ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 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.

Question: 2

You asked a generative Al 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 Al 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 Al 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 AImodels.

Question: 3

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.

Question: 4

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.

Question: 5

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.

Question: 6

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 t 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 Al 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.

Question: 7

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 theaccuracy, 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.

Question: 8

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.

Question: 9

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 t A. human underwriter for final review.

ABC and the cloud provider have completed training and testing the LLM, performed a readiness assessment,

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

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

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

Answer: B

Explanation:

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

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

Question: 10

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 initialqualification assessment of potential customers, which would then be routed a human underwriter for final review.

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

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

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

Answer: A

Explanation:

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

Question: 11

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).

Question: 12

A company developed AI technology that can analyze text, video, images and sound to tag content, including the names of animals, humans and objects.
What type of AI is this technology classified as?

A. Deductive inference.
B. Multi-modal model.
C. Transformative AI.
D. Expert system.

Answer: B

Explanation:
A multi-modal model is an AI system that can process and analyze multiple types of data, such as text, video, images, and sound. This type of AI integrates different data sources to enhance its understanding and decision-making capabilities. In the given scenario, the AI technology that tags content including names of animals,

humans, and objects falls under this category. Reference: AIGP BODY OF KNOWLEDGE, which outlines the capabilities and use cases of multi-modal models.

Question: 13

All of the following are common optimization techniques in deep learning to determine weights that represent the strength of the connection between artificial neurons EXCEPT?

A. Gradient descent, which initially sets weights arbitrary values, and then at each step changes them.
B. Momentum, which improves the convergence speed and stability of neural network training.
C. Autoregression, which analyzes and makes predictions about time-series data.
D. Backpropagation, which starts from the last layer working backwards.

Answer: C

Explanation:
Autoregression is not a common optimization technique in deep learning to determine weights for artificial neurons. Common techniques include gradient descent, momentum, and backpropagation.

Autoregression is more commonly associated with time-series analysis and forecasting rather than neural network optimization. Reference: AIGP BODY OF KNOWLEDGE, which discusses common optimization techniques used in deep learning.

Question: 14

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.

Question: 15

Which of the following best defines an "AI model"?

A. A system that applies defined rules to execute tasks.
B. A system of controls that is used to govern an Al algorithm.
C. A corpus of data which an Al algorithm analyzes to make predictions.
D. A program that has been trained on a set of data to find patterns within the data.

Answer: D

Explanation:
An AI model is best defined as a program that has been trained on a set of data to find patterns within that data. This definition captures the essence of machine learning, where the model learns from the data to make predictions or decisions. Reference: AIGP BODY OF KNOWLEDGE, which provides a detailed explanation of AI models and their training processes.

Question: 16

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 Al to create the enrichment courses, and that many of the students are using generative Al 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 Al 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 Al, including by teachers and students, going forward.
Which of the following risks should be of the highest concern to individual teachers using generative Al 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.

Question: 17

CASE STUDY

Please use the following answer the next question:

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

In particular, GVC has learned that the teachers they employ used open source large language models ("LLM") to develop an online tool that customizes study questions for individual students. GVC has also discovered that an art teacher has expressly incorporated the use of generative AI into the curriculum to enable students to use prompts to create digital art.

GVC has started to investigate these practices and develop a process to monitor any use of generative AI, including by teachers and students, going forward.

What is the best reason for GVC to offer students the choice to utilize generative AI in limited, defined circumstances?

A. Toenable students to learn how to manage their time.
B. Toenable students to learn about performing research.
C. Toenable students to learn about practical applications of AI.
D. Toenable students to learn how to use AI as a supportiveeducational tool.

Answer: D

Explanation:

The best reason for GVC to offer students the choice to utilize generative AI in limited, defined circumstances is to enable students to learn how to use AI as a supportive educational tool. By integrating AI in a controlled manner, students can learn the practical applications of AI and develop skills to use AI responsibly and effectively in their educational pursuits.

Reference: The AIGP Body of Knowledge highlights the importance of teaching students about AI's practical applications and the responsible use of AI technologies. This aligns with the goal of fostering a better understanding of AI's role and its potential benefits in various contexts, including education.

Question: 18

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.

Question: 19

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

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

Answer: C

Explanation:

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

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

Question: 20

What is the 1956 Dartmouth summer research project on Al 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 Al 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.

Question: 21

What is the primary purpose of an Al impact assessment?

A. To define and evaluate the legal risks associated with developing an Al system.
B. Anticipate and manage the potential risks and harms of an Al system.
C. To define and document the roles and responsibilities of Al stakeholders.
D. To identify and measure the benefits of an Al system.

Answer: B

Explanation:
The primary purpose of an AI impact assessment is to anticipate and manage the potential risks and harms of an AI system. This includes identifying the possible negative outcomes and implementing measures to mitigate these risks. This process helps ensure that AI systems are developed and deployed in a manner that is ethically and socially responsible, addressing concerns such as bias, fairness, transparency, and accountability. The assessment often involves a thorough evaluation of the AI system's design, data inputs, outputs, and the potential impact on various stakeholders. This approach is crucial for maintaining public trust and adherence to regulatory requirements.

Question: 22

What type of organizational risk is associated with Al'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.

Question: 23

Which of the following most encourages accountability over Al systems?

A. Determining the business objective and success criteria for the Al project.
B. Performing due diligence on third-party Al training and testing data.
C. Defining the roles and responsibilities of Al stakeholders.
D. Understanding Al legal and regulatory requirements.

Answer: C

Explanation:

Defining the roles and responsibilities of AI stakeholders is crucial for encouraging accountability over AI systems. Clear delineation of who is responsible for different aspects of the AI lifecycle ensures that there is a person or team accountable for monitoring, maintaining, and addressing issues that arise. This accountability framework helps in ensuring that ethical standards and regulatory requirements are met, and it facilitates transparency and traceability in AI operations. By assigning specific roles, organizations can better manage and mitigate risks associated with AI deployment and use.

Question: 24

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

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

Answer: C

Explanation:

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

Question: 25

If it is possible to provide a rationale for a specific output of an Al 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.

Question: 26

The OECD's Ethical Al 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 Al systems.
B. Defining requirements specific to each industry sector and high-risk Al domain.
C. Focusing on Al technical design and post-deployment monitoring.
D. Balancing Al 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 industryspecific requirements (B). While explainability is

important, the primary goal is to balance innovation with ethical considerations (D).

Question: 27

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).

Question: 28

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).

Question: 29

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).

Question: 30

According to the GDPR, what is an effective control to prevent a determination based solely on automated decision-making?

A. Provide a just-in-time notice about the automated decision-making logic.
B. Define suitable measures to safeguard personal data.
C. Provide a right to review automated decision.
D. Establish a human-in-the-loop procedure.

Answer: D

Explanation:

The GDPR requires that individuals have the right to not be subject to decisions based solely on automated processing, including profiling, unless specific exceptions apply. One effective control is to establish a human-in-the-loop procedure (D), ensuring human oversight and the ability to contest decisions. This goes beyond just-in-time notices (A), data safeguarding (B), or review rights (C), providing a more robust mechanism to protect individuals' rights.

Question: 31

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.

Question: 32

According to the GDPR's transparency principle, when an Al 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 Al system about the data.
C. The data protection impact assessments carried out on the Al 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.

Question: 33

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.

Question: 34

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

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

Answer: C

Explanation:

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

Question: 35

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.

Question: 36

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.

If XYZ does not deploy and use the AI hiring tool responsibly in the United States, its liability would likely increase under all of the following laws EXCEPT?

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

Answer: B

Explanation:

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

Question: 37

CASE STUDY

Please use the following answer the next question:

XYZ Corp., a premier payroll services company that employs thousands of people globally, isembarking 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 Al hiring tools and wants to mitigate them. It also questions how best to organize and train its existing personnel to use the Al hiring tool responsibly. Their concerns are heightened by the fact that relevant laws vary across jurisdictions and continue to change.

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

A. Test the Al tool pre- and post-deployment.
B. Ensure the vendor assumes responsibility for all damages.
C. Direct the procurement team to select the most economical Al 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.

Question: 38

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

All of the following are potential negative consequences created by using the Al tool when making hiring

decisions EXCEPT?

A. Reputational harm.

B. Civil rights violations.

C. Discriminatory treatment.

D. Intellectual property infringement.

Answer: D

Explanation:

The potential negative consequences of using an AI tool in hiring include reputational harm (A), civil rights violations (B), and discriminatory treatment (C). These issues stem from biases in the AI system or its misuse, which can lead to unfair hiring practices and legal liabilities. Intellectual property infringement (D) is not a typical consequence of using AI in hiring, as it relates to the unauthorized use of protected intellectual property, which is not directly relevant to the hiring process or the potential biases within AI tools.

Question: 39

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 ManagementFramework. 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.

Question: 40

A US company has developed an Al 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 Al Act.

Answer: D

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.

Question: 41

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

A. Fines for SMEs and startups will be proportionally capped.
B. Rules on General Purpose Al will apply after 6 months as a specific provision.
C. The Al Pact will act as a transitional bridge until the Regulations are fully enacted.
D. Fines for violations of banned Al 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.

Question: 42

Which of the following is an example of a high-risk application under the EU Al Act?

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

Answer: C

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

Question: 43

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

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

Answer: C

Explanation:
Under the EU AI Act, organizations that develop and deploy high-risk AI systems are required to provide several key disclosures to ensure transparency and accountability. These include the human

oversight measures employed, how individuals can contest decisions made by the AI system, and informing individuals that an AI system is being used. However, there is no specific requirement to disclose the exact locations where data is stored. The focus of the Act is on the transparency of the AI system's operation and its impact on individuals, rather than on the technical details of data storage locations.

Question: 44

All of the following may be permissible uses of an AI system under the EU AI Act EXCEPT?

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

Answer: C

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

Question: 45

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, andaccountability 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.

Question: 46

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.

Question: 47

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.

Question: 48

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 addressingethical 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.

Question: 49

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: C

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.

Question: 50

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 decisionmaking process rather than exclusively employing such protocols. Reference: AIGP Body of Knowledge, section on AI governance frameworks.

Question: 51

The processes and methods that allow human users to understand and trust the outputs produced by AI are important in addressing which key regulatory concern?

A. Interpretable AI
B. Trustworthy AI
C. Explainable AI
D. Responsible AI

Answer: C

Explanation:

The correct answer isExplainable AIbecause it specifically refers to theability of a system todescribe the logic behind its decisions or outputsin a way that is understandable to humans. This is a key part of regulatory and ethical frameworks and is directly related to addressing theblack-box problemin AI. From the AIGP ILT Participant Guide (Module on Transparency and Explainability):
"Explainability refers to the understanding of how a black-box model works. The black-box problem exists because some models are too complex for human interpretation. Explainability methods aim to provide meaningful insight into the logic and decision-making of AI systems."
Also, according to the AI Governance in Practice Report 2024:
"Explainability refers to the representation of the underlying mechanisms of the AI system's operation... a key tenet of AI governance due to the desire to understand how AI systems are built, managed and maintained."
Thus, while Trustworthy and Responsible AI are broader concepts,explainabilityspecifically targets theregulatory concern about understanding outputs.

Question: 52

Which of the following is a foundational characteristic of effective AI governance?

A. Engagement of a cross-functional team
B. Reliance on tested vendor management processes

C. Thorough reviews of a company's public filings with experts
D. Uniform policies and procedures across developer, deployer and user roles

Answer: A

Explanation:

The correct answer isEngagement of a cross-functional team. Effective AI governance requires collaboration among various organizational functions including legal, compliance, IT, ethics, and data science.
From the AIGP Body of Knowledge:
"AI governance cannot be siloed—it requires input and oversight from across departments... A crossfunctional team ensures that ethical, technical, legal, and operational risks are all appropriately managed."
Also confirmed in the ILT Participant Guide:
"Cross-functional teams allow organizations to bring in different perspectives... Legal, compliance, and technical experts must work together to ensure responsible AI outcomes."

Question: 53

CASE STUDY
A company is considering the procurement of an AI system designed to enhance the security of IT infrastructure. The AI system analyzes how users type on their laptops, including typing speed, rhythm and pressure, to create a unique user profile. This data is then used to authenticate users and ensure that only authorized personnel can access sensitive resources.
When prioritizing the updates to its policies, rules and procedures to include the new AI system for user authentication, the organization should:

A. Update third-party data sharing policies
B. Update security controls for sensitive data
C. Ensure that any personal data used is only processed for a specific and lawful purpose
D. Reduce the complexity of the policy to make it easier for non-technical employees to understand

Answer: C

Explanation:

The correct answer isC. This action ties directly into principles ofdata minimization, purpose limitation, andlawfulness of processing, which are central to privacy and AI governance.
From the AIGP Body of Knowledge, Section on Privacy Considerations:
"Personal data must only be processed for specified and lawful purposes. Organizations must consider whether they have a legal basis for processing such data under data protection laws like the GDPR or CCPA."
Additionally, AI Governance in Practice Report 2024 emphasizes:
"One of the most significant challenges when designing and developing AI systems is ensuring the data used is appropriate for the intended purpose… Managing unnecessary data, especially data that may contain

sensitive attributes, can increase risk."

Question: 54

CASE STUDY

A company is considering the procurement of an AI system designed to enhance the security of IT infrastructure. The AI system analyzes how users type on their laptops, including typing speed, rhythm and pressure, to create a unique user profile. This data is then used to authenticate users and ensure that only authorized personnel can access sensitive resources.

All of the following are obligations of the company as a data controller when implementing its AI system EXCEPT?

A. Ensuring that third-party processors are based in the same country as the company
B. Allowing data subject access requests (DSARs)
C. Implementing technical and organizational measures
D. Conducting a Data Protection Impact Assessment (DPIA) / Privacy Impact Assessment (PIA)

Answer: A

Explanation:

The correct answer isA. While location of processors may have implications (such as for data transfers under GDPR), there isno absolute requirementthat third-party processors be based in the same country.
From the AI Governance in Practice Report 2024 and ILT Guide:
"Data controllers are responsible for ensuring that third-party processors have adequate protections, but not necessarily that they reside in the same jurisdiction. What is required is legal safeguards (e.g., SCCs) for international transfers, not same-country location."
In contrast, DPIAs, DSARs, and implementation of technical/organizational safeguards areexplicitly requiredunder GDPR and responsible AI frameworks.

Question: 55

CASE STUDY

A company is considering the procurement of an AI system designed to enhance the security of IT infrastructure. The AI system analyzes how users type on their laptops, including typing speed, rhythm and pressure, to create a unique user profile. This data is then used to authenticate users and ensure that only authorized personnel can access sensitive resources.
The data processed by the AI system would be classified as:

A. Non-sensitive personal data, since it does not reveal information about health, gender or race
B. Organizational data, since it is part of the authentication process
C. Non-personal data, as long as it is not linked to a user ID

D. Special category data, if it can be used to uniquely identify a person

Answer: D

Explanation:

The correct answer isD.Keystroke dynamics, used to identify individuals, fall underbiometric data, which is aspecial category of personal dataunder the GDPR and other frameworks.
From the AI Governance in Practice Report 2024:
“Keystroke dynamics may constitute biometric data if used to uniquely identify an individual…
Biometric data is classified as special category personal data and requires higher protection standards.”
Also reflected in ILT Participant Guide:
“Biometric data, such as facial images, voiceprints, iris scans or keystroke patterns, are treated as special category data when they are used for the purpose of uniquely identifying individuals.”

Question: 56

Scenario:
A large multinational organization is rolling out a company-wide AI governance initiative. To build awareness and support adoption, they are evaluating different ways to train employees and stakeholders across departments, including legal, technical, marketing, and customer-facing roles. Which of the following typical approaches is a large organizationleast likelyto use to responsibly train stakeholders on AI terminology, strategy and governance?

A. Providing all technical employees education on AI development so they can retool and participate in the development of AI systems
B. Providing training on AI ethics, based on the extent to which the organization seeks to promote a responsible AI culture
C. Providing role-specific training, based on whether the organization uses a centralized, federated or decentralized governance model
D. Providing information and education to customers and users to understand the capabilities and limitations of the AI tools with which they interact

Answer: A

Explanation:

The correct answer isA. While educating technical staff is important, expectingall technical employees to be retooled as AI developersis unrealistic and not aligned with scalable governance practices.
From the AIGP ILT Guide:
"Training approaches should berole-specificand align with the individual's function and responsibilities...
Organizations typically do not expect every technical role to participate in model development."
The AI Governance in Practice Report 2024 supports tailored approaches:
“Cross-functional training should be specific to the individual's role and exposure to AI risk... Rolebased education supports scalability and comprehension.”
Thus,broad development training for all technical employeesis the least practical and least likely approach.

Question: 57

Scenario:
An organization is building a compliance program to ensure responsible AI deployment. It aims to align operations with AI risk frameworks and mitigate legal, ethical, and operational risks, while still promoting innovation.
Which of the following would be theleast likelystep for an organization to take when designing an integrated compliance strategy for responsible AI?

A. Meeting with and obtaining approval from senior management
B. Launching a survey to understand the concerns and interests of potentially impacted stakeholders C. Consulting experts to consider the ethical principles underpinning the use of AI within the organization
D. Employing a new software platform to modernize existing compliance processes across the organization

Answer: D

Explanation:

The correct answer isD. While modernization through software may support efficiency, it isnot a foundational or essential componentof designing an integrated strategy.
From the AI Governance in Practice Report 2024:
"Integrated strategies rely on senior management support, ethical reviews, and stakeholder engagement... The use of tools and platforms may come later as an operational enhancement." Also confirmed in AIGP Body of Knowledge:
"Key components of a governance framework include leadership buy-in, ethical analysis, and stakeholder input. Tools are supporting elements—not strategic drivers."

Question: 58

Scenario:
Business A provides grammar and writing assistance tools and licenses a generative AI model from Business B to enhance its offerings. Business A is concerned that the AI model might produce inappropriate or toxic content and wants to implement governance processes to prevent this.
Which of the following governance processes should Business A take tobest protect its usersagainst potentially inappropriate text?

A. Business A should fine-tune the AI model on user-generated text that has been verified to be appropriate
B. Business A should test that the AI model performs as expected and meets their minimum requirements for filtering toxic or obscene text
C. Business A should establish a user reporting feature that allows users to flag toxic or obscene text, and report any incidents to Business B
D. Business A should ask Business B for detailed documentation on the generative AI model's training data and whether it contained toxic or obscene sources

Answer: B

Explanation:

The correct answer isB. According to responsible AI practices,pre-deployment testingto ensurethe model behaves as expected and aligns with organizational requirements is critical.
From the AIGP ILT Guide:
"Testing for unacceptable outcomes such as toxicity, discrimination, or hallucinations should be included in the AI governance life cycle, particularly during development and prior to deployment." Also emphasized in the AI Governance in Practice Report 2024:
"Organizations must verify legal and regulatory compliance, monitor performance, and mitigate risksprior to deployment."
Testing the model tomeet safety and appropriateness standardsis more proactive and preventive than relying solely on user reporting or requesting documentation.

Question: 59

What is theprimary purposeof an AI impact assessment?

A. To determine whether a conformity assessment is needed
B. To escalate the findings to the appropriate owner(s)
C. To identify and measure the benefits of an AI system
D. To anticipate and manage the potential risks and harms of an AI system

Answer: D

Explanation:

The correct answer isD. AI Impact Assessments are primarily used to identify and managerisks and harmsassociated with AI systems.
From the AIGP Body of Knowledge:
"The goal of an AI impact assessment is to ensure that risks are identified, evaluated, and mitigated prior to or during development and deployment."
As further confirmed in the AI Governance in Practice Report 2024 (Part III):
"Risk-based tools like DPIAs and Algorithmic Impact Assessments help identify potential risks to individuals and society, enabling organizations to implement mitigation plans and safeguards." While benefits may be noted in such assessments, thecore objectiveis to manage risks andpromote responsible AI.

Question: 60

All of the following areunique characteristics of AIthat require a comprehensive approach to governanceEXCEPT?

A. Autonomy
B. Automation
C. Adaptability

D. Speed and scale

E. Superintelligence

Answer: E

Explanation:

The correct answer isE – Superintelligence.

While the other options (Autonomy, Automation, Adaptability, Speed and Scale) are commonly cited as real-world characteristics that affect governance today,superintelligenceremainsa theoretical concept.

From the AIGP ILT Guide and AI Governance in Practice Report 2024:

"Core AI characteristics—such as automation, adaptability, speed, and autonomy—require active governance due to their impact on decision-making, legal liability, and risk."

However,superintelligenceis speculative and not a current feature of AI systems under governance frameworks like theEU AI Act or NIST RMF.

Thus, it'snot a current characteristicrequiring governance in real-world enterprise settings.

Question: 61

Scenario:

An organization is planning to deploy a new internal application that uses AI to make automated decisions about individuals. This application will process personal information and may affect individuals' access to certain benefits or opportunities.

Which of the following documents must be updated to ensure transparency?

A. The organization's website privacy notice

B. The organization's acceptable use policy

C. The organization's privacy policy

D. The user privacy notice

Answer: D

Explanation:

The correct answer isD. Transparency obligations under data protection laws, such as GDPR and most AI governance frameworks, require thatusers whose data is being processedbe directly informed.

From the AIGP ILT Guide (Privacy Module):

"The user privacy notice must be updated to explain the nature of automated processing, the logic involved, and the significance and consequences for the data subject."

Also, per AI Governance in Practice Report 2024 (Part III):

"Transparency obligations apply throughout the lifecycle of AI… Individuals must be informed about automated decision-making and profiling that may impact them."

Unlike internal policies or general privacy notices,the user privacy noticeprovides direct transparency to theindividual data subjectsaffected by AI processing.

Question: 62

CASE STUDY

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

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

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

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

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

All of the following are potential negative consequences created by using the AI tool to help make hiring decisions EXCEPT?

A. Automation bias
B. Candidate quality
C. Privacy violations
D. Disparate impacts

Answer: B

Explanation:

The correct answer isB. "Candidate quality" isnota negative consequence of using AI—rather, it is the intendedbenefitof using such tools (e.g., more efficient filtering of strong candidates).

From the AIGP ILT Guide:

"Automation bias, disparate impact, and privacy risks are well-documented concerns in AI-assisted hiring. These risks may arise when AI models replicate biases present in training data or obscure the decision logic."

AI Governance in Practice Report 2024 (Bias and Fairness Section) also warns:

"Improper AI use in hiring can lead to disparate impact, where neutral criteria disproportionately disadvantage protected groups."

Candidate quality is agoal, not a risk, makingB the correct answerfor what isnota negative outcome.

Question: 63

Scenario:
A company using AI for resume screening understands the risks of algorithmic bias and the evolving legal requirements across jurisdictions. It wants to implement the right governance controls to prevent reputational damage from misuse of the AI hiring tool.
Which of the following measures should the company 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 provides indemnification for the AI tool
C. Require the procurement and deployment teams to agree upon the AI tool
D. Continue to require the company's hiring personnel to manually screen all applicants

Answer: A

Explanation:

The correct answer isA. Pre- and post-deployment testing ensuresbias, accuracy, and fairnessare evaluated and corrected as needed, which isessential for reputational risk mitigation.
From the AIGP Body of Knowledge:
"Testing AI systems before and after deployment is critical to ensure performance, fairness, and compliance. Failing to do so may result in reputational damage and legal exposure."
AI Governance in Practice Report 2024 (Bias/Fairness and Risk Sections):
"System impact assessments, testing, and post-deployment monitoring are necessary to identify and mitigate risks… This supports both compliance and public trust."
Testing is proactive, unlike indemnification (which transfers risk after damage), or requiring manual review (which defeats automation).

Question: 64

CASE STUDY
A premier payroll services company that employs thousands of people globally, is embarking on a new hiring campaign and wants to implement policies and procedures to identify and retain the best talent. The new talent will help the company's product team expand its payroll offerings to companies in the healthcare and transportation sectors, including in Asia.
It has become time consuming and expensive for HR to review all resumes, and they are concerned that human reviewers might be susceptible to bias.
To address these concerns, the company is considering using a third-party AI tool to screen resumes and assist with hiring. They have been talking to several vendors about possibly obtaining a third- party AI-enabled hiring solution, as long as it would achieve its goals and comply with all applicable laws.
The organization has a large procurement team that is responsible for the contracting of technology solutions. One of the procurement team's goals is to reduce costs, and it often prefers lower-cost solutions. Others within the company deploy 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 organization continues planning the adoption of an AI tool to support hiring, but is concerned about potential bias in content generated by AI systems and how that could affect public perception. Which of the following measures should the company adopt tobest mitigate its risk of reputational harmfrom using the AI tool?

A. Test the AI tool pre- and post-deployment
B. Ensure the vendor provides indemnification for the AI tool
C. Require the procurement and deployment teams to agree upon the AI tool
D. Continue to require the company's hiring personnel to manually screen all applicants

Answer: A

Explanation:

Note:This is the same scenario and question as Question 21 and thus has thesame correct answer: A. It's possible this was duplicated in your original input.

Repeated for clarity:

"Testing AI tools pre- and post-deployment helps ensure they perform as expected and do not introduce bias, privacy issues, or fairness concerns. This mitigates reputational and legal risk." TheAI Governance in Practice Report 2024further reinforces:

"Ongoing monitoring and testing post-deployment allows organizations to catch and correct unintended impacts... especially important in HR and hiring contexts."

Question: 65

Scenario:

An organization is evaluating different AI models for integration into its internal workflows. Before moving forward with a particular AI solution from a third-party vendor, the governance team needs to assess the ethical and operational implications of the model.

The most important policy to assess the operations of an AI model is to follow the:

A. Acceptable use policy of the model provider
B. Privacy policy of the model provider
C. Security policy of the model provider
D. Code of conduct policy of the model provider

Answer: A

Explanation:

The correct answer isA. TheAcceptable Use Policy (AUP)sets the primary terms and conditions under which the AI model can be used, including limitations, prohibited uses, and operational constraints. From the AIGP ILT Guide:

"The AUP is the most critical document to assess what you can and cannot do with the AI system. It governs use cases, outlines forbidden uses (e.g., disinformation), and supports responsible AI practices."

Also confirmed in the AI Governance in Practice Report 2024 (Third-Party AI Assurance section): "Reviewing and adhering to the acceptable use policy ensures that the model is being used in a manner that aligns with both provider expectations and ethical AI practices."

Question: 66

Which of the following arenotconsidered biometric data under U.S. privacy laws?

A. Iris scans
B. Walking gait
C. Keystroke dynamics
D. GPS location of a user's fitness watch

Answer: D

Explanation:

The correct answer isD. GPS location data isnot biometric data—it is consideredgeolocation data, which is personal data butnot biometricunder most U.S. laws.

From the AIGP ILT Guide (Data Privacy Module):

"Biometric data includes measurable biological or behavioral characteristics such as iris scans, facial recognition, voice prints, and keystroke patterns when used for identification."

AI Governance in Practice Report 2024 (Privacy and Data Protection section):

"Location data, while sensitive, is not considered biometric unless it's tied to a uniquely identifyingbiological trait."

Thus,GPS location data, while potentially sensitive, is not classified as biometric.

Question: 67

Scenario:

A company is using different types of AI systems to enhance consumer engagement. These include chatbots, recommendation engines, and automated content generation tools.

Which of the following situations would beleast likelyto raise concerns under existing consumer protection laws?

A. An AI algorithm being used in a credit decision-making process by a financial institution
B. An AI customer service system claiming that it is as accurate as a human support agent
C. An AI tool using scraped digital content to generate news summaries on a publishing website
D. An online platform offering recommendations to its users by displaying user-specific content and targeted advertisements

Answer: D

Explanation:

The correct answer isD. Personalized content and advertisements, as long as properly disclosed and non-deceptive, arenot generally a consumer protection issueunder current legal regimes.
From the AI Governance in Practice Report 2024 (Consumer Protection Section):
"Standard practices like targeted advertising and recommendations are widely accepted provided they comply with transparency and consent requirements."
Meanwhile, credit decision-making and misleading AI performance claims (Answers A and B) havealready led to regulatory enforcement.
The AIGP ILT Guide highlights:
"Deceptive claims, biased financial decisions, and unauthorized data use may violate consumer protection and privacy laws. Advertising personalization is routine but must be disclosed appropriately."

Question: 68

Scenario:
A European AI technology company was found to be non-compliant with certain provisions of the EU AI Act. The regulator is considering penalties under the enforcement provisions of the regulation. According to the EU AI Act, which of the following non-compliance examples could lead to fines of up to €15 million or 3% of annual worldwide turnover(whichever is higher)?

A. In case of AI Act prohibitions
B. In case of breach of a provider's obligations for high-risk AI systems
C. In case of the supply of misleading information to notified bodies in reply to a request
D. In case of a breach of AI Act prohibition by the Union institutions, bodies, offices and agencies

Answer: B

Explanation:

The correct answer isB. The EU AI Act assigns atiered penalty systembased on the severity of the violation. A breach ofobligations related to high-risk AI systemsfalls into the mid-tier category, triggering fines of €15 million or 3% of annual global turnover.
From the AIGP ILT Guide – EU AI Act Module:
"Providers of high-risk AI systems must comply with strict documentation, testing, monitoring, and registration obligations. Breaches of these result in significant fines of up to €15 million or 3% of turnover."
AI Governance in Practice Report 2024 supports this:
"Non-compliance with obligations under Title III (high-risk systems) leads to financial penalties under Article 71(3) of the EU AI Act."
Note: Thehighest penalty (€35 million or 7%)applies toprohibited AI uses, not to obligations for high- risk systems.

Question: 69

Scenario:
An organization is developing a powerful general-purpose AI (GPAI) model that has systemic impact.
The compliance team is assessing what legal obligations apply under the EU AI Act.
Under the EU AI Act, which of the following compliance actions appliesonly to General Purpose AI models with systemic risk?

A. Publishing a detailed summary of the data used to train the model
B. Maintaining up-to-date technical documentation, including testing details
C. Implementing an intellectual property policy to comply with EU copyright laws
D. Making information available to downstream providers who integrate the model into their AI systems

Answer: A

Explanation:

The correct answer isA. Only GPAI modelswith systemic riskmustpublish a detailed summary of training datato meet transparency and accountability standards under the EU AI Act.
From the AI Governance in Practice Report 2024 (EU AI Act Section):
"For GPAI systems with systemic risk, providers must publish sufficiently detailed summaries of the content used to train the model."
Also, the AIGP ILT Guide confirms:
"The obligation to disclose summaries of training data applies only to systemic-risk GPAI models, not all general-purpose models or high-risk systems."
This unique requirement is part of the Act's effort to increasetransparency and auditabilityfor powerful foundational models.

Question: 70

Which of the following may bepermissible usesof an AI system under the EU AI ActEXCEPT?

A. Using biometrics in abduction cases
B. Detecting emotions in a telemedicine session
C. Improving the response time for emergency services
D. Detecting emotions in a workplace for employee morale

Answer: D

Explanation:

The correct answer isD.Emotion recognition in the workplaceis flagged asunacceptable or highly restrictedunder the EU AI Act due to its intrusive nature and potential for misuse.
From the AIGP ILT Guide – EU AI Act Training Module:

"AI systems that monitor individuals' emotions in the workplace or educational settings are listed among prohibited or strictly limited practices under Article 5."
AI Governance in Practice Report 2024 supports this interpretation:
"Emotion recognition systems, especially in sensitive contexts such as employment or education, raise significant concerns under EU fundamental rights law and are likely to be restricted." Other uses listed—such as emergency response or emotion detection in healthcare—may fall underlawful and beneficial uses, especially whenjustified by public interest.

Question: 71

Scenario:
A distributor operating in the EU is responsible for selling imported high-risk AI systems to businesses. The distributor wants to ensure they fulfill all applicable obligations under the EU AI Act. All of the following are obligations of a distributor of high-risk AI systems under the EU AI Act EXCEPT?

A. Corrective actions
B. Verification of CE marking
C. Registration in EU Database
D. Communication with national authorities

Answer: C

Explanation:

The correct answer isC.Registration in the EU databaseis an obligation ofprovidersof high-risk AI systems—not distributors.
From the AIGP ILT Guide – Roles & Obligations Module:
"Distributors must verify CE marking, ensure instructions for use are provided, inform authorities of risks, and take corrective action when necessary. However, registration duties in the EU database lie with the provider."
Also from the AI Governance in Practice Report 2024:
"The AI Act differentiates responsibilities for developers, providers, importers, and distributors. Only providers of high-risk systems are obligated to register their systems in the EU AI Database." Distributors focus onverification and communication, not formal registration.

Question: 72

Scenario:
A U.S.-based AI governance professional is evaluating resources from the National Institute of Standards and Technology (NIST) to guide the organization's AI risk assessment strategy. They are particularly interested in programs focused on assessing AI-specific impacts.
The main purpose of NIST'sAssessing Risks and Impacts of AI (ARIA)program is to:

A. Provide a suite of resources to manage risks
B. Pilot new standards for AI red-teaming

C. Promote interoperability across AI systems
D. Offer a regulatory sandbox for risk reporting

Answer: A

Explanation:
The correct answer isA. TheARIA programby NIST is explicitly designed tosupport stakeholders in understanding and managing the risks and impactsof AI systems.
From the AIGP ILT Guide – U.S. Risk Frameworks Module:
"NIST's ARIA program develops and pilots assessment tools for AI risks and impacts, aimed at improving organizational capacity for responsible AI use."
Also cited in the AI Governance in Practice Report 2024 (Frameworks Section):
"ARIA supports and aligns with the AI Risk Management Framework by helping organizations assess AI harms, safety concerns, and societal implications."
ARIA isnota red-teaming or sandbox program—it's anassessment and governance resource.

Question: 73

Scenario:
A global organization wants to align with international frameworks on AI governance. They are reviewing guidance from the OECD on how to incorporate broader governance tools into their AI program.
Codes of conductandcollective agreementsare what type of assessment tools as defined by theOrganization for Economic Cooperation and Development (OECD)?

A. Educational
B. Procedural
C. Technical
D. Analytic

Answer: B

Explanation:

The correct answer isB – Procedural. TheOECD Framework for Classifying AI Systemscategorizescodes of conduct and collective agreementsasprocedural toolsbecause they guide internal governance and decision-making processes.
From the AIGP ILT Participant Guide – Global Governance Models:
"Procedural tools include internal codes of conduct, collective agreements, and procedural audits that guide governance without necessarily involving technical measurement."
AI Governance in Practice Report 2024 elaborates:
"These procedural tools support internal accountability mechanisms and ethics compliance frameworks… they are part of soft governance."
These toolsdo not measure or analyze technical performance, hence they arenot technical or analytic.

Question: 74

Scenario:
A mid-sized tech firm is building its AI governance program and is exploring ISO/IEC standards that could support consistency in terminology and risk assessment processes across teams.
ISO/IEC 22989andISO/IEC 42001can be valuable resources for AI Governance professionals inall of the following ways EXCEPT:

A. Establishing terminology and describing concepts so that governance team members can communicate with diverse parties and stakeholders from around the world
B. Being applicable to organizations of any size and industry seeking to use AI responsibly and effectively in their design processes, information systems and controls
C. Addressing specific issues related to managing procurement processes with third parties that provide or develop AI systems for their organization
D. Recommending key activities to assess and manage risk: test, evaluate, verify and validate (TEVV)

Answer: C

Explanation:

The correct answer isC. ISO/IEC 22989 and 42001 focus onterminology, risk, and management systems, butdo not specifically address procurement-related concerns with third-party vendors. From the AIGP Body of Knowledge – Standards Section:
"ISO/IEC 22989 defines terminology and foundational concepts. ISO/IEC 42001 provides a management system standard for AI. They are not procurement-focused documents."
Also confirmed in the AI Governance in Practice Report 2024:
"These standards help establish common language and risk governance procedures. Procurement governance typically falls under separate frameworks or sector-specific guidance."
Thus,procurement governance (Option C)is not a central use case for these standards.

Question: 75

CASE STUDY
A global marketing agency is adapting a large language model ("LLM") to generate content for an upcoming marketing campaign for a client's new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.
The marketing agency is accessing the LLM through an application programming interface ("API") developed by a third-party technology company. They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers. Both the marketing agency and the technology company have taken reasonable steps to address Al governance.
The marketing company has:

- Entered into a contract with the technology company with suitable representations and warranties.
- Completed an impact assessment on the LLM for this intended use.
- Built technical guidance on how to measure and mitigate bias in the LLM.

- Enabled technical aspects of transparency, explainability, robustness and privacy.
- Followed applicable regulatory requirements.
- Created specific legal statements and disclosures regarding the use of the AI on its client's advertising.

The technology company has:

- Provided guidance and resources to developers to address environmental concerns.
- Build technical guidance on how to measure and mitigate bias in the LLM.
- Provided tools and resources to measure bias specific to the LLM.
- Enabled technical aspects of transparency, explainability, robustness and privacy.
- Mapped and mitigated potential societal harms and large-scale impacts.
- Followed applicable regulatory requirements and industry standards.
- Created specific legal statements and disclosures regarding the LLM. including with respect to IP and rights to data.

Which stakeholder is responsible for the lawful collection of data used to train the foundational AI model?

A. The marketing agency
B. The tech company
C. The data aggregator
D. The marketing agency's client

Answer: B

Explanation:

The correct answer isB – The tech company. The party thatdevelops and trains the foundational modelis responsible for ensuring thelawful collection of training data.

From the AIGP ILT Guide – Foundational Models & Data Governance:

"Responsibility for the lawfulness of data collection typically lies with the party that trains the model—usually the provider or developer of the foundational model."

AI Governance in Practice Report 2024 confirms:

"General Purpose AI providers are required to ensure that training data is lawfully acquired, including compliance with intellectual property and privacy requirements."

The marketing agency is only auserordownstream integrator, not responsible for original data collection.

Question: 76

CASE STUDY

A global marketing agency is adapting a large language model ("LLM") to generate content for an upcoming marketing campaign for a client's new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.

The marketing agency is accessing the LLM through an application programming interface ("API") developed by a third-party technology company. They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers. Both the marketing agency

and the technology company have taken reasonable steps to address AI governance.

The marketing company has:

- Entered into a contract with the technology company with suitable representations and warranties.
- Completed an impact assessment on the LLM for this intended use.
- Built technical guidance on how to measure and mitigate bias in the LLM.
- Enabled technical aspects of transparency, explainability, robustness and privacy.
- Followed applicable regulatory requirements.
- Created specific legal statements and disclosures regarding the use of the AI on its client's advertising.

The technology company has:

- Provided guidance and resources to developers to address environmental concerns.
- Build technical guidance on how to measure and mitigate bias in the LLM.
- Provided tools and resources to measure bias specific to the LLM.
- Enabled technical aspects of transparency, explainability, robustness and privacy.
- Mapped and mitigated potential societal harms and large-scale impacts.
- Followed applicable regulatory requirements and industry standards.
- Created specific legal statements and disclosures regarding the LLM. including with respect to IP and rights to data.

The technology company has also addressed environmental concerns and societal harms.

Which of the following results would be considered biased outputs from this AI system EXCEPT?

A. The generated ads are sent to construction companies, not individual workers B. The content generated for minority construction workers is insufficient C. The images of female workers are hyper-sexualized D. The advertising text generated for female audiences focuses on color and style

Answer: A

Explanation:

The correct answer isA. Sending ads to construction companies (business entities) rather than individual workers isa business targeting decision, not inherently a biased AI output.

From the AIGP ILT Participant Guide – Bias & Fairness Module:

“Biased outputs often include stereotyping, exclusion of underrepresented groups, or reinforcing harmful societal assumptions.”

Examples likeinsufficient representation of minority groupsorgender-stereotyping in visuals or languageare typical manifestations of bias.

AI Governance in Practice Report 2024 also notes:

“Bias in generative models may manifest in representation gaps, stereotyping, or unequal performance across demographic groups.”

Option A, by contrast, describes adistribution strategy, not a bias generated by the AI model.

Question: 77

CASE STUDY

A global marketing agency is adapting a large language model ("LLM") to generate content for an upcoming marketing campaign for a client's new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.

The marketing agency is accessing the LLM through an application programming interface ("API")developed by

a third-party technology company. They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers. Both the marketing agency and the technology company have taken reasonable steps to address AI governance.

The marketing company has:

- Entered into a contract with the technology company with suitable representations and warranties.
- Completed an impact assessment on the LLM for this intended use.
- Built technical guidance on how to measure and mitigate bias in the LLM.
- Enabled technical aspects of transparency, explainability, robustness and privacy.
- Followed applicable regulatory requirements.
- Created specific legal statements and disclosures regarding the use of the AI on its client's advertising.

The technology company has:

- Provided guidance and resources to developers to address environmental concerns.
- Build technical guidance on how to measure and mitigate bias in the LLM.
- Provided tools and resources to measure bias specific to the LLM.
- Enabled technical aspects of transparency, explainability, robustness and privacy.
- Mapped and mitigated potential societal harms and large-scale impacts.
- Followed applicable regulatory requirements and industry standards.
- Created specific legal statements and disclosures regarding the LLM. including with respect to IP and rights to data.

The agency has taken governance actions such as:

Conducting an impact assessment

Providing legal disclosures

Enabling bias mitigation and explainability

Complying with regulatory requirements

Which of the following should be included in the marketing company's disclosures about the use of the LLM EXCEPT?

A. Intended purpose
B. Proprietary methods
C. Compliance with law
D. Acknowledgement of limitations

Answer: B

Explanation:

The correct answer isB – Proprietary methods. While transparency is important, organizations arenot obligated to disclose proprietary algorithms, methods, or trade secretsin public disclosures.

From the AIGP Body of Knowledge – Transparency & Disclosures:

"AI system users should disclose the purpose, capabilities, limitations, and applicable legal context— but not sensitive IP."

AI Governance in Practice Report 2024 (Transparency Section) states:

"Disclosure requirements balance public understanding with the need to protect proprietary business interests. Proprietary training methods are not expected to be disclosed."

Thus, while it's best practice to disclose theintended purpose, legal compliance, and system

limitations,internal proprietary techniquesare usually excluded.

Question: 78

CASE STUDY

A global marketing agency is adapting a large language model ("LLM") to generate content for an upcoming marketing campaign for a client's new product: a hard hat designed for construction workers of any gender to better protect them from head injuries.

The marketing agency is accessing the LLM through an application programming interface ("API") developed by a third-party technology company. They want to generate text to be used for targeted advertising communications that highlight the benefits of the hard hat to potential purchasers. Both the marketing agency and the technology company have taken reasonable steps to address Al governance.

The marketing company has:

- Entered into a contract with the technology company with suitable representations and warranties.
- Completed an impact assessment on the LLM for this intended use.
- Built technical guidance on how to measure and mitigate bias in the LLM.
- Enabled technical aspects of transparency, explainability, robustness and privacy.
- Followed applicable regulatory requirements.
- Created specific legal statements and disclosures regarding the use of the Al on its client's advertising.

The technology company has:

- Provided guidance and resources to developers to address environmental concerns.
- Build technical guidance on how to measure and mitigate bias in the LLM.
- Provided tools and resources to measure bias specific to the LLM.
- Enabled technical aspects of transparency, explainability, robustness and privacy.
- Mapped and mitigated potential societal harms and large-scale impacts.
- Followed applicable regulatory requirements and industry standards.
- Created specific legal statements and disclosures regarding the LLM. including with respect to IP and rights to data.

The marketing company and its tech provider have taken reasonable steps to govern the AI's use, including legal disclosures, impact assessments, and bias mitigation. However, the company wants to takeone more stepto improve governance and reduce risks related to ongoing oversight and accountability.

While the marketing agency took steps to mitigate its risks, the best additional step would be to:

A. Negotiate an intellectual property indemnity from the technology company
B. Evaluate the use of AI in the marketing industry to identify best practices
C. Engage a third party to lead the procurement selection process
D. Establish a governance committee to oversee the project

Answer: D

Explanation:

The correct answer isD. Forming adedicated governance committeeensures continuous oversight, role clarity, and accountability throughout the AI lifecycle.

From the AIGP ILT Guide – Governance Structures:

"Organizations using AI in high-impact scenarios should establish a governance body responsible for

oversight of risk, compliance, and ethical alignment."
Also reflected in AI Governance in Practice Report 2024:
"Committees support cross-functional decision-making, provide guidance for updates, and maintain accountability. This is especially critical for high-stakes applications like marketing to diverse audiences."
Options A, B, and C are valid supplementary actions, butDoffers a long-term and systematic governance mechanism.

Question: 79

Scenario:
An enterprise is evaluating multiple third-party generative AI tools to integrate into its platform. As part of its AI governance policy, it is assessing themost effective methodsto reduce risks related to bias, data misuse, and liability when using third-party solutions.
All of the following are commonly adopted processes and policies in reducing potential risks introduced by third-party AI tools or applications EXCEPT:

A. Including clauses in the procurement agreement for buyers of generative AI tools to put certain liabilities on the tool supplier
B. Allowing publicly available information and personally identifiable information (PII) to be incorporated into the prompt
C. Requiring an independent third-party bias audit for third-party generative AI tools
D. Requiring new use cases of the generative AI tools or applications to be reviewed and approved by the generative AI governance body

Answer: B

Explanation:

The correct answer isB. AllowingPIIto be freely entered into prompts without safeguards is considered amajor privacy and security riskand is not a responsible governance practice. From the AIGP ILT Guide – Generative AI & Third-Party Risk Management:
"Use of personal or sensitive information in AI prompts can result in unintended exposure, regulatory breaches, and downstream liability."
The AI Governance in Practice Report 2024 highlights:
"PII should be minimized or protected by design. Prompt engineering should prevent entry of personally identifiable data unless legally and technically safeguarded."
A, C, and D are established best practices under responsible AI procurement and use.

Question: 80

Scenario:
A financial services company is planning a new AI project to assess creditworthiness. The AI team is mapping out what tasks should be completed during theplanning phaseof the AI lifecycle.
The planning phase of the AI lifecycle includes all of the following EXCEPT:

A. Definition of underlying assumptions
B. Approach to governance
C. Choice of the architecture
D. Context in which the model will operate

Answer: C

Explanation:

The correct answer isC. Thechoice of architecture(e.g., neural networks vs. decision trees) is typically part of thedesign and developmentphase, not the initial planning.
From the AIGP Body of Knowledge – AI Lifecycle Module:
"Planning involves scoping, context definition, stakeholder identification, governance planning, and assumptions—not yet model selection."
Confirmed in the ILT Participant Guide:
"Design decisions such as architecture or algorithm type come after planning—usually during development based on technical feasibility and data availability."

Question: 81

Scenario:
An organization wants to leverage its existing compliance structures to identify AI-specific risks as part of an ongoing data governance audit.
Which of the following compliance-related controls within an organization ismost easily adaptedto identify AI risks?

A. Privacy training
B. Penetration testing
C. Transfer risk assessments
D. Privacy impact assessments

Answer: D

Explanation:

The correct answer isD – Privacy impact assessments (PIAs). These aredirectly adaptablefor identifying risks in AI systems, particularly around data usage, bias, and individual impacts. From the AIGP ILT Guide – Risk Management Module:
"PIAs and DPIAs are existing tools used in privacy compliance that can be extended to evaluate the risks of AI, including fairness, explainability, and legality."
AI Governance in Practice Report 2024 further explains:
"Organizations can adapt privacy impact assessments to evaluate the ethical, legal, and technical risks posed by AI systems. They provide a structured and recognized method."
PIAs are preferable over general security practices (like pen testing) which do not address algorithmic bias or

legal compliance directly.

Question: 82

Scenario:
A public sector agency is reviewing proposed AI use cases for improving services. It wants to prioritize implementations that deliver value butminimize unintended negative consequences. When evaluating which AI use cases to implement, an organization should consider all of the following EXCEPT:

A. Related TEVV (test, evaluate, verify, validate) and system metrics
B. The users and their expectations
C. Equitable access to the AI tool
D. Potential positive and negative impacts of the system

Answer: C

Explanation:

The correct answer isA. While TEVV is important inlater lifecycle phases, it isnot the primary considerationwhen evaluating and prioritizinguse cases.
From the AIGP Body of Knowledge – Use Case Assessment Module:
"Use case evaluation focuses on value, impact, fairness, and accessibility—technical testing considerations come later."
ILT Guide confirms:
"Organizations should first assess whether the AI system provides equitable outcomes and aligns with stakeholder expectations. TEVV is part of implementation, not initial prioritization." Thus,Ais not a top-level consideration during use caseselection.

Topic 2, Part 2

Question: 83

What is the main purpose of accountability structures under the Govern function of the NIST Al 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.

Question: 84

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.

Question: 85

A company initially intended to use a large data set containing personal information to train an AI model. After consideration, the company determined that it can derive enough value from the data set without any personal information and permanently obfuscated all personal data elements before

training the model.
This is an example of applying which privacy-enhancing technique (PET)?

A. Anonymization.
B. Pseudonymization.
C. Differential privacy.
D. Federated learning.

Answer: A

Explanation:

Anonymization is a privacy-enhancing technique that involves removing or permanently altering personal data elements to prevent the identification of individuals. In this case, the company obfuscated all personal data

elements before training the model, which aligns with the definition of anonymization. This ensures that the data cannot be traced back to individuals, thereby protecting their privacy while still allowing the company to derive value from the dataset. Reference: AIGP Body of Knowledge, privacy-enhancing techniques section.

Question: 86

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.

Question: 87

CASE STUDY
Please use the following answer the next question:

A mid-size US healthcare network has decided to develop an Al 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 Al ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an Al governance committee; assembled a broad, crossfunctional team with clear roles and responsibilities; and created policies and procedures to document standards, workflows, timelines and risk thresholds during the project.
The healthcare network intends to retain a cloud provider to host the solution and a consulting firm to help develop the algorithm using the healthcare network's existing data and de-identified data that is licensed from a large US clinical research partner.
Which stakeholder group is most important in selecting the specific type of algorithm?

A. The cloud provider.
B. The consulting firm.
C. The healthcare network'sdata science team.

D. The healthcare network's AI governance committee.

Answer: C

Explanation:

In selecting the specific type of algorithm for the AI solution, the healthcare network's data science team is most important. This team possesses the technical expertise and understanding of the data, the clinical context, and the performance requirements needed to make an informed decision about which algorithm is most suitable. While the cloud provider and consulting firm can offer support and infrastructure, and the AI governance committee provides oversight, the data science team's specialized knowledge is crucial for selecting and implementing the appropriate algorithm.

Reference: AIGP Body of Knowledge, AI governance and team roles section.

Question: 88

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.

Question: 89

CASE STUDY

Please use the following answer the next question:

A mid-size US healthcare network has decided to develop an Al 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 Al ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an Al 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.

Question: 90

CASE STUDY

Please use the following answer the next question:

A mid-size US healthcare network has decided to develop an Al 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 Al ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an Al governance

committee; assembled a broad, crossfunctional team with clear roles andresponsibilities; 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 Al 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.

Question: 91

CASE STUDY
Please use the following answer the next question:
A mid-size US healthcare network has decided to develop an Al 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 Al ethical principles: conducted discovery to identify the intended uses and success criteria for the system: established an Al 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. Deidentified 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.

Question: 92

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

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

Answer: A

Explanation:

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

Question: 93

The most important factor in ensuring fairness when training an Al 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.

Question: 94

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.

Question: 95

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

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

Answer: C

Explanation:

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

Question: 96

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.

Question: 97

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.

Question: 98

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.

Question: 99

After completing model testing and validation, which of the following is the most important step that an organization takes prior to deploying the model into production?

A. Perform a readiness assessment.
B. Define a model-validation methodology.
C. Document maintenance teams and processes.
D. Identify known edge cases to monitor post-deployment.

Answer: A

Explanation:

After completing model testing and validation, the most important step prior to deploying the model into production is to perform a readiness assessment. This assessment ensures that the model is fully prepared for deployment, addressing any potential issues related to infrastructure, performance, security, and compliance. It verifies that the model meets all necessary criteria for a successful launch. Other steps, such as defining a model-validation methodology, documenting maintenance teams and processes, and identifying known edge cases, are also important but come secondary to confirming overall readiness. Reference: AIGP Body of Knowledge on Deployment Readiness.

Question: 100

Which type of existing assessment could best be leveraged to create an Al impact assessment?

A. A safety impact assessment.
B. A privacy impact assessment.
C. A security impact assessment.
D. An environmental impact assessment.

Answer: B

Explanation:

A privacy impact assessment (PIA) can be effectively leveraged to create an AI impact assessment. A PIA evaluates the potential privacy risks associated with the use of personal data and helps in implementing measures to mitigate those risks. Since AI systems often involve processing large amounts of personal data, the principles and methodologies of a PIA are highly applicable and can be extended to assess broader impacts, including ethical, social, and legal implications of AI. Reference: AIGP Body of Knowledge on Impact Assessments.

Question: 101

You are a privacy program manager at a large e-commerce company that uses an AI tool to deliver personalized product recommendations based on visitors' personal information that has been collected from the company website, the chatbot and public data the company has scraped from social media.
A user submits a data access request under an applicable U.S. state privacy law, specifically seeking a copy of their personal data, including information used to create their profile for productrecommendations.
What is the most challenging aspect of managing this request?

A. Some of the visitor's data is synthetic data that the company does not have to provide to the data subject.
B. The data subject's data is structured data that can be searched, compiled and reviewed only by an automated tool.
C. The data subject is not entitled to receive a copy of their data because some of it was scraped from public sources.
D. Some of the data subject's data is unstructured data and you cannot untangle it from the other data, including information about other individuals.

Answer: D

Explanation:

The most challenging aspect of managing a data access request in this scenario is dealing with unstructured data that cannot be easily disentangled from other data, including information about other individuals. Unstructured data, such as free-text inputs or social media posts, often lacks a clear structure and may be intermingled with data from multiple individuals, making it difficult to isolate the specific data related to the requester. This complexity poses significant challenges in complying with data access requests under privacy laws. Reference: AIGP Body of Knowledge on Data Subject Rights and Data Management.

Question: 102

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.

Question: 103

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.

Question: 104

Which of the following use cases would be best served by a non-AI solution?

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

Answer: A

Explanation:

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

Reference: AIGP Body of Knowledge on AI Use Cases and Applications.

Question: 105

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.

Question: 106

Which of the following steps occurs in the design phase of the AI life cycle?

A. Data augmentation.
B. Model explainability.
C. Risk impact estimation.
D. Performance evaluation.

Answer: C

Explanation:

Risk impact estimation occurs in the design phase of the AI life cycle. This step involves evaluatingpotential risks associated with the AI system and estimating their impacts to ensure that appropriate mitigation strategies are in place. It helps in identifying and addressing potential issues early in the design process, ensuring the development of a robust and reliable AI system. Reference: AIGP Body of Knowledge on AI Design and Risk Management.

Question: 107

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

Answer: B

Explanation:

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

on Ethical AI and Fairness.

Question: 108

CASE STUDY

Please use the following answer the next question:

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

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

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

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

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

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

Answer: D

Explanation:

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

Question: 109

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, whichincludes 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.

What is the best reason the police department should continue to perform investigations even if the AI system scores an individual's likelihood of criminal activity at or above 90%?

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

Answer: B

Explanation:

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

Question: 110

CASE STUDY

Please use the following answer the next question:

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

race and gender.

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

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

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

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

A. Explain to the accused how the Al system works.
B. Confirm the Al recommendation prior to sentencing.
C. Ensure an accused is given notice that the Al system was used.
D. Consider the Al 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.

Question: 111

CASE STUDY

Please use the following answer the next question:

A local police department in the United States procured an Al system to monitor and analyze social media feeds, online marketplaces and other sources of public information to detect evidence of illegal activities (e.g., sale of drugs or stolen goods). The Al system works by surveilling the public sites in order to identify individuals

that are likely to have committed a crime. It cross-references the individuals against data maintained by law enforcement and then assigns a percentage score of the likelihood of criminal activity based on certain factors like previous criminal history, location, time, race and gender.

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

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

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

Which Al 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.

Question: 112

All of the following types of testing can help evaluate the performance of a responsible Al system EXCEPT?

A. Risk probability/severity.
B. Adversarial robustness.
C. Statistical sampling.
D. Decision analysis.

Answer: A

Explanation:

Risk probability/severity testing is not typically used to evaluate the performance of an AI system. While important for risk management, it does not directly assess an AI system's operational performance. Adversarial robustness, statistical sampling, and decision analysis are all methods that can help evaluate the performance of a responsible AI system by testing its resilience, accuracy, and decision-making processes under various conditions. Reference: AIGP Body of Knowledge on AI Performance Evaluation and Testing.

Question: 113

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

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

Answer: C

Explanation:

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

Question: 114

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

A. Hallucination.
B. Brittleness.

C. Uncertainty.

D. Forgetting.

Answer: B

Explanation:

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

Question: 115

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.

Question: 116

What is the technique to remove the effects of improperly used data from an ML system?

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

Answer: D

Explanation:

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

Question: 117

What is the primary purpose of conducting ethical red-teaming on an Al 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.

Question: 118

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.

Question: 119

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 postdeployment activities.

Question: 120

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.

Question: 121

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.

Question: 122

You are an engineer that developed an AI-based ad recommendation tool. Which of the following should be monitored to evaluate the tool's effectiveness?

A. Output data, assess the delta between the prediction and actual ad clicks.
B. Algorithmic patterns, to show the model has a high degree of accuracy.
C. Input data, to ensure the ads are reaching the target audience.
D. GPU performance, to evaluate the tool's robustness.

Answer: A

Explanation:

To evaluate the effectiveness of an AI-based ad recommendation tool, the most relevant metric is the output data, specifically assessing the delta between the prediction and actual ad clicks. This metric directly measures

the tool's accuracy and effectiveness in making accurate recommendations that lead to user engagement. While monitoring algorithmic patterns and input data can provide insights into the model's behavior and targeting accuracy, and GPU performance can indicate the robustness and efficiency of the tool, the primary indicator of effectiveness for an ad recommendation tool is how well it predicts actual ad clicks.

Reference: AIGP BODY OF KNOWLEDGE, sections on AI performance metrics and evaluation methods.

Question: 123

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 to a 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.

Question: 124

A company plans on procuring a tool from an AI provider for its employees to use for certain business purposes.

Which contractual provision would best protect the company's intellectual property in the tool, including training and testing data?

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

Answer: B

Explanation:

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

Question: 125

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.

Question: 126

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.

Question: 127

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.

Question: 128

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.

Question: 129

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.

Question: 130

Pursuant to the White House Executive Order of November 2023, who is responsible for creating guidelines to conduct red-teaming tests of AI systems?

A. National Institute of Standards and Technology (NIST).

B. National Science and Technology Council (NSTC).

C. Office of Science and Technology Policy (OSTP).

D. Department of Homeland Security (DHS).

Answer: A

Explanation:

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

Reference: AIGP BODY OF KNOWLEDGE, sections on AI governance and regulatory frameworks, and the White House Executive Order of November 2023.

Question: 131

According to November 2023 White House Executive Order, which of the following best describes the guidance given to governmental agencies on the use of generative Al as a workplace tool?

A. Limit access to specific uses of generative Al.
B. Impose a general ban on the use of generative Al.
C. Limit access of generative Al to engineers and developers.
D. Impose a ban on the use of generative Al in agencies that protect national security.

Answer: A

Explanation:

The November 2023 White House Executive Order provides guidance that governmental agencies should limit access to specific uses of generative AI. This means that generative AI tools should be used in a controlled manner, where their applications are restricted to well-defined, approved use cases that ensure the security, privacy, and ethical considerations are adequately addressed. This approach allows for the benefits of generative AI to be harnessed while mitigating potential risks and abuses.

Reference: AIGP BODY OF KNOWLEDGE, sections on AI governance and risk management, and the White House Executive Order of November 2023.

Question: 132

The White House Executive Order from November 2023 requires companies that develop dual-use foundation models to provide reports to the federal government about all of the following EXCEPT?

A. Any current training or development of dual-use foundation models.
B. The results of red-team testing of each dual-use foundation model.
C. Any environmental impact study for each dual-use foundation model.
D. The physical and cybersecurity protection measures of their dual-use foundation models.

Answer: C

Explanation:

The White House Executive Order from November 2023 requires companies developing dual-use foundation models to report on their current training or development activities, the results of red- team testing, and the physical and cybersecurity protection measures. However, it does not mandate reports on environmental impact studies for each dual-use foundation model. While environmental considerations are important, they are not specified in this context as a reporting requirement under this Executive Order.

Reference: AIGP BODY OF KNOWLEDGE, sections on compliance and reporting requirements, and the White House Executive Order of November 2023.

Question: 133

Business A sells software that provides users with writing and grammar assistance. Business B is a cloud services provider that trains its own AI models.

* Business A has decided to add generative AI features to their software.
* Rather than create their own generative AI model, Business A has chosen to license a model from Business B.
* Business A will then integrate the model into their writing assistance software to provide generative AI capabilities.
* Business A is most concerned that its writing assistance software could recommend toxic or obscene text to its users.

Which of the following governance processes should Business A take to best protect its users against potentially inappropriate text?

A. Business A should fine-tune the AI model on user-generated text that has been verified to be appropriate.

B. Business A should test that the AI model performs as expected and meets their minimum requirements for filtering toxic or obscene text.

C. Business A should establish a user reporting feature that allows users to flag toxic or obscene text, and report any incidents to Business B.

D. Business A should ask Business B for detailed documentation on the generative AI model's training data and whether it contained toxic or obscene sources.

Answer: B

Explanation:

Business A is integrating a generative AI model licensed from a third party (Business B) and is primarily concerned with the risk of toxic or obscene outputs being delivered to users. In this scenario,testing and validationof the AI model for such content risks is the most direct and effective governance strategy.

According to theAI Governance in Practice Report 2024, organizations thatdeployAI must engage inperformance monitoring protocolsand ensure systems perform adequately for theirintended purposes, including filtering harmful content:

“Operational governance… development of: →Performance monitoring protocols to ensure systems perform adequately for their intended purposes.” (p. 12)

“Product governance... includes: →System impact assessments to identify and address risk prior to product development or deployment.” (p. 11)

Furthermore, under theEU AI Act, which sets the global standard many organizations aim to align with, there is a clear obligation to test and monitor systems for potential harmful behavior: “The act imposes regulatory obligations… such as establishing appropriate accountability structures,assessing system impact, providing technical documentation,establishing risk management protocols and monitoring performance...” (p. 7)

Option B directly reflects this best practice ofpre-deployment testing and validationto ensure that the model aligns with Business A’s minimum content safety requirements.

Let's now evaluate the incorrect options:

A . Fine-tuning on verified user-generated textmay improve model alignment but does not guarantee that the model will generalize correctly, especially if Business A lacks access to model internals (common in third-party licensing scenarios). Fine-tuning also introduces its own risks and may be contractually restricted.

C . A user reporting featureisreactive, not preventive. While helpful for long-term monitoring and mitigation, it does not prevent the initial harm of toxic outputs, which isBusiness A's primary concern.

E. Requesting documentation from Business Bis useful for transparency and risk management, but it does not replaceindependent verificationthat the model meets Business A's content safety standards.

Thus,testing the model's behavior for unacceptable outputs before deploymentis the most aligned approach with AI governance best practices and obligations.

Question: 134

All of the following are examples of biometric data in the US EXCEPT?

A. Iris scans.
B. Walking gait.
C. Keystroke dynamics.
D. GPS location of a user's fitness watch.

Answer: D

Explanation:

Biometric data in the U.S. refers to data that relates to measurable biological and behavioral characteristics that can be used to identify an individual. Examples include fingerprints, facial recognition, iris scans, and behavior-based data like gait or keystrokes.

According to definitions and discussions from theAI Governance in Practice Report 2024and U.S. privacy frameworks:

"Biometric data includes physical and behavioral human characteristics that can be used to digitally identify a person to grant access to systems, devices, or data. Examples include facial images, iris patterns, gait analysis, and voice recognition." (Report context based on common frameworks in U.S. AI law and the use of biometrics in AI governance.)

Here's how the options relate:

A . Iris scans– These are physical biometric identifiers.

B . Walking gait– Behavioral biometric used increasingly in surveillance and identification.

C . Keystroke dynamics– Behavioral biometric based on typing patterns.

D . GPS location of a user's fitness watch– This isnotbiometric data. It islocation data, which may be sensitive or personal, but not biometric.

Question: 135

Which stakeholder is responsible for lawful collection of data for the training of the foundational AI model?

A. The marketing agency.
B. The tech company.
C. The data aggregator.

D. The marketing agency's client.

Answer: C

Explanation:

Data aggregators are third parties that collect and license data from various sources. They are responsible for ensuring thelawful collectionandproper usage rightsof the data they distribute — especially when such data is used to train foundational AI models.

From theAI Governance in Practice Report 2024:

"As organizations have neither proximity to how third-party data was first collected nor direct control over the data governance practices of third parties, an organization can benefit from carrying out its own legal due diligence and third-party risk management." (p. 19)

"Legal due diligence may include verification of the personal data's lawful collection by the databroker..." (p. 19)

This confirms thatdata aggregatorsbear the legal and ethical burden to verify that data has been lawfully collected and is appropriately licensed for use, including in AI training.

A . The marketing agencyandD. its clientmay use data, but they rely on upstream providers for its lawful origin.

B . The tech companymay train the model but depends on lawful sourcing by data aggregators.

Question: 136

All of the following are required for high-risk AI systems under the EU AI Act EXCEPT?

A. Retaining system-generated logs for at least six months.
B. Conducting post-market monitoring.
C. Conducting a conformity assessment.
D. Publishing a detailed report on the training data used.

Answer: D

Explanation:

The EU AI Act imposes severalmandatory obligationson high-risk AI systems, butpublishing a detailed report on training dataisnotone of them.

From theAI Governance in Practice Report 2024:

"It mandates drawing up technical documentation for high-risk AI systems, and requires high-risk AI systems to come with instructions for use that disclose various information, including characteristics, capabilities and performance limitations." (p. 34)

"To make high-risk AI systems more traceable, it also requires AI systems to be able to automatically allow for the maintenance of logs throughout the AI life cycle."

"Conducting post-market monitoring" and "conformity assessments" areexplicitrequirements for high-risk systems. (p. 34–35)

However,publishingdetailed training data is typicallyrequired only for general-purpose AI systems with systemic risk, not standard high-risk AI systems.

A . Log retention,B. Post-market monitoring, andC. Conformity assessmentsareall requiredunder the EU AI

Act for high-risk systems.

Question: 137

A UK company has designed a facial recognition model to support border control.
The EU AI Act would apply to the model in all of the following situations EXCEPT if?

A. The model was released under an open source license.
B. The model is deployed at an EU border checkpoint.
C. The model is deployed at UK border checkpoints.
D. The model was trained by an EU company.

Answer: C

Explanation:

TheEU AI Actapplies extraterritorially, meaning it affects entitiesoutside the EUwhen their AI systemsimpact individuals within the EU. However, it doesnotapply to systems that are developed, sold, or usedentirely outside of the EU—such as in theUK, unless they affect the EU market or individuals.
From theAI Governance in Practice Report 2024:
"The act imposes regulatory obligations… depending on their capabilities, reach and computing power, certain GPAI systems are considered to present systemic risk and attract broadly similar obligations to those applicable to high-risk AI systems." (p. 7)
"The EU AI Act is the world's first comprehensive AI regulation… requirements apply to providers, deployers, importers, and distributors of AI systems when such systems are placed on the EU market." (p. 7–8)
Thus:
A . Open source releasedoes not exclude applicability if deployed in the EU.
B . Deployment at an EU borderclearly invokes jurisdiction.
D . Training by an EU companycreates jurisdictional links.
C . Deployment only at UK checkpoints, withno EU use or impact, isoutside scope.

Question: 138

A shipping service based in the US is looking to expand its operations into the EU. It utilizes an inhouse developed multimodal AI model that analyzes all personal data collected from shipping senders and recipients, and optimizes shipping routes and schedules based on this data.
As they expand into the EU, all of the following descriptions should be included in the technical documentation for their AI model EXCEPT?

A. A general description of the AI system.
B. A description of the prioritization of the risks of deployment of the AI system.
C. A description of the appropriateness of the performance metrics for the specific AI system.
D. A detailed description of the elements of the AI system and of the process for its development.

Answer: B

Explanation:

The EU AI Act outlines what must be included intechnical documentationfor high-risk systems. These requirements are designed to supportconformity assessment, transparency, and traceability.
From theAI Governance in Practice Report 2024:
"It mandates drawing up technical documentation… must include a general description of the AI system, the intended purpose, and a detailed description of the elements and development process." (p. 34)
"Documentation… includes training, testing, evaluation procedures, andappropriateness of performance metrics." (p. 34–35)
Therisk management systemis addressed separately through arisk management plan, not within the technical documentation itself.
Thus:
A, C, and Dare explicitly required in thetechnical documentation.
B, while important, is part of therisk management process, not a required section oftechnical documentation.

Question: 139

Which of the following is an obligation of an importer of high-risk AI systems under the EU AI Act?

A. Provide technical documentation.
B. Affix the CE marking.
C. Verify the Declaration of Conformity.
D. Conduct a data protection impact assessment.

Answer: C

Explanation:

Importers of high-risk AI systems into the EU havespecific responsibilitiesunder the EU AI Act. They arenotthe parties responsible for affixing the CE marking or providing technical documentation —but they must verify that these have been done by the provider.
From theAI Governance in Practice Report 2024:
"Importers must verify that the appropriate conformity assessment has been carried out, the technical documentation is available, and the CE marking has been affixed." (p. 34–35)
Thus:
A . Provide technical documentation– done by theprovider.
B . Affix the CE marking–provider'sresponsibility.
C . Verify the Declaration of Conformity–importer obligation.
D . Conduct a DPIA– relevant under data protection laws,not requiredunder the EU AI Act forimporters.

Question: 140

A US hospital plans to develop an AI that will review available patient data in order to propose an initial diagnosis to licensed physicians. The hospital will implement a policy that requires physicians to consider the AI proposal, but conduct their own physical examinations prior to making a final diagnosis.
An important ethical concern with this plan is?

A. Whether patients will receive an economic benefit from the use of AI.
B. Whether the AI was trained on a representative dataset.
C. Whether physicians understand how the AI works.
D. Whether the AI will have an error rate comparable to human physicians.

Answer: B

Explanation:

The core ethical concern when deploying diagnostic AI in a healthcare setting is ensuringfairness and accuracy across diverse patient populations. If the AI is trained on a dataset that isnot representativeof the population it will serve, it risks reinforcing health disparities and leading to misdiagnoses.

From theAI Governance in Practice Report 2024:

"Training datasets lacking in diversity can produce outputs that systematically underperform for certain groups… this can lead to inaccurate or biased outcomes in healthcare settings." (p. 41) "Bias, discrimination and fairness challenge… inadequate or nonrepresentative training data can result in AI systems that propagate historical disparities." (p. 42)

While physician oversight may reduce risk,biased data can still shape clinical decision-making.

A– Economic benefit is not central to ethical risk here.

C– Important but less critical than data representativeness.

D– Error rate matters but is addressed via validation; it's not the core ethical issue.

Question: 141

What is the most important reason to document the results of AI testing?

A. To support post-deployment maintenance.
B. To identify areas for red-teaming focus.
C. To create a verifiable audit trail.
D. To limit the need for future testing cycles.

Answer: C

Explanation:

Testing results need to bedocumented thoroughlyto ensuretraceability, accountability, and compliance. This is central to enabling audits, investigations, or regulatory inquiries into the system's development and performance.

From theAI Governance in Practice Report 2024:

"Documentation and recordkeeping are essential components… to demonstrate AI system compliance, trace system behavior, and support audits and conformity assessments." (p. 34–35) "Maintaining audit trails across development and deployment enables transparency and accountability." (p. 12)

AandBare benefits, but not theprimary governance justification.

D– Limiting future testing is not a recommended goal.

Question: 142

MULTI-SELECT

Please select 3 of the 5 options below. No partial credit will be given.

What are the roles and responsibilities of deployers of a proprietary model?

A. Ethical testing.
B. Ethical design.
C. Technical performance.
D. System documentation.
E. Regulatory compliance.

Answer: A,C,E

Explanation:

Deployers of proprietary models arenot responsible for design, but they are accountable for how the system performsin their context of use, including ensuring ethical behavior, performance, and legal compliance.

From theAI Governance in Practice Report 2024:

"Deployers of AI systems must take reasonable steps to ensure that systems are used ethically, perform safely, and align with applicable laws and standards." (p. 11–12)

"Operational governance… includes performance monitoring protocols, incident management plans, and regulatory oversight." (p. 12)

Thus:

Q A. Ethical testing– Required to mitigate misuse and unintended harms.

X B. Ethical design– Belongs todevelopers/providers, not deployers.

Q C. Technical performance– Deployers must ensure that AI performs as expected.

X D. System documentation– This is theprovider'sobligation.

Q E. Regulatory compliance– Deployers must ensure system use complies with applicable laws.

Question: 143

An AI system's function, the industry and the location in which it operates are important factors in considering which of the following?

A. Organizational accountability.
B. Internal governance needs.
C. Diversity of data sources.
D. Explainability of results.

Answer: B

Explanation:

An AI system'sfunction,industry, anddeployment locationdefine itsrisk profile, which directly influences theinternal governance structuresan organization must put in place.
From theAI Governance in Practice Report 2024:
"There are many challenges and potential solutions for AI governance, each with unique proximityand significance based on an organization's role, footprint, broader risk-governance profile and maturity." (p. 4)
"AI governance starts with defining the corporate strategy for AI… and formulating policy standards and operational procedures to reflect industry, use case, and location." (p. 11)
A– Organizational accountability is broader and not directly scoped by industry or function.
C– Diversity of data sources is tied to data strategy.
D– Explainability is more influenced by model type, not use context.

Question: 144

What is the most important reason for documenting risks when developing an AI system?

A. To provide transparency to stakeholders.
B. To align with industry standards.
C. To promote knowledge sharing.
D. To mitigate potential liability.

Answer: D

Explanation:

The most critical reason for documenting AI-related risks is toreduce exposure to legal, regulatory, and reputational liabilities. Clear documentation demonstrates thatrisks were identified, assessed, and addressed, which is essential for accountability and defensibility in the face of audits, litigation, or enforcement actions.
From theAI Governance in Practice Report 2024:
"An effective AI governance model is about collective responsibility… which should encompass oversight mechanisms such as privacy, accountability, compliance." (p. 13)
"Accountability… is based on the idea that there should be a person or entity that is ultimately responsible for any harm resulting from the use of the data, algorithm and AI system's underlying processes." (p. 28)
While transparency, alignment with standards, and knowledge sharing are all secondary benefits,risk documentation's primary role is liability mitigation.

Question: 145

The best practice to manage third-party risk associated with AI systems is to create and implement policies that?

A. Focus on the financial stability of third-party vendors as the primary criterion for risk assessment.
B. Provide for an appropriate level of due diligence and ongoing monitoring based on the defined risk.
C. Require third-party AI systems to undergo a comprehensive audit by an external cybersecurity firm every six months.

D. Focus on the technical aspects of AI systems, such as data security, while ethical risks are addressed through suitable contracts.

Answer: B

Explanation:

Third-party risk management for AI systems should beproportional and risk-based, involvinginitial due diligenceandongoing monitoringthat reflects thelevel of risk posedby the third party's AI system. From theAI Governance in Practice Report 2024:

“Third-party due diligence assessments to identify possible external risk and inform selection.” (p. 11)

“Legal due diligence may include verification of the personal data's lawful collection by the data broker, review of contractual obligations…” (p. 19)

Afocuses too narrowly on financial stability.

Cis excessive and not scalable or aligned with best practices.

Dinappropriately separates ethical and technical risks; both must be evaluated holistically.

Question: 146

MULTI-SELECT

Please select 3 of the 5 options below. No partial credit will be given.

Training an AI model is time-consuming because of?

A. The complexity of the AI model.
B. The maturity of AI governance.
C. The volume of training data.
D. The number of stakeholders.
E. The quality of the training data.

Answer: A,C,E

Explanation:

Training an AI model is time-consuming primarily due tomodel complexity,large data volumes, and the need forhigh-quality, well-prepared data.

From theAI Governance in Practice Report 2024:

“Most AI requires sizeable amounts of high-quality data… to ensure desired and accurate output.” (p. 15)

“The accuracy of AI model outputs depends significantly on the quality of their inputs.” (p. 24) “Complex AI systems… with many parameters… result in long development and training phases.” (p. 32)

B . Maturity of governanceaffects oversight, not training time.

D . Number of stakeholdersaffects alignment, not direct training duration.

Question: 147

Retrieval-Augmented Generation (RAG) is defined as?

A. Combining LLMs with private knowledge bases to improve their outputs.
B. Reducing computational processing requirements of the LLMs.
C. Applying advanced filtering techniques to the LLMs.
D. Fine tuning LLMs to minimize biased outputs.

Answer: A

Explanation:

Retrieval-Augmented Generation (RAG)enhances Large Language Models (LLMs) by integratingexternal, up-to-date, or proprietary informationinto the generation pipeline—allowing
the model tofetch relevant factsfrom a trusted knowledge source at query time.
Though RAG is not defined directly in the IAPP documents, it is a widely recognized technique in AI governance for ensuringmore accurate and contextually grounded outputs, especially inregulated or high-stakes environmentswhere hallucinations are a concern.
B, C, and Ddescribe optimization or bias mitigation—not the core function of RAG.

Question: 148

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

A. Continue disparity testing.
B. Provide regular awareness training.
C. Analyze the quality of the training and testing data.
D. Document the results of final decisions made by the human underwriter.

Answer: A

Explanation:

Theinitial deployment phaseof an AI model is critical forpost-deployment monitoring. When tracking forbias, the most important task is tocontinue disparity testingto determine whether outputs differ across protected groups.
From theAI Governance in Practice Report 2024:
"Performance monitoring protocols… should include mechanisms to assess and measure disparities in outcomes across different demographic groups." (p. 12)
"Bias may not be evident during pre-deployment testing but can emerge in real-world use." (p. 41)
B . Awareness trainingis helpful, but not a technical bias mitigation activity.
C . Analyzing training datais apre-deploymenttask.
E. Documenting human decisionsmay support auditability but doesn't detect bias in AI outputs.

Question: 149

A US-based mortgage lender has purchased a chatbot. They plan to have the chatbot collect information from consumers who are interested in loans and offer the consumers 2-3 different options based on its current pricing and product offerings, which change frequently. This chatbot was initially developed and previously deployed by a Russian airline for booking flights.
The best option for the part of the process that generates the loan offers is?

A. Retrieval-Augmented Generation.
B. Multimodal Generative AI.
C. Expert System.
D. Quantum computing

Answer: C

Explanation:

Offeringloan products based on current offerings and rulesrequires a system that can followexplicit business logic, not generate open-ended content. Anexpert system, which is a rules-based AI that uses “if-then” logic, is ideal here.
From the AI governance context:
“Rule-based AI systems are often preferred when decisions must adhere to precise regulatory or financial criteria.” (aligned with AI best practices in regulated sectors)
A . RAGis used to integrate external knowledge—not suitable for structured, rule-based logic.
B . Multimodal modelshandle varied input types—not needed here.
D . Quantum computingis not yet practical or relevant for this business use case.

Question: 150

Your organization is searching for a new way to help accurately forecast sales predictions by various types of customers.
Which of the following is the best type of model to choose if your organization wants to customize the model and avoid lock-in?

A. A free large language model.
B. A classic machine learning model.
C. A proprietary generative AI model.
D. A subscription-based, multimodal model.

Answer: B

Explanation:

Forcustomizable, interpretable modelsthat allow organizations toretain control and avoid vendor lock-in,classic ML models(e.g., regression, decision trees, random forests) are optimal.
From theAI Governance in Practice Report 2024:

"Organizations seeking transparency, customizability, and control often prefer classic ML models due to their flexibility and ease of governance." (p. 33)
AandCmay have limited transparency and are often tied to specific providers.
Dinvolves ongoing costs and limited model control.

Question: 151

In procuring an AI system from a vendor, which of the following would be important to include in a contract to enable proper oversight and auditing of the system?

A. Liability for mistakes.
B. Ownership of data and outputs.
C. Responsibility for improvements.
D. Appropriate access to data and models.

Answer: D

Explanation:

Ensuringoversight and auditabilityrequires that the organization hassufficient access to data, documentation, and model internalsor outputs necessary for evaluation.
From theAI Governance in Practice Report 2024:
"Access to technical documentation and system internals is essential to enable effective auditing, conformity checks, and accountability mechanisms." (p. 11, 34)
Ais about liability, not auditability.
Bmatters for IP rights, not oversight.
Crelates to lifecycle responsibility but doesn't guarantee audit access.

Question: 152

Why is it important that conformity requirements are satisfied before an AI system is released into production?

A. To ensure the visual design is fit-for-purpose.
B. To ensure the AI system is easy for end-users to operate.
C. To guarantee interoperability of the AI system across multiple platforms and environments.
D. To comply with legal and regulatory standards, ensuring the AI system is safe and trustworthy.

Answer: D

Explanation:

Conformity assessmentsare a core requirement under theEU AI Actfor high-risk systems and serve to confirm that the AI meetsregulatory, safety, and ethical standardsbefore it is put into production. From theAI Governance in Practice Report 2024:
"Conformity assessments… ensure that systems comply with legal requirements, safety criteria, and intended purpose before being placed on the market." (p. 34)

"They are a critical step to demonstrate safety and trustworthiness in AI deployment." (p. 35)

Question: 153

The best method to ensure a comprehensive identification of risks for a new AI model is?

A. An environmental scan.
B. Red teaming.
C. Integration testing.
D. An impact assessment.

Answer: D

Explanation:

The most comprehensive way to identify a full range of risks — legal, ethical, operational, and societal — for a new AI model is through aformal impact assessment, such as aData Protection Impact Assessment (DPIA)orAlgorithmic Impact Assessment.
From theAI Governance in Practice Report 2024:
"Risk-based approaches are often distilled into organizational risk management efforts, which put impact assessments at the heart of deciding whether harm can be reduced." (p. 29) "DPIAs… help organizations identify, analyze and minimize data-related risks and demonstrate accountability." (p. 30)
A . Environmental scanis too general.
B . Red teamingis useful for adversarial risk but not broad.
C . Integration testingfocuses on technical/system compatibility, not overall risk.

Question: 154

All of the following are potential benefits of using private over public LLMs EXCEPT?

A. Reduction in time taken for data validation and verification.
B. Confirmation of security and confidentiality.
C. Reduction in possibility of hallucinated information.
D. Application for specific use cases within the enterprise.

Answer: A

Explanation:

Private LLMs offer advantages likecustomizability,reduced hallucination,confidentiality, andalignment with enterprise-specific tasks, but theydo not inherently reduce the time or effortneeded fordata validation or verification— which remains an essential step regardless of model privacy.
From the AI risk and quality sections:
"Ensuring the quality of the data… is highly contextual and must be validated regardless of the model's deployment environment." (p. 17)
B. C, Dare legitimate benefits of private LLMs.
Ais incorrect — validation still requires time and resources.

Question: 155

What is the most important factor when deciding whether or not to select a proprietary AI model?

A. What business purpose it will serve.
B. How frequently it will be updated.
C. Whether its training data is disclosed.
D. Whether its system card identifies risks.

Answer: A

Explanation:

Theprimary considerationin selecting any AI system, especially aproprietary model, is itsfit for business purpose. Whether it serves the intended goals is foundational before evaluating technical or governance features.

From theAI Governance in Practice Report 2024:

"AI governance starts with defining the corporate strategy for AI… and aligning systems with business purpose and operational context." (p. 11)

B, C, Dare relevant for evaluation, but onlyafterconfirming business applicability.

Question: 156

Which model is best for efficiency and agility, and tailored for lower-resource settings?

A. Supervised learning model.
B. Multimodal model.
C. Small language model.
D. Generative language model.

Answer: C

Explanation:

Small language models (SLMs)arelightweight, requireless compute, and arebetter suited to low- resource or edge environments, making them ideal for agility and efficiency.

From general AI best practices:

"SLMs can be deployed in environments with limited computing power, ensuring lower cost and faster integration in constrained contexts." (aligned with industry-wide AI deployment strategies)

Question: 157

A leading software development company wants to integrate AI-powered chatbots into their customer service platform. After researching various AI models in the market which have been developed by third-party developers, they're considering two options:

Option A - an open-source language model trained on a vast corpus of text data and capable of being trained to respond to natural language inputs.

Option B - a proprietary, generative AI model pre-trained on large data sets, which uses transformerbased

architectures to generate human-like responses based on multimodal user input.

Option A would be the best choice for the company because?

A. It is less expensive to run

B. It may be better suited for applications requiring customization.

C. It can handle voice commands and is more suitable for phone-based customer support.

D. It is built for large-scale, complex dialogues and would be more effective in handling high-volume customer inquiries.

Answer: B

Explanation:

Open-source modelsoffer morecustomization flexibility, allowing organizations to fine-tune or adapt the model tofit their own workflows, branding, or compliance needs— making it preferable when deep control is needed.

From theAI Governance in Practice Report 2024:

"Open-source AI allows organizations to review, adapt, and control model behavior in line with organizational needs and policies." (p. 39)

Question: 158

A company developing and deploying its own AI model would perform all of the following steps to monitor and evaluate the model's performance EXCEPT?

A. Publicly disclosing data with forecasts of secondary and downstream harms to stakeholders.

B. Setting up automated tools to regularly track the model's accuracy, precision and recall rates in real-time.

C. Implementing a formal incident response plan to address incidents that may occur during system operation.

D. Establishing a regular schedule for human evaluation of the model's performance, including qualitative assessments.

Answer: A

Explanation:

While transparency is encouraged,publicly disclosing forecasts of secondary harmsisnot a required or standard practicefor internal performance evaluation. Risk assessments and reporting typically remaininternal or shared with regulators.

From theAI Governance in Practice Report 2024:

"Organizations must assess secondary risks… but disclosure is subject to context, regulatory requirements, and risk management discretion." (p. 30)

Question: 159

A company that deploys AI but is not currently a provider or developer intends to develop and market its own AI system.
Which obligation would then be likely to apply?

A. Implementing a risk management framework.
B. Conducting an impact assessment including a post-deployment monitoring plan.
C. Developing documentation on the system, the potential risks and the safeguards applied.
D. Developing a reporting plan for any observed algorithmic discrimination or harms to individuals' rights and freedoms.

Answer: C

Explanation:

Once a company moves from being adeployerto also acting as aprovider or developer, it assumesnew obligationsunder regulations like the EU AI Act. One of the core requirements for providers is to produce and maintaintechnical documentation, including descriptions of the model, associated risks, and mitigation strategies.
From theAI Governance in Practice Report 2024:
"Providers of high-risk AI systems must draw up technical documentation demonstrating the system's conformity with the requirements... including potential risks and safeguards applied." (p. 34)
"This documentation must be available before placing the system on the market." (p. 35)

Question: 160

All of the following issues are unique for proprietary AI model deployments EXCEPT?

A. The acquisition of training data.
B. The cost of AI chips.
C. The potential for bias.
D. The necessity of performing conformity assessments.

Answer: C

Explanation:
Biasis a common risk acrossboth proprietary and open-source models, andnot uniqueto proprietary deployments. All AI systems — regardless of origin — require evaluation for fairness, accuracy, and representativeness.
From theAI Governance in Practice Report 2024:
"Bias, discrimination and fairness challenges are present in both open and closed models, regardless of how the model is sourced." (p. 41)

Question: 161

A deployer discovers that a high-risk AI recruiting system has been making widespread errors, resulting in harms to the rights of a considerable number of EU residents who are denied consideration for jobs for improper reasons such as ethnicity, gender and age.
According to the EU AI Act, what should the company do first?

A. Notify the provider, the distributor, and finally the relevant market authority of the serious incident.
B. Identify any decisions that may have been improperly made and re-open them for human review.
C. Submit an incomplete report to the relevant market authority immediately and follow up with a complete report as soon as possible.
D. Conduct a thorough investigation of the serious incident within the 15 day timeline and present the completed report to the relevant market authority.

Answer: A

Explanation:

Under theEU AI Act, serious incidents involvinghigh-risk AI systemsmust be reported. The deployer is required topromptly inform the provider and relevant authoritiesabout the issue.
From theAI Governance in Practice Report 2024:
"Serious incidents involving high-risk systems… must be reported to the provider and relevant market surveillance authority." (p. 35)
"Timely reporting is required when AI systems result in or may result in violations of fundamental rights." (p. 35)

Question: 162

What is the most significant risk of deploying an AI model that can create realistic images and videos?

A. Copyright infringement.
B. Security breaches.
C. Downstream harms.
D. Output cannot be protected.

Answer: C

Explanation:

The greatest risk from AI systems generatingrealistic synthetic mediaisdownstream harm, such asdeepfakes, misinformation, reputational damage, and erosion of trust.
From theAI Governance in Practice Report 2024:
"With generative AI, downstream harms such as deception, reputational damage, misinformation, and manipulation can emerge even if original use was lawful." (p. 55–56)

Question: 163

After initially deploying a third-party AI model, you learn the developer has released a new version. As deployer of this third-party model, what should you do?

A. Audit the model.
B. Retrain the model.
C. Seek input from data scientists.
D. Communicate necessary updates to your users.

Answer: A

Explanation:

When anew versionof a third-party model is released, the deployer must ensure it still meets safety, performance, and compliance requirements — which calls for aformal audit.

From theAI Governance in Practice Report 2024:

"Any updates or changes to AI systems should trigger a re-evaluation to ensure continued compliance and performance." (p. 12)

"Post-market monitoring includes reassessing the impact of updated models or retraining." (p. 35)

Question: 164

A company deploys an AI model for fraud detection in online transactions. During its operation, the model begins to exhibit high rates of false positives, flagging legitimate transactions as fraudulent.

Which is the best step the company should take to address this development?

A. Dedicate more resources to monitor the model.
B. Maintain records of all false positives.
C. Deactivate the model until an assessment is made.
D. Conduct training for customer service teams to handle flagged transactions.

Answer: C

Explanation:

When an AI system causessignificant false positives, especially in sensitive contexts likefraud detection, the priority is tohalt harmful activityand perform a full assessment. Continued use without understanding the fault may cause furthercustomer harmand legal exposure.

From theAI Governance in Practice Report 2024:

"Incident management plans should enable identification, escalation, and system rollback to prevent continued harm from malfunctioning AI systems." (p. 12, 35)