

Exam Questions AIGP

Artificial Intelligence Governance Professional

<https://www.2passeasy.com/dumps/AIGP/>



NEW QUESTION 1

- (Topic 1)

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: D

Explanation:

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 2

- (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 3

- (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 4

- (Topic 1)

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

Explanation:

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 5

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 6

- (Topic 1)

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: D

Explanation:

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 7

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 8

- (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 9

- (Topic 1)

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: C

Explanation:

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 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.

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

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

Answer: B

Explanation:

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

NEW QUESTION 10

- (Topic 1)

What is the key feature of Graphical Processing Units (GPUs) that makes them well-suited to running AI applications?

- A. GPUs run many tasks concurrently, resulting in faster processing.
- B. GPUs can access memory quickly, resulting in lower latency than CPUs.
- C. GPUs can run every task on a computer, making them more robust than CPUs.
- D. The number of transistors on GPUs doubles every two years, making the chips smaller and lighter.

Answer: A

Explanation:

GPUs (Graphical Processing Units) are well-suited to running AI applications due to their ability to run many tasks concurrently, which significantly enhances processing speed. This parallel processing capability makes GPUs ideal for handling the large-scale computations required in AI and deep learning tasks.

Reference: AIGP BODY OF KNOWLEDGE, which explains the importance of compute infrastructure in AI applications.

NEW QUESTION 13

- (Topic 1)

Under the Canadian Artificial Intelligence and Data Act, when must the Minister of Innovation, Science and Industry be notified about a high-impact AI system?

- A. When use of the system causes or is likely to cause material harm.
- B. When the algorithmic impact assessment has been completed.
- C. Upon release of a new version of the system.
- D. Upon initial deployment of the system.

Answer: D

Explanation:

According to the Canadian Artificial Intelligence and Data Act, high-impact AI systems must notify the Minister of Innovation, Science and Industry upon initial deployment. This requirement ensures that the authorities are aware of the deployment of significant AI systems and can monitor their impacts and compliance with regulatory standards from the outset. This initial notification is crucial for maintaining oversight and ensuring the responsible use of AI technologies.

Reference: AIGP Body of Knowledge, domain on AI laws and standards.

NEW QUESTION 16

- (Topic 1)

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: B

Explanation:

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 18

- (Topic 1)

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

Explanation:

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 21

- (Topic 2)

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: A

Explanation:

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 26

- (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 30

- (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 32

- (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.

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: D

Explanation:

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 37

- (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 41

- (Topic 2)

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: A

Explanation:

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 44

- (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.

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: A

Explanation:

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 48

- (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 53

- (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 55

- (Topic 2)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 60

- (Topic 2)

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: C

Explanation:

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 65

- (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.

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

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

Answer: D

Explanation:

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

NEW QUESTION 67

- (Topic 2)

All of the following are reasons to deploy a challenger AI model in addition a champion AI model EXCEPT to?

- A. Provide a framework to consider alternatives to the champion model.
- B. Automate real-time monitoring of the champion model.
- C. Perform testing on the champion model.
- D. Retrain the champion model.

Answer: D

Explanation:

Deploying a challenger AI model alongside a champion model is a strategy used to compare the performance of different models in a real-world environment. This

approach helps in providing a framework to consider alternatives to the champion model, automating real-time monitoring of the champion model, and performing testing on the champion model. However, retraining the champion model is not a reason to deploy a challenger model. Retraining is a separate process that involves updating the champion model with new data or techniques, which is not related to the use of a challenger model. Reference: AIGP BODY OF KNOWLEDGE, sections on model evaluation and management.

NEW QUESTION 68

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

- A. The cloud provider.
- B. The consulting firm.
- C. The healthcare network's data science team.
- D. The healthcare network's AI governance committee.

Answer: C

Explanation:

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 70

- (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 75

- (Topic 2)

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: D

Explanation:

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 76

- (Topic 2)

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: A

Explanation:

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 81

- (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 84

- (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 85

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AIGP Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AIGP Product From:

<https://www.2passeasy.com/dumps/AIGP/>

Money Back Guarantee

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