



## **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:

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 2

- (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 3

- (Topic 1)

The framework set forth in the White House Blueprint for an AI Bill of Rights addresses all of the following EXCEPT?

- A. Human alternatives, consideration and fallback.
- B. High-risk mitigation standards.
- C. Safe and effective systems.
- D. Data privacy.

**Answer:** B

#### Explanation:

The White House Blueprint for an AI Bill of Rights focuses on protecting civil rights, privacy, and ensuring AI systems are safe and effective. It includes principles like data privacy (D), human alternatives (A), and safe and effective systems (C). However, it does not specifically address high-risk mitigation standards as a distinct category (B).

### NEW QUESTION 4

- (Topic 1)

All of the following are common optimization techniques in deep learning to determine weights that represent the strength of the connection between artificial neurons EXCEPT?

- A. Gradient descent, which initially sets weights arbitrary values, and then at each step changes them.
- B. Momentum, which improves the convergence speed and stability of neural network training.
- C. Autoregression, which analyzes and makes predictions about time-series data.
- D. Backpropagation, which starts from the last layer working backwards.

**Answer:** C

#### Explanation:

Autoregression is not a common optimization technique in deep learning to determine weights for artificial neurons. Common techniques include gradient descent, momentum, and backpropagation. Autoregression is more commonly

associated with time-series analysis and forecasting rather than neural network optimization. Reference: AIGP BODY OF KNOWLEDGE, which discusses common optimization techniques used in deep learning.

#### NEW QUESTION 5

- (Topic 1)

According to the EU AI Act, providers of what kind of machine learning systems will be required to register with an EU oversight agency before placing their systems in the EU market?

- A. AI systems that are harmful based on a legal risk-utility calculation.
- B. AI systems that are "strong" general intelligence.
- C. AI systems trained on sensitive personal data.
- D. AI systems that are high-risk.

**Answer:** D

#### Explanation:

According to the EU AI Act, providers of high-risk AI systems are required to register with an EU oversight agency before these systems can be placed on the market. This requirement is part of the Act's framework to ensure that high-risk AI systems comply with stringent safety, transparency, and accountability standards. High-risk systems are those that pose significant risks to health, safety, or fundamental rights. Registration with oversight agencies helps facilitate ongoing monitoring and enforcement of compliance with the Act's provisions. Systems categorized under other criteria, such as those trained on sensitive personal data or exhibiting "strong" general intelligence, also fall under scrutiny but are primarily covered under different regulatory requirements or classifications.

#### NEW QUESTION 6

- (Topic 1)

If it is possible to provide a rationale for a specific output of an AI system, that system can best be described as?

- A. Accountable.
- B. Transparent.
- C. Explainable.
- D. Reliable.

**Answer:** C

#### Explanation:

If it is possible to provide a rationale for a specific output of an AI system, that system can best be described as explainable. Explainability in AI refers to the ability to interpret and understand the decision-making process of the AI system. This involves being able to articulate the factors and logic that led to a particular output or decision. Explainability is critical for building trust, enabling users to understand and validate the AI system's actions, and ensuring compliance with ethical and regulatory standards. It also facilitates debugging and improving the system by providing insights into its behavior.

#### NEW QUESTION 7

- (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 8

- (Topic 1)

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

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 9

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

Which of the following measures should XYZ adopt to best mitigate its risk of reputational harm from using the AI tool?

- A. Test the AI tool pre- and post-deployment.
- B. Ensure the vendor assumes responsibility for all damages.
- C. Direct the procurement team to select the most economical AI tool.
- D. Continue to require XYZ's hiring personnel to manually screen all applicants.

**Answer: A**

**Explanation:**

To mitigate the risk of reputational harm from using an AI hiring tool, XYZ Corp should rigorously test the AI tool both before and after deployment. Pre-deployment testing ensures the tool works correctly and does not introduce bias or other issues. Post-deployment testing ensures the tool continues to operate as intended and adapts to any changes in data or usage patterns. This approach helps to identify and address potential issues proactively, thereby reducing the risk of reputational harm. Ensuring the vendor assumes responsibility for damages (B) does not address the root cause of potential issues, selecting the most economical tool (C) may compromise quality, and continuing manual screening (D) defeats the purpose of using the AI tool.

**NEW QUESTION 10**

- (Topic 1)

**CASE STUDY**

Please use the following answer the next question:

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

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

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

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

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

**Answer: B**

**Explanation:**

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

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

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

**NEW QUESTION 10**

- (Topic 1)

Which of the following most encourages accountability over AI systems?

- A. Determining the business objective and success criteria for the AI project.
- B. Performing due diligence on third-party AI training and testing data.
- C. Defining the roles and responsibilities of AI stakeholders.
- D. Understanding AI legal and regulatory requirements.

**Answer: C**

**Explanation:**

Defining the roles and responsibilities of AI stakeholders is crucial for encouraging accountability over AI systems. Clear delineation of who is responsible for different aspects of the AI lifecycle ensures that there is a person or team accountable for monitoring, maintaining, and addressing issues that arise. This accountability framework helps in ensuring that ethical standards and regulatory requirements are met, and it facilitates transparency and traceability in AI operations. By assigning specific roles, organizations can better manage and mitigate risks associated with AI deployment and use.

**NEW QUESTION 15**

- (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 20**

- (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 21**

- (Topic 2)

All of the following are elements of establishing a global AI governance infrastructure EXCEPT?

- A. Providing training to foster a culture that promotes ethical behavior.
- B. Creating policies and procedures to manage third-party risk.
- C. Understanding differences in norms across countries.
- D. Publicly disclosing ethical principles.

**Answer:** D

**Explanation:**

Establishing a global AI governance infrastructure involves several key elements, including providing training to foster a culture that promotes ethical behavior, creating policies and procedures to manage third-party risk, and understanding differences in norms across countries. While publicly disclosing ethical principles can enhance transparency and trust, it is not a core element necessary for the establishment of a governance infrastructure. The focus is more on internal processes and structures rather than public disclosure. Reference: AIGP Body of Knowledge on AI Governance and Infrastructure.

**NEW QUESTION 26**

- (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 28**

- (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 33**

- (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 37**

- (Topic 2)

**CASE STUDY**

Please use the following answer the next question:

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

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

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

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

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 41**

- (Topic 2)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 45**

- (Topic 2)

You are part of your organization's ML engineering team and notice that the accuracy of a model that was recently deployed into production is deteriorating. What is the best first step address this?

- A. Replace the model with a previous version.
- B. Conduct champion/challenger testing.
- C. Perform an audit of the model.
- D. Run red-teaming exercises.

**Answer: B**

**Explanation:**

When the accuracy of a model deteriorates, the best first step is to conduct champion/challenger testing. This involves deploying a new model (challenger) alongside the current model (champion) to compare their performance. This method helps identify if the new model can perform better under current conditions without immediately discarding the existing model. It provides a controlled environment to test improvements and understand the reasons behind the deterioration. This approach is preferable to directly replacing the model, performing audits, or running red-teaming exercises, which may be subsequent steps based on the findings from the champion/challenger testing.

Reference: AIGP BODY OF KNOWLEDGE, sections on model performance management and testing strategies.

**NEW QUESTION 49**

- (Topic 2)

**CASE STUDY**

Please use the following answer the next question:

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

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

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

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

What is the best reason the police department should continue to perform investigations even if the AI system scores an individual's likelihood of criminal activity at or above 90%?

- A. Because the department did not perform an impact assessment for this intended use.
- B. Because AI systems that affect fundamental civil rights should not be fully automated.
- C. Because investigations may identify additional individuals involved in the crime.
- D. Because investigations may uncover information relevant to sentencing.

**Answer: B**

**Explanation:**

The best reason for the police department to continue performing investigations even if the

AI system scores an individual's likelihood of criminal activity at or above 90% is that AI systems affecting fundamental civil rights should not be fully automated.

Human oversight is essential to ensure that decisions impacting civil liberties are made with due consideration of context and mitigating factors that an AI might not fully appreciate. This approach ensures fairness, accountability, and adherence to legal standards. Reference: AIGP Body of Knowledge on AI Ethics and Human Oversight.

**NEW QUESTION 52**

- (Topic 2)

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 55**

- (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 59**

- (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 62**

- (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 65**

- (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 68**

- (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 70

- (Topic 2)

You are the chief privacy officer of a medical research company that would like to collect and use sensitive data about cancer patients, such as their names, addresses, race and ethnic origin, medical histories, insurance claims, pharmaceutical prescriptions, eating and drinking habits and physical activity. The company will use this sensitive data to build an AI algorithm that will spot common attributes that will help predict if seemingly healthy people are more likely to get cancer. However, the company is unable to obtain consent from enough patients to sufficiently collect the minimum data to train its model. Which of the following solutions would most efficiently balance privacy concerns with the lack of available data during the testing phase?

- A. Deploy the current model and recalibrate it over time with more data.
- B. Extend the model to multi-modal ingestion with text and images.
- C. Utilize synthetic data to offset the lack of patient data.
- D. Refocus the algorithm to patients without cancer.

**Answer: C**

#### Explanation:

Utilizing synthetic data to offset the lack of patient data is an efficient solution that balances privacy concerns with the need for sufficient data to train the model. Synthetic data can be generated to simulate real patient data while avoiding the privacy issues associated with using actual patient data. This approach allows for the development and testing of the AI algorithm without compromising patient privacy, and it can be refined with real data as it becomes available. Reference: AIGP Body of Knowledge on Data Privacy and AI Model Training.

#### NEW QUESTION 75

- (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 79

- (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 84

- (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 85

- (Topic 2)

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

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

**Answer:** A

#### **Explanation:**

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

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

### NEW QUESTION 87

- (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 90

- (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 91

- (Topic 2)

Which of the following AI uses is best described as human-centric?

- A. Pattern recognition algorithms are used to improve the accuracy of weather predictions, which benefits many industries and everyday life.
- B. Autonomous robots are used to move products within a warehouse, allowing human workers to reduce physical strain and alleviate monotony.
- C. Machine learning is used for demand forecasting and inventory management, ensuring that consumers can find products they want when they want them.
- D. Virtual assistants are used adapt educational content and teaching methods to individuals, offering personalized recommendations based on ability and needs.

**Answer:** D

**Explanation:**

Human-centric AI focuses on improving the human experience by addressing individual needs and enhancing human capabilities. Option D exemplifies this by using virtual assistants to tailor educational content to each student's unique abilities and needs, thereby supporting personalized learning and improving educational outcomes. This use case directly benefits individuals by providing customized assistance and adapting to their learning pace and style, aligning with the principles of human-centric AI.

Reference: AIGP BODY OF KNOWLEDGE, sections on trustworthy AI and human-centric AI principles.

**NEW QUESTION 94**

- (Topic 2)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 99**

- (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 104**

- (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 107**

.....

## 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/>