

IAPP

Exam Questions AIGP

Artificial Intelligence Governance Professional



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)

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

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

Answer: B

Explanation:

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

NEW QUESTION 3

- (Topic 1)

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 4

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 5

- (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 6

- (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 7

- (Topic 1)

What is the 1956 Dartmouth summer research project on AI best known as?

- A. A meeting focused on the impacts of the launch of the first mass-produced computer.
- B. A research project on the impacts of technology on society.
- C. A research project to create a test for machine intelligence.
- D. A meeting focused on the founding of the AI field.

Answer: D

Explanation:

The 1956 Dartmouth summer research project on AI is best known as a meeting focused on the founding of the AI field. This conference is historically significant because it marked the formal beginning of artificial intelligence as an academic discipline. The term "artificial intelligence" was coined during this event, and it laid the foundation for future research and development in AI.

Reference: The AIGP Body of Knowledge highlights the importance of the Dartmouth Conference as a pivotal moment in the history of AI, which established AI as a distinct field of study and research.

NEW QUESTION 8

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 9

- (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 EII 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 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.

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:

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.

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

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

Answer: B

Explanation:

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

NEW QUESTION 11

- (Topic 1)

Under the Canadian Artificial Intelligence and Data Act, when must the Minister of Innovation, Science and Industry be notified about a high-impact AI system?

- A. When use of the system causes or is likely to cause material harm.
- B. When the algorithmic impact assessment has been completed.
- C. Upon release of a new version of the system.
- D. Upon initial deployment of the system.

Answer: D

Explanation:

According to the Canadian Artificial Intelligence and Data Act, high-impact AI systems must notify the Minister of Innovation, Science and Industry upon initial deployment. This requirement ensures that the authorities are aware of the deployment of significant AI systems and can monitor their impacts and compliance with regulatory standards from the outset. This initial notification is crucial for maintaining oversight and ensuring the responsible use of AI technologies. Reference: AIGP Body of Knowledge, domain on AI laws and standards.

NEW QUESTION 16

- (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)

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

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

Answer: B

Explanation:

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

NEW QUESTION 23

- (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 27

- (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 32

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 36

- (Topic 2)

Training data is best defined as a subset of data that is used to?

- A. Enable a model to detect and learn patterns.
- B. Fine-tune a model to improve accuracy and prevent overfitting.
- C. Detect the initial sources of biases to mitigate prior to deployment.
- D. Resemble the structure and statistical properties of production data.

Answer: A

Explanation:

Training data is used to enable a model to detect and learn patterns. During the training phase, the model learns from the labeled data, identifying patterns and relationships that it will later use to make predictions on new, unseen data. This process is fundamental in building an AI model's capability to perform tasks accurately. Reference: AIGP Body of Knowledge on Model Training and Pattern Recognition.

NEW QUESTION 40

- (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 45

- (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 48

- (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)

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 57

- (Topic 2)

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

- A. Output data, assess the delta between the prediction and actual ad clicks.
- B. Algorithmic patterns, to show the model has a high degree of accuracy.
- C. Input data, to ensure the ads are reaching the target audience.
- D. GPU performance, to evaluate the tool's robustness.

Answer: A

Explanation:

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

NEW QUESTION 61

- (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 63

- (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 68

- (Topic 2)

A company has trained an ML model primarily using synthetic data, and now intends to use live personal data to test the model. Which of the following is NOT a best practice apply during the testing?

- A. The test data should be representative of the expected operational data.
- B. Testing should minimize human involvement to the extent practicable.
- C. The test data should be anonymized to the extent practicable.
- D. Testing should be performed specific to the intended uses.

Answer: B

Explanation:

Minimizing human involvement to the extent practicable is not a best practice during the testing of an ML model. Human oversight is crucial during testing to ensure that the model performs correctly and ethically, and to interpret any anomalies or issues that arise. Best practices include using representative test data, anonymizing data to the extent practicable, and performing testing specific to the intended uses of the model. Reference: AIGP Body of Knowledge on AI Model Testing and Human Oversight.

NEW QUESTION 71

- (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 75

- (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 77

- (Topic 2)

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

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

Answer: C

Explanation:

Ensuring fairness when training an AI system largely depends on the data attributes and variability. This involves having a diverse and representative dataset that accurately reflects the population the AI system will serve. Fairness can be compromised if the data is biased or lacks variability, as the model may learn and perpetuate these biases. Diverse data attributes ensure that the model learns from a wide range of examples, reducing the risk of biased predictions. Reference: AIGP Body of Knowledge on Ethical AI Principles and Data Management.

NEW QUESTION 81

- (Topic 2)

Which of the following steps occurs in the design phase of the AI life cycle?

- A. Data augmentation.
- B. Model explainability.
- C. Risk impact estimation.
- D. Performance evaluation.

Answer: C

Explanation:

Risk impact estimation occurs in the design phase of the AI life cycle. This step involves evaluating potential risks associated with the AI system and estimating their impacts to ensure that appropriate mitigation strategies are in place. It helps in identifying and addressing potential issues early in the design process, ensuring the development of a robust and reliable AI system. Reference: AIGP Body of Knowledge on AI Design and Risk Management.

NEW QUESTION 85

- (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 89

- (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 94

- (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 99

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

The healthcare network intends to retain a cloud provider to host the solution and a

consulting firm to help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.

Which of the following steps can best mitigate the possibility of discrimination prior to training and testing the AI solution?

- A. Procure more data from clinical research partners.
- B. Engage a third party to perform an audit.
- C. Perform an impact assessment.
- D. Create a bias bounty program.

Answer: C

Explanation:

Performing an impact assessment is the best step to mitigate the possibility of discrimination before training and testing the AI solution. An impact assessment, such as a Data Protection Impact Assessment (DPIA) or Algorithmic Impact Assessment (AIA), helps identify potential biases and discriminatory outcomes that could arise from the AI system. This process involves evaluating the data and the algorithm for fairness, accountability, and transparency. It ensures that any biases in the data are detected and addressed, thus preventing discriminatory practices and promoting ethical AI deployment. Reference: AIGP Body of

Knowledge on Ethical AI and Impact Assessments.

NEW QUESTION 100

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

The healthcare network intends to retain a cloud provider to host the solution and a consulting firm to help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.

Which stakeholder group is most important in selecting the specific type of algorithm?

- A. The cloud provider.
- B. The consulting firm.
- C. The healthcare network's data science team.
- D. The healthcare network's AI governance committee.

Answer: C

Explanation:

In selecting the specific type of algorithm for the AI solution, the healthcare network's data science team is most important. This team possesses the technical expertise and understanding of the data, the clinical context, and the performance requirements needed to make an informed decision about which algorithm is most suitable. While the cloud provider and consulting firm can offer support and infrastructure, and the AI governance committee provides oversight, the data science team's specialized knowledge is crucial for selecting and implementing the appropriate algorithm. Reference: AIGP Body of Knowledge, AI governance and team roles section.

NEW QUESTION 105

- (Topic 2)

What is the technique to remove the effects of improperly used data from an ML system?

- A. Data cleansing.
- B. Model inversion.
- C. Data de-duplication.
- D. Model disgorgement.

Answer: D

Explanation:

Model disgorgement is the technique used to remove the effects of improperly used data from an ML system. This process involves retraining or adjusting the model to eliminate any biases or inaccuracies introduced by the inappropriate data. It ensures that the model's outputs are not influenced by data that was not meant to be used or was used incorrectly. Reference: AIGP Body of Knowledge on Data Management and Model Integrity.

NEW QUESTION 108

- (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 113

- (Topic 2)

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

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

Answer: B

Explanation:

Respecting intellectual property rights means adhering to licensing terms and ensuring that generated content complies with these terms. A system that categorizes and applies filters based on licensing terms ensures that content is used legally and ethically, respecting the rights of content creators. While providing

attribution is important, categorization and application of filters based on licensing terms are more directly tied to compliance with intellectual property laws. This principle is elaborated in the IAPP AIGP Body of Knowledge sections on intellectual property and compliance.

NEW QUESTION 117

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AIGP Practice Exam Features:

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

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