



**IAPP**

## **Exam Questions AIGP**

Artificial Intelligence Governance Professional

### NEW QUESTION 1

- (Topic 1)

#### CASE STUDY

Please use the following answer the next question:

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

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

ABC and the cloud provider have completed training and testing the LLM, performed a readiness assessment, and made the decision to deploy the LLM into production. ABC has designated an internal compliance team to monitor the model during the first month, specifically to evaluate the accuracy, fairness, and reliability of its output. After the first month in production, ABC realizes that the LLM declines a higher percentage of women's loan applications due primarily to women historically receiving lower salaries than men.

The best approach to enable a customer who wants information on the AI model's parameters for underwriting purposes is to provide?

- A. A transparency notice.
- B. An opt-out mechanism.
- C. Detailed terms of service.
- D. Customer service support.

**Answer:** A

#### Explanation:

The best approach to enable a customer who wants information on the AI model's parameters for underwriting purposes is to provide a transparency notice. This notice should explain the nature of the AI system, how it uses customer data, and the decision-making process it follows. Providing a transparency notice is crucial for maintaining trust and compliance with regulatory requirements regarding the transparency and accountability of AI systems.

Reference: According to the AIGP Body of Knowledge, transparency in AI systems is essential to ensure that stakeholders, including customers, understand how their data is being used and how decisions are made. This aligns with ethical principles of AI governance, ensuring that customers are informed and can make knowledgeable decisions regarding their interactions with AI systems.

### NEW QUESTION 2

- (Topic 1)

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 3

- (Topic 1)

All of the following are penalties and enforcements outlined in the EU AI Act EXCEPT?

- A. Fines for SMEs and startups will be proportionally capped.
- B. Rules on General Purpose AI will apply after 6 months as a specific provision.
- C. The AI Pact will act as a transitional bridge until the Regulations are fully enacted.
- D. Fines for violations of banned AI applications will be €35 million or 7% global annual turnover (whichever is higher).

**Answer:** C

#### Explanation:

The EU AI Act outlines specific penalties and enforcement mechanisms to ensure compliance with its regulations. Among these, fines for violations of banned AI applications can be as high as €35 million or 7% of the global annual turnover of the offending organization, whichever is higher. Proportional caps on fines are applied to SMEs and startups to ensure fairness. General Purpose AI rules are to apply after a 6-month period as a specific provision to ensure that stakeholders have adequate time to comply. However, there is no provision for an "AI Pact" acting as a transitional bridge until the regulations are fully enacted, making option C the correct answer.

### 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 to a human underwriter for final review.

ABC and the cloud provider have completed training and testing the LLM, performed a readiness assessment, and made the decision to deploy the LLM into production. ABC has designated an internal compliance team to monitor the model during the first month, specifically to evaluate the accuracy, fairness, and reliability of its output. After the first month in production, ABC realizes that the LLM declines a higher percentage of women's loan applications due primarily to women historically receiving lower salaries than men.

What is the best strategy to mitigate the bias uncovered in the loan applications?

- A. Retrain the model with data that reflects demographic parity.
- B. Procure a third-party statistical bias assessment tool.
- C. Document all instances of bias in the data set.
- D. Delete all gender-based data in the data set.

**Answer:** A

**Explanation:**

Retraining the model with data that reflects demographic parity is the best strategy to mitigate the bias uncovered in the loan applications. This approach addresses the root cause of the bias by ensuring that the training data is representative and balanced, leading to more equitable decision-making by the AI model. Reference: The AIGP Body of Knowledge stresses the importance of using high-quality, unbiased training data to develop fair and reliable AI systems. Retraining the model with balanced data helps correct biases that arise from historical inequalities, ensuring that the AI system makes decisions based on equitable criteria.

**NEW QUESTION 5**

- (Topic 1)

An EU bank intends to launch a multi-modal AI platform for customer engagement and automated decision-making assist with the opening of bank accounts. The platform has been subject to thorough risk assessments and testing, where it proves to be effective in not discriminating against any individual on the basis of a protected class.

What additional obligations must the bank fulfill prior to deployment?

- A. The bank must obtain explicit consent from users under the privacy Directive.
- B. The bank must disclose how the AI system works under the EII Digital Services Act.
- C. The bank must subject the AI system an adequacy decision and publish its appropriate safeguards.
- D. The bank must disclose the use of the AI system and implement suitable measures for users to contest automated decision-making.

**Answer:** D

**Explanation:**

Under the EU regulations, particularly the GDPR, banks using AI for decision-making must inform users about the use of AI and provide mechanisms for users to contest decisions. This is part of ensuring transparency and accountability in automated processing. Explicit consent under the privacy directive (A) and disclosing under the Digital Services Act (B) are not specifically required in this context. An adequacy decision is related to data transfers outside the EU (C).

**NEW QUESTION 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 explain ability 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)

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 9**

- (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 10**

- (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:

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.

The frameworks that would be most appropriate for XYZ's governance needs would be the NIST AI Risk Management Framework and?

- A. NIST Information Security Risk (NIST SP 800-39).
- B. NIST Cyber Security Risk Management Framework (CSF 2.0).
- C. IEEE Ethical System Design Risk Management Framework (IEEE 7000-21).
- D. Human Rights, Democracy, and Rule of Law Impact Assessment (HUDERIA).

**Answer:** C

**Explanation:**

The IEEE Ethical System Design Risk Management Framework (IEEE 7000-21) would be most appropriate for XYZ Corp's governance needs in addition to the NIST AI Risk Management Framework. The IEEE framework specifically addresses ethical concerns during system design, which is crucial for ensuring the responsible use of AI in hiring. It complements the NIST framework by focusing on ethical risk management, aligning well with XYZ Corp's goals of deploying AI responsibly and mitigating associated risks.

**NEW QUESTION 15**

- (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 19**

- (Topic 1)

**CASE STUDY**

Please use the following answer the next question:

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

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

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

What is the best reason for GVC to offer students the choice to utilize generative AI in limited, defined circumstances?

- A. To enable students to learn how to manage their time.
- B. To enable students to learn about performing research.
- C. To enable students to learn about practical applications of AI.
- D. To enable students to learn how to use AI as a supportive educational tool.

**Answer:** D

**Explanation:**

The best reason for GVC to offer students the choice to utilize generative AI in limited, defined circumstances is to enable students to learn how to use AI as a supportive educational tool. By integrating AI in a controlled manner, students can learn the practical applications of AI and develop skills to use AI responsibly and effectively in their educational pursuits.

Reference: The AIGP Body of Knowledge highlights the importance of teaching students about AI's practical applications and the responsible use of AI technologies. This aligns with the goal of fostering a better understanding of AI's role and its potential benefits in various contexts, including education.

**NEW QUESTION 23**

- (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.

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.

The organization has a large procurement team that is responsible for the contracting of technology solutions. One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company are responsible for integrating and deploying technology solutions into the organization's operations in a responsible, cost-effective manner.

The organization is aware of the risks presented by AI hiring tools and wants to mitigate

them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 25**

- (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 30**

- (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 31**

- (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 32**

- (Topic 1)

An AI system that maintains its level of performance within defined acceptable limits despite real world or adversarial conditions would be described as?

- A. Robust.
- B. Reliable.
- C. Resilient.
- D. Reinforced.

**Answer:** C

**Explanation:**

An AI system that maintains its level of performance within defined acceptable limits despite real-world or adversarial conditions is described as resilient. Resilience in AI refers to the system's ability to withstand and recover from unexpected challenges, such as cyber-attacks, hardware failures, or unusual input data. This characteristic ensures that the AI system can continue to function effectively and reliably in various conditions, maintaining performance and integrity. Robustness, on the other hand, focuses on the system's strength against errors, while reliability ensures consistent performance over time. Resilience combines these aspects with the capacity to adapt and recover.

**NEW QUESTION 35**

- (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 40**

- (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 42**

- (Topic 1)

Which of the following is a subcategory of AI and machine learning that uses labeled datasets to train algorithms?

- A. Segmentation.
- B. Generative AI.
- C. Expert systems.
- D. Supervised learning.

**Answer:** D

**Explanation:**

Supervised learning is a subcategory of AI and machine learning where labeled datasets are used to train algorithms. This process involves feeding the algorithm a dataset where the input-output pairs are known, allowing the algorithm to learn and make predictions or decisions based on new, unseen data. Reference: AIGP BODY OF KNOWLEDGE, which describes supervised learning as a model trained on labeled data (e.g., text recognition, detecting spam in emails).

**NEW QUESTION 43**

- (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 47**

- (Topic 2)

What is the best reason for a company adopt a policy that prohibits the use of generative AI?

- A. Avoid using technology that cannot be monetized.
- B. Avoid needing to identify and hire qualified resources.
- C. Avoid the time necessary to train employees on acceptable use.
- D. Avoid accidental disclosure to its confidential and proprietary information.

**Answer:** D

**Explanation:**

The primary concern for a company adopting a policy prohibiting the use of generative AI is the risk of accidental disclosure of confidential and proprietary information. Generative AI tools can inadvertently leak sensitive data during the creation process or through data sharing. This risk outweighs the other reasons listed, as protecting sensitive information is critical to maintaining the company's competitive edge and legal compliance. This rationale is discussed in the sections on risk management and data privacy in the IAPP AIGP Body of Knowledge.

**NEW QUESTION 48**

- (Topic 2)

To maintain fairness in a deployed system, it is most important to?

- A. Protect against loss of personal data in the model.
- B. Monitor for data drift that may affect performance and accuracy.
- C. Detect anomalies outside established metrics that require new training data.
- D. Optimize computational resources and data to ensure efficiency and scalability.

**Answer:** B

**Explanation:**

To maintain fairness in a deployed system, it is crucial to monitor for data drift that may affect performance and accuracy. Data drift occurs when the statistical properties of the input data change over time, which can lead to a decline in model performance. Continuous monitoring and updating of the model with new data

ensure that it remains fair and accurate, adapting to any changes in the data distribution. Reference: AIGP Body of Knowledge on Post-Deployment Monitoring and Model Maintenance.

#### NEW QUESTION 53

- (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 58

- (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 62

- (Topic 2)

Retraining an LLM can be necessary for all of the following reasons EXCEPT?

- A. To minimize degradation in prediction accuracy due to changes in data.
- B. Adjust the model's hyper parameters specific use case.
- C. Account for new interpretations of the same data.
- D. To ensure interpretability of the model's predictions.

**Answer: D**

#### Explanation:

Retraining an LLM (Large Language Model) is primarily done to improve or maintain its performance as data changes over time, to fine-tune it for specific use cases, and to incorporate new data interpretations to enhance accuracy and relevance. However, ensuring interpretability of the model's predictions is not typically a reason for retraining. Interpretability relates to how easily the outputs of the model can be understood and explained, which is generally addressed through different techniques or methods rather than through the retraining process itself. References to this can be found in the IAPP AIGP Body of Knowledge discussing model retraining and interpretability as separate concepts.

#### NEW QUESTION 64

- (Topic 2)

What is the best method to proactively train an LLM so that there is mathematical proof that no specific piece of training data has more than a negligible effect on the model or its output?

- A. Clustering.



- B. Transfer learning.
- C. Differential privacy.
- D. Data compartmentalization.

**Answer:** C

**Explanation:**

Differential privacy is a technique used to ensure that the inclusion or exclusion of a single data point does not significantly affect the outcome of any analysis, providing a way to mathematically prove that no specific piece of training data has more than a negligible effect on the model or its output. This is achieved by introducing randomness into the data or the algorithms processing the data. In the context of training large language models (LLMs), differential privacy helps in protecting individual data points while still enabling the model to learn effectively. By adding noise to the training process, differential privacy provides strong guarantees about the privacy of the training data.

Reference: AIGP BODY OF KNOWLEDGE, pages related to data privacy and security in model training.

**NEW QUESTION 67**

- (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 70**

- (Topic 2)

According to November 2023 White House Executive Order, which of the following best describes the guidance given to governmental agencies on the use of generative AI as a workplace tool?

- A. Limit access to specific uses of generative AI.
- B. Impose a general ban on the use of generative AI.
- C. Limit access of generative AI to engineers and developers.
- D. Impose a ban on the use of generative AI in agencies that protect national security.

**Answer:** A

**Explanation:**

The November 2023 White House Executive Order provides guidance that governmental agencies should limit access to specific uses of generative AI. This means that generative AI tools should be used in a controlled manner, where their applications are restricted to well-defined, approved use cases that ensure the security, privacy, and ethical considerations are adequately addressed. This approach allows for the benefits of generative AI to be harnessed while mitigating potential risks and abuses.

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

**NEW QUESTION 74**

- (Topic 2)

The White House Executive Order from November 2023 requires companies that develop dual-use foundation models to provide reports to the federal government about all of the following EXCEPT?

- A. Any current training or development of dual-use foundation models.
- B. The results of red-team testing of each dual-use foundation model.
- C. Any environmental impact study for each dual-use foundation model.
- D. The physical and cybersecurity protection measures of their dual-use foundation models.

**Answer:** C

**Explanation:**

The White House Executive Order from November 2023 requires companies developing dual-use foundation models to report on their current training or development activities, the results of red-team testing, and the physical and cybersecurity protection measures. However, it does not mandate reports on environmental impact studies for each dual-use foundation model. While environmental considerations are important, they are not specified in this context as a reporting requirement under this Executive Order.

Reference: AIGP BODY OF KNOWLEDGE, sections on compliance and reporting requirements, and the White House Executive Order of November 2023.

**NEW QUESTION 79**

- (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 80**

- (Topic 2)

The most important factor in ensuring fairness when training an AI system is?

- A. The architecture and model selection.
- B. The data labeling and classification.
- C. The data attributes and variability.
- D. The model accuracy and scale.

**Answer:** C

**Explanation:**

Ensuring fairness when training an AI system largely depends on the data attributes and variability. This involves having a diverse and representative dataset that accurately reflects the population the AI system will serve. Fairness can be compromised if the data is biased or lacks variability, as the model may learn and perpetuate these biases.

Diverse data attributes ensure that the model learns from a wide range of examples, reducing the risk of biased predictions. Reference: AIGP Body of Knowledge on Ethical AI Principles and Data Management.

**NEW QUESTION 81**

- (Topic 2)

During the planning and design phases of the AI development life cycle, bias can be reduced by all of the following EXCEPT?

- A. Stakeholder involvement.
- B. Feature selection.
- C. Human oversight.
- D. Data collection.

**Answer:** B

**Explanation:**

Bias in AI can be reduced during the planning and design phases through stakeholder involvement, human oversight, and careful data collection. While feature selection is critical in the development phase, it does not specifically occur during planning and design. Ensuring diverse stakeholder involvement and human oversight helps identify and mitigate potential biases early, and data collection ensures a representative dataset. Reference: AIGP Body of Knowledge on AI Development Lifecycle and Bias Mitigation.

**NEW QUESTION 86**

- (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 90**

- (Topic 2)

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

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

**Answer: D**

**Explanation:**

Deploying a challenger AI model alongside a champion model is a strategy used to compare the performance of different models in a real-world environment. This approach helps in providing a framework to consider alternatives to the champion model, automating real-time monitoring of the champion model, and performing testing on the champion model. However, retraining the champion model is not a reason to deploy a challenger model. Retraining is a separate process that involves updating the champion model with new data or techniques, which is not related to the use of a challenger model.

Reference: AIGP BODY OF KNOWLEDGE, sections on model evaluation and management.

**NEW QUESTION 91**

- (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 94**

- (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 98**

- (Topic 2)

Which of the following deployments of generative AI best respects intellectual property rights?

- A. The system produces content that is modified to closely resemble copyrighted work.
- B. The system categorizes and applies filters to content based on licensing terms.
- C. The system provides attribution to creators of publicly available information.
- D. The system produces content that includes trademarks and copyrights.

**Answer: B**

**Explanation:**

Respecting intellectual property rights means adhering to licensing terms and ensuring that generated content complies with these terms. A system that categorizes and applies filters based on licensing terms ensures that content is used legally and ethically, respecting the rights of content creators. While providing attribution is important, categorization and application of filters based on licensing terms are more directly tied to compliance with intellectual property laws. This principle is elaborated in the IAPP AIGP Body of Knowledge sections on intellectual property and compliance.

NEW QUESTION 103  
.....



## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### AIGP Practice Exam Features:

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

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