



IAPP

Exam Questions AIGP

Artificial Intelligence Governance Professional

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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.S. intellectual property law.
- B. Generative AI is generally trained using intellectual property owned by third parties.
- C. Students must expressly consent to this use of generative AI.
- D. 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.

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 3

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

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

D. Litigation and Product Development.

Answer: A

Explanation:

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

NEW QUESTION 4

- (Topic 1)

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)

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.

Which of the following is the most important reason to train the underwriters on the model prior to deployment?

- A. To provide a reminder of a right appeal.
- B. To solicit on-going feedback on model performance.
- C. To apply their own judgment to the initial assessment.
- D. To ensure they provide transparency applicants on the model.

Answer: C

Explanation:

Training underwriters on the model prior to deployment is crucial so they can apply their own judgment to the initial assessment. While AI models can streamline the process, human judgment is still essential to catch nuances that the model might miss or to account for any biases or errors in the model's decision-making process.

Reference: The AIGP Body of Knowledge emphasizes the importance of human oversight

in AI systems, particularly in high-stakes areas such as underwriting and loan approvals. Human underwriters can provide a critical review and ensure that the model's assessments are accurate and fair, integrating their expertise and understanding of complex cases.

NEW QUESTION 7

- (Topic 1)

The OECD's Ethical AI Governance Framework is a self-regulation model that proposes to prevent societal harms by?

- A. Establishing explain ability criteria to responsibly source and use data to train AI systems.
- B. Defining requirements specific to each industry sector and high-risk AI domain.
- C. Focusing on AI technical design and post-deployment monitoring.
- D. Balancing AI innovation with ethical considerations.

Answer: D

Explanation:

The OECD's Ethical AI Governance Framework aims to ensure that AI development and deployment are carried out ethically while fostering innovation. The framework includes principles like transparency, accountability, and human rights protections to prevent societal harm. It does not focus solely on technical design or post-deployment monitoring (C), nor does it establish industry-specific requirements (B). While explainability is important, the primary goal is to balance innovation with ethical considerations (D).

NEW QUESTION 8

- (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 9

- (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 10

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 10

- (Topic 1)

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 .. 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.

During the first month when ABC monitors the model for bias, it is most important to?

- A. Continue disparity testing.
- B. Analyze the quality of the training and testing data.
- C. Compare the results to human decisions prior to deployment.
- D. Seek approval from management for any changes to the model.

Answer: A

Explanation:

During the first month of monitoring the model for bias, it is most important to continue disparity testing. Disparity testing involves regularly evaluating the model's decisions to identify and address any biases, ensuring that the model operates fairly across different demographic groups.

Reference: Regular disparity testing is highlighted in the AIGP Body of Knowledge as a critical practice for maintaining the fairness and reliability of AI models. By continuously monitoring for and addressing disparities, organizations can ensure their AI systems remain compliant with ethical and legal standards, and mitigate any unintended biases that may arise in production.

NEW QUESTION 14

- (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 16

- (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 21

- (Topic 1)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 25

- (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 27

- (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 31

- (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 35

- (Topic 1)

Which risk management framework/guide/standard focuses on value-based engineering methodology?

- A. ISO/IEC Guide 51 (Safety).
- B. ISO 31000 Guidelines (Risk Management).
- C. IEEE 7000-2021 Standard Model Process for Addressing Ethical Concerns during System Design.
- D. Council of Europe Human Rights, Democracy, and the Rule of Law Assurance Framework (HUDERIA) for AI Systems.

Answer: C

Explanation:

The IEEE 7000-2021 Standard focuses on a value-based engineering methodology for addressing ethical concerns during system design. This standard guides

engineers and organizations in integrating ethical considerations into the design and development processes of AI systems, ensuring that these technologies are developed responsibly and align with human values. Reference: AIGP Study Material, section on risk management frameworks and standards.

NEW QUESTION 38

- (Topic 1)

A Canadian company is developing an AI solution to evaluate candidates in the course of job interviews.

Before offering the AI solution in the EU market, the company must take all of the following steps EXCEPT?

- A. Register the AI 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 39

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 41

- (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 43

- (Topic 1)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 45

- (Topic 1)

A U.S. mortgage company developed an AI platform that was trained using anonymized details from mortgage applications, including the applicant's education, employment and demographic information, as well as from subsequent payment or default information. The AI platform will be used automatically grant or deny new mortgage applications, depending on whether the platform views an applicant as presenting a likely risk of default.

Which of the following laws is NOT relevant to this use case?

- A. Fair Housing Act.
- B. Fair Credit Reporting Act.
- C. Equal Credit Opportunity Act.
- D. Title VII of the Civil Rights Act of 1964.

Answer: D

Explanation:

The U.S. mortgage company's AI platform relates to housing and credit, making the Fair Housing Act (A), Fair Credit Reporting Act (B), and Equal Credit Opportunity Act (C) relevant. Title VII of the Civil Rights Act of 1964 deals with employment discrimination and is not directly relevant to the mortgage application context (D).

NEW QUESTION 50

- (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 52

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 54

- (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 59

- (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 60

- (Topic 2)

To maintain fairness in a deployed system, it is most important to?

- A. Protect against loss of personal data in the model.
- B. Monitor for data drift that may affect performance and accuracy.
- C. Detect anomalies outside established metrics that require new training data.
- D. Optimize computational resources and data to ensure efficiency and scalability.

Answer: B

Explanation:

To maintain fairness in a deployed system, it is crucial to monitor for data drift that may affect performance and accuracy. Data drift occurs when the statistical properties of the input data change over time, which can lead to a decline in model performance. Continuous monitoring and updating of the model with new data ensure that it remains fair and accurate, adapting to any changes in the data distribution. Reference: AIGP Body of Knowledge on Post-Deployment Monitoring and Model Maintenance.

NEW QUESTION 61

- (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 65

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 67

- (Topic 2)

What is the primary purpose of conducting ethical red-teaming on an AI system?

- A. To improve the model's accuracy.
- B. To simulate model risk scenarios.
- C. To identify security vulnerabilities.
- D. To ensure compliance with applicable law.

Answer: B

Explanation:

The primary purpose of conducting ethical red-teaming on an AI system is to simulate model risk scenarios. Ethical red-teaming involves rigorously testing the AI system to identify potential weaknesses, biases, and vulnerabilities by simulating real-world attack or failure scenarios. This helps in proactively addressing issues that could compromise the system's reliability, fairness, and security. Reference: AIGP Body of Knowledge on AI Risk Management and Ethical AI Practices.

NEW QUESTION 72

- (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 75

- (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 78

- (Topic 2)

A company plans on procuring a tool from an AI provider for its employees to use for certain business purposes. Which contractual provision would best protect the company's intellectual property in the tool, including training and testing data?

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

Answer: B

Explanation:

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

NEW QUESTION 83

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

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

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

During the procurement process, what is the most likely reason that the third-party consultant asked each vendor for information about the diversity of their datasets?

- A. To comply with applicable law.
- B. To assist the fairness of the AI system.
- C. To evaluate the reliability of the AI system.

D. To determine the explainability of the AI system.

Answer: B

Explanation:

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

NEW QUESTION 88

- (Topic 2)

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

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

Answer: C

Explanation:

Ensuring fairness when training an AI system largely depends on the data attributes and variability. This involves having a diverse and representative dataset that accurately reflects the population the AI system will serve. Fairness can be compromised if the data is biased or lacks variability, as the model may learn and perpetuate these biases.

Diverse data attributes ensure that the model learns from a wide range of examples, reducing the risk of biased predictions. Reference: AIGP Body of Knowledge on Ethical AI Principles and Data Management.

NEW QUESTION 92

- (Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 95

- (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 97

- (Topic 2)

CASE STUDY

Please use the following answer the next question:

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

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

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

The police department has now been using the AI system for a year. An internal review has found that every time the system scored a likelihood of criminal activity at or above 90%, the police investigation subsequently confirmed that the individual had, in fact, committed a crime. Based on these results, the police department

wants to forego investigations for cases where the AI system gives a score of at least 90% and proceed directly with an arrest.

The best human oversight mechanism for the police department to implement is that a police officer should?

- A. Explain to the accused how the AI system works.
- B. Confirm the AI recommendation prior to sentencing.
- C. Ensure an accused is given notice that the AI system was used.
- D. Consider the AI recommendation as part of the criminal investigation.

Answer: D

Explanation:

The best human oversight mechanism for the police department to implement is for a police officer to consider the AI recommendation as part of the criminal investigation. This ensures that the AI system's output is used as a tool to aid human decision-making rather than replace it. The police officer should integrate the AI's insights with other evidence and contextual information to make informed decisions, maintaining a balance between technological aid and human judgment.

Reference: AIGP Body of Knowledge on AI Integration and Human Oversight.

NEW QUESTION 102

- (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 104

- (Topic 2)

All of the following are reasons to deploy a challenger AI model in addition a champion AI model EXCEPT to?

- A. Provide a framework to consider alternatives to the champion model.
- B. Automate real-time monitoring of the champion model.
- C. Perform testing on the champion model.
- D. Retrain the champion model.

Answer: D

Explanation:

Deploying a challenger AI model alongside a champion model is a strategy used to compare the performance of different models in a real-world environment. This approach helps in providing a framework to consider alternatives to the champion model, automating real-time monitoring of the champion model, and performing testing on the champion model. However, retraining the champion model is not a reason to deploy a challenger model. Retraining is a separate process that involves updating the champion model with new data or techniques, which is not related to the use of a challenger model.

Reference: AIGP BODY OF KNOWLEDGE, sections on model evaluation and management.

NEW QUESTION 106

- (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 109

- (Topic 2)

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

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

Answer: B

Explanation:

Respecting intellectual property rights means adhering to licensing terms and ensuring that generated content complies with these terms. A system that categorizes and applies filters based on licensing terms ensures that content is used legally and ethically, respecting the rights of content creators. While providing attribution is important, categorization and application of filters based on licensing terms are more directly tied to compliance with intellectual property laws. This principle is elaborated in the IAPP AIGP Body of Knowledge sections on intellectual property and compliance.

NEW QUESTION 114

.....

Relate Links

100% Pass Your AIGP Exam with ExamBible Prep Materials

<https://www.exambible.com/AIGP-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>