

# 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:

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.

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 3**

- (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 4**

- (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 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)

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 7**

- (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 8**

- (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 9**

- (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 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)

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 14**

- (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 16**

- (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 18**

- (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 20**

- (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 23**

- (Topic 2)

What is the term for an algorithm that focuses on making the best choice achieve an immediate objective at a particular step or decision point, based on the available information and without regard for the longer-term best solutions?

- A. Single-lane.
- B. Optimized.
- C. Efficient.
- D. Greedy.

**Answer:** D

**Explanation:**

A greedy algorithm is one that makes the best choice at each step to achieve an immediate objective, without considering the longer-term consequences. It focuses on local optimization at each decision point with the hope that these local solutions will lead to an optimal global solution. However, greedy algorithms do not always produce the best overall solution for certain problems, but they are useful when an immediate, locally optimal solution is desired. Reference: AIGP Body of Knowledge, algorithm types section.

**NEW QUESTION 25**

- (Topic 2)

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

- A. A non-profit wants to develop a social media presence.O
- B. An e-commerce provider wants to make personalized recommendations.
- C. A business analyst wants to forecast future cost overruns and underruns.
- D. A customer service agency wants automate answers to common questions.

**Answer:** A

**Explanation:**

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

**NEW QUESTION 30**

- (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 32**

- (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 willgive privacy notice to individuals before using their personal data to train or test the tool.
- B. The provider willdefend and indemnify the company against infringement claims.
- C. The provider willobtain and maintain insurance to cover potential claims.
- D. The provider willwarrant 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 33

- (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 36

- (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 41

- (Topic 2)

The planning phase of the AI life cycle articulates all of the following EXCEPT the?

- A. Objective of the model.
- B. Approach to governance.
- C. Choice of the architecture.
- D. Context in which the model will operate.

**Answer: B**

#### Explanation:

The planning phase of the AI life cycle typically includes defining the objective of the model, choosing the appropriate architecture, and understanding the context in which the model will operate. However, the approach to governance is usually established as part of the overall AI governance framework, not specifically within the planning phase. Governance encompasses broader organizational policies and procedures that ensure AI development and deployment align with legal, ethical, and operational standards. Reference: AIGP Body of Knowledge, AI lifecycle planning phase section.

### NEW QUESTION 42

- (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 45

- (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 49

- (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 50

.....



## 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 First Try
- \* AIGP Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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