# 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:
XYZ Corp., a premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent. The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors,
including in Asia.
It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias. Address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third-party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.
The organization has a large procurement team that is responsible for the contracting of technology solutions. One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company are responsible for integrating and deploying technology solutions into the organization's operations in a responsible, cost- effective manner.
The organization is aware of the risks presented by AI hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the AI hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.
Which other stakeholder groups should be involved in the selection and implementation of the AI hiring tool?

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

**Answer:** A

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

**NEW QUESTION 2**
- (Topic 1)
Each of the following actors are typically engaged in the AI development life cycle EXCEPT?

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

**Answer:** B

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

**NEW QUESTION 3**
- (Topic 1)
An EU bank intends to launch a multi-modal AI platform for customer engagement and automated decision-making assist with the opening of bank accounts. The platform has been subject to thorough risk assessments and testing, where it proves to be effective in not discriminating against any individual on the basis of a protected class.
What additional obligations must the bank fulfill prior to deployment?

A. The bank must obtain explicit consent from users under the privacy Directive.
B. The bank must disclose how the AI system works under the Ell Digital Services Act.
C. The bank must subject the AI system an adequacy decision and publish its appropriate safeguards.
D. The bank must disclose the use of the AI system and implement suitable measures for users to contest automated decision-making.

**Answer:** D

**Explanation:**
Under the EU regulations, particularly the GDPR, banks using AI for decision-making must inform users about the use of AI and provide mechanisms for users to contest decisions. This is part of ensuring transparency and accountability in automated processing. Explicit consent under the privacy directive (A) and disclosing under the Digital Services Act (B) are not specifically required in this context. An adequacy decision is related to data transfers outside the EU (C).

**NEW QUESTION 4**
- (Topic 1)
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. Toprovide a reminder of a right appeal.
B. Tosolicit on-going feedback on model performance.
C. Toapply their own judgment to the initial assessment.
D. Toensure 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 5**
- (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 6**
- (Topic 1)
According to the GDPR's transparency principle, when an AI system processes personal data in automated decision-making, controllers are required to provide data subjects specific information on?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 7**
- (Topic 1)
The framework set forth in the White House Blueprint for an AI Bill of Rights addresses all of the following EXCEPT?

A. Human alternatives, consideration and fallback.
B. High-risk mitigation standards.
C. Safe and effective systems.
D. Data privacy.

**Answer:** B

**Explanation:**

The White House Blueprint for an AI Bill of Rights focuses on protecting civil rights, privacy, and ensuring AI systems are safe and effective. It includes principles like data privacy (D), human alternatives (A), and safe and effective systems (C). However, it does not specifically address high-risk mitigation standards as a distinct category (B).

**NEW QUESTION 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)
What is the 1956 Dartmouth summer research project on AI best known as?

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

**Answer:** D

**Explanation:**

The 1956 Dartmouth summer research project on AI is best known as a meeting focused on the founding of the AI field. This conference is historically significant because it marked the formal beginning of artificial intelligence as an academic discipline. The term "artificial intelligence" was coined during this event, and it laid the foundation for future research and development in AI.
Reference: The AIGP Body of Knowledge highlights the importance of the Dartmouth
Conference as a pivotal moment in the history of AI, which established AI as a distinct field of study and research.

## NEW QUESTION 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 18**
- (Topic 1)
A US company has developed an AI system, CrimeBuster 9619, that collects information about incarcerated individuals to help parole boards predict whether someone is likely to commit another crime if released from prison.
When considering expanding to the EU market, this type of technology would?

A. Require the company to register the tool with the EU database.
B. Be subject approval by the relevant EU authority.
C. Require a detailed conformity assessment.
D. Be banned under the EU AI Act.

**Answer:** C

**Explanation:**
Under the EU AI Act, high-risk AI systems like CrimeBuster 9619 would require a detailed conformity assessment before being deployed in the EU market. This assessment ensures that the AI system complies with all relevant regulations and standards, addressing potential risks related to privacy, security, and discrimination. The company would not need to register the tool with the EU database (A), seek approval from an EU authority (B), or face a ban (D) as long as it meets the necessary conformity requirements.

**NEW QUESTION 19**
- (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 20**
- (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 22**
- (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)
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 26**
- (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 27**
- (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. Thatthe AI has full human cognitive abilities that can independently decide where to take the driver.
B. That they have obtained appropriate intellectual property (IP) licenses to use data for training the AI.
C. That the AI has strong cybersecurity to prevent malicious actors from taking control of the car.
D. That the AI can differentiate among ethnic backgrounds of pedestrians.

**Answer:** A

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


**NEW QUESTION 32**
- (Topic 1)
Which of the following is a subcategory of AI and machine learning that uses labeled datasets to train algorithms?

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

**Answer:**

D

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

**NEW QUESTION 36**
- (Topic 1)
A company is creating a mobile app to enable individuals to upload images and videos, and analyze this data using ML to provide lifestyle improvement recommendations. The signup form has the following data fields:
* 1.First name 2.Last name 3.Mobile number 4.Email ID 5.New password 6.Date of birth 7.Gender
In addition, the app obtains a device's IP address and location information while in use. What GDPR privacy principles does this violate?

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

**Answer:** A

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

**NEW QUESTION 41**
- (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)
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 52**
- (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 57**
- (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 59**
- (Topic 2)
Which of the following is the least relevant consideration in assessing whether users should be given the right to opt out from an AI system?

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

**Answer:** D

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

**NEW QUESTION 60**
- (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.
Which AI risk would NOT have been identified during the procurement process based on the categories of information requested by the third-party consultant?

A. Security.
B. Accuracy.
C. Explainability.
D. Discrimination.

**Answer:** A

**Explanation:**
The AI risk that would not have been identified during the procurement process based on the categories of information requested by the third-party consultant is security. The consultant focused on accuracy rates, diversity of training data, and system functionality, which pertain to performance and fairness but do not directly address the security aspects of the AI system. Security risks involve ensuring that the system is protected against unauthorized access, data breaches, and other vulnerabilities that could compromise its integrity. Reference: AIGP Body of Knowledge on AI Security and Risk Management.

**NEW QUESTION 62**
- (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 64**
- (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 68**
- (Topic 2)
Retraining an LLM can be necessary for all of the following reasons EXCEPT?

A. To minimize degradation in prediction accuracy due tochanges in data.
B. Adjust the model's hyper parameters specific use case.
C. Account for new interpretations of the same data.
D. To ensure interpretability of the model's predictions.

**Answer:** D

**Explanation:**
Retraining an LLM (Large Language Model) is primarily done to improve or maintain its performance as data changes over time, to fine-tune it for specific use cases, and to incorporate new data interpretations to enhance accuracy and relevance. However, ensuring interpretability of the model's predictions is not typically a reason for retraining. Interpretability relates to how easily the outputs of the model can be understood and explained, which is generally addressed through different techniques or methods rather than through the retraining process itself. References to this can be found in the IAPP AIGP Body of Knowledge discussing model retraining and interpretability as separate concepts.

**NEW QUESTION 72**
- (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 77**
- (Topic 2)
CASE STUDY
Please use the following answer the next question:
A mid-size US healthcare network has decided to develop an AI solution to detect a type of cancer that is most likely arise in adults. Specifically, the healthcare network intends to create a recognition algorithm that will perform an initial review of all imaging and then route records a radiologist for secondary review pursuant agreed-upon criteria (e.g., a confidence score below a threshold).
To date, the healthcare network has taken the following steps: defined its AI ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an AI governance committee; assembled a broad, crossfunctional team with clear roles and responsibilities; and created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.
The healthcare network intends to retain a cloud provider to host the solution and a consulting firm to help develop the algorithm using the healthcare network's existing data
and de-identified data that is licensed from a large US clinical research partner.
In the design phase, which of the following steps is most important in gathering the data from the clinical research partner?

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

**Answer:** D

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

**NEW QUESTION 80**
- (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 82**
- (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 87**
- (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 90**
- (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 93**
- (Topic 2)
During the planning and design phases of the AI development life cycle, bias can be reduced by all of the following EXCEPT?

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

**Answer:** B

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

**NEW QUESTION 95**
- (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 97
- (Topic 2)
When monitoring the functional performance of a model that has been deployed into
production, all of the following are concerns EXCEPT?

A. Feature drift.
B. System cost.
C. Model drift.
D. Data loss.

**Answer:** B

**Explanation:**
When monitoring the functional performance of a model deployed into production, concerns typically include feature drift, model drift, and data loss. Feature drift refers to changes in the input features that can affect the model's predictions. Model drift is when the model's performance degrades over time due to changes in the data or environment. Data loss can impact the accuracy and reliability of the model. However, system cost, while important for budgeting and financial planning, is not a direct concern when monitoring the functional performance of a deployed model. Reference: AIGP Body of Knowledge on Model Monitoring and Maintenance.

## NEW QUESTION 102
- (Topic 2)
An artist has been using an AI tool to create digital art and would like to ensure that it has copyright protection in the United States.
Which of the following is most likely to enable the artist to receive copyright protection?

A. Ensure the tool was trained using publicly available content.
B. Obtain a representation from the AI provider on how the tool works.
C. Provide a log of the prompts the artist used to generate the images.
D. Update the images in a creative way to demonstrate that it is the artist's.

**Answer:** D

**Explanation:**
For the artist to receive copyright protection, the most effective approach is to demonstrate that the final artwork includes sufficient creative input by the artist. By updating or altering the images in a way that reflects the artist's personal creativity, the artist can claim originality, which is a core requirement for copyright protection under U.S. law. The other options do not directly address the originality and creative input required for copyright. This is highlighted in the sections on copyright protection in the IAPP AIGP Body of Knowledge.

## NEW QUESTION 107
- (Topic 2)
What is the main purpose of accountability structures under the Govern function of the NIST AI Risk Management Framework?

A. To empower and train appropriate cross-functional teams.
B. To establish diverse, equitable and inclusive processes.
C. To determine responsibility for allocating budgetary resources.
D. To enable and encourage participation by external stakeholders.

**Answer:** A

**Explanation:**
The NIST AI Risk Management Framework's Govern function emphasizes the importance of establishing accountability structures that empower and train cross-functional teams. This is crucial because cross-functional teams bring diverse perspectives and expertise, which are essential for effective AI governance and risk management. Training these teams ensures that they are well-equipped to handle their responsibilities and can make informed decisions that align with the organization's AI principles and ethical standards. Reference: NIST AI Risk Management Framework documentation, Govern function section.

## NEW QUESTION 108
- (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 copyrightedwork.
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 112**
- (Topic 2)
CASE STUDY
Please use the following answer the next question:
A mid-size US healthcare network has decided to develop an AI solution to detect a type of cancer that is most likely arise in adults. Specifically, the healthcare network intends to create a recognition algorithm that will perform an initial review of all imaging and then route records a radiologist for secondary review pursuant Agreed-upon criteria (e.g., a confidence score below a threshold).
To date, the healthcare network has taken the following steps: defined its AI ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an AI governance committee; assembled a broad, crossfunctional team with clear roles and responsibilities; and created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.
The healthcare network intends to retain a cloud provider to host the solution and a consulting firm to help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.
The most significant risk from combining the healthcare network's existing data with the clinical research partner data is?

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

**Answer:** A

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

**NEW QUESTION 116**
......

# 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/