% and proceed directly with an arrest.

During the procurement process, what is the most likely reason that the third-party consultant asked each vendor for information about the diversity of their datasets?

- A.** To comply with applicable law.
- B.** To assist the fairness of the AI system.
- C.** To evaluate the reliability of the AI system.
- D.** To determine the explainability of the AI system.

Answer: ([SHOW ANSWER](#))

The third-party consultant asked each vendor for information about the diversity of their datasets to assist in ensuring the fairness of the AI system. Diverse datasets help prevent biases and ensure that the AI system performs equitably across different demographic groups. This is crucial for a law enforcement application, where fairness and avoiding discriminatory practices are of paramount importance. Ensuring diversity in training data helps in building a more just and unbiased AI system. Reference: AIGP Body of Knowledge on Ethical AI and Fairness.

NEW QUESTION: 95

During the planning and design phases of the AI development life cycle, bias can be reduced by all of the following EXCEPT?

- A. Stakeholder involvement.
- B. Feature selection.
- C. Human oversight.
- D. Data collection.

Answer: ([SHOW ANSWER](#))

Bias in AI can be reduced during the planning and design phases through stakeholder involvement, human oversight, and careful data collection. While feature selection is critical in the development phase, it does not specifically occur during planning and design. Ensuring diverse stakeholder involvement and human oversight helps identify and mitigate potential biases early, and data collection ensures a representative dataset. Reference: AIGP Body of Knowledge on AI Development Lifecycle and Bias Mitigation.

NEW QUESTION: 96**CASE STUDY**

Please use the following answer the next question:

XYZ Corp., a premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent.

The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.

It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.

The organization has a large procurement team that is responsible for the contracting of technology solutions.

One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company are responsible for integrating and deploying technology solutions into the organization's operations in a responsible, cost-effective manner.

The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

Which of the following measures should XYZ adopt to best mitigate its risk of reputational harm from using the AI tool?

- A. Test the AI tool pre- and post-deployment.

- B. Ensure the vendor assumes responsibility for all damages.
- C. Direct the procurement team to select the most economical AI tool.
- D. Continue to require XYZ's hiring personnel to manually screen all applicants.

Answer: (SHOW ANSWER)

To mitigate the risk of reputational harm from using an AI hiring tool, XYZ Corp should rigorously test the AI tool both before and after deployment. Pre-deployment testing ensures the tool works correctly and does not introduce bias or other issues. Post-deployment testing ensures the tool continues to operate as intended and adapts to any changes in data or usage patterns. This approach helps to identify and address potential issues proactively, thereby reducing the risk of reputational harm. Ensuring the vendor assumes responsibility for damages (B) does not address the root cause of potential issues, selecting the most economical tool (C) may compromise quality, and continuing manual screening (D) defeats the purpose of using the AI tool.

Valid AIGP Dumps shared by EduDump.com for Helping Passing AIGP Exam! EduDump.com now offer the **newest AIGP exam dumps**, the EduDump.com AIGP exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AIGP dumps with Test Engine here: <https://www.edudump.com/exams/IAPP/AIGP/premium/> (219 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)