

ISACA AI FUNDAMENTALS PRACTICE TEST (SAMPLE)

STUDY GUIDE



Prepare with structure. Pass with confidence



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	15

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 term best describes a field that integrates rule-based systems with data-driven learning to improve reliability?**
 - A. Hybrid approaches
 - B. Explainable AI (XAI)
 - C. Data privacy issues
 - D. Deep learning
- 2. Which concept relates to shaping statistical language models and advances in natural language processing?**
 - A. Markov Chains
 - B. Reinforcement Learning
 - C. Symbolic AI
 - D. Probabilistic Reasoning
- 3. Which of the following is NOT a common bias type in AI?**
 - A. Representation bias
 - B. Measurement bias
 - C. Data drift bias
 - D. Historical bias
- 4. Which term denotes a strategic approach to AI development that promotes fairness and reduces bias?**
 - A. Responsible AI
 - B. Privacy challenges in AI
 - C. Differential privacy
 - D. EU AI Act of 2024
- 5. What process creates new sets of features to aid the training model in performing its task?**
 - A. Feature engineering
 - B. Feature selection
 - C. Dimensionality reduction
 - D. Data normalization

- 6. Which learning approach works by adjusting behavior to labels and features in a test set generated automatically?**
- A. Supervised learning
 - B. Self-supervised learning
 - C. Mask learning
 - D. Generative AI
- 7. Used for image recognition tasks?**
- A. Convolutional neural networks (CNNs)
 - B. Generative AI
 - C. Hybrid approaches
 - D. Mask learning
- 8. Which role ensures that data accurately reflects the task for which the model is trained?**
- A. Subject matter experts
 - B. Data scientists
 - C. Data engineers
 - D. Product owners
- 9. Which option best describes a common drawback of machine learning approaches?**
- A. They are less predictable, require substantial data and computational resources
 - B. They are less predictable but require minimal data
 - C. They are completely deterministic
 - D. They never require data
- 10. What term describes toolkits for deep learning that offer pre-built modules and optimization algorithms?**
- A. Cloud Platforms
 - B. Toolkits for Deep Learning Models
 - C. Specialized Frameworks
 - D. Training Loop

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which term best describes a field that integrates rule-based systems with data-driven learning to improve reliability?

- A. Hybrid approaches**
- B. Explainable AI (XAI)
- C. Data privacy issues
- D. Deep learning

Hybrid approaches blend rule-based systems, which encode explicit, hand-crafted rules, with data-driven learning, which derives patterns from data. This combination leverages the predictability and transparency of rules alongside the adaptability and accuracy of learning from real-world data. By using rules to guide decisions and learning to refine them, these systems can achieve greater reliability, especially in environments where formal constraints matter but data can reveal nuances the rules alone don't capture. Explainable AI relates to making decisions understandable and can be a feature of hybrid systems, but it describes interpretability rather than the integration concept itself. Data privacy issues and deep learning are not the concept described by the question; deep learning is a technique that could be part of a hybrid approach but doesn't specify the integration with rule-based logic.

2. Which concept relates to shaping statistical language models and advances in natural language processing?

- A. Markov Chains**
- B. Reinforcement Learning
- C. Symbolic AI
- D. Probabilistic Reasoning

Modeling text as a Markov process where the next word is determined by a limited, recent history is what shaped statistical language models and NLP advances. In a Markov chain for language, you treat each word (or token) as a state, and the probability of the next word depends on a small number of preceding words (often just the previous one or two). This memory-limited assumption keeps the problem tractable and lets us estimate these transition probabilities from large text data, giving rise to classic n-gram models like bigrams and trigrams. This approach provides a practical way to predict or generate text by multiplying the observed conditional probabilities. It also forms the basis for extensions like hidden Markov models used in tagging and recognition tasks, where observed words depend on hidden states, all within the same chain-and-probability framework. Other options involve different ideas: reinforcement learning focuses on learning decisions and policies through interaction, symbolic AI emphasizes explicit rules and logic, and probabilistic reasoning covers uncertainty in broader contexts but does not specifically describe the sequential, state-based modeling that Markov chains formalize for language.

3. Which of the following is NOT a common bias type in AI?

- A. Representation bias
- B. Measurement bias
- C. Data drift bias**
- D. Historical bias

Understanding biases in AI starts with recognizing where unfair outcomes come from. Data drift describes a change in the data distribution over time, which can cause a model to become less accurate as contexts shift. It's about evolving inputs, not a systematic tendency baked into how data are collected or labeled. Because of that, it isn't categorized as a common bias type. The idea behind the other terms is different: representation bias occurs when the training data don't reflect the real population, leading to skewed performance; measurement bias happens when data collection or measurement tools introduce consistent errors; historical bias arises when past decisions or societal prejudices are embedded in the data, biasing outcomes. So, data drift stands apart from these predefined bias types, making it the correct choice for NOT being a common bias type.

4. Which term denotes a strategic approach to AI development that promotes fairness and reduces bias?

- A. Responsible AI**
- B. Privacy challenges in AI
- C. Differential privacy
- D. EU AI Act of 2024

Responsible AI is a strategic approach to building and deploying AI that prioritizes fairness, transparency, accountability, and governance throughout the system's lifecycle. It means designing processes that actively identify and reduce bias—from data collection and labeling to model training and evaluation—using diverse and representative data, applying fairness metrics, and implementing bias mitigation techniques. It also involves making models explainable, establishing clear accountability and decision-making processes, and continuously monitoring performance to catch and address unfair outcomes over time. This holistic, ethics-driven framework aims to prevent harm and ensure AI decisions are justifiable and trustworthy. Privacy challenges in AI focuses on protecting individuals' data and preventing leakage or misuse, which is related but distinct from a broad, strategy-based emphasis on fairness. Differential privacy is a technique that adds controlled noise to data or results to protect privacy, not a general approach to fairness across all AI systems. The EU AI Act of 2024 is a regulatory framework governing how AI systems should be developed and used, setting obligations and compliance requirements rather than defining the overarching practice of building fair AI.

5. What process creates new sets of features to aid the training model in performing its task?

- A. Feature engineering**
- B. Feature selection
- C. Dimensionality reduction
- D. Data normalization

Feature engineering is the process of creating new features from raw data to improve model performance. By deriving new attributes—such as combining two measurements into a ratio, applying a log transformation to handle skew, binning continuous variables, creating interaction terms, or aggregating information over time or groups—you provide the model with signals that can better capture patterns in the data. This step is guided by domain knowledge and data understanding, aiming to reveal relationships the model might not discover from the original features alone. In contrast, selecting features means choosing which of the existing attributes to use, not creating new ones. Dimensionality reduction transforms features into a smaller set of components to reduce complexity, which can produce new representations but its primary goal is compactness rather than directly crafting signals from domain knowledge. Data normalization scales features for consistency but does not generate new features.

6. Which learning approach works by adjusting behavior to labels and features in a test set generated automatically?

- A. Supervised learning
- B. Self-supervised learning**
- C. Mask learning
- D. Generative AI

Self-supervised learning uses supervisory signals created from the data itself rather than relying on human-labeled examples. It designs a task where part of the input is transformed, masked, or shuffled, and the model must predict the original or the missing pieces. The labels come from the data generation process, so the model learns by adjusting its behavior to these automatically produced targets. This matches the idea of a test set generated automatically, since the targets are created without external labeling. Supervised learning needs human-provided labels, so it doesn't fit the automatic-label scenario. Mask learning is a specific self-supervised technique, not the broad approach itself, and generative AI focuses more on producing new data than on learning from automatically generated labels.

7. Used for image recognition tasks?

- A. Convolutional neural networks (CNNs)**
- B. Generative AI
- C. Hybrid approaches
- D. Mask learning

Image recognition relies on understanding spatial relationships in pixel data. Convolutional neural networks are built for grid-structured inputs like images: they use small, learnable filters that slide across the image to detect simple patterns such as edges, then stack many layers to build up increasingly complex features. The weight-sharing and local connectivity mean the same detector works across the image, and pooling helps reduce dimensionality while preserving essential information, giving translation invariance. This combination enables robust recognition even in large, diverse datasets. Generative AI focuses on creating new images or content rather than classification, hybrid approaches mix methods but the most direct, effective mechanism for recognizing visual content is provided by CNNs, and mask learning typically pertains to segmentation tasks rather than general recognition.

8. Which role ensures that data accurately reflects the task for which the model is trained?

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

Aligning data with the task through domain knowledge is essential. Subject matter experts provide the deep understanding of the real-world problem, define what counts as correct answers, and establish realistic scenarios and labeling standards. Their input ensures the ground-truth labels, categories, and edge cases truly reflect how the task should be performed in practice, so the model learns from data that represents the intended use. While other roles contribute significantly—data scientists design models and evaluate performance, data engineers build and maintain the data pipelines to deliver clean data at scale, and product owners set goals and priorities—they do not inherently anchor the data to the actual domain and task requirements as SME input does.

9. Which option best describes a common drawback of machine learning approaches?

- A. They are less predictable, require substantial data and computational resources**
- B. They are less predictable but require minimal data
- C. They are completely deterministic
- D. They never require data

A common drawback of machine learning approaches is that they can be less predictable and require substantial data and computational resources. In practice, models often behave differently when exposed to new or shifted data, because performance depends on how well the training data represent real-world situations. This unpredictability is tied to the need for large, varied datasets to learn robust patterns and to the inherent randomness in training processes (like initialization and stochastic optimization). Additionally, many powerful models demand significant computing power and storage, both for training and sometimes for deployment, which can be costly and time-consuming. Other statements don't fit because machine learning is not completely deterministic—randomness in training can lead to different outcomes between runs. It also isn't data-free; learning requires data to identify patterns. And training effectively typically requires more than a minimal amount of data to generalize well.

10. What term describes toolkits for deep learning that offer pre-built modules and optimization algorithms?

- A. Cloud Platforms
- B. Toolkits for Deep Learning Models**
- C. Specialized Frameworks
- D. Training Loop

Toolkits for deep learning models describe a collection of ready-to-use components and optimization routines that scientists and developers can assemble into neural networks quickly. These toolkits provide pre-built modules such as layers, activation functions, loss functions, optimizers, and data handling utilities, enabling experimentation without implementing common building blocks from scratch. That direct phrasing captures the idea of a modular, reusable set of deep learning components. Cloud platforms focus on infrastructure rather than bundled model-building components. Specialized frameworks refer to broader ecosystems for building models, which can include more than just modular toolkits. A training loop is the iterative process of training a model, not a collection of reusable components.

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