

AIGP Dumps

Artificial Intelligence Governance Professional

<https://www.certleader.com/AIGP-dumps.html>



NEW QUESTION 1

- (Topic 1)

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

- A. Data architects.
- B. Government regulators.
- C. Socio-cultural and technical experts.
- D. Legal and privacy governance experts.

Answer: B

Explanation:

Typically, actors involved in the AI development life cycle include data architects (who design the data frameworks), socio-cultural and technical experts (who ensure the AI system is socio-culturally aware and technically sound), and legal and privacy governance experts (who handle the legal and privacy aspects). Government regulators, while important, are not directly engaged in the development process but rather oversee and regulate the industry. Reference: AIGP BODY OF KNOWLEDGE and AI development frameworks.

NEW QUESTION 2

- (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 3

- (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 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)

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 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)

A US company has developed an AI system, CrimeBuster 9619, that collects information about incarcerated individuals to help parole boards predict whether someone is likely to commit another crime if released from prison.

When considering expanding to the EU market, this type of technology would?

- A. Require the company to register the tool with the EU database.
- B. Be subject approval by the relevant EU authority.
- C. Require a detailed conformity assessment.
- D. Be banned under the EU AI Act.

Answer: C

Explanation:

Under the EU AI Act, high-risk AI systems like CrimeBuster 9619 would require a detailed conformity assessment before being deployed in the EU market. This assessment ensures that the AI system complies with all relevant regulations and standards, addressing potential risks related to privacy, security, and discrimination. The company would not need to register the tool with the EU database (A), seek approval from an EU authority (B), or face a ban (D) as long as it meets the necessary conformity requirements.

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)

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 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 12

- (Topic 1)

Under the NIST AI Risk Management Framework, all of the following are defined as characteristics of trustworthy AI EXCEPT?

- A. Tested and Effective.
- B. Secure and Resilient.
- C. Explainable and Interpretable.
- D. Accountable and Transparent.

Answer: A

Explanation:

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

NEW QUESTION 17

- (Topic 1)

A U.S. mortgage company developed an AI platform that was trained using anonymized details from mortgage applications, including the applicant's education, employment and demographic information, as well as from subsequent payment or default information. The AI platform will be used automatically grant or deny new mortgage applications, depending on whether the platform views an applicant as presenting a likely risk of default.

Which of the following laws is NOT relevant to this use case?

- A. Fair Housing Act.
- B. Fair Credit Reporting Act.
- C. Equal Credit Opportunity Act.
- D. Title VII of the Civil Rights Act of 1964.

Answer: D

Explanation:

The U.S. mortgage company's AI platform relates to housing and credit, making the Fair Housing Act (A), Fair Credit Reporting Act (B), and Equal Credit

Opportunity Act (C) relevant. Title VII of the Civil Rights Act of 1964 deals with employment discrimination and is not directly relevant to the mortgage application context (D).

NEW QUESTION 22

- (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 27

- (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 31

- (Topic 2)

All of the following types of testing can help evaluate the performance of a responsible AI system EXCEPT?

- A. Risk probability/severity.
- B. Adversarial robustness.
- C. Statistical sampling.
- D. Decision analysis.

Answer: A

Explanation:

Risk probability/severity testing is not typically used to evaluate the performance of an AI system. While important for risk management, it does not directly assess an AI system's operational performance. Adversarial robustness, statistical sampling, and decision analysis are all methods that can help evaluate the performance of a responsible AI system by testing its resilience, accuracy, and decision-making processes under various conditions. Reference: AIGP Body of Knowledge on AI Performance Evaluation and Testing.

NEW QUESTION 33

- (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 34

- (Topic 2)

What is the primary purpose of conducting ethical red-teaming on an AI system?

- A. To improve the model's accuracy.
- B. To simulate model risk scenarios.
- C. To identify security vulnerabilities.
- D. To ensure compliance with applicable law.

Answer: B

Explanation:

The primary purpose of conducting ethical red-teaming on an AI system is to simulate model risk scenarios. Ethical red-teaming involves rigorously testing the AI system to identify potential weaknesses, biases, and vulnerabilities by simulating real-world attack or failure scenarios. This helps in proactively addressing issues that could compromise the system's reliability, fairness, and security. Reference: AIGP Body of Knowledge on AI Risk Management and Ethical AI Practices.

NEW QUESTION 37

- (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 42

- (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 44

- (Topic 2)

You are an engineer that developed an AI-based ad recommendation tool. Which of the following should be monitored to evaluate the tool's effectiveness?

- A. Output data, assess the delta between the prediction and actual ad clicks.

- B. Algorithmic patterns, to show the model has a high degree of accuracy.
- C. Input data, to ensure the ads are reaching the target audience.
- D. GPU performance, to evaluate the tool's robustness.

Answer: A

Explanation:

To evaluate the effectiveness of an AI-based ad recommendation tool, the most relevant metric is the output data, specifically assessing the delta between the prediction and actual ad clicks. This metric directly measures the tool's accuracy and effectiveness in making accurate recommendations that lead to user engagement. While monitoring algorithmic patterns and input data can provide insights into the model's behavior and targeting accuracy, and GPU performance can indicate the robustness and efficiency of the tool, the primary indicator of effectiveness for an ad recommendation tool is how well it predicts actual ad clicks. Reference: AIGP BODY OF KNOWLEDGE, sections on AI performance metrics and evaluation methods.

NEW QUESTION 49

- (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 54

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 58

- (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 61

- (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 66

- (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 71

- (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 75

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AIGP Exam with Our Prep Materials Via below:

<https://www.certleader.com/AIGP-dumps.html>