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 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. Toprovide a reminder of a right appeal.
B. Tosolicit on-going feedback on model performance.
C. Toapply their own judgment to the initial assessment.
D. Toensure 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 2
- (Topic 1)
Which of the following best defines an "Al model"?

A. A system that applies defined rules to execute tasks.
B. A system of controls that is used to govern an Al algorithm.
C. A corpus of data which an Al 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 3
- (Topic 1)
The framework set forth in the White House Blueprint for an Al 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 4
- (Topic 1)
All of the following are common optimization techniques in deep learning to determine weights that represent the strength of the connection between artificial neurons EXCEPT?

A. Gradient descent, which initially sets weights arbitrary values, and then at each step changes them.
B. Momentum, which improves the convergence speed and stability of neural network training.
C. Autoregression, which analyzes and makes predictions about time-series data.
D. Backpropagation, which starts from the last layer working backwards.

Answer: C

Explanation:
Autoregression is not a common optimization technique in deep learning to
determine weights for artificial neurons. Common techniques include gradient descent, momentum, and backpropagation. Autoregression is more commonly

associated with time- series analysis and forecasting rather than neural network optimization. Reference: AIGP BODY OF KNOWLEDGE, which discusses common optimization techniques used in deep learning.

NEW QUESTION 5
- (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 6
- (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 .. human underwriter for final review.
ABC and the cloud provider have completed training and testing the LLM, performed a readiness assessment, and made the decision to deploy the LLM into production. ABC has designated an internal compliance team to monitor the model during the first month, specifically to evaluate the accuracy, fairness, and reliability of its output. After the first month in production, ABC realizes that the LLM declines a higher percentage of women's loan applications due primarily to women historically receiving lower salaries than men.
During the first month when ABC monitors the model for bias, it is most important to?

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

Answer: A

Explanation:
During the first month of monitoring the model for bias, it is most important to continue disparity testing. Disparity testing involves regularly evaluating the model's decisions to identify and address any biases, ensuring that the model operates fairly across different demographic groups.
Reference: Regular disparity testing is highlighted in the AIGP Body of Knowledge as a
critical practice for maintaining the fairness and reliability of AI models. By continuously monitoring for and addressing disparities, organizations can ensure their AI systems remain compliant with ethical and legal standards, and mitigate any unintended biases that may arise in production.

NEW QUESTION 7
- (Topic 1)
According to the GDPR, what is an effective control to prevent a determination based solely on automated decision-making?

A. Provide a just-in-time notice about the automated decision-making logic.
B. Define suitable measures to safeguard personal data.
C. Provide a right to review automated decision.
D. Establish a human-in-the-loop procedure.

Answer: D

Explanation:
The GDPR requires that individuals have the right to not be subject to decisions based solely on automated processing, including profiling, unless specific exceptions apply. One effective control is to establish a human-in-the-loop procedure (D), ensuring human oversight and the ability to contest decisions. This goes beyond just-in- time notices (A), data safeguarding (B), or review rights (C), providing a more robust mechanism to protect individuals' rights.

NEW QUESTION 8
- (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.
What is the best reason for GVC to offer students the choice to utilize generative AI in limited, defined circumstances?

A. Toenable students to learn how to manage their time.

B. Toenable students to learn about performing research.
C. Toenable students to learn about practical applications of AI.
D. Toenable students to learn how to use AI as a supportive educational tool.

Answer: D

Explanation:
The best reason for GVC to offer students the choice to utilize generative AI in limited, defined circumstances is to enable students to learn how to use AI as a supportive educational tool. By integrating AI in a controlled manner, students can learn the practical applications of AI and develop skills to use AI responsibly and effectively in their educational pursuits.
Reference: The AIGP Body of Knowledge highlights the importance of teaching students about AI's practical applications and the responsible use of AI technologies. This aligns with the goal of fostering a better understanding of AI's role and its potential benefits in various contexts, including education.

NEW QUESTION 9
- (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 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.
Which of the following measures should XYZ adopt to best mitigate its risk of reputational harm from using the AI tool?

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

Answer: A

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

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-discriminationlaws.

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 11
- (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 12
- (Topic 1)
What is the primary purpose of an AI impact assessment?

A. To define and evaluate the legal risks associated with developing an AI system.
B. Anticipate and manage the potential risks and harms of an AI system.
C. To define and document the roles and responsibilities of AI stakeholders.
D. To identify and measure the benefits of an AI system.

Answer: B

Explanation:
The primary purpose of an AI impact assessment is to anticipate and manage the potential risks and harms of an AI system. This includes identifying the possible negative outcomes and implementing measures to mitigate these risks. This process helps ensure that AI systems are developed and deployed in a manner that is ethically and socially responsible, addressing concerns such as bias, fairness, transparency, and accountability. The assessment often involves a thorough evaluation of the AI system's design, data inputs, outputs, and the potential impact on various stakeholders. This approach is crucial for maintaining public trust and adherence to regulatory requirements.

NEW QUESTION 14
- (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 17
- (Topic 2)
All of the following are elements of establishing a global Al governance infrastructure EXCEPT?

A. Providing training to foster a culture that promotes ethical behavior.
B. Creating policies and procedures to manage third-partyrisk.
C. Understanding differences in norms across countries.
D. Publicly disclosing ethical principles.

Answer: D

Explanation:
Establishing a global AI governance infrastructure involves several key elements, including providing training to foster a culture that promotes ethical behavior, creating policies and procedures to manage third-party risk, and understanding differences in norms across countries. While publicly disclosing ethical principles can enhance transparency and trust, it is not a core element necessary for the establishment of a governance infrastructure. The focus is more on internal processes and structures rather than public disclosure. Reference: AIGP Body of Knowledge on AI Governance and Infrastructure.

NEW QUESTION 19
- (Topic 2)
In the machine learning context, feature engineering is the process of?

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

Answer: D

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

NEW QUESTION 22
- (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 Al?

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 Al 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 24
- (Topic 2)
You are part of your organization’s ML engineering team and notice that the accuracy of a model that was recently deployed into production is deteriorating. What is the best first step address this?

A. Replace the model with a previous version.
B. Conduct champion/challenger testing.
C. Perform an audit of the model.
D. Run red-teaming exercises.

Answer: B

Explanation:
When the accuracy of a model deteriorates, the best first step is to conduct champion/challenger testing. This involves deploying a new model (challenger) alongside the current model (champion) to compare their performance. This method helps identify if the new model can perform better under current conditions without immediately discarding the existing model. It provides a controlled environment to test improvements and understand the reasons behind the deterioration. This approach is preferable to directly replacing the model, performing audits, or running red-teaming exercises, which may be subsequent steps based on the findings from the champion/challenger testing.
Reference: AIGP BODY OF KNOWLEDGE, sections on model performance management and testing strategies.

NEW QUESTION 29
- (Topic 2)
A company initially intended to use a large data set containing personal information to train an Al model. After consideration, the company determined that it can derive enough value from the data set without any personal information and permanently obfuscated all personal data elements before training the model. This is an example of applying which privacy-enhancing technique (PET)?

A. Anonymization.

B. Pseudonymization.
C. Differential privacy.
D. Federated learning.

Answer: A

Explanation:
Anonymization is a privacy-enhancing technique that involves removing or permanently altering personal data elements to prevent the identification of individuals. In this case, the company obfuscated all personal data elements before training the model, which aligns with the definition of anonymization. This ensures that the data cannot be traced back to individuals, thereby protecting their privacy while still allowing the company to derive value from the dataset. Reference: AIGP Body of Knowledge, privacy-enhancing techniques section.

NEW QUESTION 30
- (Topic 2)
Which of the following is the least relevant consideration in assessing whether users should be given the right to opt out from an AI system?

A. Feasibility.
B. Risk to users.
C. Industry practice.
D. Cost of alternative mechanisms.

Answer: D

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

NEW QUESTION 32
- (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 34
- (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 35
- (Topic 2)
You are the chief privacy officer of a medical research company that would like to collect and use sensitive data about cancer patients, such as their names, addresses, race and ethnic origin, medical histories, insurance claims, pharmaceutical prescriptions, eating and drinking habits and physical activity.
The company will use this sensitive data to build an AI algorithm that will spot common attributes that will help predict if seemingly healthy people are more likely to get cancer. However, the company is unable to obtain consent from enough patients to sufficiently collect the minimum data to train its model.
Which of the following solutions would most efficiently balance privacy concerns with the lack of available data during the testing phase?

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

Answer: C

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

NEW QUESTION 36
- (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 41
- (Topic 2)
CASE STUDY
Please use the following answer the next question:
A mid-size US healthcare network has decided to develop an Al 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 Al ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an Al 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 46
- (Topic 2)
Which of the following steps occurs in the design phase of the Al 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 49

- (Topic 2)
Which of the following AI uses is best described as human-centric?

A. Pattern recognition algorithms are used to improve the accuracy of weather predictions, which benefits many industries and everyday life.
B. Autonomous robots are used to move products within a warehouse, allowing human workers to reduce physical strain and alleviate monotony.
C. Machine learning is used for demand forecasting and inventory management, ensuring that consumers can find products they want when they want them.
D. Virtual assistants are used adapt educational content and teaching methods to individuals, offering personalized recommendations based on ability and needs.

Answer: D

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

NEW QUESTION 53
- (Topic 2)
All of the following are reasons to deploy a challenger Al model in addition a champion Al 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 57
- (Topic 2)
When monitoring the functional performance of a model that has been deployed into production, all of the following are concerns EXCEPT?

A. Feature drift.
B. System cost.
C. Model drift.
D. Data loss.

Answer: B

Explanation:
When monitoring the functional performance of a model deployed into production, concerns typically include feature drift, model drift, and data loss. Feature drift refers to changes in the input features that can affect the model's predictions. Model drift is when the model's performance degrades over time due to changes in the data or environment. Data loss can impact the accuracy and reliability of the model. However, system cost, while important for budgeting and financial planning, is not a direct concern when monitoring the functional performance of a deployed model. Reference: AIGP Body of Knowledge on Model Monitoring and Maintenance.

NEW QUESTION 61
- (Topic 2)
What is the main purpose of accountability structures under the Govern function of the NIST Al 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 62
- (Topic 2)
CASE STUDY
Please use the following answer the next question:
A mid-size US healthcare network has decided to develop an Al 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 Al 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 65
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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