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 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.
What is the best strategy to mitigate the bias uncovered in the loan applications?

A. Retrain the model with data that reflects demographic parity.
B. Procure a third-party statistical bias assessment tool.
C. Document all instances of bias in the data set.
D. Delete all gender-based data in the data set.

Answer: A

Explanation:
Retraining the model with data that reflects demographic parity is the best strategy to mitigate the bias uncovered in the loan applications. This approach addresses the root cause of the bias by ensuring that the training data is representative and balanced, leading to more equitable decision-making by the AI model.
Reference: The AIGP Body of Knowledge stresses the importance of using high-quality,
unbiased training data to develop fair and reliable AI systems. Retraining the model with balanced data helps correct biases that arise from historical inequalities, ensuring that the AI system makes decisions based on equitable criteria.

NEW QUESTION 3
- (Topic 1)
According to the Singapore Model AI Governance Framework, all of the following are recommended measures to promote the responsible use of AI EXCEPT?

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

Answer: C

Explanation:
The Singapore Model AI Governance Framework recommends several measures to promote the responsible use of AI, such as determining the level of human involvement in decision-making, adapting governance structures, and establishing communications and collaboration among stakeholders. However, employing human-over-the-loop protocols is not specifically mentioned in this framework. The focus is more on integrating human oversight appropriately within the decision-making process rather than exclusively employing such protocols. Reference: AIGP Body of Knowledge, section on AI governance frameworks.

NEW QUESTION 4
- (Topic 1)
Which of the following is an example of a high-risk application under the EU AI Act?

A. A resume scanning tool that ranks applicants.
B. An AI-enabled inventory management tool.
C. A government-run social scoring tool.
D. A customer service chatbot tool.

Answer: C

Explanation:
The EU AI Act categorizes certain applications of AI as high-risk due to their potential impact on fundamental rights and safety. High-risk applications include those
used in critical areas such as employment, education, and essential public services. A government-run social scoring tool, which assesses individuals based on their social behavior or perceived trustworthiness, falls under this category because of its profound implications for privacy, fairness, and individual rights. This contrasts with other AI applications like resume scanning tools or customer service chatbots, which are generally not classified as high-risk under the EU AI Act.

NEW QUESTION 5
- (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 6
- (Topic 1)
If it is possible to provide a rationale for a specific output of an AI system, that system can best be described as?

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

Answer: C

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

NEW QUESTION 7
- (Topic 1)
All of the following may be permissible uses of an AI system under the EU AI Act EXCEPT?

A. To detect an individual's intent for law enforcement purposes.
B. To promote equitable distribution of welfare benefits.
C. To implement social scoring.
D. To manage border control.

Answer: C

Explanation:
The EU AI Act explicitly prohibits the use of AI systems for social scoring by public authorities, as it can lead to discrimination and unfair treatment of individuals based on their social behavior or perceived trustworthiness. While AI can be used to promote equitable distribution of welfare benefits, manage border control, and even detect an individual's intent for law enforcement purposes (within strict regulatory and ethical boundaries), implementing social scoring systems is not permissible under the Act due to the significant risks to fundamental rights and freedoms.

NEW QUESTION 8
- (Topic 1)
You asked a generative AI tool to recommend new restaurants to explore in Boston, Massachusetts that have a specialty Italian dish made in a traditional fashion without spinach and wine. The generative AI tool recommended five restaurants for you to visit.
After looking up the restaurants, you discovered one restaurant did not exist and two others did not have the dish.
This information provided by the generative AI tool is an example of what is commonly called?

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

Answer: C

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

NEW QUESTION 9
- (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)
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 Ell law and includes suitable safeguards.
C. Is deemed to solely benefit the individual and includes documented legitimate interests.
D. Is necessary for entering into or performing under a contract between the data subject and data controller.

Answer: A

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

NEW QUESTION 10
- (Topic 1)
Which of the following disclosures is NOT required for an EU organization that developed and deployed a high-risk Al 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 Al 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 15
- (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 Al tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party Al-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 Al hiring tools and wants to mitigate
them. It also questions how best to organize and train its existing personnel to use the Al 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 Al hiring tool responsibly in the United States, its liability would likely increase under all of the following laws EXCEPT?

A. Anti-discriminationlaws.
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 19
- (Topic 1)
CASE STUDY
Please use the following answer the next question:
Good Values Corporation (GVC) is a U.S. educational services provider that employs teachers to create and deliver enrichment courses for high school students. GVC has learned that many of its teacher employees are using generative AI to create the enrichment courses, and that many of the students are using generative AI to complete their assignments.
In particular, GVC has learned that the teachers they employ used open source large language models ("LLM") to develop an online tool that customizes study questions for individual students. GVC has also discovered that an art teacher has expressly incorporated the use of generative AI into the curriculum to enable students to use prompts to create digital art.
GVC has started to investigate these practices and develop a process to monitor any use
of generative AI, including by teachers and students, going forward.
Which of the following risks should be of the highest concern to individual teachers using generative AI to ensure students learn the course material?

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

Answer: B

Explanation:
The highest concern for individual teachers using generative AI to ensure students learn the course material is model accuracy. Ensuring that the AI-generated content is accurate and relevant to the curriculum is crucial for effective learning. If the AI model produces inaccurate or irrelevant content, it can mislead students and hinder their understanding of the subject matter.
Reference: According to the AIGP Body of Knowledge, one of the core risks posed by AI
systems is the accuracy of the data and models used. Ensuring the accuracy of AI- generated content is essential for maintaining the integrity of the educational material and achieving the desired learning outcomes.

NEW QUESTION 20
- (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 thechips 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 24
- (Topic 1)
Under the Canadian Artificial Intelligence and Data Act, when must the Minister of Innovation, Science and Industry be notified about a high-impact AI system?

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

Answer: D

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

NEW QUESTION 27
- (Topic 1)
Which of the following is NOT a common type of machine learning?

A. Deep learning.
B. Cognitive learning.
C. Unsupervised learning.
D. Reinforcement learning.

Answer: B

Explanation:
The common types of machine learning include supervised learning, unsupervised learning, reinforcement learning, and deep learning. Cognitive learning is not a type of machine learning; rather, it is a term often associated with the broader field of cognitive science and psychology. Reference: AIGP BODY OF KNOWLEDGE and standard AI/ML literature.

NEW QUESTION 28
- (Topic 1)
A Canadian company is developing an Al solution to evaluate candidates in the course of job interviews.
Before offering the Al solution in the EU market, the company must take all of the following steps EXCEPT?

A. Register the Al solution in a public EU database.
B. Establish a risk and quality management system.
C. Engage a third-party auditor to perform a bias audit.
D. Draw up technical documentation and instructions for use.

Answer: A

Explanation:
Before offering an AI solution in the EU market, a Canadian company must take several steps to comply with the EU AI Act. These steps include establishing a risk and quality management system (B), engaging a third-party auditor to perform a bias audit (C), and drawing up technical documentation and instructions for use (D). However, there is no requirement to register the AI solution in a public EU database (A). This registration step is not specified as part of the compliance requirements under the EU AI Act for such solutions.

NEW QUESTION 33
- (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 Al tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party Al-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 Al hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the Al 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 Al 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 35
- (Topic 1)
Which of the following is a subcategory of Al and machine learning that uses labeled datasets to train algorithms?

A. Segmentation.
B. Generative Al.
C. Expert systems.
D. Supervised learning.

Answer: D

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

NEW QUESTION 39
- (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 43
- (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-partyrisk.
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 45
- (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 50
- (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 51
- (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 55

- (Topic 2)
CASE STUDY
Please use the following answer the next question:
A local police department in the United States procured an Al 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 Al 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 Al 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 Al 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 Al system gives a score of at least 90% and proceed directly with an arrest.
What is the best reason the police department should continue to perform investigations even if the Al system scores an individual's likelihood of criminal activity at or above 90%?

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

Answer: B

Explanation:
The best reason for the police department to continue performing investigations even if the
AI system scores an individual's likelihood of criminal activity at or above 90% is that AI systems affecting fundamental civil rights should not be fully automated. Human oversight is essential to ensure that decisions impacting civil liberties are made with due consideration of context and mitigating factors that an AI might not fully appreciate. This approach ensures fairness, accountability, and adherence to legal standards. Reference: AIGP Body of Knowledge on AI Ethics and Human Oversight.

NEW QUESTION 56
- (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 60
- (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 Al 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 61
- (Topic 2)
Which of the following use cases would be best served by a non-AI solution?

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

Answer: A

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

NEW QUESTION 66
- (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, what is the most important step for the healthcare network to take when mapping its existing data to the clinical research partner data?

A. Apply privacy-enhancing technologies to the data.
B. Identify fits and gaps in the combined data.
C. Ensure the data is labeled and formatted.
D. Evaluate the country of origin of the data.

Answer: B

Explanation:
In the design phase of integrating data from different sources, identifying fits and gaps is crucial. This process involves understanding how well the data from the clinical research partner aligns with the healthcare network's existing data. It ensures that the combined data set is coherent and can be effectively used for training the AI algorithm. This step helps in spotting any discrepancies, inconsistencies, or missing data that might affect the performance and accuracy of the AI model. It directly addresses the integrity and compatibility of the data, which is foundational before applying any privacy-enhancing technologies, labeling, or evaluating the origin of the data. Reference: AIGP Body of Knowledge on Data Integration and Quality.

NEW QUESTION 71
- (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 72
- (Topic 2)
You are the chief privacy officer of a medical research company that would like to collect and use sensitive data about cancer patients, such as their names, addresses, race and ethnic origin, medical histories, insurance claims, pharmaceutical prescriptions, eating and drinking habits and physical activity.
The company will use this sensitive data to build an AI algorithm that will spot common attributes that will help predict if seemingly healthy people are more likely to get cancer. However, the company is unable to obtain consent from enough patients to sufficiently collect the minimum data to train its model.
Which of the following solutions would most efficiently balance privacy concerns with the lack of available data during the testing phase?

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

Answer: C

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

NEW QUESTION 76
- (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 80
- (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 operationaldata.
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 83
- (Topic 2)
A company plans on procuring a tool from an AI provider for its employees to use for certain business purposes.
Which contractual provision would best protect the company's intellectual property in the tool, including training and testing data?

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

Answer: B

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

NEW QUESTION 88
- (Topic 2)
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 93
- (Topic 2)
Pursuant to the White House Executive Order of November 2023, who is responsible for creating guidelines to conduct red-teaming tests of AI systems?

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

Answer: A

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

NEW QUESTION 97
- (Topic 2)
CASE STUDY
Please use the following answer the next question:
A local police department in the United States procured an Al 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 Al 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 Al 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 Al 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 Al system gives a score of at least 90% and proceed directly with an arrest.
During the procurement process, what is the most likely reason that the third-party consultant asked each vendor for information about the diversity of their datasets?

A. To comply with applicable law.
B. To assist the fairness of the Al system.
C. To evaluate the reliability of the Al system.
D. To determine the explainability of the Al system.

Answer: B

Explanation:
The third-party consultant asked each vendor for information about the diversity of their datasets to assist in ensuring the fairness of the AI system. Diverse datasets help prevent biases and ensure that the AI system performs equitably across different demographic groups. This is crucial for a law enforcement application, where fairness and avoiding discriminatory practices are of paramount importance. Ensuring diversity in training data helps in building a more just and unbiased AI system. Reference: AIGP Body of Knowledge on Ethical AI and Fairness.

NEW QUESTION 99
- (Topic 2)
Which of the following steps occurs in the design phase of the Al 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 102
- (Topic 2)
Which of the following elements of feature engineering is most important to mitigate the potential bias in an Al 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 105
- (Topic 2)
What is the technique to remove the effects of improperly used data from an ML system?

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

Answer: D

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

NEW QUESTION 106
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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