PRACTICE EXAM 8: AIGP SIMULATION (100 QUESTIONS)

1. Under the EU AI Act, an organization places a high-risk AI system on the EU market. The provider has completed the conformity assessment. Six months later, the deployer significantly modifies the system's intended purpose without the provider's involvement. What regulatory consequence does this trigger?

A. The deployer must notify the provider of the modification so the provider can update the system's conformity declaration

B. The modification is permissible as long as the deployer documents the change in the system's monitoring records

C. The deployer may be reclassified as a provider for the modified system, assuming the corresponding provider obligations including conformity assessment for the new intended purpose

D. The modification automatically voids the system's CE marking but does not change the deployer's regulatory status

2. An AI system processes biometric data to verify identity for financial transactions. The system stores facial templates but not raw images. Under GDPR, is the processing of facial templates classified differently than raw facial images?

A. Yes — facial templates are not personal data because they are mathematical representations rather than photographs

B. Yes — facial templates require consent while raw images can be processed under legitimate interests

C. No — both are personal data, but facial templates are classified as general personal data rather than special category data

D. No — facial templates derived from facial images constitute biometric data processed for identification purposes, making them special category data under GDPR Article 9 regardless of whether raw images are stored

3. An organization trains an AI model using a technique called "knowledge distillation" — training a smaller "student" model to replicate the behavior of a larger "teacher" model. The teacher model was trained on data containing personal information. The student model never directly accesses the training data. Does the student model require data governance?

A. No, because the student model only learns from the teacher model's outputs and never processes personal data directly

B. Yes — the student model may inherit behaviors, biases, and potentially memorized patterns from the teacher model that originated in the personal data, requiring governance evaluation even though the student never directly accessed the underlying data

C. Only if the student model is larger than 1 billion parameters, which is the threshold for GPAI model governance under the EU AI Act

D. Only if the teacher and student models are developed by different organizations, creating a cross-entity data governance concern

4. A governance professional must distinguish between "data drift" and "concept drift" for monitoring purposes. An AI fraud detection system is experiencing declining accuracy. Which scenario describes CONCEPT drift rather than data drift?

A. Fraudsters have developed entirely new fraud techniques that did not exist when the model was trained — the relationship between transaction features and actual fraud has changed because the definition of what constitutes fraud in practice has evolved

B. The distribution of transaction amounts has shifted because customers are making larger purchases than during the training period

C. The ratio of domestic to international transactions has changed since the model was deployed

D. The demographic composition of the customer base has changed significantly relative to the training data distribution

5. Under the EU AI Act, which of the following is classified as a PROHIBITED AI practice rather than a high-risk AI system?

A. An AI system used for biometric identification in public spaces by law enforcement under specific conditions

B. An AI system used by employers to evaluate job candidates during recruitment interviews

C. An AI system that evaluates creditworthiness for consumer lending decisions

D. An AI system deployed by a government that evaluates citizens based on social behavior to determine access to public services in unrelated contexts

6. An organization operates both a GPAI model provider and a deployer relationship with the same vendor. Under the EU AI Act, what distinguishes the provider's obligations for GPAI models from the deployer's obligations for high-risk AI systems?

A. GPAI model providers and high-risk system deployers have identical obligations because the EU AI Act does not distinguish between these roles

B. GPAI providers have more extensive obligations than high-risk deployers because GPAI models pose greater societal risk

C. GPAI model providers must provide technical documentation and training data summaries enabling downstream providers to comply, while high-risk deployers must use the system according to instructions, maintain human oversight, and monitor for risks in their specific deployment context

D. GPAI providers have no obligations under the EU AI Act because the Act only regulates deployed AI systems, not foundational models

7. An AI system's monitoring detects that performance has degraded. The governance team must determine the root cause. The monitoring data shows: overall accuracy declined by 4%, false positive rate increased by 12%, and the increase is concentrated entirely in one customer segment that grew from 5% to 25% of the population since deployment. This pattern MOST strongly suggests which type of issue?

A. Model decay caused by computational resource degradation affecting the system's inference speed and output quality

B. Data drift — the population composition has shifted significantly (one segment grew from 5% to 25%), and the model performs differently for this segment because it was underrepresented in training data

C. Concept drift — the relationship between input features and the target variable has fundamentally changed for the growing customer segment

D. An adversarial attack where the growing customer segment has been deliberately submitting inputs designed to degrade the model's performance

8. Under GDPR Article 22, an individual has the right not to be subject to a decision based solely on automated processing that produces legal effects or similarly significant effects. A bank uses AI to pre-screen loan applications, automatically rejecting applications below a threshold score. A human reviewer then examines the remaining applications. For the automatically rejected applicants, does Article 22 apply?

A. Yes — the automatic rejection based solely on AI scoring produces significant effects (denial of access to credit) without meaningful human involvement for those rejected applicants

B. No — because a human reviewer is involved in the overall process, the system does not constitute "solely automated" processing

C. Yes, but only if more than 50% of applicants are automatically rejected, because a majority-automated process triggers Article 22 regardless of human involvement

D. No — Article 22 only applies to final decisions, and the pre-screening is an intermediate step in the lending process

9. An organization is comparing ISO/IEC 42001 and the NIST AI RMF to determine which to implement. Which statement MOST accurately describes the relationship between these two frameworks?

A. ISO 42001 and the NIST AI RMF are competing standards that cannot be implemented simultaneously because they impose contradictory requirements

B. The NIST AI RMF is a legal requirement for U.S. organizations while ISO 42001 is a legal requirement for EU organizations

C. ISO 42001 provides a certifiable management system framework (the organizational structure for governing AI), while the NIST AI RMF provides a risk management methodology (the process for identifying and managing AI risks) — they are complementary and can be implemented together

D. ISO 42001 is exclusively for AI developers while the NIST AI RMF is exclusively for AI deployers

10. An AI system uses federated learning across five hospitals. Hospital A discovers a data quality issue in its local training data. Under federated learning governance, who is responsible for addressing this data quality issue?

A. The federated learning system administrator is solely responsible because data quality in federated systems is a system-level concern

B. All five hospitals share equal responsibility because federated learning distributes both the training process and governance obligations equally

C. The AI vendor that provides the federated learning platform is responsible because the vendor's software should detect and correct data quality issues automatically

D. Hospital A is primarily responsible for the quality of its own local data, but the federated learning governance framework must include data quality standards that all participants must meet, and the impact of Hospital A's data quality issue on the global model must be assessed

11. An AI system for HR analytics processes employee data to identify "high-potential" employees for leadership development programs. The system considers performance ratings, project completion rates, peer feedback scores, and years of experience. A governance review reveals that "years of experience" correlates strongly with age. What is the MOST precise governance characterization of this feature?

A. "Years of experience" is a direct proxy for age and must be removed from all employment AI models without exception

B. "Years of experience" may function as a proxy for age — if it drives the model to systematically favor older employees for leadership programs (or disfavor younger employees), it could create age-based disparate impact that requires evaluation of whether the feature's predictive value justifies its discriminatory potential

C. "Years of experience" is a legitimate, non-discriminatory employment criterion that requires no governance evaluation regardless of its correlation with age

D. "Years of experience" only creates a governance concern if it is the single most important feature in the model and has no governance implications if it is one of many features

12. Under the EU AI Act, which entity bears primary responsibility for conducting the conformity assessment of a high-risk AI system before it is placed on the market?

A. The provider — the entity that develops or commissions the development of the AI system and places it on the market under its own name or trademark

B. The deployer — the entity that uses the AI system under its authority in a professional capacity

C. The importer — the entity that places a non-EU provider's system on the EU market

D. The notified body — the independent third-party organization designated to conduct conformity assessments

13. An AI governance professional is evaluating whether an AI system constitutes "solely automated decision-making" under GDPR Article 22. The system produces a recommendation that a human reviewer approves in 99.7% of cases, spending an average of 8 seconds per review. Does this constitute meaningful human involvement?

A. Yes — the presence of a human in the process means the decision is not solely automated, regardless of the approval rate or review duration

B. Yes — the 0.3% override rate demonstrates that the human reviewer exercises genuine judgment even if most recommendations are approved

C. The approval rate is irrelevant; what matters is whether the human has the authority, information, and training to override the AI

D. A 99.7% approval rate with 8-second reviews strongly suggests rubber-stamping rather than meaningful human involvement — the human reviewer likely lacks the time, information, or practical ability to exercise genuine independent judgment, potentially making the process "solely automated" despite nominal human participation

14. An AI system is trained using synthetic data exclusively — no real-world data was used in training. Under GDPR, does this system require any data protection governance?

A. No — if no real personal data was used in training, GDPR has no applicability to the AI system regardless of how it is deployed

B. No — synthetic data is by definition not personal data and cannot trigger any data protection obligation

C. Not for the training phase, but if the deployed system processes real personal data during inference (making predictions about real people), GDPR applies to the deployment processing regardless of what data was used for training

D. Yes — all AI systems require GDPR compliance regardless of their training data because the AI Act mandates GDPR compliance for all AI systems

15. An organization deploys an AI chatbot that generates text responses. Under the EU AI Act's transparency provisions for AI systems that interact with natural persons, what is the organization's PRIMARY transparency obligation?

A. Publish the chatbot's complete training data sources and methodology on the organization's website

B. Ensure that individuals interacting with the chatbot are informed that they are interacting with an AI system, unless this is obvious from the circumstances

C. Provide each user with a copy of the chatbot's model card before the conversation begins

D. Disclose the chatbot's confidence score alongside every response so users can evaluate the reliability of each statement

16. A GPAI model provider releases a model that the provider classifies as not posing systemic risk. A downstream deployer integrates the model into a high-risk use case (medical diagnostics). Under the EU AI Act, what happens to the GPAI provider's obligations?

A. The GPAI provider's obligations remain unchanged — the provider must comply with the general GPAI transparency and documentation requirements, but the high-risk compliance obligations fall on the deployer (or the entity integrating the GPAI into the high-risk system)

B. The GPAI provider automatically becomes subject to the systemic risk provisions because a downstream use case is high-risk

C. The GPAI provider's obligations increase to match the high-risk deployer's obligations because the provider's model is now part of a high-risk system

D. The GPAI provider bears no obligations once the model is released because all downstream obligations transfer to the deployer

17. An AI system for insurance pricing uses a gradient boosting model. When a customer asks why their premium is higher than a neighbor's, the insurer provides a list of the model's top 10 global feature importances. The customer argues this does not explain their individual premium. From a governance perspective, is the customer correct?

A. No — global feature importances provide the most comprehensive explanation of model behavior and satisfy all explainability requirements

B. No — insurers are exempt from individual explainability requirements under GDPR because insurance pricing is classified as a legitimate business interest

C. Yes, but only if the customer can demonstrate that their premium was calculated using automated processing with no human involvement whatsoever

D. Yes — global feature importances explain average model behavior across all customers but do not explain why this specific customer received this specific premium, and individual explainability requires local explanation methods that identify the factors driving this particular decision

18. An organization uses an AI system to monitor network traffic for cybersecurity threats. The system processes metadata (sender, recipient, timestamps, packet sizes) but not message content. Under GDPR, does the metadata processing require a lawful basis?

A. No — network metadata is technical operational data and is not personal data under GDPR

B. Yes — metadata about individuals' communications (who communicated with whom, when, and how much data was exchanged) constitutes personal data that can reveal patterns about individuals' behavior, relationships, and activities, requiring a lawful basis for processing

C. Only if the metadata is stored for more than 30 days, because short-term technical processing of metadata is exempt from GDPR

D. Only if the metadata is processed in the EU, because network traffic data generated in transit through EU infrastructure is not subject to GDPR territorial scope

19. Under the NIST AI RMF, the "Map" function includes contextualizing AI system risks. What does "contextualizing" mean in practical governance terms?

A. Creating a visual diagram (map) of the AI system's technical architecture and data flows

B. Comparing the AI system's risk profile against a standardized industry risk taxonomy

C. Understanding the specific environment in which the AI system operates — including the affected stakeholders, the social context, the regulatory landscape, the deployment conditions, and the potential for harm in this particular setting — because the same system poses different risks in different contexts

D. Mapping the AI system's features to the corresponding EU AI Act Annex III high-risk use cases

20. A governance professional must determine whether an AI system's post-deployment performance change qualifies as a "serious incident" under the EU AI Act. The system, classified as high-risk, has experienced a 3% decline in overall accuracy but the decline is entirely concentrated in one demographic group, whose accuracy dropped by 27%. Does this constitute a serious incident?

A. The concentrated 27% accuracy decline for one demographic group likely constitutes a serious incident because it may result in harm to individuals in that group through systematically incorrect decisions — the seriousness of an incident is determined by its impact on affected individuals, not by the aggregate performance change

B. No — a 3% overall accuracy decline is within normal operational variance and does not meet the threshold for serious incident reporting

C. Only if the affected demographic group constitutes more than 10% of the system's total user population

D. Only if the demographic group experiencing the decline has filed formal complaints about the system's performance

21. An AI vendor's contract states that the vendor retains "all rights to model improvements derived from customer data." A governance professional identifies this as problematic. Under which governance principle is this clause MOST concerning?

A. The principle of safety, because model improvements derived from customer data may introduce instability

B. The principle of transparency, because the vendor should disclose exactly which improvements were derived from each customer's data

C. The principle of fairness, because model improvements may benefit some customers more than others

D. The principle of purpose limitation and data rights — the clause effectively grants the vendor ongoing rights to use the organization's data for the vendor's commercial benefit, potentially training models sold to competitors, without clear limitation on how the derived improvements will be used

22. An organization deploys an AI system for medical triage. During the first month, the system correctly identifies a critically ill patient that a human triage nurse had initially assessed as lower-acuity. This "save" is documented. Over the same month, the system misclassifies three moderately ill patients as low-acuity, causing delayed treatment. These misclassifications are not documented because the monitoring system only flags missed critical cases. What governance gap does this asymmetric documentation create?

A. The documentation gap is acceptable because monitoring should focus on the most severe outcomes (missed critical cases) rather than moderate misclassifications

B. The asymmetric documentation creates survivorship bias in governance assessment — the organization sees the system's successes (documented saves) but not its failures (undocumented misclassifications), producing an artificially positive view of system performance that may mask systematic errors in non-critical classifications

C. The documentation gap only matters for regulatory reporting purposes and has no impact on the system's actual clinical governance

D. The documentation gap can be resolved by reducing the monitoring scope to focus exclusively on critical cases, eliminating the inconsistency

23. Under the EU AI Act, real-time biometric identification in publicly accessible spaces by law enforcement is generally prohibited, with limited exceptions. Which of the following is one of the enumerated exceptions that may permit such use?

A. Targeted search for specific victims of abduction, trafficking, or sexual exploitation, or for specific missing persons, subject to prior judicial authorization

B. General crime prevention in areas with historically high crime rates, subject to police commander authorization

C. Border security screening at airports for all arriving international passengers, subject to interior ministry authorization

D. Identity verification for access to government buildings, subject to building security manager authorization

24. An organization's AI system for content moderation classifies user-generated content into categories: "safe," "borderline," and "prohibited." The system automatically removes "prohibited" content and sends "borderline" content to human review. A governance professional notes that the definition of "borderline" has shifted over time as the model has been retrained on human reviewers' decisions. This creates a feedback loop. What is the governance significance?

A. The feedback loop improves the system's alignment with human judgment and is a desirable feature of iterative model improvement

B. The feedback loop has no governance significance because human reviewers make the final determination for borderline content

C. The feedback loop only matters if the human reviewers' decisions are inconsistent with the organization's stated content policies

D. The feedback loop means the boundary between "borderline" and "prohibited" may be shifting without deliberate policy decision — the system learns from reviewers who may have implicit biases or inconsistent standards, and those patterns become the new classification norm, potentially changing what content is removed versus reviewed without any explicit policy change

25. A governance professional must advise on the appropriate lawful basis for processing employee data to train an AI attrition prediction model. The model would analyze work patterns, communication frequency, and performance metrics. The HR team proposes using "legitimate interests" as the lawful basis. What is the MOST important governance consideration before accepting this proposal?

A. Whether the organization has published a privacy policy that mentions AI processing of employee data

B. Whether the AI vendor has certified that the model complies with GDPR requirements for employee data processing

C. Whether the HR team has obtained written approval from the CEO to use employee data for AI training purposes

D. Whether a legitimate interests assessment has been conducted that balances the organization's interest in predicting attrition against employees' reasonable expectations, privacy rights, and the potential consequences of being profiled — particularly given the power imbalance between employer and employee that may undermine the legitimacy of "legitimate interests" as a basis

26. An organization uses AI to generate personalized pricing for customers — offering different prices to different customers for the same product based on predicted willingness to pay. A governance professional raises a concern. Which governance principle is MOST directly implicated?

A. The principle of transparency, because customers should be informed that they may be seeing different prices than other customers

B. The principle of fairness — if willingness-to-pay predictions correlate with demographic characteristics (income, location, device type as socioeconomic proxies), the pricing system may produce outcomes where protected groups systematically pay more, constituting discriminatory pricing even without using demographic data directly

C. The principle of safety, because incorrect pricing predictions could cause financial harm to customers who overpay

D. The principle of accuracy, because the willingness-to-pay predictions may not be precise enough to justify personalized pricing

27. Under the NIST AI RMF, the "Govern" function is described as cross-cutting — informing and being informed by the Map, Measure, and Manage functions. What does "cross-cutting" mean in this context?

A. The Govern function must be completed before Map, Measure, and Manage can begin, creating a sequential dependency

B. The Govern function operates independently from the other three functions and does not interact with them

C. The Govern function establishes the organizational context, policies, roles, and culture that shape how Map, Measure, and Manage are performed — and findings from those functions feed back to refine governance policies and priorities, creating a continuous interaction rather than a one-time setup

D. "Cross-cutting" indicates that the Govern function applies only to systems that span multiple organizational departments

28. An organization receives a data subject access request asking what personal data is held in an AI model's training dataset. The model was trained on 10 million records and identifying one individual's specific contribution is technically extremely difficult. Under GDPR, what is the organization's obligation?

A. The organization must make reasonable efforts to identify and provide the data subject's information from the training data — if identification is genuinely impossible despite reasonable effort, the organization should document this impossibility and inform the data subject, while still providing any other personal data held about them outside the training dataset

B. The organization is completely exempt from responding to the access request because AI training data is automatically exempt from GDPR data subject access rights

C. The organization must recreate the entire training dataset in a searchable format to fulfill the access request, regardless of the technical effort required

D. The organization may refuse the request entirely because the data subject should have exercised their access rights before the data was used for training

29. An organization uses an AI system for automated document classification. The system was designed for English-language documents. An employee begins using it to classify documents in French without governance review. The system produces classifications for the French documents, but their accuracy has not been validated. What type of governance issue is this?

A. A technical performance issue that the AI system's monitoring will automatically detect and flag

B. Unauthorized use outside the system's validated scope — the system was validated for English documents only, and applying it to French documents without validation means every French classification is unverified, creating unknown risk proportional to the consequences of misclassification

C. A training deficiency because the employee should have known the system was designed only for English

D. Not a governance issue because the system is producing outputs for French documents, indicating it can process them

30. An AI system used in criminal justice has been challenged for discriminatory outcomes. The vendor argues the system was trained on "objective" criminal justice data. A governance professional identifies the flaw in this defense. What is the core problem with characterizing criminal justice data as "objective"?

A. Criminal justice data cannot be used for AI training because it is classified as sensitive personal data under GDPR Article 9

B. Criminal justice data is only "objective" if it has been validated by an independent third party before use in AI training

C. The term "objective" only applies to data collected through randomized controlled trials, which criminal justice data never involves

D. Criminal justice data reflects the outcomes of a system with well-documented disparities — arrest rates, charging decisions, sentencing patterns, and incarceration rates are influenced by enforcement priorities, resource allocation, and systemic biases, making this data a record of the justice system's behavior rather than an objective measure of criminal activity

31. A governance professional must determine whether the EU AI Act's transparency obligation for emotion recognition systems applies to a customer satisfaction survey tool. The tool analyzes text responses to open-ended survey questions and categorizes sentiment as positive, neutral, or negative. Is this tool an "emotion recognition system" under the Act?

A. Yes — any system that analyzes human-generated text for sentiment is an emotion recognition system under the Act

B. No — because the tool analyzes text rather than biometric data (facial expressions, voice patterns), and the Act's emotion recognition provisions apply specifically to systems that detect emotions from biometric data

C. Basic text sentiment analysis (positive/neutral/negative) that categorizes expressed opinions from survey text likely differs from emotion recognition as defined in the Act, which refers to identifying or inferring emotions or intentions of natural persons from biometric data — though the precise boundary may depend on regulatory guidance and the depth of the emotional inference

D. Yes — all sentiment analysis tools are classified as emotion recognition systems regardless of input modality

32. An organization's AI system for fraud detection flags a transaction as potentially fraudulent. The transaction is frozen, and the customer is unable to access their funds for 72 hours while the investigation is completed. The investigation concludes the transaction was legitimate. Under which GDPR right might the customer seek remedy for the 72-hour fund freeze?

A. The right to rectification (Article 16) if the AI system held inaccurate data about the customer, the right to object (Article 21) to the automated fraud monitoring, and potentially the right to compensation (Article 82) for material or non-material damage caused by the wrongful freeze — multiple GDPR provisions may apply depending on the circumstances

B. Only the right to erasure (Article 17), because the customer can request deletion of the fraud flag from their account records

C. Only the right to data portability (Article 20), because the customer can request transfer of their transaction data to a different financial institution

D. No GDPR right applies because fraud detection is a legal obligation that overrides all data subject rights without exception

33. An AI governance committee debates whether to require "explainability by design" — building interpretability into AI systems from the architecture stage — versus "post-hoc explainability" — applying explanation tools after the model is built. For a high-risk credit scoring system, which approach is MOST governance-appropriate?

A. Post-hoc explainability is always preferred because it can be applied to any model architecture without constraining the development team's choices

B. The choice is purely technical and has no governance implications because both approaches produce equivalent explanations

C. Neither approach is necessary because the EU AI Act does not require explainability for credit scoring systems

D. Explainability by design is preferred for high-risk applications because it ensures the model's decision process is inherently interpretable rather than approximated by a separate explanation tool — post-hoc methods may produce explanations that do not faithfully represent the model's actual reasoning, which is particularly problematic for consequential decisions like credit scoring

34. An AI system provider issues a model update that changes the model's behavior in ways the deployer did not anticipate. Under the EU AI Act, what obligation does the provider have regarding model updates for high-risk systems?

A. No specific obligation — providers may update models at any time without notice because model improvement is a standard maintenance activity

B. The provider must inform deployers of updates that materially affect the system's performance or behavior, and the update documentation must enable deployers to assess whether their compliance status is affected

C. The provider must obtain written consent from every deployer before implementing any model update

D. The provider must submit model updates to the national competent authority for approval before distributing them to deployers

35. An AI system generates a decision that harms an individual. The individual seeks accountability. The AI system was developed by Provider X, deployed by Organization Y, and makes decisions using data supplied by Data Broker Z. Under AI governance accountability frameworks, who is accountable?

A. Only Provider X, because the provider is always solely accountable for all harms produced by AI systems they develop

B. Only Organization Y, because the deployer has exclusive accountability for all outcomes of AI systems used under its authority

C. Accountability is distributed — Provider X is accountable for the system's design, training, and documentation; Organization Y is accountable for the deployment context, human oversight, and monitoring; and Data Broker Z is accountable for the data quality and lawfulness of the data supplied — each entity bears accountability proportionate to its contribution to the harm

D. No one is accountable because the distributed nature of AI system development creates an accountability gap that current governance frameworks cannot resolve

36. An organization uses differential privacy to protect personal data in its AI training dataset. After applying differential privacy, the organization claims the training data is "fully anonymized" and therefore no longer subject to GDPR. Is this claim correct?

A. Yes — differential privacy mathematically guarantees that individual records cannot be identified, making the data fully anonymized and exempt from GDPR

B. Yes, but only if the differential privacy epsilon parameter is set below 1.0, which is the threshold recognized by GDPR for anonymization

C. The claim is correct but only for the specific AI model trained on the differentially private data, not for any future models trained on the same data

D. The claim may not be correct — differential privacy reduces re-identification risk but whether it achieves full anonymization (as opposed to pseudonymization) depends on the implementation parameters, the specific techniques used, and whether the resulting data truly cannot be related to an identifiable person by any reasonable means, which must be assessed case by case

37. An organization implements an AI system for employee scheduling. Workers complain that the system consistently assigns night shifts to the same group of employees. Investigation reveals the system

optimizes for minimizing schedule conflicts, and these employees have fewer scheduling constraints (no childcare obligations, no second jobs, no class schedules). The system is technically functioning as designed. Under responsible AI principles, what governance concept applies?

A. Human-centricity — the system optimizes for operational efficiency without considering employee welfare, and the employees who accommodate the system's needs most easily (those with fewer external obligations) are penalized with the least desirable shifts, creating a system that treats employee flexibility as a resource to be exploited rather than a contribution to be reciprocated

B. Only employment law applies, and the scheduling pattern is lawful because night shift assignment based on scheduling flexibility is not discrimination

C. Only the principle of accuracy applies, because the system should more accurately predict which employees would prefer night shifts

D. No governance concept applies because the system is functioning as designed and meeting its stated optimization objective

38. An AI system processes insurance claims and generates automated denial letters. A governance audit discovers that the denial letters contain technical model output language that policyholders cannot understand — phrases like "claim confidence score below threshold" and "feature vector analysis indicates policy exclusion." Under governance principles, what is the most relevant concern?

A. The denial letters need only technical accuracy because policyholders can request human clarification through the standard appeals process

B. The automated denial letters fail the principle of meaningful transparency — affected individuals cannot understand the basis for their denial if explanations use technical AI terminology, and policyholders need plain-language explanations that identify the specific reasons for denial in terms they can act upon

C. The technical language is only a concern for policyholders with limited education and does not affect the governance assessment for the general population

D. The concern is limited to customer satisfaction metrics rather than governance compliance

39. An organization's AI governance committee is evaluating whether its AI-powered spam filter requires governance review. The committee debates whether a spam filter constitutes an "AI system" or

merely a traditional rule-based software tool. The spam filter uses machine learning to classify emails based on patterns learned from training data. Under the EU AI Act's definition, is this an AI system?

A. No — spam filters are explicitly excluded from the EU AI Act's scope because email filtering is a standard IT function

B. No — an AI system must produce outputs that affect individuals' rights or safety, and spam filtering does not meet this threshold

C. Yes — the EU AI Act defines AI systems broadly as machine-based systems that process inputs to generate outputs such as predictions, recommendations, or decisions, and a machine learning spam classifier that generates classification predictions falls within this definition

D. The classification depends on whether the spam filter uses deep learning or shallow learning techniques

40. A healthcare AI system was validated on a dataset of 50,000 patients at a single hospital. The system is deployed at 200 hospitals serving diverse populations. A governance professional argues this validation is insufficient. On what SPECIFIC basis?

A. The validation dataset may not represent the demographic diversity, clinical practices, disease prevalence, and data recording standards across 200 different hospitals — a system validated at one institution cannot be assumed to perform equivalently at sites with materially different patient populations and clinical contexts

B. The validation dataset of 50,000 is too small for any AI system deployment regardless of the number of deployment sites

C. The validation is only insufficient if the 200 hospitals are in different countries, because hospitals within the same country have standardized practices

D. The validation is sufficient because 50,000 patients provides statistical power that exceeds the minimum required for healthcare AI validation under the EU AI Act

41. An organization deploys an AI system and the monitoring system generates an alert indicating potential bias. The operations team investigates and determines the alert is a false positive — no actual bias exists. The governance professional argues this false positive should still be documented. Why?

A. The documentation is unnecessary because false positive alerts are routine operational noise that do not warrant governance records

B. False positive alerts should be documented because they contribute to understanding the monitoring system's own performance — tracking false positive rates over time reveals whether the monitoring thresholds are appropriately calibrated, and undocumented false positives prevent this analysis

C. False positives should only be documented if they exceed 10% of all alerts, which is the standard threshold for monitoring recalibration

D. The documentation is only necessary for regulatory compliance and has no practical governance value

42. Under the EU AI Act, what is the purpose of the EU database for high-risk AI systems?

A. To provide a publicly accessible registry where providers and deployers register high-risk AI systems before they are placed on the market or put into service, enabling transparency and regulatory oversight

B. To serve as a technical repository where providers upload their AI models' source code for regulatory inspection

C. To function as a marketplace where approved high-risk AI systems can be purchased by deployers

D. To store the complete training datasets of all registered high-risk AI systems for audit purposes

43. An organization uses an AI system for hiring that was developed in-house. The system's performance metrics are strong. A governance professional recommends conducting an external algorithmic audit despite strong internal metrics. What governance value does an EXTERNAL audit provide that internal assessment cannot?

A. External auditors always have more technical expertise than internal data science teams and will produce more accurate performance evaluations

B. External audits are legally required under the EU AI Act for all high-risk AI systems used in employment decisions

C. External auditors provide independence — they are not influenced by organizational incentives to confirm that systems are performing well, and their fresh perspective may identify issues that internal teams have normalized, rationalized, or failed to recognize due to familiarity with the system

D. External audits generate marketing-quality compliance certificates that demonstrate governance maturity to customers and investors

44. An AI system for predicting recidivism assigns risk scores that inform parole decisions. A defense attorney argues that the defendant cannot challenge the risk score because the system's methodology is proprietary. The prosecution argues that the vendor validated the system and the validation results are public. Under due process governance principles, what is the MOST significant gap in the prosecution's argument?

A. The prosecution's argument addresses the system's general validity but not the specific application to this defendant — public validation results show the system works on average, but they do not enable the defense to understand why this specific defendant received this specific score, challenge whether the inputs were accurate, or identify whether the system has known limitations for the defendant's demographic profile

B. The prosecution's argument is fully adequate because public validation results provide all the information needed to assess the system's reliability

C. The gap is only that the prosecution did not also provide the vendor's model card, which would supplement the validation results

D. The gap is only that the system should have been validated by a government-appointed auditor rather than the vendor's own validation team

45. A governance professional reviews an organization's AI governance documentation and finds that the same "template" impact assessment has been used for five different AI systems — the systems' names are changed but the risk analysis, affected populations, and mitigation measures are identical across all five documents. What governance failure does this represent?

A. Template-based impact assessments that are not tailored to each system's specific risks, affected populations, and deployment context provide no meaningful governance value — they create the appearance of compliance without the substance of risk analysis, and each system requires an assessment that reflects its unique characteristics

B. Using templates for impact assessments is a governance best practice that ensures consistency and completeness across all assessments

C. The template approach is acceptable as long as the five AI systems belong to the same risk classification tier

D. The template approach is only problematic if the five systems process different types of data

46. An organization's AI system for sentiment analysis processes customer reviews to identify product issues. The system assigns sentiment scores (positive, neutral, negative). A governance review discovers the system assigns more negative sentiment scores to reviews written in African American Vernacular English (AAVE) compared to reviews with identical sentiment written in Standard American English. What type of bias is this?

A. Selection bias, because the training data underrepresents reviews written in AAVE

B. Measurement bias — the system's sentiment analysis tool systematically misinterprets the linguistic features of AAVE, producing inaccurate sentiment measurements for this linguistic group that reflect the system's inability to process the dialect rather than actual differences in customer sentiment

C. Concept drift, because customer sentiment patterns have changed since the model was trained

D. Label bias, because the human annotators who labeled the training data assigned more negative sentiment to AAVE text

47. An organization implements the NIST AI RMF's Measure function by establishing metrics for a deployed AI system. The metrics include: overall accuracy, precision, recall, F1 score, and inference latency. A governance professional argues this metric set is incomplete. What category of metric is MOST critically missing?

A. User satisfaction metrics measuring how end users perceive the AI system's performance

B. Financial metrics measuring the AI system's return on investment and cost per prediction

C. Fairness metrics disaggregated across relevant protected groups — without fairness metrics, the organization can demonstrate that the system performs well overall but cannot demonstrate that it performs equitably for all populations it affects

D. Security metrics measuring the system's resistance to adversarial attacks and data breaches

48. An organization develops an AI system using training data that includes personal data collected under a privacy notice authorizing processing for "providing our services." The AI system is deployed for a different purpose — internal workforce analytics. Under GDPR's purpose limitation principle, is the workforce analytics use permissible?

A. Yes — any internal business use is compatible with the original purpose of "providing services" because workforce analytics supports service delivery

B. Yes — purpose limitation does not apply to AI training data because the data has been transformed into statistical patterns during the training process

C. Only if the organization updates its privacy notice within 30 days of deploying the system for workforce analytics

D. No — workforce analytics is a materially different purpose from "providing services," and using data collected for one purpose to train an AI system deployed for an incompatible purpose violates purpose limitation, requiring either a compatible purpose assessment or a new lawful basis

49. An AI system used by a hospital to predict patient deterioration generates alerts for nursing staff. A critical care nurse receives 47 alerts during a 12-hour shift. Only 3 of the 47 alerts result in clinical intervention. After six months, the nurse has developed a pattern of responding slower to all alerts. This behavioral change is an example of which governance-relevant phenomenon?

A. Alert fatigue — the high ratio of non-actionable to actionable alerts has degraded the nurse's responsiveness to all alerts including genuine emergencies, demonstrating that the alert system's sensitivity threshold creates a human factors risk that the governance framework must address

B. Alert fatigue is a clinical training issue for the nursing department to address through education about the importance of responding to every alert

C. The nurse's slower response time indicates they need additional technical training on interpreting the AI system's alerts

D. The 3-out-of-47 actionable alert ratio demonstrates that the system is performing well and the nurse's behavior simply reflects appropriate triage of low-priority alerts

50. An organization's AI governance committee receives a proposal for an AI system that would predict customer churn. The committee classifies the system as "minimal risk" because churn prediction does not directly affect individuals' rights. A governance professional challenges this classification. What is the MOST persuasive basis for the challenge?

A. Churn prediction systems are automatically classified as high-risk under the EU AI Act's Annex III

B. All AI systems that process customer data must be classified as high-risk regardless of their purpose

C. The churn predictions may be used to make decisions that DO affect individuals — such as offering retention discounts to high-value predicted churners while providing no retention efforts for low-value customers, effectively creating differentiated service quality based on AI profiling

D. The system should be classified as limited-risk because it interacts with customers through personalized offers

51. An organization deploys an AI system for predictive maintenance in a factory. Three months after deployment, the factory introduces new machinery that uses different materials and operates under different conditions. The AI system continues to operate normally. What is the MOST important governance question the organization should ask?

A. Whether the factory's insurance policy covers damages caused by AI monitoring failures for the new machinery type

B. Whether the AI system's performance metrics have changed since the new machinery was introduced

C. Whether the new machinery's maintenance schedule is compatible with the AI system's alert frequency

D. Whether the AI system was ever trained on data from machinery of the new type — because if not, the system cannot detect the new machinery's failure modes and its "normal" operation may reflect inability to monitor rather than confirmation that the machinery is healthy

52. An AI governance professional discovers that the organization's AI system for loan underwriting has been processing applications from a customer segment not represented in the original training data for six months. The system has been producing decisions for this segment, but no validation has been performed. The segment constitutes 15% of total applications. What is the MOST immediate governance action?

A. Wait for the quarterly performance review to evaluate the system's accuracy for the new segment before taking any action

B. Immediately implement human review for all applications from the unvalidated segment while conducting emergency validation testing — the system has been making unvalidated decisions affecting 15% of applicants for six months, creating both compliance and fairness risk

C. Retrain the model immediately using data from the new segment to correct the validation gap

D. Notify the AI vendor that the system is processing an unvalidated population and request the vendor's assessment

53. Under the NIST AI RMF, what is the relationship between risk tolerance and risk management?

A. Risk tolerance defines the level of risk the organization is willing to accept — and this threshold determines which risks require mitigation, which can be accepted with documentation, and which are unacceptable regardless of mitigation, making risk tolerance a prerequisite governance decision that shapes all subsequent risk management activities

B. Risk tolerance and risk management are identical concepts that refer to the same organizational process

C. Risk tolerance is determined by the NIST AI RMF Playbook for each industry category and cannot be customized by individual organizations

D. Risk tolerance only applies to high-risk AI systems and is not a relevant concept for minimal or limited-risk systems

54. An organization uses AI-generated synthetic data to augment an underrepresented group in its training dataset. After augmentation, the model's fairness metrics improve significantly. However, a governance professional raises a concern about the synthetic data. What is the MOST significant concern?

A. Synthetic data always introduces noise that degrades model performance regardless of its representation benefits

B. Using synthetic data to improve fairness metrics is prohibited under the EU AI Act because it constitutes artificial manipulation of compliance evidence

C. If the synthetic data does not accurately represent the real-world characteristics and variation within the underrepresented group, the model may learn artificial patterns rather than genuine ones — appearing to satisfy fairness metrics while actually making poorly calibrated decisions for the group it was designed to serve better

D. Synthetic data augmentation only requires governance review if it increases the training dataset by more than 50%

55. An AI system generates automated medical reports. The system is designed to flag abnormal findings for physician review. A governance audit reveals that the system generates reports with different levels of detail depending on the complexity of the case — complex cases receive shorter, less detailed reports. This happens because the model has lower confidence on complex cases and generates less text when uncertain. What governance concern does this create?

A. The concern is limited to the system's natural language generation quality and can be addressed by improving the report generation algorithm

B. The system should generate reports of identical length regardless of case complexity to ensure consistency

C. The concern is that physicians cannot see the confidence level in the reports and therefore cannot distinguish between concise reports for simple cases and inadequately detailed reports for complex ones

D. The most complex cases — where detailed reporting is MOST clinically important — receive the LEAST detailed reports, creating an inverse relationship between clinical need and report quality that could lead to important findings being understated or omitted precisely when thorough documentation matters most

56. A governance professional is evaluating two AI vendors for a high-risk deployment. Vendor A provides a model with 95% accuracy and a 200-page technical documentation package. Vendor B provides a model with 91% accuracy and complete transparency including training data access, disaggregated metrics, audit rights, and human-readable model explanations. For a high-risk deployment, which vendor better satisfies governance requirements?

A. Vendor B — for high-risk deployments, the deployer needs transparency, audit capability, and disaggregated metrics to fulfill independent governance obligations, and these capabilities cannot be satisfied by aggregate accuracy alone, regardless of how high the accuracy is

B. Vendor A — the higher accuracy means fewer errors affecting individuals, which is the primary governance concern for high-risk systems

C. Both vendors equally satisfy governance requirements because the EU AI Act does not specify a minimum accuracy threshold

D. Neither vendor satisfies requirements because high-risk AI systems must achieve 99% accuracy before deployment

57. An AI system for automated insurance claim processing denies a claim. The policyholder appeals and requests an explanation. The insurer provides: "Your claim was denied based on analysis of multiple policy factors." Under GDPR and consumer protection principles, is this explanation adequate?

A. Yes — the explanation confirms the denial was based on policy factors, which provides sufficient information for the policyholder

B. No — the explanation must identify the specific factors that contributed to the denial and how they influenced the decision, because a generic reference to "multiple policy factors" provides no actionable information that would enable the policyholder to understand, challenge, or address the basis for the denial

C. Yes, but only if the insurer also provides the policyholder with a copy of the complete insurance policy

D. The adequacy of the explanation depends on whether the claim value exceeds €5,000, because GDPR's explanation requirements scale with financial significance

58. An organization operates an AI system that was developed by an external vendor. The system processes personal data. Under GDPR, what data protection role does the organization typically fulfill?

A. Data processor, because the vendor developed the system and determines the means of processing

B. Joint controller with the vendor, because both parties influence how personal data is processed

C. Data controller, because the organization determines the purposes and means of processing personal data through the AI system, regardless of who developed the system

D. No GDPR role, because the vendor bears all data protection obligations for systems it develops

59. An organization trains an AI model on employee performance data. An employee exercises their GDPR right to erasure, requesting deletion of all their personal data. The organization deletes the employee's records from its databases. However, the AI model was trained on data that included this employee's records. Must the organization retrain the model without the employee's data?

A. Yes — the organization must immediately retrain the model to completely remove the employee's data influence from the model's learned parameters

B. No — the right to erasure applies only to identifiable data, and if the employee's data is no longer identifiable within the model's parameters, the model does not need to be retrained

C. Yes, but only if the employee specifically requests model retraining in their erasure request

D. The answer depends on whether the employee's data can be considered "personal data" within the model — if the model has generalized the employee's data into statistical patterns that cannot be attributed to any individual, retraining may not be required, but if the model memorizes or can reproduce the specific data, the right to erasure may extend to the model itself

60. An AI system for patient risk stratification assigns priority levels for preventive care programs. The system was trained on data from a health system that historically provided more preventive care to privately insured patients. A governance audit reveals the system assigns higher priority to patients whose profiles resemble those who historically received preventive care — effectively predicting access to care rather than medical need. What governance action is required?

A. Evaluate whether the system is measuring healthcare access rather than genuine medical risk, test whether the system's priority assignments correlate more strongly with insurance status than with clinical indicators, and if confirmed, redesign the model to predict health outcomes directly rather than using historical care patterns as the target variable

B. No action is required because the system is accurately predicting which patients will benefit from preventive care based on historical evidence

C. The only action needed is to add insurance status as a control variable in the model to adjust for insurance-related patterns

D. The action is limited to disclosing to patients that the priority system is based on historical care patterns

61. An organization's AI governance committee is reviewing a new AI vendor agreement. The agreement contains a limitation of liability clause that caps the vendor's liability at the total fees paid in the prior 12 months. For a high-risk AI system where a single incident could cause millions of euros in harm to affected individuals, what governance concern does this clause create?

A. Limitation of liability clauses are standard commercial terms that do not create governance concerns regardless of the AI system's risk classification

B. The clause is only problematic if it explicitly excludes liability for bias-related harms, which should always be uncapped

C. The liability cap creates a governance concern because it limits the organization's ability to recover damages if the vendor's system causes harm — for a high-risk system, the potential harm may vastly exceed the contractual liability cap, leaving the deploying organization bearing the financial responsibility for vendor-caused harms

D. The clause is acceptable as long as the organization purchases supplemental AI liability insurance

62. An AI system for credit scoring uses an ensemble model. The governance professional discovers that when the individual models in the ensemble disagree, the system defaults to the most conservative model's prediction (denial). This design choice was made by the development team without governance input. What governance issue does this design decision represent?

A. The design decision is purely technical and has no governance implications because it concerns model architecture rather than deployment policy

B. The conservative default means the system systematically favors denial over approval when models disagree — this is a value-laden design choice with significant implications for applicants (more denials) that should have been made through governance review rather than unilateral developer decision, because it determines how the system treats uncertain cases

C. The conservative default is always the correct design choice for credit scoring because false approvals (granting credit to unqualified applicants) pose greater risk than false denials

D. The design choice only requires governance review if the ensemble disagrees on more than 20% of applications

63. An AI system for employee monitoring tracks keyboard activity, application usage, and screen time. The system generates a "productivity score" displayed on the employee's manager's dashboard. An employee with a chronic pain condition takes frequent short breaks (2-3 minutes every 30 minutes) to manage their condition. The system scores this employee as "low productivity" due to the breaks. Under disability nondiscrimination and AI governance frameworks, what is the governance failure?

A. The system is correctly measuring productivity and the employee's low score reflects their actual output regardless of the medical reason

B. The only governance requirement is to inform the employee that their productivity is being monitored by an AI system

C. The governance failure is limited to the manager's use of the productivity score and does not relate to the AI system's design

D. The AI system's productivity metric penalizes behavior related to a disability accommodation without incorporating the accommodation into its scoring logic — creating discriminatory measurement that disadvantages employees with disabilities, which the governance framework should have identified during the impact assessment

64. An organization uses an AI system to generate marketing email subject lines. The system was trained on historical email performance data. A governance review discovers the system generates subject lines that perform well but frequently use manipulative psychological techniques — artificial urgency ("Last chance!"), false scarcity ("Only 3 left!"), and fear-based messaging ("Don't miss out or you'll regret it!"). The marketing team argues the system is performing its optimization objective. What governance principle applies?

A. The optimization objective (maximizing email open rates) has produced outputs that may violate consumer protection principles prohibiting unfair or deceptive practices — the system learned that manipulative techniques maximize engagement, and governance must evaluate whether the optimization objective should be constrained to exclude deceptive or manipulative content

B. The system is performing correctly because email marketing subject lines are not regulated by consumer protection law

C. The governance concern is limited to ensuring the marketing emails include an unsubscribe link

D. The manipulative techniques are only a governance concern if they are used for financial products, not for general marketing

65. An AI governance professional must determine whether a system that generates "risk scores" for insurance applicants constitutes "profiling" under GDPR Article 4(4). The system processes personal data to evaluate aspects relating to the natural person's health, financial situation, and lifestyle to predict insurance risk. Does this constitute profiling?

A. No — profiling only occurs when the processing leads to automated decisions, and risk scoring is merely an input to human decision-making

B. No — insurance underwriting is a regulated financial activity exempt from GDPR's profiling provisions

C. Yes — GDPR defines profiling as automated processing of personal data to evaluate certain aspects relating to a natural person, including analyzing or predicting aspects concerning health, economic situation, and personal preferences, which describes exactly what the insurance risk scoring system does

D. Only if the risk scores are stored in a structured filing system for more than six months

66. An organization's AI system for automated recruitment screening produces a shortlist. A rejected applicant files a discrimination complaint. During investigation, the governance team discovers the system's training data, test results, and model documentation were all created by the same three-person data science team with no external review or validation. What governance weakness does this reveal?

A. A three-person team is too small to develop an AI system for recruitment, which requires a minimum of five developers

B. The absence of independent review or validation creates an objectivity gap — the same team that built the system evaluated its own work, potentially missing biases they unconsciously built in or normalized during development, and high-risk employment AI requires independent assessment to verify the development team's conclusions

C. The governance weakness is limited to the team's failure to document their work in the organization's official documentation template

D. The weakness is only relevant if the three team members share the same demographic background

67. An AI system for loan processing has been operating for two years. The organization notices that the model's predictions have become increasingly accurate over time. Normally this would be positive, but

the governance professional is concerned. What could explain INCREASING accuracy that is actually a governance concern?

A. Increasing accuracy is always positive and cannot represent a governance concern under any circumstances

B. The model may be overfitting to the specific characteristics of the current customer population, making it increasingly good at predicting outcomes for the current population but increasingly fragile — any population shift could cause sudden, dramatic performance collapse

C. Increasing accuracy indicates concept drift and the model should be immediately retrained

D. The increasing accuracy may reflect a feedback loop — if the model's approved loans are the only ones that generate outcome data (repayment or default), the model is increasingly validated only on cases it already approves, creating an echo chamber where the model becomes more "accurate" by confirming its own previous decisions rather than by genuinely improving its predictive capability

68. Under the EU AI Act, what obligation does a deployer of a high-risk AI system have regarding the system's instructions for use?

A. The deployer must use the high-risk AI system in accordance with the instructions for use accompanying the system — including implementing the specified human oversight measures, monitoring for the risks identified by the provider, and reporting serious incidents

B. The deployer may use the system in any manner it deems appropriate as long as the system's CE marking is valid

C. The deployer is only required to retain the instructions for use on file and produce them upon request from a regulatory authority

D. The deployer must translate the instructions for use into all official languages of the EU member states where the system is deployed

69. An organization uses an AI system for credit decisioning. The system was validated and found to be fair across racial groups at the time of deployment. Eighteen months later, monitoring reveals a racial disparity has emerged. Investigation shows the disparity developed gradually as the model processed more applications from a demographic that was underrepresented in the original training data. What does this demonstrate about pre-deployment fairness validation?

A. Pre-deployment validation is sufficient if it is thorough and the disparity reflects a monitoring failure rather than a validation failure

B. Pre-deployment fairness validation captures a snapshot of the system's behavior at deployment but does not guarantee ongoing fairness — population shifts, data drift, and real-world dynamics can introduce disparities over time that were not present during validation, demonstrating that continuous post-deployment fairness monitoring is essential

C. The disparity proves the original validation was flawed and the system should never have been deployed

D. Pre-deployment validation is unnecessary because all fairness issues can be detected through post-deployment monitoring

70. An organization uses an AI system to generate performance evaluation summaries for employees. The system processes quantitative metrics and qualitative peer feedback. A governance audit reveals that the system generates more positive language in summaries for employees whose names are commonly associated with the majority ethnic group and more neutral or critical language for employees with names associated with minority ethnic groups — even when the underlying performance data is identical. What type of bias is this?

A. Selection bias, because the system selects different data points to include in summaries based on the employee's name

B. Historical bias, because the training data reflects historical performance evaluation patterns that favored majority-group employees

C. The system exhibits name-based bias in natural language generation — the model has learned associations between names and evaluation tone from training data, producing systematically different language for identical performance based on the employee's name, which functions as a proxy for ethnicity

D. Label bias, because the peer feedback data used for training was labeled differently for majority and minority group employees

71. An AI system for medical diagnosis achieves 95% accuracy overall. A governance professional requests disaggregated accuracy by patient age group. The results show: ages 18-40: 97%, ages 41-65: 96%, ages 66+: 84%. The system serves a hospital where 40% of patients are over 66. What does this disaggregation reveal that the aggregate metric conceals?

A. The aggregate metric is sufficient and the age-disaggregated results do not change the governance assessment

B. The disaggregation reveals that the system's marketing claim of "95% accuracy" should be replaced with "97% accuracy" to reflect its best-performing segment

C. The disaggregation only matters if the accuracy differences correlate with protected characteristics beyond age

D. The 95% aggregate masks a significant performance gap — the system is substantially less accurate for patients over 66, who constitute 40% of the hospital's population, meaning the patients most likely to have complex medical conditions receive the least accurate AI diagnostic assistance

72. An organization is transitioning an AI system from development to deployment. The governance professional must verify "deployment readiness." Which element is MOST commonly overlooked in deployment readiness assessments?

A. Verification that the human oversight personnel designated for the system have actually completed the required training on the system's capabilities, limitations, and override procedures — systems are often deployed with oversight "planned" but not yet implemented, rendering the oversight mechanism ineffective from day one

B. Verification that the AI system's source code has been code-reviewed by at least two senior developers

C. Verification that the AI vendor has provided a written guarantee of the system's performance for the first 12 months

D. Verification that the system's user interface has passed usability testing with a representative sample of end users

73. An organization operates an AI system for insurance claim fraud detection. The system flags 5% of claims for investigation. A governance audit reveals that of the flagged claims, only 2% are ultimately confirmed as fraudulent. This means 98% of flagged claimants are subjected to fraud investigation despite being legitimate. What governance metric describes this situation?

A. The system has a 2% accuracy rate, which is unacceptably low for any AI application

B. The system has a 98% recall rate, indicating it captures nearly all genuine fraud cases

C. The system has a 98% false positive rate within flagged claims — meaning the overwhelming majority of flagged claimants are legitimate customers subjected to the burden, delay, and stigma of fraud investigation, and governance must evaluate whether this ratio represents acceptable performance or creates disproportionate harm to innocent policyholders

D. The system has a 5% false negative rate, indicating it misses 5% of actual fraud cases

74. An AI system for automated essay scoring is used in a college entrance exam. Analysis reveals the system assigns higher scores to longer essays regardless of content quality. Students who write longer, lower-quality essays receive higher scores than students who write concise, higher-quality essays. What governance concept does this represent?

A. A training data quality issue that can be resolved by curating a training dataset with more variation in essay length

B. The system has learned to use essay length as a proxy for quality — a shortcut that produces the appearance of grading without measuring the construct (writing quality) it was designed to assess, creating a construct validity failure where the system measures word count rather than the intended assessment criteria

C. The system is functioning correctly because longer essays demonstrate greater student effort and should receive higher scores

D. A post-deployment monitoring gap because the correlation between length and score should have been detected during the first week of monitoring

75. An organization's AI system produces an output that causes harm. The organization's incident response plan is activated. The plan calls for: (1) containment, (2) investigation, (3) remediation, (4) notification, (5) lessons learned. A governance professional argues this sequence has a critical gap. What is missing?

A. Regulatory reporting to the national competent authority, which should occur early in the sequence rather than only after full investigation

B. Customer compensation, which should be included as step 2 before investigation begins

C. Third-party forensic analysis, which should replace the internal investigation to ensure objectivity

D. The incident response plan addresses affected individuals only through "notification" (step 4) but does not include active remediation of harm to affected individuals — notification informs people about the incident but does not actively assess and restore their situation, which should be a distinct step

76. A governance committee reviews an AI system and determines it is compliant with all EU AI Act requirements at the time of assessment. The committee issues an approval valid for 24 months. A governance professional argues that time-bounded compliance approvals are problematic for AI systems. What is the basis for this argument?

A. AI systems can change between assessments through data drift, model retraining, population shifts, or regulatory changes — a 24-month approval may not reflect the system's actual compliance status at month 18 if material changes have occurred, and governance should combine periodic reassessment with trigger-based review when material changes are detected

B. Time-bounded approvals are appropriate because the EU AI Act requires compliance reassessment every 24 months

C. The professional's argument has no basis because compliance status does not change between assessments if the system is not modified

D. Time-bounded approvals should be replaced with perpetual approvals that remain valid until a specific compliance violation is documented

77. An organization uses AI to screen job applicants. The system filters candidates based on skills matching, experience requirements, and education credentials. A governance review reveals the system filters out candidates who have employment gaps. The development team argues this is a neutral, non-discriminatory criterion. A governance professional disagrees. On what basis?

A. Employment gaps are a protected characteristic under the EU AI Act and cannot be used as a screening criterion

B. Employment gaps disproportionately affect women (maternity leave), people with disabilities (medical leave), caregivers, veterans (deployment), and formerly incarcerated individuals — making a facially neutral criterion produce disparate impact on multiple protected groups

C. Employment gaps only create a governance concern if they exceed 24 months in duration

D. The governance professional's concern is valid only if the organization operates in a jurisdiction with specific "ban-the-gap" legislation

78. An AI vendor's model card states the system was tested for fairness using "equalized odds." A governance professional must verify this claim. What SPECIFIC metric does equalized odds require?

A. Equal approval rates across demographic groups, meaning the same percentage of each group receives favorable outcomes

B. Equal precision across demographic groups, meaning the system is equally reliable when it predicts a positive outcome for members of any group

C. Equal individual treatment, meaning any two individuals with identical features receive identical predictions regardless of demographic group membership

D. Equal true positive rates AND equal false positive rates across demographic groups — meaning the system correctly identifies positive cases at the same rate AND incorrectly flags negative cases at the same rate for all groups

79. An organization deploys an AI content moderation system. Users appeal content removal decisions. The appeals process involves human review. Analysis reveals that the human appeals reviewers overturn the AI's removal decision in 31% of cases. What governance question should this overturn rate prompt?

A. The overturn rate is irrelevant because human reviewers are always more accurate than AI systems in content moderation decisions

B. The 31% overturn rate should prompt automatic recalibration of the AI system to match the human reviewers' decisions exactly

C. The overturn rate should prompt investigation into whether the 31% reversed removals share common characteristics — identifying systematic patterns in what the AI incorrectly removes helps diagnose specific failure modes that can be addressed through model improvement, and also indicates whether certain types of content or user populations are disproportionately affected by incorrect removals

D. A 31% overturn rate is within acceptable parameters for content moderation systems and does not require governance investigation

80. An AI governance professional is asked to identify the SINGLE most important question an organization should ask before deploying any AI system. Across all four AIGP domains and all governance frameworks studied, what question is it?

A. "Who is affected by this system, how might it harm them, and do we have adequate controls to prevent that harm?" — this question integrates risk identification (Domain I), legal compliance (Domain II), development governance (Domain III), and deployment oversight (Domain IV) into a single inquiry that drives every governance decision

B. "Does this system comply with the EU AI Act's risk classification requirements?"

C. "Has this system passed all technical performance benchmarks?"

D. "Does this system have a complete model card and impact assessment on file?"

81. An organization uses an AI model trained on customer data from the United States. The organization wants to deploy the same model in Japan. Beyond translation and localization, what is the MOST critical cross-border governance consideration?

A. The model must be retrained from scratch using exclusively Japanese customer data because U.S. training data cannot produce valid predictions for Japanese customers

B. The only consideration is ensuring the model's user interface supports the Japanese language

C. The consideration is limited to registering the model with the Japanese AI regulatory authority

D. The model was trained on U.S. customer behavior patterns that may not transfer to Japan — cultural differences in purchasing behavior, financial practices, communication styles, and demographic patterns mean the model's learned associations may produce inaccurate or biased predictions for Japanese customers, requiring validation on Japanese population data before deployment

82. An organization's AI governance committee discovers that the IT department deployed a machine learning model three months ago without any governance review. The model analyzes customer purchasing patterns to optimize inventory. The IT team argues it is "just a statistics tool, not AI." Under the EU AI Act's definition, does this model constitute an AI system?

A. The classification depends on whether the model was developed using a named AI framework (TensorFlow, PyTorch) versus a general statistical software package

B. A machine learning model that processes inputs to generate predictions constitutes an AI system under the EU AI Act's broad definition, regardless of whether the IT team characterizes it as "just statistics" — the Act's definition encompasses machine-based systems that generate outputs like predictions from inputs, which describes this model

C. The model is not an AI system because inventory optimization does not directly affect individual consumers' rights

D. The model only constitutes an AI system if it uses deep learning with more than three hidden layers

83. An organization operates an AI system for medical imaging analysis. The system flags potential abnormalities for radiologist review. A governance audit reveals that the system's false negative rate (missed abnormalities) is 3% overall but 11% for images taken on a specific older scanner model used primarily by a community health center serving a low-income population. What governance principle is MOST directly at issue?

A. The principle of transparency, because patients at the community health center should be informed of the scanner-related performance limitation

B. The principle of data minimization, because the system should not process images from scanners it cannot accurately analyze

C. The principle of fairness and safety — the nearly 4x higher false negative rate for images from the scanner used by the low-income community creates both an equity concern (disparate quality of AI-assisted diagnosis based on which facility patients attend) and a safety concern (missed abnormalities in a vulnerable population)

D. The principle of accountability, because no individual has been assigned responsibility for monitoring per-scanner performance metrics

84. An organization's AI system for processing benefit applications was deployed five years ago. The system has been retrained three times, expanded to two new use cases, and modified to incorporate three new data sources. The governance documentation reflects only the original deployment. An external audit is announced. What is the organization's MOST significant governance exposure?

A. The organization cannot demonstrate the governance history of the system's evolution — it cannot show that retraining, expansion, and modification were governance-reviewed, that impact assessments were updated, or that the current system operates within governed parameters

B. The exposure is limited to potential fines for documentation gaps, which are minor administrative matters

C. The exposure is limited to the two new use cases, which require retroactive impact assessments

D. The organization's only exposure is to the external auditor's fee, which will be higher due to the documentation gap

85. An AI governance professional must explain to a non-technical audience why AI governance differs from traditional software quality assurance. What is the MOST compelling distinction?

A. Traditional software is deterministic — it executes coded rules that produce the same output every time — while AI systems learn patterns from data that may contain biases, evolve through retraining, produce probabilistic outputs that can vary, and may behave differently for different populations in ways that the developers did not explicitly program

B. AI systems are more expensive to develop than traditional software and therefore require more rigorous oversight of the development budget

C. AI governance requires specialized software tools while traditional QA can be performed with standard testing frameworks

D. The distinction is purely regulatory — AI systems are subject to the EU AI Act while traditional software is not

86. An AI system for automated tax filing assistance provides tax advice to millions of users. The system occasionally provides advice that is technically correct under federal tax law but incorrect under the user's state tax law. Users who follow the advice file incorrect state returns. Under consumer protection and professional responsibility frameworks, what governance control should have been in place?

A. A disclaimer informing users that the system provides federal tax guidance only and state-specific advice should be sought from a tax professional

B. The system should be restricted to providing federal tax guidance only and should refuse to answer any state-specific tax questions

C. The governance concern is limited to ensuring the system's training data includes the most recent federal tax code updates

D. The system should detect the user's state, incorporate state-specific tax rules into its analysis, and clearly disclose when its advice may not account for all state-level requirements — because users who rely on AI tax guidance expect it to be comprehensive, and providing partially correct advice that users assume is complete creates foreseeable harm

87. An AI vendor presents benchmark results showing their model ranked #1 on a widely-used evaluation benchmark. A governance professional evaluating the vendor for a specific deployment is skeptical of relying on benchmark rankings. What is the SINGLE most important reason for this skepticism?

A. Benchmarks evaluate performance on standardized tasks under controlled conditions that may not reflect the deployer's specific data characteristics, population demographics, use case requirements, and operational environment — a model that excels on a benchmark may underperform in the specific deployment context that matters

B. Benchmark rankings are unreliable because vendors can overfit their models to specific benchmarks

C. Benchmark rankings change frequently and the #1 position may have changed since the vendor produced its marketing materials

D. Benchmark results are only valid for the specific hardware configuration on which the benchmark was conducted

88. An AI system deployed for customer service generates responses. A monitoring review reveals the system's response quality has remained stable but the system has developed a pattern of responses that subtly discourage customers from pursuing legitimate complaints — using phrases like "many customers find that the issue resolves on its own" and "processing your complaint may take 6-8 weeks." The system learned these patterns from training data containing historical responses where agents were incentivized to reduce complaint escalations. What governance concept describes this?

A. Concept drift, because the relationship between customer queries and appropriate responses has changed since the model was trained

B. A technical hallucination issue where the model generates fabricated information about processing timelines

C. The system has learned operationally dysfunctional organizational patterns from its training data — the historical incentive to reduce escalations produced training examples that subtly discouraged complaints, and the model has generalized these patterns as the "correct" way to respond to customers

D. A monitoring failure because the system's response quality metrics should have detected the discouraging language patterns

89. An organization operates 25 AI systems across its enterprise. A governance maturity assessment reveals that all systems have governance documentation, monitoring, and assigned ownership. However, the assessment also reveals that governance insights from one system's experiences are never shared with teams managing other systems. The same issues are discovered repeatedly across different systems. What specific governance capability is missing?

A. A dedicated AI ethics board that reviews all AI systems collectively rather than individually

B. An automated governance monitoring platform that correlates findings across all 25 systems in real-time

C. A formal regulatory scanning function that distributes regulatory updates to all AI system owners simultaneously

D. Organizational learning — the systematic capture, analysis, and sharing of governance insights across the AI portfolio so that lessons learned from one system's governance experiences improve governance for all systems

90. An organization's AI governance program has been operational for four years. Every governance metric is trending positively. All systems are documented, monitored, and compliant. Incident response is tested. Vendor management is systematic. The governance professional must identify the ONE remaining gap that prevents the program from achieving the highest maturity level. What is this gap?

A. The organization needs to increase its governance budget to exceed the industry benchmark

B. The organization has not systematically captured governance insights across its AI portfolio, shared lessons across teams, or used accumulated evidence to proactively improve governance practices — it

performs governance well but does not learn from governance, which is the defining characteristic of the highest maturity level

C. The organization needs to achieve ISO/IEC 42001 certification to reach the highest maturity level

D. The organization needs to hire additional governance staff to increase the team's capacity for review

91. An AI governance professional is preparing for the AIGP certification exam. Based on the complete Body of Knowledge, practice examinations, and governance principles studied, what is the SINGLE most important concept the professional should carry into the exam?

A. AI governance is about applied judgment — recognizing which principles apply in context, understanding why they matter, and synthesizing knowledge from across all four domains to determine the most appropriate governance response in situations where no single textbook answer exists

B. Memorize the EU AI Act's article numbers and penalty tiers because the exam tests specific regulatory citations

C. Focus on the technical aspects of AI model development because the exam primarily tests machine learning knowledge

D. Remember that the NIST AI RMF's four functions (Govern, Map, Measure, Manage) are the answer to every governance question

92. An AI system generates automated performance reviews for call center employees based on call duration, resolution rates, customer satisfaction scores, and adherence to scripts. The system consistently rates employees who spend longer on calls as lower-performing. However, these employees often handle the most complex and emotionally sensitive customer issues. The system does not account for call complexity in its scoring. What is the governance failure?

A. The system should only use customer satisfaction scores as the sole performance metric because it is the only metric that captures the full complexity of customer interactions

B. The system generates reviews that should only be used as conversation starters for manager-employee discussions and should never directly influence compensation decisions

C. The AI system's performance metric is misaligned with actual job performance — it penalizes behaviors (longer calls) that may reflect higher-quality service for complex cases, creating a

measurement that systematically disadvantages employees who handle the most difficult work rather than rewarding them for it

D. The system should weight all metrics equally to eliminate the bias toward call duration

93. An organization deploys an AI system with comprehensive governance controls. Five years later, the system has been retrained multiple times, applied to new populations, and modified with new features — all without governance review because the original approval was interpreted as a blanket authorization. What governance mechanism should have prevented this?

A. A contractual provision requiring the AI vendor to prevent any modifications to the deployed system

B. A change management process that defines which types of modifications (retraining, scope expansion, feature addition) require governance review, specifying thresholds that trigger reassessment rather than allowing the original approval to serve as perpetual authorization for all future changes

C. An annual renewal process requiring the governance committee to re-approve the system every year regardless of whether changes have occurred

D. A technical freeze that prevents any changes to the AI system's model, data, or features after the initial governance approval

94. An AI system for medical imaging produces a prediction with a confidence score of 52%. The clinical protocol requires human radiologist review for all predictions below 70% confidence. The radiologist reviews the image and agrees with the AI's prediction. A governance professional argues that agreeing with a 52% confidence prediction requires additional documentation. Why?

A. The AI system should be programmed to suppress predictions below 60% confidence because they are too unreliable to display

B. Below 60% confidence, the AI system's contribution to the diagnostic process is minimal and the radiologist should ignore the AI output entirely

C. The radiologist should disagree with the AI system whenever the confidence is below 70% to demonstrate meaningful human oversight

D. At 52% confidence, the AI's prediction is only slightly better than a coin flip — the radiologist's agreement should be based on their independent clinical assessment rather than influenced by the AI

output, and documentation should confirm that the radiologist conducted an independent evaluation rather than anchoring to the low-confidence AI prediction

95. An organization's AI governance program has achieved full maturity across all dimensions: policies, monitoring, incident response, vendor management, and organizational learning. A new challenge emerges: the organization begins deploying agentic AI systems that autonomously make multi-step decisions and take actions in the real world without human approval between steps. What NEW governance challenge do agentic systems create that the existing framework may not adequately address?

A. Agentic systems require governance that addresses autonomous multi-step action chains where errors can compound across steps, decisions build on previous decisions, and the speed of autonomous operation may outpace human ability to intervene — requiring new governance mechanisms like operational boundaries, automatic rollback capabilities, and intervention checkpoints that the existing framework's human-in-the-loop approach was not designed for

B. Agentic AI systems do not create any new governance challenges because all AI governance principles apply equally to agentic and non-agentic systems

C. The only new challenge is that agentic systems require more computational resources, increasing the organization's cloud computing costs

D. Agentic systems are prohibited under the EU AI Act and the organization should not deploy them

96. A governance professional reviews the organization's AI governance metrics. The dashboard shows: 100% of systems documented, 100% of impact assessments completed, 95% employee training completion, 48 governance committee meetings held. The professional notes these are all activity metrics. What OUTCOME metric would MOST directly demonstrate the governance program's effectiveness?

A. The total number of AI governance policies published per year

B. The average time to complete a governance review

C. The trend in AI-related incidents, complaints, and near-misses over time — a declining trend demonstrates that governance controls are actually preventing or catching issues before they cause harm, directly measuring whether the program reduces risk rather than just completing activities

D. The total governance budget expressed as a percentage of total AI development spending

97. An organization's AI governance committee has been operational for three years. The committee has approved every AI deployment proposal presented to it — a 100% approval rate. A governance professional argues this pattern is itself a governance concern. Why?

A. Governance committees must reject at least 25% of proposals to demonstrate independence and rigor

B. A 100% approval rate suggests the committee may not be functioning as an effective governance check — it may be rubber-stamping proposals, organizational culture may discourage objections, or proposals may not be presenting sufficient detail for the committee to identify genuine concerns

C. The approval rate is appropriate because it demonstrates the organization's AI development teams consistently produce well-governed proposals

D. The 100% approval rate is only a concern if the committee has fewer than five members

98. An AI system for automated résumé screening is deployed globally. The system was trained primarily on résumés from the United States and United Kingdom. A governance professional identifies that the system may systematically disadvantage candidates from countries where résumé conventions differ significantly — different formatting, different credential systems, different professional norms. The development team argues the system evaluates "content not format." What is the flaw in this argument?

A. The development team is correct because machine learning models process semantic content rather than formatting conventions

B. NLP models learn patterns that include formatting, structure, document organization, and linguistic conventions — not just semantic content — meaning the system may penalize candidates whose résumés follow conventions the model has rarely seen, even when the underlying qualifications are equivalent

C. The flaw is limited to the system's inability to process non-Latin character sets

D. The argument is only flawed if the system uses optical character recognition to process résumé images rather than parsing text directly

99. An AI governance professional must articulate the ONE principle that connects all governance activities across all four AIGP domains. What principle is it?

A. AI governance exists to ensure that AI systems serve human values throughout their lifecycle — by understanding what they are, knowing what constrains them, governing how they are built, and overseeing how they operate — because the ultimate purpose of every governance activity is to ensure that the benefits of AI are realized while its risks to individuals and society are responsibly managed

B. AI governance exists primarily to ensure compliance with the EU AI Act

C. AI governance exists primarily to protect organizations from regulatory penalties and litigation

D. AI governance exists primarily to ensure AI systems achieve maximum technical performance

100. Having completed Practice Exam 8, a candidate has now completed 800 of 1,000 practice questions. What study strategy would MOST effectively prepare the candidate for the final two practice exams and the actual AIGP certification examination?

A. Retake Exams 1-8 repeatedly until achieving a perfect score on each exam

B. Focus exclusively on reading and rereading the Part One learning chapters without further practice testing

C. Analyze error patterns across all eight exams — identifying whether persistent weaknesses reflect knowledge gaps, application errors, or reading errors — and target the specific type of error in the specific domain where it occurs, using exam explanations as focused study material rather than restudying all content uniformly

D. Focus exclusively on memorizing the answer keys from Exams 1-8 in case similar questions appear on the remaining exams

Practice Exam 8: Answer Key and Explanations

1. C — Under the EU AI Act, a deployer that significantly modifies a high-risk AI system's intended purpose without the provider's involvement may be reclassified as a provider for the modified system. This reclassification carries the full suite of provider obligations, including conducting a new conformity assessment for the modified intended purpose.

2. D — Facial templates derived from facial images constitute biometric data processed for the purpose of uniquely identifying a natural person, making them special category data under GDPR Article 9. The mathematical transformation from raw image to template does not change the data's classification — the template serves the same identification function as the original image.

3. B — Knowledge distillation transfers the teacher model's learned behavior to the student, including biases, memorized patterns, and behaviors originating from personal data in the teacher's training set. The student model inherits governance-relevant characteristics even without directly accessing the underlying data, requiring governance evaluation of the distilled model.

4. A — Concept drift occurs when the relationship between inputs and the correct output changes. New fraud techniques that did not exist during training change what "fraud" looks like in practice — the same transaction features now have different meanings. Data drift, by contrast, is a shift in the input distribution without a change in the underlying relationship.

5. D — Social scoring by public authorities that evaluates citizens based on social behavior to determine access to public services in unrelated contexts is classified as a prohibited practice under the EU AI Act. This differs from high-risk systems, which are permitted with appropriate governance controls, because the harm from government social scoring is deemed unacceptable regardless of safeguards.

6. C — GPAI model providers must supply technical documentation and training data summaries enabling downstream integration. High-risk deployers must use the system per instructions, maintain human oversight, and monitor deployment risks. These are distinct, complementary obligations reflecting different roles in the AI value chain.

7. B — The performance degradation concentrates in a segment that grew from 5% to 25% of the population. This is data drift — the population composition has shifted, and the model performs differently for the underrepresented segment. The model's training did not include sufficient representation of this group to maintain accuracy as their proportion grew.

8. A — For applicants automatically rejected below the threshold score, the decision is solely automated with no human involvement, and it produces significant effects (denial of credit access). Both conditions for Article 22 are met for these specific applicants, regardless of human involvement elsewhere in the process.

9. C — ISO 42001 provides a certifiable management system framework — the organizational structure for governing AI. The NIST AI RMF provides a risk management methodology — the process for identifying and managing AI risks. They address different governance needs and can be implemented together as complementary frameworks.

10. D — Hospital A is primarily responsible for its own local data quality, but the federated governance framework must establish data quality standards for all participants because poor data from one node affects the global model. The impact on all five hospitals must be assessed since federated model updates propagate across the entire network.

11. B — "Years of experience" may function as a proxy for age. If it drives the model to systematically favor or disfavor employees based on seniority in ways that correlate with age, it creates potential age-based disparate impact. The governance evaluation must assess whether the feature's predictive value justifies its discriminatory potential.

12. A — The provider — the entity that develops or commissions the AI system and places it on the market under its own name — bears primary responsibility for the conformity assessment. This includes ensuring the system meets all applicable requirements before it enters the market.

13. D — A 99.7% approval rate with 8-second average reviews strongly indicates the human reviewer is not exercising genuine independent judgment. Meaningful human involvement requires adequate time, relevant information, and practical ability to reach an independent conclusion — conditions that an 8-second review almost certainly cannot satisfy.

14. C — Synthetic-only training data means GDPR has no applicability to the training phase. However, if the deployed system processes real personal data during inference — making predictions about real individuals — GDPR applies fully to that deployment processing regardless of the training data's synthetic nature.

15. B — The EU AI Act's transparency provision for AI systems interacting with natural persons requires that individuals be informed they are interacting with an AI system, unless this is obvious from the circumstances. This is the primary transparency obligation — ensuring people know they are communicating with AI.

16. A — The GPAI provider's obligations remain at the general transparency and documentation level. When a downstream entity integrates the GPAI model into a high-risk system, the high-risk compliance

obligations fall on that integrating entity, not on the original GPAI provider whose classification is independent of downstream uses.

17. D — Global feature importances explain average model behavior across all predictions but do not explain the specific factors driving any individual customer's premium. Individual explainability requires local methods (SHAP values, LIME) that identify which features contributed most to this particular decision for this particular customer.

18. B — Network metadata — who communicated with whom, when, how frequently, and how much data was exchanged — constitutes personal data under GDPR because it reveals patterns about individuals' behavior, relationships, and activities. Even without content, metadata can profile individuals, requiring a lawful basis for processing.

19. C — "Contextualizing" means understanding the AI system's specific operating environment: affected stakeholders, social context, regulatory landscape, deployment conditions, and harm potential. The same AI model poses different risks in different contexts, and the Map function ensures risks are evaluated in the specific deployment setting rather than in the abstract.

20. A — The seriousness of an incident is determined by its impact on affected individuals, not by aggregate statistics. A 27% accuracy decline concentrated in one demographic group likely causes systematic harm to individuals in that group through incorrect decisions, meeting the threshold for a serious incident despite the modest 3% overall decline.

21. D — The clause grants the vendor ongoing commercial rights to data derived from the organization's operations without clear limitations. This raises purpose limitation concerns (customer data used for the vendor's commercial model improvement) and competitive risk (the vendor may train models sold to the organization's competitors using the organization's own data).

22. B — The organization documents the AI system's successes (saves) but not its failures (misclassifications), creating survivorship bias. The governance assessment sees only positive outcomes, producing an artificially favorable view that masks systematic errors in non-critical classifications invisible to the monitoring framework.

23. A — The EU AI Act permits real-time biometric identification in public spaces by law enforcement for the targeted search for specific victims of abduction, trafficking, or sexual exploitation, or for

specific missing persons — subject to prior judicial or administrative authorization. This is a narrow, specified exception to the general prohibition.

24. D — The feedback loop means the boundary between "borderline" and "prohibited" content may shift without any deliberate policy decision. Human reviewers' implicit biases and inconsistent standards become the new classification norm through retraining, potentially changing content moderation outcomes without explicit policy authorization.

25. D — A legitimate interests assessment must balance the organization's interest against employees' rights and reasonable expectations. The employer-employee power imbalance is particularly relevant — employees may have limited practical ability to object to processing by their employer, which may undermine the legitimacy of "legitimate interests" as a basis for invasive profiling like attrition prediction.

26. B — If willingness-to-pay predictions correlate with demographic characteristics through proxies like income, location, or device type, the pricing system may produce outcomes where protected groups systematically pay higher prices. This constitutes discriminatory pricing through AI profiling even without using demographic data directly.

27. C — "Cross-cutting" means the Govern function establishes the organizational context — policies, roles, culture, and commitment — that shapes how Map, Measure, and Manage are performed. Findings from those functions feed back to refine governance priorities, creating continuous interaction rather than sequential dependency.

28. A — The organization must make reasonable efforts to identify the data subject's information within the training data. If identification is genuinely impossible despite reasonable effort, the organization should document this impossibility and inform the data subject, while still providing any other personal data held about them outside the training dataset.

29. B — The system was validated for English documents only. Applying it to French documents without validation means every French classification is unverified — the system may produce outputs that appear reasonable but are actually inaccurate, creating unknown risk proportional to the consequences of misclassification in the specific use context.

30. D — Criminal justice data reflects the outcomes of a system with documented disparities. Arrest rates, charging decisions, and sentencing patterns are influenced by enforcement priorities, resource

allocation, and systemic biases. This data records the justice system's behavior, not an objective measure of criminal activity, making "objective" a misleading characterization.

31. C — Basic text sentiment analysis categorizing expressed opinions likely differs from emotion recognition as defined in the EU AI Act, which refers to identifying emotions from biometric data such as facial expressions and voice patterns. However, the boundary depends on the depth of emotional inference and evolving regulatory guidance.

32. A — Multiple GDPR rights may apply depending on circumstances: rectification (Article 16) if inaccurate data caused the flag, the right to object (Article 21) to the automated monitoring, and compensation (Article 82) for material or non-material damage from the wrongful freeze. The specific applicability depends on the facts.

33. D — Explainability by design ensures the model's decision process is inherently interpretable rather than approximated by a separate tool. Post-hoc explanation methods may produce explanations that do not faithfully represent the model's actual reasoning — a particularly problematic risk for consequential decisions like credit scoring where explanation accuracy matters.

34. B — The provider must inform deployers of updates that materially affect the system's performance or behavior. Update documentation must enable deployers to assess whether their compliance status is affected. This ensures deployers can evaluate updates against their deployment context before relying on modified system behavior.

35. C — Accountability is distributed across the value chain. The provider is accountable for design, training, and documentation. The deployer is accountable for deployment context, oversight, and monitoring. The data broker is accountable for data quality and lawfulness. Each entity's accountability is proportionate to its contribution to the harm.

36. D — Whether differential privacy achieves full anonymization (versus pseudonymization) depends on implementation parameters, techniques used, and whether the data truly cannot be related to identifiable individuals by any reasonable means. This must be assessed case by case rather than assumed from the application of the technique.

37. A — Human-centricity requires AI systems to serve human welfare. The scheduling system treats employee flexibility as a resource to exploit — those who accommodate most easily receive the worst

shifts repeatedly, while those who constrain the system receive better schedules. The system penalizes cooperative behavior rather than reciprocating it.

38. B — Denial letters using phrases like "feature vector analysis" and "confidence score below threshold" fail meaningful transparency because policyholders cannot understand or act on technical AI terminology. Governance requires plain-language explanations identifying specific denial reasons in terms the affected individual can comprehend and address.

39. C — The EU AI Act defines AI systems broadly as machine-based systems that process inputs to generate outputs such as predictions, recommendations, or decisions. A machine learning spam classifier that learns patterns from training data and generates classification predictions falls within this definition regardless of perceived simplicity.

40. A — Validation at a single hospital cannot represent the demographic diversity, clinical practices, disease prevalence, and data recording standards across 200 different hospitals. Different institutions serve different populations with different characteristics, and performance validated at one site cannot be assumed to transfer to materially different deployment contexts.

41. B — False positive alerts should be documented to track the monitoring system's own performance over time. Monitoring false positive rates reveals whether thresholds need recalibration, and undocumented false positives prevent this meta-monitoring analysis. Understanding monitoring system behavior is part of governance oversight.

42. A — The EU database provides a publicly accessible registry where providers and deployers register high-risk AI systems, enabling regulatory oversight and public transparency about which high-risk AI systems are in operation. This supports the Act's transparency and accountability objectives.

43. C — External auditors provide independence — they are not influenced by organizational incentives to confirm systems are performing well. Their fresh perspective may identify issues that internal teams have normalized or rationalized through familiarity. Independence is the core governance value of external assessment.

44. A — Public validation results show the system works on average but do not enable the defense to challenge this specific defendant's score. The defense cannot verify whether inputs were accurate, whether the system has known limitations for the defendant's demographic profile, or which factors drove the individual prediction — preventing meaningful adversarial challenge.

45. A — Template impact assessments with identical risk analysis, affected populations, and mitigations across different systems provide no meaningful governance value. Each system has unique characteristics requiring tailored assessment. Templates create compliance appearance without governance substance.

46. B — Measurement bias occurs when the measurement instrument systematically produces different results for different groups. The sentiment analysis tool misinterprets AAVE linguistic features, producing inaccurate negative sentiment scores that reflect the tool's limitation rather than actual differences in customer sentiment.

47. C — Without fairness metrics disaggregated across relevant protected groups, the organization can demonstrate overall performance but cannot verify equitable treatment for all affected populations. Fairness metrics are essential for any system that makes or influences decisions affecting individuals.

48. D — Workforce analytics is a materially different purpose from "providing services." Using data collected for one purpose to train an AI system deployed for an incompatible purpose violates GDPR's purpose limitation principle, requiring either a compatible purpose assessment demonstrating consistency or a new, independent lawful basis.

49. A — Alert fatigue is a governance-relevant phenomenon where excessive non-actionable alerts degrade human response to all alerts, including genuine emergencies. The 3-out-of-47 actionable ratio has systematically reduced the nurse's responsiveness, creating a patient safety risk that the governance framework's alert threshold design must address.

50. C — The churn predictions may drive decisions that directly affect individuals — retention discounts for predicted churners, reduced service for predicted non-churners. The system's risk classification must consider not just its direct outputs but the downstream decisions those outputs enable, which may create differential treatment based on AI profiling.

51. D — If the AI system was never trained on data from the new machinery type, it cannot detect that machinery's failure modes. "Normal" operation may reflect the system's inability to monitor rather than confirmation that the equipment is healthy — creating a dangerous false sense of security for unmonitored equipment.

52. B — The most immediate action is implementing human review for the unvalidated segment while conducting emergency validation. Six months of unvalidated decisions affecting 15% of applicants

creates both compliance risk (decisions made outside the system's validated scope) and fairness risk (unknown performance for this population).

53. A — Risk tolerance defines the level of risk the organization accepts. This threshold determines which risks require mitigation, which can be accepted with documentation, and which are unacceptable regardless of mitigation. Risk tolerance is a prerequisite governance decision that shapes all subsequent risk management activities.

54. C — If the synthetic data does not accurately represent the real-world characteristics of the underrepresented group, the model may learn artificial patterns. It would appear to satisfy fairness metrics while making poorly calibrated decisions for the group it was intended to serve better — creating the illusion of fairness without the substance.

55. D — The system generates less detail for more complex cases — precisely when thorough documentation is most clinically important. This creates an inverse relationship between clinical need and report quality, potentially causing important findings to be understated or omitted for the patients who need the most careful analysis.

56. A — For high-risk deployments, the deployer needs transparency, disaggregated metrics, and audit capability to fulfill independent governance obligations. Vendor B provides these essential governance tools. Vendor A's higher accuracy cannot compensate for the inability to verify compliance, evaluate fairness, or conduct meaningful oversight.

57. B — A generic reference to "multiple policy factors" provides no actionable information. The explanation must identify specific contributing factors and how they influenced the denial, enabling the policyholder to understand, evaluate, and potentially challenge the basis for the decision or address the identified factors.

58. C — The organization is the data controller because it determines the purposes and means of processing personal data through the AI system. The vendor's role as system developer does not transfer controller status — the entity that decides why and how personal data is processed bears controller obligations under GDPR.

59. D — The answer depends on whether the employee's data is identifiable within the model. If the model has generalized data into statistical patterns that cannot be attributed to any individual, retraining

may not be required. If the model memorizes or can reproduce specific data, the right to erasure may extend to the model itself.

60. A — The governance action must evaluate whether the system predicts healthcare access rather than medical need. If priority assignments correlate more strongly with insurance status than clinical indicators, the model should be redesigned to predict health outcomes directly rather than using historical care patterns as the prediction target.

61. C — For a high-risk system where a single incident could cause millions in harm, a liability cap at 12 months of fees leaves the deploying organization bearing the financial responsibility for vendor-caused harms that exceed the cap. The gap between potential harm and contractual recovery creates significant unmitigated financial risk.

62. B — Defaulting to the most conservative prediction when models disagree is a value-laden design choice that systematically favors denial. This decision determines how the system treats uncertain cases and has significant implications for applicants. Value-laden design choices with consequential impacts require governance review, not unilateral developer decision.

63. D — The AI system's productivity metric penalizes disability-related behavior (frequent short breaks) without incorporating the documented accommodation. This creates discriminatory measurement that systematically disadvantages employees with disabilities, which the governance framework should have identified during impact assessment.

64. A — The optimization objective (maximizing open rates) produced outputs that may violate consumer protection principles. The system learned that manipulative techniques work, and governance must evaluate whether the objective should be constrained to exclude deceptive content. A system working as designed can still produce harmful outcomes requiring governance intervention.

65. C — GDPR Article 4(4) defines profiling as automated processing of personal data to evaluate aspects relating to a natural person, including health, economic situation, and personal preferences. The insurance risk scoring system processes personal data to evaluate exactly these aspects, fitting squarely within the GDPR profiling definition.

66. B — The same team that built the system evaluated its own work, creating an objectivity gap. Internal teams may miss biases they unconsciously built in or normalized during development. High-risk

employment AI requires independent assessment to provide the objectivity that self-evaluation cannot guarantee.

67. D — Increasing accuracy may reflect a feedback loop: the model only generates outcome data for approved cases (it cannot observe defaults for denied applicants). Over time, the model confirms its own previous decisions, becoming more "accurate" within its self-selected validation set rather than genuinely improving predictive capability.

68. A — Deployers must use high-risk AI systems in accordance with the provider's instructions for use. This includes implementing specified human oversight measures, monitoring for identified risks, maintaining logs, and reporting serious incidents. The instructions are binding operational requirements, not optional guidelines.

69. B — Pre-deployment validation captures a snapshot of fairness at deployment but cannot guarantee ongoing fairness. Population shifts, data drift, and real-world dynamics introduce disparities over time that were not present during validation. This demonstrates that continuous post-deployment fairness monitoring is essential, not optional.

70. C — The system exhibits name-based bias in natural language generation. The model learned associations between names and evaluation tone from training data, producing systematically different language for identical performance data based on the employee's name — which functions as a proxy for ethnicity.

71. D — The 95% aggregate conceals that patients over 66 — constituting 40% of the hospital's population and likely having the most complex conditions — receive substantially less accurate diagnostic assistance (84% vs. 96-97% for younger groups). The patients most dependent on accurate diagnosis receive the least accurate AI support.

72. A — Systems are often deployed with oversight "planned" but not yet implemented. Verifying that designated human oversight personnel have actually completed required training on the system's capabilities, limitations, and override procedures ensures the oversight mechanism is functional from day one rather than aspirational.

73. C — A 98% false positive rate within flagged claims means the overwhelming majority of flagged claimants are legitimate customers subjected to investigation burden, delay, and stigma. Governance

must evaluate whether this ratio represents acceptable fraud detection performance or creates disproportionate harm to innocent policyholders.

74. B — The system uses essay length as a proxy for quality — a shortcut that produces the appearance of grading without measuring writing quality. This is a construct validity failure where the system measures word count rather than the intended assessment criteria, systematically advantaging verbose writers over concise, high-quality writers.

75. D — "Notification" informs affected individuals about the incident but does not actively assess and restore their situation. Active remediation — evaluating the harm to each affected individual and taking steps to make them whole — should be a distinct step in the incident response plan, not conflated with mere notification.

76. A — AI systems change between assessments through drift, retraining, population shifts, or regulatory changes. A 24-month approval may not reflect actual compliance at month 18 if material changes occurred. Governance should combine periodic reassessment with trigger-based review when material changes are detected.

77. B — Employment gaps disproportionately affect women (maternity), people with disabilities (medical leave), caregivers, veterans (deployment), and formerly incarcerated individuals. A facially neutral criterion produces disparate impact on multiple protected groups, making it a significant governance concern despite its appearance of neutrality.

78. D — Equalized odds requires equal true positive rates AND equal false positive rates across demographic groups. This means the system must correctly identify positive cases at the same rate AND incorrectly flag negative cases at the same rate for all groups — a dual requirement that is more stringent than demographic parity or predictive parity alone.

79. C — The 31% overturn rate should prompt investigation into whether reversed removals share common characteristics. Identifying patterns in what the AI incorrectly removes reveals specific failure modes, diagnoses which content types or user populations are disproportionately affected, and guides targeted model improvement.

80. A — "Who is affected, how might they be harmed, and do we have adequate controls?" integrates all four AIGP domains: risk identification (Domain I), legal compliance (Domain II), development

governance (Domain III), and deployment oversight (Domain IV). Every governance decision flows from understanding the system's potential impact on people.

81. D — U.S. customer behavior patterns may not transfer to Japan. Cultural differences in purchasing behavior, financial practices, communication styles, and demographic patterns mean the model's learned associations may produce inaccurate or biased predictions for Japanese customers. Validation on Japanese population data is essential before deployment.

82. B — The EU AI Act's definition encompasses machine-based systems that process inputs to generate outputs like predictions. A machine learning model that analyzes purchasing patterns to generate inventory predictions fits this definition regardless of whether the IT team labels it "AI" or "statistics." Functional characteristics determine classification.

83. C — The nearly 4x higher false negative rate for images from the older scanner creates both an equity concern (patients at the community health center receive less accurate AI diagnosis) and a safety concern (missed abnormalities in a population that may have less access to follow-up care). The disparity correlates with socioeconomic status through scanner quality.

84. A — The organization cannot demonstrate governance history for five years of system evolution. It cannot show that retraining, expansion, and modification were governance-reviewed, that assessments were updated, or that the current system operates within approved parameters. This creates the appearance of sustained ungoverned operation.

85. A — Traditional software executes coded rules deterministically. AI systems learn from data that may contain biases, evolve through retraining, produce probabilistic outputs, and may behave differently for different populations in ways developers did not explicitly program. This fundamental difference requires governance beyond traditional QA.

86. D — The system should detect the user's state, incorporate state-specific rules, and disclose limitations. Users relying on AI tax guidance expect comprehensive advice, and providing partially correct advice that users assume is complete creates foreseeable harm. The governance control must address the gap between user expectations and system capabilities.

87. A — Benchmarks evaluate performance on standardized tasks under controlled conditions that may not reflect the deployer's specific data, population, use case, and operational environment. A benchmark-

leading model may underperform in the deployment context that matters, making deployment-specific evaluation essential.

88. C — The system learned operationally dysfunctional patterns from training data where agents were incentivized to reduce escalations. The model generalized complaint-discouraging language as the "correct" response style, embedding organizational dysfunction into automated customer interactions without anyone explicitly programming this behavior.

89. D — Organizational learning — systematically capturing, analyzing, and sharing governance insights across the AI portfolio — is the missing capability. Without it, teams managing different systems independently discover the same issues without benefiting from each other's experiences.

90. B — The organization performs governance well but does not learn from governance. Systematically capturing insights, sharing lessons, and using accumulated evidence to proactively improve practices — the defining characteristic of the highest maturity level — is the one remaining gap preventing the program from achieving optimizing status.

91. A — AI governance is about applied judgment: recognizing which principles apply, understanding why they matter, and synthesizing knowledge across all four domains. The exam tests this practical synthesis capability, not isolated fact recall, making deep principle understanding the most valuable preparation.

92. C — The metric penalizes longer calls without accounting for complexity, systematically disadvantaging employees who handle the most difficult work. The measurement is misaligned with actual job performance — rewarding speed over quality and punishing the employees who provide the most skilled service.

93. B — A change management process defines which modifications require governance review and specifies thresholds triggering reassessment. Without explicit change management, the original approval becomes an unintended blanket authorization for all future changes, allowing the system to evolve beyond its governed parameters.

94. D — At 52% confidence, the AI's contribution is marginally better than chance. The radiologist's agreement must be based on independent clinical assessment rather than anchoring to the AI output. Documentation should confirm independent evaluation occurred, because low-confidence AI predictions can anchor human judgment without providing genuine diagnostic value.

95. A — Agentic systems create new governance challenges: autonomous multi-step action chains where errors compound, decisions build on previous decisions, and speed outpaces human intervention. Existing human-in-the-loop frameworks were not designed for this operational model, requiring new mechanisms like operational boundaries, automatic rollback, and intervention checkpoints.

96. C — A declining trend in incidents, complaints, and near-misses directly demonstrates that governance controls prevent or catch issues before harm occurs. Activity metrics measure governance effort; outcome trends measure governance effect — which is what the program exists to achieve.

97. B — A 100% approval rate over three years suggests the committee may not be functioning as an effective governance check. It may be rubber-stamping proposals, organizational culture may discourage objections, or proposals may lack sufficient detail for meaningful critical evaluation. Genuine oversight occasionally produces modifications or deferrals.

98. B — NLP models learn patterns that include formatting, structure, organization, and linguistic conventions — not just semantic content. Résumés following conventions the model rarely saw during training may be penalized even when the underlying qualifications are equivalent to those presented in familiar formats.

99. A — AI governance ensures that AI systems serve human values throughout their lifecycle. Every governance activity — from understanding AI capabilities through legal compliance to development governance and deployment oversight — ultimately serves the purpose of realizing AI's benefits while responsibly managing its risks to individuals and society.

100. C — Analyzing error patterns across all eight exams reveals whether persistent weaknesses stem from knowledge gaps, application errors, or reading errors. Each error type requires different remediation, making targeted correction more efficient than uniform restudying. Exam explanations serve as focused study material for each specific weakness.