



**IAPP**

## **Exam Questions AIGP**

Artificial Intelligence Governance Professional

### NEW QUESTION 1

- (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.

All of the following may be copyright risks from teachers using generative AI to create course content EXCEPT?

- A. Content created by an LLM may be protectable under U.S. intellectual property law.
- B. Generative AI is generally trained using intellectual property owned by third parties.
- C. Students must expressly consent to this use of generative AI.
- D. Generative AI often creates content without attribution.

**Answer:** C

#### Explanation:

All of the options listed may pose copyright risks when teachers use generative AI to create course content, except for students must expressly consent to this use of generative AI. While obtaining student consent is essential for ethical and privacy reasons, it does not directly relate to copyright risks associated with the creation and use of AI-generated content.

Reference: The AIGP Body of Knowledge discusses the importance of addressing intellectual property (IP) risks when using AI-generated content. Copyright risks are typically associated with the use of third-party data and the lack of attribution, rather than the consent of users.

### NEW QUESTION 2

- (Topic 1)

Random forest algorithms are in what type of machine learning model?

- A. Symbolic.
- B. Generative.
- C. Discriminative.
- D. Natural language processing.

**Answer:** C

#### Explanation:

Random forest algorithms are classified as discriminative models. Discriminative models are used to classify data by learning the boundaries between classes, which is the core functionality of random forest algorithms. They are used for classification and regression tasks by aggregating the results of multiple decision trees to make accurate predictions.

Reference: The AIGP Body of Knowledge explains that discriminative models, including random forest algorithms, are designed to distinguish between different classes in the data, making them effective for various predictive modeling tasks.

### NEW QUESTION 3

- (Topic 1)

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

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

**Answer:** B

#### Explanation:

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

### NEW QUESTION 4

- (Topic 1)

#### CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

**Answer: C**

**Explanation:**

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

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

**NEW QUESTION 5**

- (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 6**

- (Topic 1)

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

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

**Answer: A**

**Explanation:**

The GDPR's transparency principle requires that when personal data is processed for automated decision-making, including profiling, data subjects must be informed about the existence of such automated decision-making. Additionally, they must be provided with meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for them. This requirement ensures that data subjects are fully aware of how their personal data is being used and the potential impacts, thereby promoting transparency and trust in the processing activities.

**NEW QUESTION 7**

- (Topic 1)

The OECD's Ethical AI Governance Framework is a self-regulation model that proposes to prevent societal harms by?

- A. Establishing explainability criteria to responsibly source and use data to train AI systems.
- B. Defining requirements specific to each industry sector and high-risk AI domain.
- C. Focusing on AI technical design and post-deployment monitoring.
- D. Balancing AI innovation with ethical considerations.

**Answer: D**

**Explanation:**

The OECD's Ethical AI Governance Framework aims to ensure that AI development and deployment are carried out ethically while fostering innovation. The framework includes principles like transparency, accountability, and human rights protections to prevent societal harm. It does not focus solely on technical design or post-deployment monitoring (C), nor does it establish industry-specific requirements (B). While explainability is important, the primary goal is to balance innovation with ethical considerations (D).

**NEW QUESTION 8**

- (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 9

- (Topic 1)

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

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

**Answer:** C

#### Explanation:

The EU AI Act explicitly prohibits the use of AI systems for social scoring by public authorities, as it can lead to discrimination and unfair treatment of individuals based on their social behavior or perceived trustworthiness. While AI can be used to promote equitable distribution of welfare benefits, manage border control, and even detect an individual's intent for law enforcement purposes (within strict regulatory and ethical boundaries), implementing social scoring systems is not permissible under the Act due to the significant risks to fundamental rights and freedoms.

#### NEW QUESTION 10

- (Topic 1)

##### CASE STUDY

Please use the following answer the next question:

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 10

- (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 14

- (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 19**

- (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 23**

- (Topic 1)

**CASE STUDY**

Please use the following answer the next question:

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

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

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

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

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

**Answer: B**

**Explanation:**

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

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

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

**NEW QUESTION 26**

- (Topic 1)

Machine learning is best described as a type of algorithm by which?

- A. Systems can mimic human intelligence with the goal of replacing humans.
- B. Systems can automatically improve from experience through predictive patterns.
- C. Statistical inferences are drawn from a sample with the goal of predicting human intelligence.
- D. Previously unknown properties are discovered in data and used to predict and make improvements in the data.

**Answer: B**

**Explanation:**

Machine learning (ML) is a subset of artificial intelligence (AI) where systems use data to learn and improve over time without being explicitly programmed. Option B accurately describes machine learning by stating that systems can automatically improve from

experience through predictive patterns. This aligns with the fundamental concept of ML where algorithms analyze data, recognize patterns, and make decisions with minimal human intervention. Reference: AIGP BODY OF KNOWLEDGE, which covers the basics of AI and machine learning concepts.

### NEW QUESTION 30

- (Topic 1)

A company developed AI technology that can analyze text, video, images and sound to tag content, including the names of animals, humans and objects. What type of AI is this technology classified as?

- A. Deductive inference.
- B. Multi-modal model.
- C. Transformative AI.
- D. Expert system.

**Answer:** B

#### Explanation:

A multi-modal model is an AI system that can process and analyze multiple types of data, such as text, video, images, and sound. This type of AI integrates different data sources to enhance its understanding and decision-making capabilities. In the given scenario, the AI technology that tags content including names of animals, humans, and objects falls under this category. Reference: AIGP BODY OF KNOWLEDGE, which outlines the capabilities and use cases of multi-modal models.

### NEW QUESTION 35

- (Topic 1)

A company is working to develop a self-driving car that can independently decide the appropriate route to take the driver after the driver provides an address. If they want to make this self-driving car "strong" AI, as opposed to "weak," the engineers would also need to ensure?

- A. That the AI has full human cognitive abilities that can independently decide where to take the driver.
- B. That they have obtained appropriate intellectual property (IP) licenses to use data for training the AI.
- C. That the AI has strong cybersecurity to prevent malicious actors from taking control of the car.
- D. That the AI can differentiate among ethnic backgrounds of pedestrians.

**Answer:** A

#### Explanation:

Strong AI, also known as artificial general intelligence (AGI), refers to AI that possesses the ability to understand, learn, and apply intelligence across a broad range of tasks, similar to human cognitive abilities. For the self-driving car to be classified as "strong" AI, it would need to possess full human cognitive abilities to make independent decisions beyond pre-programmed instructions. Reference: AIGP BODY OF KNOWLEDGE and AI classifications.

### NEW QUESTION 39

- (Topic 1)

A company is creating a mobile app to enable individuals to upload images and videos, and analyze this data using ML to provide lifestyle improvement recommendations. The signup form has the following data fields:

\* 1.First name 2.Last name 3.Mobile number 4.Email ID 5.New password 6.Date of birth 7.Gender

In addition, the app obtains a device's IP address and location information while in use. What GDPR privacy principles does this violate?

- A. Purpose Limitation and Data Minimization.
- B. Accountability and Lawfulness.
- C. Transparency and Accuracy.
- D. Integrity and Confidentiality.

**Answer:** A

#### Explanation:

The GDPR privacy principles that this scenario violates are Purpose Limitation and Data Minimization. Purpose Limitation requires that personal data be collected for specified, explicit, and legitimate purposes and not further processed in a manner that is incompatible with those purposes. Data Minimization mandates that personal data collected should be adequate, relevant, and limited to what is necessary in relation to the purposes for which they are processed. In this case, collecting extensive personal information (e.g., IP address, location, gender) and potentially using it beyond the necessary scope for the app's functionality could violate these principles by collecting more data than needed and possibly using it for purposes not originally intended.

### NEW QUESTION 40

- (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 41

- (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 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.

When notifying an accused perpetrator, what additional information should a police officer provide about the use of the AI system?

- A. Information about the accuracy of the AI system.
- B. Information about how the accused can oppose the charges.
- C. Information about the composition of the training data of the system.
- D. Information about how the individual was identified by the AI system.

**Answer: D**

**Explanation:**

When notifying an accused perpetrator, the police officer should provide information about how the individual was identified by the AI system. This transparency is crucial for maintaining trust and ensuring that the accused understands the basis of the charges against them. Information about the accuracy, how to oppose the charges, and the composition of the training data, while potentially relevant, do not directly address the immediate need for the accused to understand the specific process that led to their identification. Reference: AIGP Body of Knowledge on AI Transparency and Explainability.

**NEW QUESTION 44**

- (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 48**

- (Topic 2)

During the development of semi-autonomous vehicles, various failures occurred as a result of the sensors misinterpreting environmental surroundings, such as sunlight.

These failures are an example of?

- A. Hallucination.
- B. Brittleness.
- C. Uncertainty.
- D. Forgetting.

**Answer: B**

**Explanation:**

The failures in semi-autonomous vehicles due to sensors misinterpreting environmental surroundings, such as sunlight, are examples of brittleness. Brittleness in AI systems refers to their inability to handle variations in input data or unexpected conditions, leading to failures when the system encounters situations that were not adequately covered during training. These systems perform well under specific conditions but fail when those conditions change. Reference: AIGP Body of Knowledge on AI System Robustness and Failures.

**NEW QUESTION 50**

- (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 52**

- (Topic 2)

A company initially intended to use a large data set containing personal information to train an AI 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 53**

- (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 58**

- (Topic 2)

All of the following are included within the scope of post-deployment AI maintenance EXCEPT?

- A. Ensuring that all model components are subject a control framework.
- B. Dedicating experts to continually monitor the model output.
- C. Evaluating the need for an audit under certain standards.
- D. Defining thresholds to conduct new impact assessments.

**Answer:** D

**Explanation:**

Post-deployment AI maintenance typically includes ensuring that all model components are subject to a control framework, dedicating experts to continually monitor the model output, and evaluating the need for audits under certain standards. However, defining thresholds to conduct new impact assessments is usually part of the initial deployment and ongoing governance processes rather than a maintenance activity. Maintenance focuses more on the operational aspects of the AI system rather than setting new thresholds for impact assessments.

Reference: AIGP BODY OF KNOWLEDGE, sections discussing AI lifecycle management and post-deployment activities.

**NEW QUESTION 62**

- (Topic 2)

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

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

**Answer:** A

**Explanation:**

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

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

**NEW QUESTION 63**

- (Topic 2)

**CASE STUDY**

Please use the following answer the next question:

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

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

The healthcare network intends to retain a cloud provider to host the solution and a consulting firm to help develop the algorithm using the healthcare network's existing data

and de-identified data that is licensed from a large US clinical research partner.

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 67**

- (Topic 2)

**CASE STUDY**

Please use the following answer the next question:

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

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

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

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

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 70**

- (Topic 2)

Which of the following 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 71**

- (Topic 2)

Which of the following elements of feature engineering is most important to mitigate the potential bias in an AI system?

- A. Feature selection.
- B. Feature validation.
- C. Feature transformation.
- D. Feature importance analysis.

**Answer:** A

**Explanation:**

Feature selection is the most important element of feature engineering to mitigate potential bias in an AI system. This process involves choosing the most relevant and representative features from the data set, which directly affects the model's performance and fairness. By carefully selecting features, data scientists can reduce the influence of biased or irrelevant attributes, ensuring that the AI system is more accurate and equitable. Proper feature selection helps in eliminating biases that might stem from socio-demographic factors or other sensitive variables, leading to a more balanced and fair AI model. Reference: AIGP Body of Knowledge on Fairness in AI and Feature Engineering.

**NEW QUESTION 72**

- (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.

The best human oversight mechanism for the police department to implement is that a police officer should?

- A. Explain to the accused how the AI system works.
- B. Confirm the AI recommendation prior to sentencing.
- C. Ensure an accused is given notice that the AI system was used.
- D. Consider the AI recommendation as part of the criminal investigation.

**Answer:** D

**Explanation:**

The best human oversight mechanism for the police department to implement is for a police officer to consider the AI recommendation as part of the criminal investigation. This ensures that the AI system's output is used as a tool to aid human decision-making rather than replace it. The police officer should integrate the AI's insights with other evidence and contextual information to make informed decisions, maintaining a balance between technological aid and human judgment.

Reference: AIGP Body of Knowledge on AI Integration and Human Oversight.

### NEW QUESTION 77

- (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 81

- (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 83

- (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 86

- (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 87**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### AIGP Practice Exam Features:

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

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