



IAPP

Exam Questions AIGP

Artificial Intelligence Governance Professional

NEW QUESTION 1

- (Topic 1)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 2

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 3

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 4

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 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)

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 7

- (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 8

- (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 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.

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

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

Answer: C

Explanation:

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

NEW QUESTION 10

- (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 10

- (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 13

- (Topic 1)

What is the primary reason the EU is considering updates to its Product Liability Directive?

- A. To increase the minimum warranty level for defective goods.
- B. To define new liability exemptions for defective products.
- C. Address digital services and connected products.
- D. Address free and open-source software.

Answer: C

Explanation:

The primary reason the EU is considering updates to its Product Liability Directive is to address digital services and connected products. The current directive does not adequately cover the complexities and challenges posed by modern digital and connected technologies. By updating the directive, the EU aims to ensure that it remains relevant and effective in addressing the liabilities associated with these advanced products, ensuring consumer protection and fair market practices in the digital age.

NEW QUESTION 14

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 17

- (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 19

- (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 20

- (Topic 1)

Under the NIST AI Risk Management Framework, all of the following are defined as characteristics of trustworthy AI EXCEPT?

- A. Tested and Effective.
- B. Secure and Resilient.
- C. Explainable and Interpretable.
- D. Accountable and Transparent.

Answer: A

Explanation:

The NIST AI Risk Management Framework outlines several characteristics of trustworthy AI, including being secure and resilient, explainable and interpretable, and accountable and transparent. While being tested and effective is important, it is not explicitly listed as a characteristic of trustworthy AI in the NIST framework. The focus is more on the system's ability to function safely, securely, and transparently in a way that stakeholders can understand and trust. Reference: AIGP Body of Knowledge, NIST AI RMF section.

NEW QUESTION 24

- (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 28

- (Topic 1)

A company is working to develop a self-driving car that can independently decide the appropriate route to take the driver after the driver provides an address.

If they want to make this self-driving car "strong" AI, as opposed to "weak," the engineers would also need to ensure?

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

Answer: A

Explanation:

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

NEW QUESTION 32

- (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 37

- (Topic 2)

Which of the following would be the least likely step for an organization to take when designing an integrated compliance strategy for responsible AI?

- A. Conducting an assessment of existing compliance programs to determine overlaps and integration points.
- B. Employing a new software platform to modernize existing compliance processes across the organization.
- C. Consulting experts to consider the ethical principles underpinning the use of AI within the organization.
- D. Launching a survey to understand the concerns and interests of potentially impacted stakeholders.

Answer: B

Explanation:

When designing an integrated compliance strategy for responsible AI, the least likely step would be employing a new software platform to modernize existing compliance processes. While modernizing compliance processes is beneficial, it is not as directly related to the strategic integration of ethical principles and stakeholder concerns. More critical steps include conducting assessments of existing compliance programs to identify overlaps and integration points, consulting experts on ethical principles, and launching surveys to understand stakeholder concerns. These steps ensure that the compliance strategy is comprehensive and aligned with responsible AI principles. Reference: AIGP Body of Knowledge on AI Governance and Compliance Integration.

NEW QUESTION 42

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 43

- (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 44

- (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 47

- (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 49

- (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 51

- (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 55

- (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 57

- (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 61

- (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 64

- (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 67

- (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 69

- (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 70

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

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

The most significant risk from combining the healthcare network's existing data with the clinical research partner data is?

- A. Privacy risk.
- B. Security risk.
- C. Operational risk.
- D. Reputational risk.

Answer: A

Explanation:

The most significant risk from combining the healthcare network's existing data with the clinical research partner data is privacy risk. Combining data sets, especially in healthcare, often involves handling sensitive information that could lead to privacy breaches if not managed properly. De-identified data can still pose re-identification risks when combined with other data sets. Ensuring privacy involves implementing robust data protection measures, maintaining compliance with privacy regulations such as HIPAA, and conducting thorough privacy impact assessments. Reference: AIGP Body of Knowledge on Data Privacy and Security.

NEW QUESTION 74

.....

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](#)