

ISACA AI Fundamentals Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2026 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain accurate, complete, and timely information about this product from reliable sources.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	8
Explanations	10
Next Steps	16

SAMPLE

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

- Practice answering questions under realistic conditions,
- Improve accuracy and speed,
- Review explanations to strengthen weak areas, and
- Approach the exam with greater confidence.

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 - 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

SAMPLE

- 1. Which model is trained to assign a class to a sample based on the sample's features?**
 - A. Support vector machine**
 - B. Decision tree**
 - C. Random forest**
 - D. Naive Bayes**

- 2. Which concept is a tactical approach to developing and using AI tools ensuring diversity and reducing bias?**
 - A. Responsible AI**
 - B. Differential privacy**
 - C. GDPR (EU)**
 - D. EU AI Act of 2024**

- 3. Which statement best describes privacy-preserving inference?**
 - A. Data scraping from public web pages.**
 - B. Visualizing model predictions.**
 - C. Training models on distributed datasets with no security.**
 - D. Techniques that allow performing inference on encrypted or isolated data without exposing raw data (e.g., homomorphic encryption, secure enclaves).**

- 4. Which of the following is NOT a privacy-preserving ML technique listed?**
 - A. Differential privacy**
 - B. Data anonymization**
 - C. Federated learning**
 - D. Secure multiparty computation**

- 5. What term describes the settings that determine how the learning algorithm processes data to populate the model?**
 - A. Hyperparameters**
 - B. Model training parameters**
 - C. Epochs**
 - D. Confusion matrix**

- 6. Which regulation prohibits apps that pose unacceptable risks like real-time biometric identification?**
- A. EU AI Act of 2024**
 - B. GDPR (EU)**
 - C. CCPA 2020**
 - D. Privacy challenges in AI**
- 7. What are the layers between input and output that are not directly visible called?**
- A. Input layer**
 - B. Output layer**
 - C. Hidden layers**
 - D. Data layer**
- 8. What term describes the overall field of machines performing tasks that typically require human intelligence?**
- A. Artificial Intelligence**
 - B. Artificial Narrow Intelligence**
 - C. Artificial General Intelligence**
 - D. Symbolic AI**
- 9. Which group provides domain knowledge to ensure data and modeling tasks align with real-world requirements?**
- A. Subject matter experts**
 - B. Data scientists**
 - C. Project managers**
 - D. Product owners**
- 10. Which algorithm is commonly used for classification and can operate with margin-based optimization?**
- A. Support vector machine**
 - B. K-means**
 - C. Naive Bayes**
 - D. Logistic regression**

Answers

SAMPLE

1. A
2. A
3. D
4. B
5. A
6. A
7. C
8. A
9. A
10. A

SAMPLE

Explanations

SAMPLE

1. Which model is trained to assign a class to a sample based on the sample's features?

- A. Support vector machine**
- B. Decision tree**
- C. Random forest**
- D. Naive Bayes**

Classifying a sample by its features is the job of a classifier, and the support vector machine is a classic example designed specifically for this purpose. It learns a boundary that separates different classes in the feature space and assigns a class label to new samples based on which side of that boundary they fall on. During training, it aims to maximize the margin between the classes, which helps with better generalization to unseen data when making predictions. The other methods also produce class labels but work differently. A decision tree splits the space with simple rules and assigns the label at a leaf. A random forest builds many trees and combines their outputs for a final label. Naive Bayes uses probability calculations to pick the most likely class given the features, using independence assumptions to simplify the math. The key idea here is that the model in question is a classifier that maps feature vectors directly to class labels, and the support vector machine is a primary example of that approach.

2. Which concept is a tactical approach to developing and using AI tools ensuring diversity and reducing bias?

- A. Responsible AI**
- B. Differential privacy**
- C. GDPR (EU)**
- D. EU AI Act of 2024**

Responsible AI focuses on the practical steps and controls that guide the end-to-end creation and use of AI systems to be fair, transparent, and accountable. It emphasizes reducing bias and promoting diversity through concrete actions such as using representative and diverse data, applying fairness metrics, conducting bias testing, designing inclusively, ensuring explainability, and implementing ongoing monitoring and governance. These are the tactical, day-to-day practices that directly address diversity and bias in AI tools. Differential privacy centers on protecting individual data privacy, which is important but not primarily about bias reduction or diversity in AI outputs. Regulatory frameworks like GDPR or the EU AI Act set legal obligations and governance requirements, but they do not by themselves prescribe the hands-on tactics for developing and deploying AI to minimize bias and ensure diverse representation.

3. Which statement best describes privacy-preserving inference?

- A. Data scraping from public web pages.**
- B. Visualizing model predictions.**
- C. Training models on distributed datasets with no security.**
- D. Techniques that allow performing inference on encrypted or isolated data without exposing raw data (e.g., homomorphic encryption, secure enclaves).**

Privacy-preserving inference is about getting predictions from a trained model without exposing the sensitive data used to produce those predictions. This is achieved by performing computations on encrypted data or inside protected hardware environments so that raw inputs (and sometimes outputs) remain confidential. Techniques like homomorphic encryption allow operations on ciphertexts, while secure enclaves or trusted execution environments keep data isolated during processing. This is crucial when handling sensitive information (health, finance, personal data) and you want to use models without revealing the underlying data. The other statements don't describe this protective approach: scraping data from public pages doesn't involve keeping data private during inference; visualizing model predictions is simply about showing results; training on distributed data without security pertains to the data's protection during training, not inference.

4. Which of the following is NOT a privacy-preserving ML technique listed?

- A. Differential privacy**
- B. Data anonymization**
- C. Federated learning**
- D. Secure multiparty computation**

In privacy-preserving ML, the goal is to protect individuals' data while still enabling useful learning. The techniques that provide formal or practical privacy protections include adding carefully calibrated noise to outputs or data (differential privacy), keeping raw data on devices and sharing only model updates to reduce data exposure (federated learning), and enabling joint computations without revealing inputs (secure multiparty computation). Data anonymization, while it removes direct identifiers, often leaves behind indirect identifiers and patterns that can be linked back to individuals, especially when combined with other data sources. It also lacks formal guarantees against re-identification in real-world scenarios and with high-dimensional data. Because of these weaknesses, data anonymization does not meet the same robust privacy standards as the other techniques, making it the one that does not fit as a privacy-preserving ML technique in this list.

5. What term describes the settings that determine how the learning algorithm processes data to populate the model?

- A. Hyperparameters**
- B. Model training parameters**
- C. Epochs**
- D. Confusion matrix**

Hyperparameters are the settings that determine how the learning algorithm processes data to populate the model. They are configured before training and control aspects such as learning rate, network depth or number of layers, batch size, regularization strength, and the optimizer choice. These settings shape the learning dynamics, influencing how quickly and how well the model learns from the data. They are not learned from the data itself; the actual parameters the model ends up with (weights and biases) are learned during training. An epoch, in contrast, is simply one full pass through the training dataset, which is about iteration count rather than configuration. A confusion matrix is an evaluation tool used after training to assess performance, not a training configuration. Therefore, the term that best describes the settings guiding how the algorithm processes data to build the model is hyperparameters.

6. Which regulation prohibits apps that pose unacceptable risks like real-time biometric identification?

- A. EU AI Act of 2024**
- B. GDPR (EU)**
- C. CCPA 2020**
- D. Privacy challenges in AI**

AI regulation treats some uses as unacceptable risk and bans them to prevent harm before it happens. The EU AI Act of 2024 creates a prohibited category for certain AI practices, including real-time biometric identification in public spaces for surveillance. This prohibition reflects the serious privacy and civil liberties concerns such uses raise, signaling that some AI applications are not acceptable regardless of how they're deployed. In contrast, GDPR focuses on how personal data, including biometric data, can be collected and processed—establishing lawful bases, transparency, and rights for individuals—rather than issuing a blanket prohibition on real-time biometric identification. CCPA centers on consumer privacy rights and business obligations around personal data, not on prohibiting specific AI practices. "Privacy challenges in AI" is a topic area, not a regulation. So the regulation that prohibits apps posing unacceptable risks like real-time biometric identification is the EU AI Act of 2024.

7. What are the layers between input and output that are not directly visible called?

- A. Input layer**
- B. Output layer**
- C. Hidden layers**
- D. Data layer**

Hidden layers are the layers between input and output that aren't directly visible. In a neural network, the input layer passes the raw data in, the output layer delivers the final prediction, and the hidden layers in between transform that data through weighted connections and nonlinear activations. Each hidden layer learns intermediate representations—progressively extracting features and patterns from the data—so deeper networks can model more complex relationships. The term “hidden” reflects that these internal computations aren't observed directly from the outside. The input layer and the output layer are the visible ends, and while “data layer” isn't standard terminology for this concept, it doesn't describe the internal processing layers.

8. What term describes the overall field of machines performing tasks that typically require human intelligence?

- A. Artificial Intelligence**
- B. Artificial Narrow Intelligence**
- C. Artificial General Intelligence**
- D. Symbolic AI**

The main idea here is recognizing the broad field that covers machines performing tasks that typically require human intelligence. That field is Artificial Intelligence. It encompasses systems that learn, reason, perceive, and understand language, among other capabilities. Within AI, there are narrower areas like Artificial Narrow Intelligence, which handles specific tasks; Artificial General Intelligence, which would have broad, human-like cognitive abilities; and Symbolic AI, which focuses on explicit rule-based approaches. Artificial Intelligence is the correct term because it is the overarching discipline that includes all these forms and pursuits.

9. Which group provides domain knowledge to ensure data and modeling tasks align with real-world requirements?

- A. Subject matter experts**
- B. Data scientists**
- C. Project managers**
- D. Product owners**

Subject matter experts provide the domain knowledge needed to ensure data work reflects real-world requirements. They deeply understand the actual processes, terminology, constraints, and stakeholder needs of a given field, so they can define relevant data, labeling schemes, and what success looks like in practical terms. This context helps guide feature selection, data quality checks, and interpretation of model results so the analytics deliver actionable, real-world value. Data scientists contribute the technical methods and modeling skills, while product owners and project managers shape requirements and delivery, but the precise domain knowledge comes from those who work directly in the domain.

10. Which algorithm is commonly used for classification and can operate with margin-based optimization?

A. Support vector machine

B. K-means

C. Naive Bayes

D. Logistic regression

Margin-based optimization means choosing a decision boundary that maximizes the separation between classes. A support vector machine does exactly this: it finds a hyperplane that separates the classes with the largest possible margin, with the closest points on either side—the support vectors—defining that boundary. This focus on maximizing margin often leads to better generalization, and it can be extended to nonlinear boundaries using kernel tricks, giving both linear and nonlinear SVMs. K-means is an unsupervised clustering method, not a classifier. Naive Bayes is a probabilistic classifier that relies on conditional independence assumptions, not margin optimization. Logistic regression estimates probabilities via maximum likelihood and does not optimize a geometric margin, even though it yields a linear decision boundary. So the algorithm that uses margin-based optimization for classification is the one based on support vector machines.

Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at hello@examzify.com.

Or visit your dedicated course page for more study tools and resources:

<https://isacaaifundamentals.examzify.com>

We wish you the very best on your exam journey. You've got this!

SAMPLE