

IAPP

Exam Questions AIGP

Artificial Intelligence Governance Professional



NEW QUESTION 1

- (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.

All of the following may be copyright risks from teachers using generative AI to create course content EXCEPT?

- A. Content created by an LLM may be protectable under U.
- B. intellectual property law.
- C. Generative AI is generally trained using intellectual property owned by third parties.
- D. Students must expressly consent to this use of generative AI.
- E. Generative AI often creates content without attribution.

Answer: C

Explanation:

All of the options listed may pose copyright risks when teachers use generative AI to create course content, except for students must expressly consent to this use of generative AI. While obtaining student consent is essential for ethical and privacy reasons, it does not directly relate to copyright risks associated with the creation and use of AI-generated content.

Reference: The AIGP Body of Knowledge discusses the importance of addressing intellectual property (IP) risks when using AI-generated content. Copyright risks are typically associated with the use of third-party data and the lack of attribution, rather than the consent of users.

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.

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 3

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

- A. Finance and Legal.
- B. Marketing and Compliance.
- C. Supply Chain and Marketing.

D. Litigation and Product Development.

Answer: A

Explanation:

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

NEW QUESTION 4

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 5

- (Topic 1)

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 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)

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 9

- (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)

What type of organizational risk is associated with AI's resource-intensive computing demands?

- A. People risk.
- B. Security risk.
- C. Third-party risk.
- D. Environmental risk.

Answer: D

Explanation:

AI's resource-intensive computing demands pose significant environmental risks. High-performance computing required for training and deploying AI models often leads to substantial energy consumption, which can result in increased carbon emissions and other environmental impacts. This is particularly relevant given the growing concern over climate change and the environmental footprint of technology. Organizations need to consider these environmental risks when developing AI systems, potentially exploring more energy-efficient methods and renewable energy sources to mitigate the environmental impact.

NEW QUESTION 10

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

In particular, GVC has learned that the teachers they employ used open source large language models ("LLM") to develop an online tool that customizes study questions for individual students. GVC has also discovered that an art teacher has expressly

incorporated the use of generative AI into the curriculum to enable students to use prompts to create digital art. GVC has started to investigate these practices and develop a process to monitor any use of generative AI, including by teachers and students, going forward. 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 15

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 19

- (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 21

- (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 23

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several

vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws. The organization has a large procurement team that is responsible for the contracting of technology solutions. One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company are responsible for integrating and deploying technology solutions into the organization's operations in a responsible, cost-effective manner. The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change. If XYZ does not deploy and use the AI hiring tool responsibly in the United States, its liability would likely increase under all of the following laws EXCEPT?

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

Answer: B

Explanation:

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

NEW QUESTION 26

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 31

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 34

- (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 37

- (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 42

- (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 47

- (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 52

- (Topic 1)

A company is creating a mobile app to enable individuals to upload images and videos, and analyze this data using ML to provide lifestyle improvement recommendations. The signup form has the following data fields:

* 1.First name 2.Last name 3.Mobile number 4.Email ID 5.New password 6.Date of birth 7.Gender

In addition, the app obtains a device's IP address and location information while in use. What GDPR privacy principles does this violate?

- A. Purpose Limitation and Data Minimization.
- B. Accountability and Lawfulness.
- C. Transparency and Accuracy.
- D. Integrity and Confidentiality.

Answer: A

Explanation:

The GDPR privacy principles that this scenario violates are Purpose Limitation and Data Minimization. Purpose Limitation requires that personal data be collected for specified, explicit, and legitimate purposes and not further processed in a manner that is incompatible with those purposes. Data Minimization mandates that personal data collected should be adequate, relevant, and limited to what is necessary in relation to the purposes for which they are processed. In this case, collecting extensive personal information (e.g., IP address, location, gender) and potentially using it beyond the necessary scope for the app's functionality could

violate these principles by collecting more data than needed and possibly using it for purposes not originally intended.

NEW QUESTION 55

- (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 56

- (Topic 2)

Testing data is defined as a subset of data that is used to?

- A. Assess a model's on-going performance in production.
- B. Enable a model to discover and learn patterns.
- C. Provide a robust evaluation of a final model.
- D. Evaluate a model's handling of randomized edge cases.

Answer: C

Explanation:

Testing data is a subset of data used to provide a robust evaluation of a final model. After training the model on training data, it is essential to test its performance on unseen data (testing data) to ensure it generalizes well to new, real-world scenarios. This step helps in assessing the model's accuracy, reliability, and ability to handle various data inputs. Reference: AIGP Body of Knowledge on Model Validation and Testing.

NEW QUESTION 61

- (Topic 2)

In the machine learning context, feature engineering is the process of?

- A. Converting raw data into clean data.
- B. Creating learning schema for a model apply.
- C. Developing guidelines to train and test a model.
- D. Extracting attributes and variables from raw data.

Answer: D

Explanation:

In the machine learning context, feature engineering is the process of extracting attributes and variables from raw data to make it suitable for training an AI model. This step is crucial as it transforms raw data into meaningful features that can improve the model's accuracy and performance. Feature engineering involves selecting, modifying, and creating new features that help the model learn more effectively. Reference: AIGP Body of Knowledge on AI Model Development and Feature Engineering.

NEW QUESTION 62

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 66

- (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 70

- (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 73

- (Topic 2)

During the development of semi-autonomous vehicles, various failures occurred as a result of the sensors misinterpreting environmental surroundings, such as sunlight.

These failures are an example of?

- A. Hallucination.
- B. Brittleness.
- C. Uncertainty.
- D. Forgetting.

Answer: B**Explanation:**

The failures in semi-autonomous vehicles due to sensors misinterpreting environmental surroundings, such as sunlight, are examples of brittleness. Brittleness in AI systems refers to their inability to handle variations in input data or unexpected conditions, leading to failures when the system encounters situations that were not adequately covered during training. These systems perform well under specific conditions but fail when those conditions change. Reference: AIGP Body of Knowledge on AI System Robustness and Failures.

NEW QUESTION 77

- (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 81

- (Topic 2)

What is the term for an algorithm that focuses on making the best choice achieve an immediate objective at a particular step or decision point, based on the available information and without regard for the longer-term best solutions?

- A. Single-lane.
- B. Optimized.
- C. Efficient.
- D. Greedy.

Answer: D**Explanation:**

A greedy algorithm is one that makes the best choice at each step to achieve an immediate objective, without considering the longer-term consequences. It focuses on local optimization at each decision point with the hope that these local solutions will lead to an optimal global solution. However, greedy algorithms do not always produce the best overall solution for certain problems, but they are useful when an immediate, locally optimal solution is desired. Reference: AIGP Body

of Knowledge, algorithm types section.

NEW QUESTION 83

- (Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 87

- (Topic 2)

All of the following are included within the scope of post-deployment AI maintenance EXCEPT?

- A. Ensuring that all model components are subject a control framework.
- B. Dedicating experts to continually monitor the model output.
- C. Evaluating the need for an audit under certain standards.
- D. Defining thresholds to conduct new impact assessments.

Answer: D

Explanation:

Post-deployment AI maintenance typically includes ensuring that all model components are subject to a control framework, dedicating experts to continually monitor the model output, and evaluating the need for audits under certain standards. However, defining thresholds to conduct new impact assessments is usually part of the initial deployment and ongoing governance processes rather than a maintenance activity. Maintenance focuses more on the operational aspects of the AI system rather than setting new thresholds for impact assessments.

Reference: AIGP BODY OF KNOWLEDGE, sections discussing AI lifecycle management and post-deployment activities.

NEW QUESTION 88

- (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 92

- (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 97

- (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 102

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

The healthcare network intends to retain a cloud provider to host the solution and a consulting firm to help develop the algorithm using the healthcare network's existing data

and de-identified data that is licensed from a large US clinical research partner.

In the design phase, which of the following steps is most important in gathering the data from the clinical research partner?

- A. Perform a privacy impact assessment.
- B. Combine only anonymized data.
- C. Segregate the data sets.
- D. Review the terms of use.

Answer: D

Explanation:

Reviewing the terms of use is essential when gathering data from a clinical research partner. This step ensures that the healthcare network complies with all legal and contractual obligations related to data usage. It addresses data ownership, usage limitations, consent requirements, and privacy obligations, which are critical to maintaining ethical standards and avoiding legal repercussions. This review helps ensure that the data is used in a manner consistent with the agreements made and the regulatory environment, which is fundamental for lawful and ethical AI development. Reference: AIGP Body of Knowledge on Legal and Regulatory Considerations.

NEW QUESTION 107

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 108

- (Topic 2)

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

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

Answer: B

Explanation:

Bias in AI can be reduced during the planning and design phases through stakeholder involvement, human oversight, and careful data collection. While feature

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

NEW QUESTION 113

- (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 114

- (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 118

- (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 120

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AIGP Practice Exam Features:

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

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