IAPP.AIGP.v2026-01-29.q82

Exam Code:	AIGP
Exam Name:	IAPP Certified Artificial Intelligence Governance Professional
Certification Provider:	IAPP
Free Question Number:	82
Version:	v2026-01-29
# of views:	102
# of Questions views:	820
https://www.freecram.net/torrent/IAPP.AIGP.v2026-01-29.q82.html	

NEW QUESTION: 1
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 toreduce exposure to legal, regulatory, and reputational liabilities. Clear documentation demonstrates thatrisks were identified, assessed, and addressed, which is essential for accountability and defensibility in the face of audits, litigation, or enforcement actions.
From theAI Governance in Practice Report 2024:
"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.

NEW QUESTION: 2
According to the GDPR, an individual has the right to have a human confirm or replace an automated decision unless that automated decision?
A. Is authorized with the data subject s explicit consent.
B. Is authorized by applicable Ell law and includes suitable safeguards.
C. Is deemed to solely benefit the individual and includes documented legitimate interests.
D. Is necessary for entering into or performing under a contract between the data subject and data controller.

Answer: (SHOW ANSWER)
According to the GDPR, individuals have the right to not be subject to a decision based solely on automated processing, including profiling, which produces legal effects or similarly significantly affects them. However, there are exceptions to this right, one of which is when the decision is based on the data subject's explicit consent. This means that if an individual explicitly consents to the automated decision-making process, there is no requirement for human intervention to confirm or replace the decision. This exception ensures that individuals can have control over automated decisions that affect them, provided they have given clear and informed consent.

NEW QUESTION: 3
Machine learning is best described as a type of algorithm by which?
A. Systems can mimic human intelligence with the goal of replacing humans.
B. Systems can automatically improve from experience through predictive patterns.
C. Statistical inferences are drawn from a sample with the goal of predicting human intelligence.
D. Previously unknown properties are discovered in data and used to predict and make improvements in the data.
Answer: (SHOW ANSWER)
Machine learning (ML) is a subset of artificial intelligence (AI) where systems use data to learn and improve over time without being explicitly programmed. Option B accurately describes machine learning by stating that systems can automatically improve from experience through predictive patterns. This aligns with the fundamental concept of ML where algorithms analyze data, recognize patterns, and make decisions with minimal human intervention. Reference: AIGP BODY OF KNOWLEDGE, which covers the basics of AI and machine learning concepts.

NEW QUESTION: 4
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 bedocumented thoroughlyto ensuretraceability, accountability, and compliance.
This is central to enabling audits, investigations, or regulatory inquiries into the system's development and performance.
From theAI Governance in Practice Report 2024:
"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)

* AandBare benefits, but not theprimary governance justification.
* D- Limiting future testing is not a recommended goal.

NEW QUESTION: 5
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: 6
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: 7
During the development of semi-autonomous vehicles, various failures occurred as a result of the sensors misinterpreting environmental surroundings, such as sunlight.

These failures are an example of?
A. Hallucination.
B. Brittleness.
C. Uncertainty.
D. Forgetting.
Answer: B (LEAVE A REPLY)
The failures in semi-autonomous vehicles due to sensors misinterpreting environmental surroundings, such as sunlight, are examples of brittleness. Brittleness in AI systems refers to their inability to handle variations in input data or unexpected conditions, leading to failures when the system encounters situations that were not adequately covered during training. These systems perform well under specific conditions but fail when those conditions change. Reference: AIGP Body of Knowledge on AI System Robustness and Failures.

NEW QUESTION: 8
A company is working to develop a self-driving car that can independently decide the appropriate route to take the driver after the driver provides an address.
If they want to make this self-driving car "strong" AI, as opposed to "weak," the engineers would also need to ensure?
A. Thatthe AI has full human cognitive abilities that can independently decide where to take the driver.
B. That they have obtained appropriate intellectual property (IP) licenses to use data for training the AI.
C. That the AI has strong cybersecurity to prevent malicious actors from taking control of the car.
D. That the AI can differentiate among ethnic backgrounds of pedestrians.
Answer: (SHOW ANSWER)
Strong AI, also known as artificial general intelligence (AGI), refers to AI that possesses the ability to understand, learn, and apply intelligence across a broad range of tasks, similar to human cognitive abilities.
For the self-driving car to be classified as "strong" AI, it would need to possess full human cognitive abilities to make independent decisions beyond pre-programmed instructions.
Reference: AIGP BODY OF KNOWLEDGE and AI classifications.

NEW QUESTION: 9
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.

Which stakeholder group is most important in selecting the specific type of algorithm?

A. The cloud provider.

B. The consulting firm.

C. The healthcare network'sdata science team.

D. The healthcare network's AI governance committee.

Answer: (SHOW ANSWER)

In selecting the specific type of algorithm for the AI solution, the healthcare network's data science team is most important. This team possesses the technical expertise and understanding of the data, the clinical context, and the performance requirements needed to make an informed decision about which algorithm is most suitable. While the cloud provider and consulting firm can offer support and infrastructure, and the AI governance committee provides oversight, the data science team's specialized knowledge is crucial for selecting and implementing the appropriate algorithm. Reference: AIGP Body of Knowledge, AI governance and team roles section.

NEW QUESTION: 10

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 2024 (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: 11
Which type of existing assessment could best be leveraged to create an AI impact assessment?
A. A safety impact assessment.
B. A privacy impact assessment.
C. A security impact assessment.
D. An environmental impact assessment.
Answer: (SHOW ANSWER)
A privacy impact assessment (PIA) can be effectively leveraged to create an AI impact assessment. A PIA evaluates the potential privacy risks associated with the use of personal data and helps in implementing measures to mitigate those risks. Since AI systems often involve processing large amounts of personal data, the principles and methodologies of a PIA are highly applicable and can be extended to assess broader impacts, including ethical, social, and legal implications of AI. Reference: AIGP Body of Knowledge on Impact Assessments.

NEW QUESTION: 12
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.
The data processed by the AI system would be classified as:
A. Non-sensitive personal data, since it does not reveal information about health, gender or race
B. Organizational data, since it is part of the authentication process
C. Non-personal data, as long as it is not linked to a user ID
D. Special category data, if it can be used to uniquely identify a person

Answer: (SHOW ANSWER)
The correct answer isD.Keystroke dynamics, used to identify individuals, fall underbiometric data, which is aspecial category of personal dataunder the GDPR and other frameworks.
From the AI Governance in Practice Report 2024:
"Keystroke dynamics may constitute biometric data if used to uniquely identify an individual... Biometric data is classified as special category personal data and requires higher protection standards." Also reflected in ILT Participant Guide:
"Biometric data, such as facial images, voiceprints, iris scans or keystroke patterns, are treated as special category data when they are used for the purpose of uniquely identifying individuals."

NEW QUESTION: 13
An AI system's function, the industry and the location in which it operates are important factors in considering which of the following?
A. Organizational accountability.
B. Internal governance needs.
C. Diversity of data sources.
D. Explainability of results.
Answer: (SHOW ANSWER)
An AI system'sfunction,industry, anddeployment locationdefine itsrisk profile, which directly influences theinternal governance structuresan organization must put in place.
From theAI Governance in Practice Report 2024:
"There are many challenges and potential solutions for AI governance, each with unique proximityand significance based on an organization's role, footprint, broader risk-governance profile and maturity." (p. 4)
"AI governance starts with defining the corporate strategy for AI... and formulating policy standards and operational procedures to reflect industry, use case, and location." (p. 11)
* A- Organizational accountability is broader and not directly scoped by industry or function.
* C- Diversity of data sources is tied to data strategy.
* D- Explainability is more influenced by model type, not use context.

NEW QUESTION: 14
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 causessignificant false positives, especially in sensitive contexts likefraud detection, the priority is tohalt harmful activityand perform a full assessment. Continued use without understanding the fault may cause furthercustomer harmand legal exposure.
From theAI Governance in Practice Report 2024:
"Incident management plans should enable identification, escalation, and system rollback to prevent continued harm from malfunctioning AI systems." (p. 12, 35)

NEW QUESTION: 15
You are an engineer that developed an AI-based ad recommendation tool.
Which of the following should be monitored to evaluate the tool's effectiveness?
A. Output data, assess the delta between the prediction and actual ad clicks.
B. Algorithmic patterns, to show the model has a high degree of accuracy.
C. Input data, to ensure the ads are reaching the target audience.
D. GPU performance, to evaluate the tool's robustness.
Answer: (SHOW ANSWER)
To evaluate the effectiveness of an AI-based ad recommendation tool, the most relevant metric is the output data, specifically assessing the delta between the prediction and actual ad clicks. This metric directly measures the tool's accuracy and effectiveness in making accurate recommendations that lead to user engagement. While monitoring algorithmic patterns and input data can provide insights into the model's behavior and targeting accuracy, and GPU performance can indicate the robustness and efficiency of the tool, the primary indicator of effectiveness for an ad recommendation tool is how well it predicts actual ad clicks.
Reference: AIGP BODY OF KNOWLEDGE, sections on AI performance metrics and evaluation methods.

NEW QUESTION: 16
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.
The best human oversight mechanism for the police department to implement is that a police officer should?
A. Explain to the accused how the AI system works.
B. Confirm the AI recommendation prior to sentencing.
C. Ensure an accused is given notice that the AI system was used.
D. Consider the AI recommendation as part of the criminal investigation.
Answer: (SHOW ANSWER)
The best human oversight mechanism for the police department to implement is for a police officer to consider the AI recommendation as part of the criminal investigation. This ensures that the AI system's output is used as a tool to aid human decision-making rather than replace it. The police officer should integrate the AI's insights with other evidence and contextual information to make informed decisions, maintaining a balance between technological aid and human judgment.
Reference: AIGP Body of Knowledge on AI Integration and Human Oversight.

NEW QUESTION: 17
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 harmsisnot a required or standard practicefor internal performance evaluation. Risk assessments and reporting typically remain internal or shared with regulators.

From theAI Governance in Practice Report 2024:

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

NEW QUESTION: 18

You are a privacy program manager at a large e-commerce company that uses an AI tool to deliver personalized product recommendations based on visitors' personal information that has been collected from the company website, the chatbot and public data the company has scraped from social media.

A user submits a data access request under an applicable U.S. state privacy law, specifically seeking a copy of their personal data, including information used to create their profile for product recommendations.

What is the most challenging aspect of managing this request?

A. Some of the visitor's data is synthetic data that the company does not have to provide to the data subject.

B. The data subject's data is structured data that can be searched, compiled and reviewed only by an automated tool.

C. The data subject is not entitled to receive a copy of their data because some of it was scraped from public sources.

D. Some of the data subject's data is unstructured data and you cannot untangle it from the other data, including information about other individuals.

Answer: (SHOW ANSWER)

The most challenging aspect of managing a data access request in this scenario is dealing with unstructured data that cannot be easily disentangled from other data, including information about other individuals.

Unstructured data, such as free-text inputs or social media posts, often lacks a clear structure and may be intermingled with data from multiple individuals, making it difficult to isolate the specific data related to the requester. This complexity poses significant challenges in complying with data access requests under privacy laws. Reference: AIGP Body of Knowledge on Data Subject Rights and Data Management.

NEW QUESTION: 19

The processes and methods that allow human users to understand and trust the outputs produced by AI are important in addressing which key regulatory concern?

A. Interpretable AI

B. Trustworthy AI
C. Explainable AI
D. Responsible AI
Answer: (SHOW ANSWER)
The correct answer isExplainable AIbecause it specifically refers to theability of a system todescribe the logic behind its decisions or outputsin a way that is understandable to humans. This is a key part of regulatory and ethical frameworks and is directly related to addressing theblack-box problemin AI.
From the AIGP ILT Participant Guide (Module on Transparency and Explainability):
"Explainability refers to the understanding of how a black-box model works. The black-box problem exists because some models are too complex for human interpretation. Explainability methods aim to provide meaningful insight into the logic and decision-making of AI systems."
Also, according to the AI Governance in Practice Report 2024:
"Explainability refers to the representation of the underlying mechanisms of the AI system's operation... a key tenet of AI governance due to the desire to understand how AI systems are built, managed and maintained." Thus, while Trustworthy and Responsible AI are broader concepts,explainabilityspecifically targets the regulatory concern about understanding outputs.

NEW QUESTION: 20

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 2024 (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: 21
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 2024 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: 22
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 2024 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: 23
What is the term for an algorithm that focuses on making the best choice achieve an immediate objective at a particular step or decision point, based on the available information and without regard for the longer-term best solutions?
A. Single-lane.
B. Optimized.
C. Efficient.
D. Greedy.
Answer: D (LEAVE A REPLY)
A greedy algorithm is one that makes the best choice at each step to achieve an immediate objective, without considering the longer-term consequences. It focuses on local optimization at each decision point with the hope that these local solutions will lead to an optimal global solution. However, greedy algorithms do not always produce the best overall solution for certain problems, but they are useful when an immediate, locally optimal solution is desired. Reference: AIGP Body of Knowledge, algorithm types section.

NEW QUESTION: 24
All of the following are elements of establishing a global AI governance infrastructure EXCEPT?
A. Providing training to foster a culture that promotes ethical behavior.
B. Creating policies and procedures to manage third-partyrisk.
C. Understanding differences in norms across countries.
D. Publicly disclosing ethical principles.
Answer: D (LEAVE A REPLY)
Establishing a global AI governance infrastructure involves several key elements, including providing training to foster a culture that promotes ethical behavior, creating policies and procedures to manage third-party risk, and understanding differences in norms across countries. While publicly disclosing ethical principles can enhance transparency and trust, it is not a core element necessary for the establishment of a governance infrastructure. The focus is more on internal processes and structures rather than public disclosure. Reference:
AIGP Body of Knowledge on AI Governance and Infrastructure.

NEW QUESTION: 25
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: C (LEAVE A REPLY)

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: 26
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 adeployerto also acting as aprovider or developer, it assumesnew obligationsunder regulations like the EU AI Act. One of the core requirements for providers is to produce and maintaintechnical documentation, including descriptions of the model, associated risks, and mitigation strategies.
From theAI Governance in Practice Report 2024:
"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: 27
Why is it important that conformity requirements are satisfied before an AI system is released into production?
A. To ensure the visual design is fit-for-purpose.
B. To ensure the AI system is easy for end-users to operate.
C. To guarantee interoperability of the AI system across multiple platforms and environments.
D. To comply with legal and regulatory standards, ensuring the AI system is safe and trustworthy.
Answer: (SHOW ANSWER)
Conformity assessmentsare a core requirement under theEU AI Actfor high-risk systems and serve to confirm that the AI meetsregulatory, safety, and ethical standardsbefore it is put into production.
From theAI Governance in Practice Report 2024:

"Conformity assessments... ensure that systems comply with legal requirements, safety criteria, and intended purpose before being placed on the market." (p. 34)
"They are a critical step to demonstrate safety and trustworthiness in AI deployment." (p. 35)

NEW QUESTION: 28
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 2024:
"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: 29
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 operationaldata.
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: 30

All of the following are included within the scope of post-deployment AI maintenance EXCEPT?

A. Ensuring that all model components are subject a control framework.

B. Dedicating experts to continually monitor the model output.

C. Evaluating the need for an audit under certain standards.

D. Defining thresholds to conduct new impact assessments.

Answer: (SHOW ANSWER)

Post-deployment AI maintenance typically includes ensuring that all model components are subject to a control framework, dedicating experts to continually monitor the model output, and evaluating the need for audits under certain standards. However, defining thresholds to conduct new impact assessments is usually part of the initial deployment and ongoing governance processes rather than a maintenance activity.

Maintenance focuses more on the operational aspects of the AI system rather than setting new thresholds for impact assessments.

Reference: AIGP BODY OF KNOWLEDGE, sections discussing AI lifecycle management and post-deployment activities.

NEW QUESTION: 31

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 isEngagement 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: 32
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.
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.
The frameworks that would be most appropriate for XYZ's governance needs would be the NIST AI Risk Management Framework and?
A. NIST Information Security Risk (NIST SP 800-39).
B. NIST Cyber Security Risk Management Framework (CSF 2.0).
C. IEEE Ethical System Design Risk Management Framework (IEEE 7000-21).
D. Human Rights, Democracy, and Rule of Law Impact Assessment (HUDERIA).
Answer: (SHOW ANSWER)
The IEEE Ethical System Design Risk Management Framework (IEEE 7000-21) would be most appropriate for XYZ Corp's governance needs in addition to the NIST AI Risk ManagementFramework. The IEEE framework specifically addresses ethical concerns during system design, which is crucial for ensuring the responsible use of AI in hiring. It complements the NIST framework by focusing on ethical risk management, aligning well with XYZ Corp's goals of deploying AI responsibly and mitigating associated risks.

NEW QUESTION: 33

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: C (LEAVE A REPLY)
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: 34
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)
Theinitial deployment phaseof an AI model is critical forpost-deployment monitoring. When tracking for bias, the most important task is tocontinue disparity testingto determine whether outputs differ across protected groups.
From theAI Governance in Practice Report 2024:
"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 trainingis helpful, but not a technical bias mitigation activity.
* C. Analyzing training datais apre-deploymenttask.
* D. Documenting human decisionsmay support auditability but doesn't detect bias in AI outputs.

NEW QUESTION: 35
Which of the following would be the least likely step for an organization to take when designing an integrated compliance strategy for responsible AI?
A. Conducting an assessment of existing compliance programs to determine overlaps and integration points.
B. Employing a new software platform to modernize existing compliance processes across the organization.
C. Consulting experts to consider the ethical principles underpinning the use of AI within the organization.

D. Launching a survey to understand the concerns and interests of potentially impacted stakeholders.

Answer: (SHOW ANSWER)

When designing an integrated compliance strategy for responsible AI, the least likely step would be employing a new software platform to modernize existing compliance processes. While modernizing compliance processes is beneficial, it is not as directly related to the strategic integration of ethical principles and stakeholder concerns. More critical steps include conducting assessments of existing compliance programs to identify overlaps and integration points, consulting experts on ethical principles, and launching surveys to understand stakeholder concerns. These steps ensure that the compliance strategy is comprehensive and aligned with responsible AI principles. Reference: AIGP Body of Knowledge on AI Governance and Compliance Integration.

NEW QUESTION: 36

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: (SHOW ANSWER)

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: 37

What is the primary reason the EU is considering updates to its Product Liability Directive?

A. To increase the minimum warranty level for defective goods.

B. To define new liability exemptions for defective products.

C. Address digital services and connected products.

D. Address free and open-source software.

Answer: (SHOW ANSWER)

The primary reason the EU is considering updates to its Product Liability Directive is to address digital services and connected products. The current directive does not adequately cover the

complexities and challenges posed by modern digital and connected technologies. By updating the directive, the EU aims to ensure that it remains relevant and effective in addressing the liabilities associated with these advanced products, ensuring consumer protection and fair market practices in the digital age.

NEW QUESTION: 38
A UK company has designed a facial recognition model to support border control.
The EU AI Act would apply to the model in all of the following situations EXCEPT if?
A. The model was released under an open source license.
B. The model is deployed at an EU border checkpoint.
C. The model is deployed at UK border checkpoints.
D. The model was trained by an EU company.
Answer: (SHOW ANSWER)
TheEU AI Actapplies extraterritorially, meaning it affects entitiesoutside the EUwhen their AI systems impact individuals within the EU. However, it doesnotapply to systems that are developed, sold, or used entirely outside of the EU-such as in theUK, unless they affect the EU market or individuals.
From theAI Governance in Practice Report 2024:
"The act imposes regulatory obligations... depending on their capabilities, reach and computing power, certain GPAI systems are considered to present systemic risk and attract broadly similar obligations to those applicable to high-risk AI systems." (p. 7)
"The EU AI Act is the world's first comprehensive AI regulation... requirements apply to providers, deployers, importers, and distributors of AI systems when such systems are placed on the EU market." (p. 7-8) Thus:
* A. Open source releasedoes not exclude applicability if deployed in the EU.
* B. Deployment at an EU borderclearly invokes jurisdiction.
* D. Training by an EU companycreates jurisdictional links.
* C. Deployment only at UK checkpoints, withno EU use or impact, isoutside scope.

NEW QUESTION: 39
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.
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 other stakeholder groups should be involved in the selection and implementation of the AI hiring tool?
A. Finance and Legal.
B. Marketing and Compliance.
C. Supply Chain and Marketing.
D. Litigation and Product Development.
Answer: (SHOW ANSWER)
In the selection and implementation of the AI hiring tool, involving Finance and Legal is crucial. The Finance team is essential for assessing cost implications, budget considerations, and financial risks. The Legal team is necessary to ensure compliance with applicable laws and regulations, including those related to data privacy, employment, and anti-discrimination. Involving these stakeholders ensures a comprehensive evaluation of both the financial viability and legal compliance of the AI tool, mitigating potential risks and aligning with organizational objectives and regulatory requirements.

NEW QUESTION: 40
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 thelawful collectionandproper usage rightsof the data they distribute - especially when such data is used to train foundational AI models.
From theAI Governance in Practice Report 2024:
"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 databroker..." (p.
19)
This confirms thatdata aggregatorsbear 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 agencyandD. its clientmay use data, but they rely on upstream providers for its lawful origin.
* B. The tech companymay train the model but depends on lawful sourcing by data aggregators.

NEW QUESTION: 41
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.
What is the best reason for GVC to offer students the choice to utilize generative AI in limited, defined circumstances?
A. Toenable students to learn how to manage their time.
B. Toenable students to learn about performing research.
C. Toenable students to learn about practical applications of AI.
D. Toenable students to learn how to use AI as a supportive educational tool.
Answer: (SHOW ANSWER)
The best reason for GVC to offer students the choice to utilize generative AI in limited, defined circumstances is to enable students to learn how to use AI as a supportive educational tool. By integrating AI in a controlled manner, students can learn the practical applications of AI and develop skills to use AI responsibly and effectively in their educational pursuits.
Reference: The AIGP Body of Knowledge highlights the importance of teaching students about AI's practical applications and the responsible use of AI technologies. This aligns with the goal of fostering a better understanding of AI's role and its potential benefits in various contexts, including education.

NEW QUESTION: 42
Scenario:

A company using AI for resume screening understands the risks of algorithmic bias and the evolving legal requirements across jurisdictions. It wants to implement the right governance controls to prevent reputational damage from misuse of the AI hiring tool.
Which of the following measures should the company 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 provides indemnification for the AI tool
C. Require the procurement and deployment teams to agree upon the AI tool
D. Continue to require the company's hiring personnel to manually screen all applicants
Answer: (SHOW ANSWER)
The correct answer is A. Pre- and post-deployment testing ensures bias, accuracy, and fairness are evaluated and corrected as needed, which is essential for reputational risk mitigation.
From the AIGP Body of Knowledge:
"Testing AI systems before and after deployment is critical to ensure performance, fairness, and compliance.
Failing to do so may result in reputational damage and legal exposure." AI Governance in Practice Report 2024 (Bias/Fairness and Risk Sections):
"System impact assessments, testing, and post-deployment monitoring are necessary to identify and mitigate risks... This supports both compliance and public trust." Testing is proactive, unlike indemnification (which transfers risk after damage), or requiring manual review (which defeats automation).

NEW QUESTION: 43
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: 44
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: 45
All of the following are examples of biometric data in the US EXCEPT?
A. Iris scans.
B. Walking gait.
C. Keystroke dynamics.
D. GPS location of a user's fitness watch.
Answer: (SHOW ANSWER)
Biometric data in the U.S. refers to data that relates to measurable biological and behavioral characteristics that can be used to identify an individual. Examples include fingerprints, facial recognition, iris scans, and behavior-based data like gait or keystrokes.
According to definitions and discussions from theAI Governance in Practice Report 2024and U.S. privacy frameworks:
"Biometric data includes physical and behavioral human characteristics that can be used to digitally identify a person to grant access to systems, devices, or data. Examples include facial images, iris patterns, gait analysis, and voice recognition." (Report context based on common frameworks in U.S. AI law and the use of biometrics in AI governance.) Here's how the options relate:
* A. Iris scans- These are physical biometric identifiers.
* B. Walking gait- Behavioral biometric used increasingly in surveillance and identification.

* C. Keystroke dynamics- Behavioral biometric based on typing patterns.
* D. GPS location of a user's fitness watch- This isnotbiometric data. It islocation data, which may be sensitive or personal, but not biometric.

NEW QUESTION: 46
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: 47
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 isA. TheAcceptable 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 2024 (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: 48
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, what is the most important step for the healthcare network to take when mapping its existing data to the clinical research partner data?
A. Apply privacy-enhancing technologies to the data.
B. Identify fits and gaps in the combined data.
C. Ensure the data is labeled and formatted.
D. Evaluate the country of origin of the data.
Answer: (SHOW ANSWER)
In the design phase of integrating data from different sources, identifying fits and gaps is crucial. This process involves understanding how well the data from the clinical research partner aligns with the healthcare network's existing data. It ensures that the combined data set is coherent and can be effectively used for training the AI algorithm. This step helps in spotting any discrepancies, inconsistencies, or missing data that might affect the performance and accuracy of the AI model. It directly addresses the integrity and compatibility of the data, which is foundational before applying any privacy-enhancing technologies, labeling, or evaluating the origin of the data.
Reference: AIGP Body of Knowledge on Data Integration and Quality.

NEW QUESTION: 49

Which of the following arenotconsidered biometric data under U.S. privacy laws?

A. Iris scans

B. Walking gait

C. Keystroke dynamics

D. GPS location of a user's fitness watch

Answer: (SHOW ANSWER)

The correct answer isD. GPS location data isnot biometric data-it is consideredgeolocation data, which is personal data butnot biometricunder most U.S. laws.

From the AIGP ILT Guide (Data Privacy Module):

"Biometric data includes measurable biological or behavioral characteristics such as iris scans, facial recognition, voice prints, and keystroke patterns when used for identification." AI Governance in Practice Report 2024 (Privacy and Data Protection section):

"Location data, while sensitive, is not considered biometric unless it's tied to a uniquely identifyingbiological trait." Thus,GPS location data, while potentially sensitive, is not classified as biometric.

NEW QUESTION: 50

Each of the following actors are typically engaged in the AI development life cycle EXCEPT?

A. Data architects.

B. Government regulators.

C. Socio-cultural and technical experts.

D. Legal and privacy governance experts.

Answer: (SHOW ANSWER)

Typically, actors involved in the AI development life cycle include data architects (who design the data frameworks), socio-cultural and technical experts (who ensure the AI system is socio-culturally aware and technically sound), and legal and privacy governance experts (who handle the legal and privacy aspects).

Government regulators, while important, are not directly engaged in the development process but rather oversee and regulate the industry. Reference: AIGP BODY OF KNOWLEDGE and AI development frameworks.

NEW QUESTION: 51

You are a privacy program manager at a large e-commerce company that uses an AI tool to deliver personalized product recommendations based on visitors' personal information that has been collected from the company website, the chatbot and public data the company has scraped from social media.

A user submits a data access request under an applicable U.S. state privacy law, specifically seeking a copy of their personal data, including information used to create their profile for productrecommendations.

What is the most challenging aspect of managing this request?

A. Some of the data subject's data is unstructured data and you cannot untangle it from the other data, including information about other individuals.
B. The data subject is not entitled to receive a copy of their data because some of it was scraped from public sources.
C. Some of the visitor's data is synthetic data that the company does not have to provide to the data subject.
D. The data subject's data is structured data that can be searched, compiled and reviewed only by an automated tool.
Answer: (SHOW ANSWER)
The most challenging aspect of managing a data access request in this scenario is dealing with unstructured data that cannot be easily disentangled from other data, including information about other individuals.
Unstructured data, such as free-text inputs or social media posts, often lacks a clear structure and may be intermingled with data from multiple individuals, making it difficult to isolate the specific data related to the requester. This complexity poses significant challenges in complying with data access requests under privacy laws. Reference: AIGP Body of Knowledge on Data Subject Rights and Data Management.

NEW QUESTION: 52
The White House Executive Order from November 2023 requires companies that develop dual-use foundation models to provide reports to the federal government about all of the following EXCEPT?
A. Any current training or development of dual-use foundation models.
B. The results of red-team testing of each dual-use foundation model.
C. Any environmental impact study for each dual-use foundation model.
D. The physical and cybersecurity protection measures of their dual-use foundation models.
Answer: (SHOW ANSWER)
The White House Executive Order from November 2023 requires companies developing dual-use foundation models to report on their current training or development activities, the results of red-team testing, and the physical and cybersecurity protection measures. However, it does not mandate reports on environmental impact studies for each dual-use foundation model. While environmental considerations are important, they are not specified in this context as a reporting requirement under this Executive Order.
Reference: AIGP BODY OF KNOWLEDGE, sections on compliance and reporting requirements, and the White House Executive Order of November 2023.

NEW QUESTION: 53
After completing model testing and validation, which of the following is the most important step that an organization takes prior to deploying the model into production?
A. Perform a readiness assessment.
B. Define a model-validation methodology.

C. Document maintenance teams and processes.
D. Identify known edge cases to monitor post-deployment.
Answer: (SHOW ANSWER)
After completing model testing and validation, the most important step prior to deploying the model into production is to perform a readiness assessment. This assessment ensures that the model is fully prepared for deployment, addressing any potential issues related to infrastructure, performance, security, and compliance.
It verifies that the model meets all necessary criteria for a successful launch. Other steps, such as defining a model-validation methodology, documenting maintenance teams and processes, and identifying known edge cases, are also important but come secondary to confirming overall readiness. Reference: AIGP Body of Knowledge on Deployment Readiness.

NEW QUESTION: 54
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.
Which AI risk would NOT have been identified during the procurement process based on the categories of information requested by the third-party consultant?
A. Security.
B. Accuracy.
C. Explainability.

D. Discrimination.

Answer: (SHOW ANSWER)

The AI risk that would not have been identified during the procurement process based on the categories of information requested by the third-party consultant is security. The consultant focused on accuracy rates, diversity of training data, and system functionality, which pertain to performance and fairness but do not directly address the security aspects of the AI system. Security risks involve ensuring that the system is protected against unauthorized access, data breaches, and other vulnerabilities that could compromise its integrity. Reference: AIGP Body of Knowledge on AI Security and Risk Management.

NEW QUESTION: 55

Which of the following steps occurs in the design phase of the AI life cycle?

A. Data augmentation.

B. Model explainability.

C. Risk impact estimation.

D. Performance evaluation.

Answer: (SHOW ANSWER)

Risk impact estimation occurs in the design phase of the AI life cycle. This step involves evaluating potential risks associated with the AI system and estimating their impacts to ensure that appropriate mitigation strategies are in place. It helps in identifying and addressing potential issues early in the design process, ensuring the development of a robust and reliable AI system. Reference: AIGP Body of Knowledge on AI Design and Risk Management.

NEW QUESTION: 56

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: 57

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 tA. 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.
Each of the following steps would support fairness testing by the compliance team during the first month in production EXCEPT?
A. Validating a similar level of decision-making across different demographic groups.
B. Providing the loan applicants with information about the model capabilities and limitations.
C. Identifying if additional training data should be collected for specific demographic groups.
D. Using tools to help understand factors that may account for differences in decision-making.
Answer: (SHOW ANSWER)
Providing the loan applicants with information about the model capabilities and limitations would not directly support fairness testing by the compliance team. Fairness testing focuses on evaluating the model's decisions for biases and ensuring equitable treatment across different demographic groups, rather than informing applicants about the model.
Reference: The AIGP Body of Knowledge outlines that fairness testing involves technical assessments such as validating decision-making consistency across demographics and using tools to understand decision factors. While transparency to applicants is important for ethical AI use, it does not contribute directly to the technical process of fairness testing.

NEW QUESTION: 58
What type of organizational risk is associated with AI's resource-intensive computing demands?
A. People risk.
B. Security risk.
C. Third-party risk.
D. Environmental risk.
Answer: (SHOW ANSWER)
AI's resource-intensive computing demands pose significant environmental risks. High-performance computing required for training and deploying AI models often leads to substantial energy consumption, which can result in increased carbon emissions and other environmental

impacts. This is particularly relevant given the growing concern over climate change and the environmental footprint of technology. Organizations need to consider these environmental risks when developing AI systems, potentially exploring more energy-efficient methods and renewable energy sources to mitigate the environmental impact.

NEW QUESTION: 59
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: 60
You are part of your organization's ML engineering team and notice that the accuracy of a model that was recently deployed into production is deteriorating.
What is the best first step address this?

A. Replace the model with a previous version.
B. Conduct champion/challenger testing.
C. Perform an audit of the model.
D. Run red-teaming exercises.
Answer: (SHOW ANSWER)
When the accuracy of a model deteriorates, the best first step is to conduct champion/challenger testing. This involves deploying a new model (challenger) alongside the current model (champion) to compare their performance. This method helps identify if the new model can perform better under current conditions without immediately discarding the existing model. It provides a controlled environment to test improvements and understand the reasons behind the deterioration. This approach is preferable to directly replacing the model, performing audits, or running red-teaming exercises, which may be subsequent steps based on the findings from the champion/challenger testing.

Reference: AIGP BODY OF KNOWLEDGE, sections on model performance management and testing strategies.

NEW QUESTION: 61
An EU bank intends to launch a multi-modal AI platform for customer engagement and automated decision-making assist with the opening of bank accounts. The platform has been subject to thorough risk assessments and testing, where it proves to be effective in not discriminating against any individual on the basis of a protected class.
What additional obligations must the bank fulfill prior to deployment?
A. The bank must obtain explicit consent from users under the privacy Directive.
B. The bank must disclose how the AI system works under the EII Digital Services Act.
C. The bank must subject the AI system an adequacy decision and publish its appropriate safeguards.
D. The bank must disclose the use of the AI system and implement suitable measures for users to contest automated decision-making.
Answer: (SHOW ANSWER)
Under the EU regulations, particularly the GDPR, banks using AI for decision-making must inform users about the use of AI and provide mechanisms for users to contest decisions. This is part of ensuring transparency and accountability in automated processing. Explicit consent under the privacy directive (A) and disclosing under the Digital Services Act (B) are not specifically required in this context. An adequacy decision is related to data transfers outside the EU (C).

NEW QUESTION: 62
Business A sells software that provides users with writing and grammar assistance. Business B is a cloud services provider that trains its own AI models.
* Business A has decided to add generative AI features to their software.
* Rather than create their own generative AI model, Business A has chosen to license a model from Business B:
* Business A will then integrate the model into their writing assistance software to provide generative AI capabilities.
* Business A is most concerned that its writing assistance software could recommend toxic or obscene text to its users.

Which of the following governance processes should Business A take to best protect its users against potentially inappropriate text?

A. Business A should fine-tune the AI model on user-generated text that has been verified to be appropriate.

B. Business A should test that the AI model performs as expected and meets their minimum requirements for filtering toxic or obscene text.

C. Business A should establish a user reporting feature that allows users to flag toxic or obscene text, and report any incidents to Business B.

D. Business A should ask Business B for detailed documentation on the generative AI model's training data and whether it contained toxic or obscene sources.

Answer: B (LEAVE A REPLY)

Business A is integrating a generative AI model licensed from a third party (Business B) and is primarily concerned with the risk of toxic or obscene outputs being delivered to users. In this scenario,testing and validationof the AI model for such content risks is the most direct and effective governance strategy.

According to theAI Governance in Practice Report 2024, organizations thatdeployAI must engage in performance monitoring protocolsand ensure systems perform adequately for theirintended purposes, including filtering harmful content:

"Operational governance... development of: #Performance monitoring protocols to ensure systems perform adequately for their intended purposes." (p. 12)

"Product governance... includes: #System impact assessments to identify and address risk prior to product development or deployment." (p. 11) Furthermore, under theEU AI Act, which sets the global standard many organizations aim to align with, there is a clear obligation to test and monitor systems for potential harmful behavior:

"The act imposes regulatory obligations... such as establishing appropriate accountability structures,assessing system impact, providing technical documentation,establishing risk management protocols and monitoring performance..." (p. 7) Option B directly reflects this best practice ofpre-deployment testing and validationto ensure that the model aligns with Business A's minimum content safety requirements.

Let's now evaluate the incorrect options:

* A. Fine-tuning on verified user-generated textmay improve model alignment but does not guarantee that the model will generalize correctly, especially if Business A lacks access to model internals (common in third-party licensing scenarios). Fine-tuning also introduces its own risks and may be contractually restricted.

* C. A user reporting featureisreactive, not preventive. While helpful for long-term monitoring and mitigation, it does not prevent the initial harm of toxic outputs, which isBusiness A's primary concern.

* D. Requesting documentation from Business Bis useful for transparency and risk management, but it does not replaceindependent verificationthat the model meets Business A's content safety standards.

Thus,testing the model's behavior for unacceptable outputs before deploymentis the most aligned approach with AI governance best practices and obligations.

NEW QUESTION: 63
All of the following may be permissible uses of an AI system under the EU AI Act EXCEPT?
A. To implement social scoring.
B. To detect an individual's intent for law enforcement purposes.
C. To manage border control.
D. To promote equitable distribution of welfare benefits.
Answer: (SHOW ANSWER)
The EU AI Act explicitly prohibits the use of AI systems for social scoring by public authorities, as it can lead to discrimination and unfair treatment of individuals based on their social behavior or perceived trustworthiness. While AI can be used to promote equitable distribution of welfare benefits, manage border control, and even detect an individual's intent for law enforcement purposes (within strict regulatory and ethical boundaries), implementing social scoring systems is not permissible under the Act due to the significant risks to fundamental rights and freedoms.

NEW QUESTION: 64
Retraining an LLM can be necessary for all of the following reasons EXCEPT?
A. To minimize degradation in prediction accuracy due tochanges 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: 65
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: (SHOW ANSWER)
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: 66
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)
Model disgorgement is the technique used to remove the effects of improperly used data from an ML system.
This process involves retraining or adjusting the model to eliminate any biases or inaccuracies introduced by the inappropriate data. It ensures that the model's outputs are not influenced by data that was not meant to be used or was used incorrectly. Reference: AIGP Body of Knowledge on Data Management and Model Integrity.

NEW QUESTION: 67
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.
The organization continues planning the adoption of an AI tool to support hiring, but is concerned about potential bias in content generated by AI systems and how that could affect public perception.
Which of the following measures should the company adopt tobest mitigate its risk of reputational harm from using the AI tool?
A. Test the AI tool pre- and post-deployment
B. Ensure the vendor provides indemnification for the AI tool
C. Require the procurement and deployment teams to agree upon the AI tool
D. Continue to require the company's hiring personnel to manually screen all applicants
Answer: (SHOW ANSWER)
Note:This is the same scenario and question as Question 21 and thus has thesame correct answer: A. It's possible this was duplicated in your original input.
Repeated for clarity:
"Testing AI tools pre- and post-deployment helps ensure they perform as expected and do not introduce bias, privacy issues, or fairness concerns. This mitigates reputational and legal risk."
TheAI Governance in Practice Report 2024further reinforces:
"Ongoing monitoring and testing post-deployment allows organizations to catch and correct unintended impacts... especially important in HR and hiring contexts."

NEW QUESTION: 68
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 22989andISO/IEC 42001can be valuable resources for AI Governance professionals inall 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: (SHOW ANSWER)

The correct answer isC. ISO/IEC 22989 and 42001 focus onterminology, risk, and management systems, butdo 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 2024:
"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: 69
Which of the following is a subcategory of AI and machine learning that uses labeled datasets to train algorithms?
A. Segmentation.
B. Generative AI.
C. Expert systems.
D. Supervised learning.
Answer: D (LEAVE A REPLY)
Supervised learning is a subcategory of AI and machine learning where labeled datasets are used to train algorithms. This process involves feeding the algorithm a dataset where the input-output pairs are known, allowing the algorithm to learn and make predictions or decisions based on new, unseen data. Reference:
AIGP BODY OF KNOWLEDGE, which describes supervised learning as a model trained on labeled data (e.
g., text recognition, detecting spam in emails).

NEW QUESTION: 70
The planning phase of the AI life cycle articulates all of the following EXCEPT the?
A. Objective of the model.
B. Approach to governance.
C. Choice of the architecture.
D. Context in which the model will operate.
Answer: B (LEAVE A REPLY)
The planning phase of the AI life cycle typically includes defining the objective of the model, choosing the appropriate architecture, and understanding the context in which the model will operate. However, the approach to governance is usually established as part of the overall AI governance framework, not specifically within the planning phase. Governance encompasses broader organizational policies and procedures that ensure AI development and deployment align with legal, ethical, and operational standards. Reference: AIGP Body of Knowledge, AI lifecycle planning phase section.

NEW QUESTION: 71
An artist has been using an AI tool to create digital art and would like to ensure that it has copyright protection in the United States.
Which of the following is most likely to enable the artist to receive copyright protection?
A. Ensure the tool was trained using publicly available content.
B. Obtain a representation from the AI provider on how the tool works.
C. Provide a log of the prompts the artist used to generate the images.
D. Update the images in a creative way to demonstrate that it is the artist's.
Answer: (SHOW ANSWER)
For the artist to receive copyright protection, the most effective approach is to demonstrate that the final artwork includes sufficient creative input by the artist. By updating or altering the images in a way that reflects the artist's personal creativity, the artist can claim originality, which is a core requirement for copyright protection under U.S. law. The other options do not directly address the originality and creative input required for copyright. This is highlighted in the sections on copyright protection in the IAPP AIGP Body of Knowledge.

NEW QUESTION: 72
Training data is best defined as a subset of data that is used to?
A. Enable a model to detect and learn patterns.
B. Fine-tune a model to improve accuracy and prevent overfitting.
C. Detect the initial sources of biases to mitigate prior to deployment.
D. Resemble the structure and statistical properties of production data.
Answer: (SHOW ANSWER)
Training data is used to enable a model to detect and learn patterns. During the training phase, the model learns from the labeled data, identifying patterns and relationships that it will later use to make predictions on new, unseen data. This process is fundamental in building an AI model's capability to perform tasks accurately. Reference: AIGP Body of Knowledge on Model Training and Pattern Recognition.

NEW QUESTION: 73
All of the following are penalties and enforcements outlined in the EU AI Act EXCEPT?
A. Fines for SMEs and startups will be proportionally capped.
B. Rules on General Purpose AI will apply after 6 months as a specific provision.
C. The AI Pact will act as a transitional bridge until the Regulations are fully enacted.
D. Fines for violations of banned AI applications will be €35 million or 7% global annual turnover (whichever is higher).
Answer: (SHOW ANSWER)
The EU AI Act outlines specific penalties and enforcement mechanisms to ensure compliance with its regulations. Among these, fines for violations of banned AI applications can be as high as €35 million or 7% of the global annual turnover of the offending organization, whichever is higher. Proportional caps on fines are applied to SMEs and startups to ensure fairness. General Purpose

AI rules are to apply after a 6-month period as a specific provision to ensure that stakeholders have adequate time to comply. However, there is no provision for an "AI Pact" acting as a transitional bridge until the regulations are fully enacted, making option C the correct answer.

NEW QUESTION: 74
What is the main purpose of accountability structures under the Govern function of the NIST AI Risk Management Framework?
A. To empower and train appropriate cross-functional teams.
B. To establish diverse, equitable and inclusive processes.
C. To determine responsibility for allocating budgetary resources.
D. To enable and encourage participation by external stakeholders.
Answer: (SHOW ANSWER)
The NIST AI Risk Management Framework's Govern function emphasizes the importance of establishing accountability structures that empower and train cross-functional teams. This is crucial because cross- functional teams bring diverse perspectives and expertise, which are essential for effective AI governance and risk management. Training these teams ensures that they are well-equipped to handle their responsibilities and can make informed decisions that align with the organization's AI principles and ethical standards. Reference:
NIST AI Risk Management Framework documentation, Govern function section.

NEW QUESTION: 75
Scenario:
Business A provides grammar and writing assistance tools and licenses a generative AI model from Business B to enhance its offerings. Business A is concerned that the AI model might produce inappropriate or toxic content and wants to implement governance processes to prevent this.
Which of the following governance processes should Business A take tobest protect its usersagainst potentially inappropriate text?
A. Business A should fine-tune the AI model on user-generated text that has been verified to be appropriate
B. Business A should test that the AI model performs as expected and meets their minimum requirements for filtering toxic or obscene text
C. Business A should establish a user reporting feature that allows users to flag toxic or obscene text, and report any incidents to Business B
D. Business A should ask Business B for detailed documentation on the generative AI model's training data and whether it contained toxic or obscene sources
Answer: (SHOW ANSWER)
The correct answer isB. According to responsible AI practices,pre-deployment testingto ensurethe model behaves as expected and aligns with organizational requirements is critical.
From the AIGP ILT Guide:

"Testing for unacceptable outcomes such as toxicity, discrimination, or hallucinations should be included in the AI governance life cycle, particularly during development and prior to deployment."
Also emphasized in the AI Governance in Practice Report 2024:
"Organizations must verify legal and regulatory compliance, monitor performance, and mitigate riskspriorto deployment." Testing the model tomeet safety and appropriateness standardsis more proactive and preventive than relying solely on user reporting or requesting documentation.

NEW QUESTION: 76
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: 77
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 havespecific responsibilitiesunder the EU AI Act. They arenot 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 theAI Governance in Practice Report 2024:
"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 theprovider.
* B. Affix the CE marking-provider'sresponsibility.
* C. Verify the Declaration of Conformity-importer obligation.
* D. Conduct a DPIA- relevant under data protection laws,not requiredunder the EU AI Act for importers.

NEW QUESTION: 78
Under the NIST AI Risk Management Framework, all of the following are defined as characteristics of trustworthy AI EXCEPT?
A. Tested and Effective.
B. Secure and Resilient.
C. Explainable and Interpretable.
D. Accountable and Transparent.
Answer: (SHOW ANSWER)

The NIST AI Risk Management Framework outlines several characteristics of trustworthy AI, including being secure and resilient, explainable and interpretable, and accountable and transparent. While being tested and effective is important, it is not explicitly listed as a characteristic of trustworthy AI in the NIST framework.
The focus is more on the system's ability to function safely, securely, and transparently in a way that stakeholders can understand and trust. Reference: AIGP Body of Knowledge, NIST AI RMF section.

NEW QUESTION: 79
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 integratingexternal, up-to-date, or proprietary informationinto the generation pipeline-allowing the model tofetch relevant factsfrom 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 ensuringmore accurate and contextually grounded outputs, especially inregulated or high- stakes environmentswhere hallucinations are a concern.
* B, C, and Ddescribe optimization or bias mitigation-not the core function of RAG.

NEW QUESTION: 80
Scenario:
A global organization wants to align with international frameworks on AI governance. They are reviewing guidance from the OECD on how to incorporate broader governance tools into their AI program.

Codes of conduct and collective agreements are what type of assessment tools as defined by the Organization for Economic Cooperation and Development (OECD)?
A. Educational
B. Procedural
C. Technical
D. Analytic
Answer: (SHOW ANSWER)
The correct answer is B - Procedural. The OECD Framework for Classifying AI Systems categorizes codes of conduct and collective agreements as procedural tools because they guide internal governance and decision-making processes.
From the AIGP ILT Participant Guide - Global Governance Models:
"Procedural tools include internal codes of conduct, collective agreements, and procedural audits that guide governance without necessarily involving technical measurement." AI Governance in Practice Report 2024 elaborates:

"These procedural tools support internal accountability mechanisms and ethics compliance frameworks...
they are part of soft governance."
These tools do not measure or analyze technical performance, hence they are not technical or analytic.

NEW QUESTION: 81
In procuring an AI system from a vendor, which of the following would be important to include in a contract to enable proper oversight and auditing of the system?
A. Liability for mistakes.
B. Ownership of data and outputs.
C. Responsibility for improvements.
D. Appropriate access to data and models.
Answer: (SHOW ANSWER)
Ensuringoversight and auditabilityrequires that the organization hassufficient access to data, documentation, and model internalsor outputs necessary for evaluation.
From theAI Governance in Practice Report 2024:
"Access to technical documentation and system internals is essential to enable effective auditing, conformity checks, and accountability mechanisms." (p. 11, 34)
* Ais about liability, not auditability.
* Bmatters for IP rights, not oversight.
* Crelates to lifecycle responsibility but doesn't guarantee audit access.

NEW QUESTION: 82
Which of the following is the least relevant consideration in assessing whether users should be given the right to opt out from an AI system?
A. Feasibility.
B. Risk to users.
C. Industry practice.
D. Cost of alternative mechanisms.
Answer: (SHOW ANSWER)
When assessing whether users should be given the right to opt out from an AI system, the primary considerations are feasibility, risk to users, and industry practice. Feasibility addresses whether the opt-out mechanism can be practically implemented. Risk to users assesses the potential harm or benefits users might face if they cannot opt out. Industry practice considers the norms and standards within the industry. However, the cost of alternative mechanisms, while important in the broader context of implementation, is not directly relevant to the ethical consideration of whether users should have the right to opt out. The focus should be on protecting user rights and ensuring ethical AI practices.
Reference: AIGP BODY OF KNOWLEDGE, sections discussing user rights and ethical considerations in AI.