



# IAPP

## Exam Questions AIGP

Artificial Intelligence Governance Professional

## About ExamBible

*[Your Partner of IT Exam](#)*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

## 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.

Each of the following steps would support fairness testing by the compliance team during the first month in production EXCEPT?

- A. Validating a similar level of decision-making across different demographic groups.
- B. Providing the loan applicants with information about the model capabilities and limitations.
- C. Identifying if additional training data should be collected for specific demographic groups.
- D. Using tools to help understand factors that may account for differences in decision-making.

**Answer: B**

### Explanation:

Providing the loan applicants with information about the model capabilities and limitations would not directly support fairness testing by the compliance team. Fairness testing focuses on evaluating the model's decisions for biases and ensuring equitable treatment across different demographic groups, rather than informing applicants about the model.

Reference: The AIGP Body of Knowledge outlines that fairness testing involves technical assessments such as validating decision-making consistency across demographics and using tools to understand decision factors. While transparency to applicants is important for ethical AI use, it does not contribute directly to the technical process of fairness testing.

## NEW QUESTION 2

- (Topic 1)

### CASE STUDY

Please use the following answer the next question:

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

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

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

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 3

- (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 other stakeholder groups should be involved in the selection and implementation of the AI hiring tool?

- A. Finance and Legal.

- B. Marketing and Compliance.
- C. Supply Chain and Marketing.
- D. Litigation and Product Development.

**Answer:** A

**Explanation:**

In the selection and implementation of the AI hiring tool, involving Finance and Legal is crucial. The Finance team is essential for assessing cost implications, budget considerations, and financial risks. The Legal team is necessary to ensure compliance with applicable laws and regulations, including those related to data privacy, employment, and anti-discrimination. Involving these stakeholders ensures a comprehensive evaluation of both the financial viability and legal compliance of the AI tool, mitigating potential risks and aligning with organizational objectives and regulatory requirements.

**NEW QUESTION 4**

- (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 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)

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 7**

- (Topic 1)

According to the Singapore Model AI Governance Framework, all of the following are recommended measures to promote the responsible use of AI EXCEPT?

- A. Determining the level of human involvement in algorithmic decision-making.
- B. Adapting the existing governance structure algorithmic decision-making.
- C. Employing human-over-the-loop protocols for high-risk systems.
- D. Establishing communications and collaboration among stakeholders.

**Answer:** C

**Explanation:**

The Singapore Model AI Governance Framework recommends several measures to promote the responsible use of AI, such as determining the level of human involvement in decision-making, adapting governance structures, and establishing communications and collaboration among stakeholders. However, employing

human-over-the-loop protocols is not specifically mentioned in this framework. The focus is more on integrating human oversight appropriately within the decision-making process rather than exclusively employing such protocols. Reference: AIGP Body of Knowledge, section on AI governance frameworks.

#### NEW QUESTION 8

- (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 9

- (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 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 13

- (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 14

- (Topic 1)

You asked a generative AI tool to recommend new restaurants to explore in Boston, Massachusetts that have a specialty Italian dish made in a traditional fashion without spinach and wine. The generative AI tool recommended five restaurants for you to visit.

After looking up the restaurants, you discovered one restaurant did not exist and two others did not have the dish.

This information provided by the generative AI tool is an example of what is commonly called?

- A. Prompt injection.
- B. Model collapse.
- C. Hallucination.
- D. Overfitting.

**Answer:** C

#### Explanation:

In the context of AI, particularly generative models, "hallucination" refers to the generation of outputs that are not based on the training data and are factually incorrect or non-existent. The scenario described involves the generative AI tool providing incorrect and non-existent information about restaurants, which fits the definition of hallucination. Reference: AIGP BODY OF KNOWLEDGE and various AI literature discussing the limitations and challenges of generative AI models.

### NEW QUESTION 16

- (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 18

- (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 23**

- (Topic 1)

According to the GDPR, an individual has the right to have a human confirm or replace an automated decision unless that automated decision?

- A. Is authorized with the data subject's explicit consent.
- B. Is authorized by applicable E.U. law and includes suitable safeguards.
- C. Is deemed to solely benefit the individual and includes documented legitimate interests.
- D. Is necessary for entering into or performing under a contract between the data subject and data controller.

**Answer:** A

**Explanation:**

According to the GDPR, individuals have the right to not be subject to a decision based solely on automated processing, including profiling, which produces legal effects or similarly significantly affects them. However, there are exceptions to this right, one of which is when the decision is based on the data subject's explicit consent. This means that if an individual explicitly consents to the automated decision-making process, there is no requirement for human intervention to confirm or replace the decision. This exception ensures that individuals can have control over automated decisions that affect them, provided they have given clear and informed consent.

**NEW QUESTION 26**

- (Topic 1)

Which of the following disclosures is NOT required for an EU organization that developed and deployed a high-risk AI system?

- A. The human oversight measures employed.
- B. How an individual may contest a decision.
- C. The location(s) where data is stored.
- D. The fact that an AI system is being used.

**Answer:** C

**Explanation:**

Under the EU AI Act, organizations that develop and deploy high-risk AI systems are required to provide several key disclosures to ensure transparency and accountability. These include the human oversight measures employed, how individuals can contest decisions made by the AI system, and informing individuals that an AI system is being used. However, there is no specific requirement to disclose the exact locations where data is stored. The focus of the Act is on the transparency of the AI system's operation and its impact on individuals, rather than on the technical details of data storage locations.

**NEW QUESTION 27**

- (Topic 1)

What is the key feature of Graphical Processing Units (GPUs) that makes them well-suited to running AI applications?

- A. GPUs run many tasks concurrently, resulting in faster processing.
- B. GPUs can access memory quickly, resulting in lower latency than CPUs.
- C. GPUs can run every task on a computer, making them more robust than CPUs.
- D. The number of transistors on GPUs doubles every two years, making the chips smaller and lighter.

**Answer:** A

**Explanation:**

GPUs (Graphical Processing Units) are well-suited to running AI applications due to their ability to run many tasks concurrently, which significantly enhances processing speed. This parallel processing capability makes GPUs ideal for handling the large-scale computations required in AI and deep learning tasks. Reference: AIGP BODY OF KNOWLEDGE, which explains the importance of compute infrastructure in AI applications.

**NEW QUESTION 28**

- (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 29**

- (Topic 1)

What is the primary purpose of an AI impact assessment?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 30**

- (Topic 1)

**CASE STUDY**

Please use the following answer the next question:

XYZ Corp., a premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent. The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.

It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.

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

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

The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

All of the following are potential negative consequences created by using the AI tool when making hiring decisions EXCEPT?

- A. Reputational harm.
- B. Civil rights violations.
- C. Discriminatory treatment.
- D. Intellectual property infringement.

**Answer:** D

**Explanation:**

The potential negative consequences of using an AI tool in hiring include reputational harm (A), civil rights violations (B), and discriminatory treatment (C). These issues stem from biases in the AI system or its misuse, which can lead to unfair hiring practices and legal liabilities. Intellectual property infringement (D) is not a typical consequence of using AI in hiring, as it relates to the unauthorized use of protected intellectual property, which is not directly relevant to the hiring process or the potential biases within AI tools.

**NEW QUESTION 34**

- (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 36**

- (Topic 1)

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

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

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

**Answer:** D

**Explanation:**

The U.S. mortgage company's AI platform relates to housing and credit, making the Fair Housing Act (A), Fair Credit Reporting Act (B), and Equal Credit Opportunity Act (C) relevant. Title VII of the Civil Rights Act of 1964 deals with employment discrimination and is not directly relevant to the mortgage application context (D).



### NEW QUESTION 39

- (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 43

- (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 48

- (Topic 2)

Which type of existing assessment could best be leveraged to create an AI impact assessment?

- A. A safety impact assessment.
- B. A privacy impact assessment.
- C. A security impact assessment.
- D. An environmental impact assessment.

**Answer:** B

#### **Explanation:**

A privacy impact assessment (PIA) can be effectively leveraged to create an AI impact assessment. A PIA evaluates the potential privacy risks associated with the use of personal data and helps in implementing measures to mitigate those risks. Since AI systems often involve processing large amounts of personal data, the principles and methodologies of a PIA are highly applicable and can be extended to assess broader impacts, including ethical, social, and legal implications of AI. Reference: AIGP Body of Knowledge on Impact Assessments.

### NEW QUESTION 53

- (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 57

- (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 59**

- (Topic 2)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 60**

- (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 64**

- (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 66**

- (Topic 2)

Which of the following is the least relevant consideration in assessing whether users should be given the right to opt out from an AI system?

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

**Answer:** D

**Explanation:**

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

Reference: AIGP BODY OF KNOWLEDGE, sections discussing user rights and ethical considerations in AI.

**NEW QUESTION 70**

- (Topic 2)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 71**

- (Topic 2)

After completing model testing and validation, which of the following is the most important step that an organization takes prior to deploying the model into production?

- A. Perform a readiness assessment.
- B. Define a model-validation methodology.
- C. Document maintenance teams and processes.
- D. Identify known edge cases to monitor post-deployment.

**Answer:** A

**Explanation:**

After completing model testing and validation, the most important step prior to deploying the model into production is to perform a readiness assessment. This assessment ensures that the model is fully prepared for deployment, addressing any potential issues related to infrastructure, performance, security, and compliance. It verifies that the model meets all necessary criteria for a successful launch. Other steps, such as defining a model-validation methodology, documenting maintenance teams and processes, and identifying known edge cases, are also important but come secondary to confirming overall readiness.

Reference: AIGP Body of Knowledge on Deployment Readiness.

**NEW QUESTION 74**

- (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 79**

- (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 83**

- (Topic 2)

A company plans on procuring a tool from an AI provider for its employees to use for certain business purposes.

Which contractual provision would best protect the company's intellectual property in the tool, including training and testing data?

- A. The provider will give privacy notice to individuals before using their personal data to train or test the tool.
- B. The provider will defend and indemnify the company against infringement claims.
- C. The provider will obtain and maintain insurance to cover potential claims.
- D. The provider will warrant that the tool will work as intended.

**Answer:** B

**Explanation:**

To protect the company's intellectual property, the most pertinent contractual provision is ensuring that the AI provider will defend and indemnify the company against infringement claims. This clause means the provider will take responsibility for any intellectual property disputes that arise, thereby safeguarding the company from potential legal and financial repercussions related to the use of the tool. Other options, while beneficial, do not directly address the protection of intellectual property. This concept is detailed in the contractual best practices section of the IAPP AIGP Body of Knowledge.

**NEW QUESTION 88**

- (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 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.

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 92**

- (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 95**

- (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 96**

- (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 98**

- (Topic 2)

An artist has been using an AI tool to create digital art and would like to ensure that it has copyright protection in the United States. Which of the following is most likely to enable the artist to receive copyright protection?

- A. Ensure the tool was trained using publicly available content.
- B. Obtain a representation from the AI provider on how the tool works.
- C. Provide a log of the prompts the artist used to generate the images.
- D. Update the images in a creative way to demonstrate that it is the artist's.

**Answer: D**

**Explanation:**

For the artist to receive copyright protection, the most effective approach is to demonstrate that the final artwork includes sufficient creative input by the artist. By updating or altering the images in a way that reflects the artist's personal creativity, the artist can claim originality, which is a core requirement for copyright protection under U.S. law. The other options do not directly address the originality and creative input required for copyright. This is highlighted in the sections on copyright protection in the IAPP AIGP Body of Knowledge.

**NEW QUESTION 103**

- (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 104**

.....

## Relate Links

**100% Pass Your AIGP Exam with ExamBible Prep Materials**

<https://www.exambible.com/AIGP-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>