

IAPP

Exam Questions AIGP

Artificial Intelligence Governance Professional



NEW QUESTION 1

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

Answer: A

Explanation:

The best approach to enable a customer who wants information on the AI model's parameters for underwriting purposes is to provide a transparency notice. This notice should explain the nature of the AI system, how it uses customer data, and the decision-making process it follows. Providing a transparency notice is crucial for maintaining trust and compliance with regulatory requirements regarding the transparency and accountability of AI systems.

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

NEW QUESTION 2

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

Which 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: A

Explanation:

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 3

- (Topic 1)

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: B

Explanation:

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 4

- (Topic 1)

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 a human underwriter for final review.

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

What is the best strategy to mitigate the bias uncovered in the loan applications?

- A. Retrain the model with data that reflects demographic parity.
- B. Procure a third-party statistical bias assessment tool.
- C. Document all instances of bias in the data set.
- D. Delete all gender-based data in the data set.

Answer: A

Explanation:

Retraining the model with data that reflects demographic parity is the best strategy to mitigate the bias uncovered in the loan applications. This approach addresses the root cause of the bias by ensuring that the training data is representative and balanced, leading to more equitable decision-making by the AI model.

Reference: The AIGP Body of Knowledge stresses the importance of using high-quality, unbiased training data to develop fair and reliable AI systems. Retraining the model with balanced data helps correct biases that arise from historical inequalities, ensuring that the AI system makes decisions based on equitable criteria.

NEW QUESTION 5

- (Topic 1)

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 a human underwriter for final review.

ABC and the cloud provider have completed training and testing the LLM, performed a readiness assessment, and made the decision to deploy the LLM into production. ABC has designated an internal compliance team to monitor the model during the first month, specifically to evaluate the accuracy, fairness, and reliability of its output. After the first

month in production, ABC realizes that the LLM declines a higher percentage of women's loan applications due primarily to women historically receiving lower salaries than men.

Which of the following is the most important reason to train the underwriters on the model prior to deployment?

- A. To provide a reminder of a right appeal.
- B. To solicit on-going feedback on model performance.
- C. To apply their own judgment to the initial assessment.
- D. To ensure they provide transparency applicants on the model.

Answer: C

Explanation:

Training underwriters on the model prior to deployment is crucial so they can apply their own judgment to the initial assessment. While AI models can streamline the process, human judgment is still essential to catch nuances that the model might miss or to account for any biases or errors in the model's decision-making process.

Reference: The AIGP Body of Knowledge emphasizes the importance of human oversight in AI systems, particularly in high-stakes areas such as underwriting and loan approvals. Human underwriters can provide a critical review and ensure that the model's assessments are accurate and fair, integrating their expertise and understanding of complex cases.

NEW QUESTION 6

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 7

- (Topic 1)

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: C

Explanation:

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 8

- (Topic 1)

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

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

Answer: A

Explanation:

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 9

- (Topic 1)

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

- A. To detect an individual's intent for law enforcement purposes.
- B. To promote equitable distribution of welfare benefits.
- C. To implement social scoring.
- D. To manage border control.

Answer: C

Explanation:

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 10

- (Topic 1)

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: C

Explanation:

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 Management Framework. 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 10

- (Topic 1)

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: C

Explanation:

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 15

- (Topic 1)

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: D

Explanation:

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 17

- (Topic 1)

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 EII 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: A

Explanation:

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 22

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 27

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

students to use prompts to create digital art.

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

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

- A. Financial cost.
- B. Model accuracy.
- C. Technical complexity.
- D. Copyright infringement.

Answer: B

Explanation:

The highest concern for individual teachers using generative AI to ensure students learn the course material is model accuracy. Ensuring that the AI-generated content is accurate and relevant to the curriculum is crucial for effective learning. If the AI model produces inaccurate or irrelevant content, it can mislead students and hinder their understanding of the subject matter.

Reference: According to the AIGP Body of Knowledge, one of the core risks posed by AI

systems is the accuracy of the data and models used. Ensuring the accuracy of AI-generated content is essential for maintaining the integrity of the educational material and achieving the desired learning outcomes.

NEW QUESTION 28

- (Topic 1)

A Canadian company is developing an AI solution to evaluate candidates in the course of job interviews.

Before offering the AI solution in the EU market, the company must take all of the following steps EXCEPT?

- A. Register the AI solution in a public EU database.
- B. Establish a risk and quality management system.
- C. Engage a third-party auditor to perform a bias audit.
- D. Draw up technical documentation and instructions for use.

Answer: A

Explanation:

Before offering an AI solution in the EU market, a Canadian company must take several steps to comply with the EU AI Act. These steps include establishing a risk and quality management system (B), engaging a third-party auditor to perform a bias audit (C), and drawing up technical documentation and instructions for use (D). However, there is no requirement to register the AI solution in a public EU database (A). This registration step is not specified as part of the compliance requirements under the EU AI Act for such solutions.

NEW QUESTION 29

- (Topic 1)

An AI system that maintains its level of performance within defined acceptable limits despite real world or adversarial conditions would be described as?

- A. Robust.
- B. Reliable.
- C. Resilient.
- D. Reinforced.

Answer: C

Explanation:

An AI system that maintains its level of performance within defined acceptable limits despite real-world or adversarial conditions is described as resilient. Resilience in AI refers to the system's ability to withstand and recover from unexpected challenges, such as cyber-attacks, hardware failures, or unusual input data. This characteristic ensures that the AI system can continue to function effectively and reliably in various conditions, maintaining performance and integrity. Robustness, on the other hand, focuses on the system's strength against errors, while reliability ensures consistent performance over time. Resilience combines these aspects with the capacity to adapt and recover.

NEW QUESTION 34

- (Topic 1)

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: B

Explanation:

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 37

- (Topic 1)

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.

All of the following are potential negative consequences created by using the AI tool when making hiring decisions EXCEPT?

- A. Reputational harm.
- B. Civil rights violations.
- C. Discriminatory treatment.
- D. Intellectual property infringement.

Answer: D

Explanation:

The potential negative consequences of using an AI tool in hiring include reputational harm (A), civil rights violations (B), and discriminatory treatment (C). These issues stem from biases in the AI system or its misuse, which can lead to unfair hiring practices and legal liabilities. Intellectual property infringement (D) is not a typical consequence of using AI in hiring, as it relates to the unauthorized use of protected intellectual property, which is not directly relevant to the hiring process or the potential biases within AI tools.

NEW QUESTION 38

- (Topic 2)

Testing data is defined as a subset of data that is used to?

- A. Assess a model's on-going performance in production.
- B. Enable a model to discover and learn patterns.
- C. Provide a robust evaluation of a final model.
- D. Evaluate a model's handling of randomized edge cases.

Answer: C

Explanation:

Testing data is a subset of data used to provide a robust evaluation of a final model. After training the model on training data, it is essential to test its performance on unseen data (testing data) to ensure it generalizes well to new, real-world scenarios. This step helps in assessing the model's accuracy, reliability, and ability to handle various data inputs. Reference: AIGP Body of Knowledge on Model Validation and Testing.

NEW QUESTION 39

- (Topic 2)

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: B

Explanation:

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 42

- (Topic 2)

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: B

Explanation:

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 44

- (Topic 2)

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

Explanation:

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 46

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 50

- (Topic 2)

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: D

Explanation:

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

- (Topic 2)

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: D

Explanation:

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 56

- (Topic 2)

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: A

Explanation:

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 57

- (Topic 2)

A company has trained an ML model primarily using synthetic data, and now intends to use live personal data to test the model. Which of the following is NOT a best practice apply during the testing?

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

Answer: B

Explanation:

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 62

- (Topic 2)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 67

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

datasets?

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

Answer: B

Explanation:

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

NEW QUESTION 70

- (Topic 2)

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: C

Explanation:

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 75

- (Topic 2)

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 of the following steps can best mitigate the possibility of discrimination prior to training and testing the AI solution?

- A. Procure more data from clinical research partners.
- B. Engage a third party to perform an audit.
- C. Perform an impact assessment.
- D. Create a bias bounty program.

Answer: C

Explanation:

Performing an impact assessment is the best step to mitigate the possibility of discrimination before training and testing the AI solution. An impact assessment, such as a Data Protection Impact Assessment (DPIA) or Algorithmic Impact Assessment (AIA), helps identify potential biases and discriminatory outcomes that could arise from the AI system. This process involves evaluating the data and the algorithm for fairness, accountability, and transparency. It ensures that any biases in the data are detected and addressed, thus preventing discriminatory practices and promoting ethical AI deployment. Reference: AIGP Body of Knowledge on Ethical AI and Impact Assessments.

NEW QUESTION 79

- (Topic 2)

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: D

Explanation:

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 84

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 87

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 88

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AIGP Practice Exam Features:

- * AIGP Questions and Answers Updated Frequently
- * AIGP Practice Questions Verified by Expert Senior Certified Staff
- * AIGP Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AIGP Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AIGP Practice Test Here](#)