

ISACA AI FUNDAMENTALS PRACTICE TEST (SAMPLE)

STUDY GUIDE



**SAMPLE STUDY GUIDE
WITH LIMITED QUESTIONS**

**VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE**

<https://isacaaifundamentals.examzify.com>



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 metric is the harmonic mean of precision and recall?**
 - A. F-score
 - B. Accuracy
 - C. Confusion matrix
 - D. Hyperparameters
- 2. Which of the following is NOT listed as a major stage in the AI lifecycle?**
 - A. Data labeling
 - B. Problem definition
 - C. Data collection/curation
 - D. Model deployment
- 3. 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
- 4. Which component specifically enables explanations of AI decision processes to users?**
 - A. Explainable AI (XAI)
 - B. Convolutional neural networks (CNNs)
 - C. Data privacy issues
 - D. Generative AI
- 5. Which architecture is especially suitable for modeling sequential data and maintaining context across steps?**
 - A. Transformer
 - B. Convolutional Neural Network
 - C. Recurrent Neural Network
 - D. Generative Adversarial Network

6. What term describes media that show real or imagined individuals, often used for fake news?
- A. Deepfakes
 - B. Synthetic media
 - C. Media manipulation
 - D. Impersonation videos
7. Which regulation right to access, correct, erase data, and object to processing in the EU?
- A. GDPR (EU)
 - B. CCPA 2020
 - C. Personal Information Protection Law (PIPL - China)
 - D. EU AI Act of 2024
8. Which learning paradigm is used to detect novel patterns in data without relying on labeled examples?
- A. Unsupervised Learning
 - B. Reinforcement Learning
 - C. Probabilistic Reasoning
 - D. Symbolic AI
9. If a metric specifies the outcomes to be achieved, it is best described as a KGI rather than a KPI. Which option reflects this idea?
- A. KPI
 - B. KGI
 - C. ROI analysis
 - D. Process metric
10. Which set of elements are essential components of an AI governance policy?
- A. Data handling, privacy, fairness, accountability, risk management.
 - B. Data handling, privacy, fairness, accountability, risk management, roles, and approval workflows.
 - C. Roles and approval workflows only.
 - D. Data handling, privacy, fairness, accountability, risk management.

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which metric is the harmonic mean of precision and recall?

- A. F-score**
- B. Accuracy
- C. Confusion matrix
- D. Hyperparameters

The metric that combines precision and recall using a harmonic mean is the F-score, commonly called the F1 score. It is defined as $F1 = 2 \times (\text{precision} \times \text{recall}) / (\text{precision} + \text{recall})$. The harmonic mean brings both numbers into balance, so if either precision or recall is low, the F-score drops significantly, reflecting poor overall performance on the positive class. For example, with precision 0.9 and recall 0.5, the F1 score is about 0.64, showing that high precision alone doesn't compensate for low recall. This makes the F-score preferable when you want a single measure that penalizes imbalanced performance between correctly identifying positives and avoiding false positives. Accuracy, on the other hand, looks at overall correct predictions and can be misleading if the class distribution is imbalanced; a confusion matrix is just the table of true/false positives and negatives; hyperparameters are model settings, not a performance metric.

2. Which of the following is NOT listed as a major stage in the AI lifecycle?

- A. Data labeling**
- B. Problem definition
- C. Data collection/curation
- D. Model deployment

In AI project lifecycles, the major stages typically include defining the problem, collecting and curating data, training and validating the model, and deploying the model into production. Data labeling is a crucial activity for supervised learning, but it's usually treated as part of data preparation rather than a standalone major stage. It involves adding annotations to data so the model can learn from it, and it often happens within the data preparation step after data collection. Therefore, data labeling is not listed as a major stage, while problem definition, data collection/curation, and model deployment are.

3. 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.

4. Which component specifically enables explanations of AI decision processes to users?

- A. Explainable AI (XAI)**
- B. Convolutional neural networks (CNNs)
- C. Data privacy issues
- D. Generative AI

Explainable AI (XAI) enables users to understand how a model reaches its decisions. It focuses on transparency and provides ways to see which inputs most influenced a result, along with simple surrogate models or local reasoning for individual predictions. This visibility builds trust, allows oversight, and supports governance and compliance in contexts where decisions matter. The other options don't fit as well. A convolutional neural network is a powerful pattern-recognition model but isn't designed to explain its outputs to users. Data privacy concerns focus on protecting data rather than revealing the decision process. Generative AI aims to create new content and doesn't inherently provide explanations of how decisions are made.

5. Which architecture is especially suitable for modeling sequential data and maintaining context across steps?

- A. Transformer
- B. Convolutional Neural Network
- C. Recurrent Neural Network**
- D. Generative Adversarial Network

Modeling sequences requires memory of what happened earlier. A recurrent neural network provides that by carrying a hidden state through the sequence, so each step combines the current input with the previous context to produce an output and an updated context. This setup lets the model handle variable-length sequences and learn dependencies that span many steps, which is exactly what you need when what comes next depends on past elements. Training through time adjusts how much influence earlier steps have on later predictions, reinforcing long-range consistency. Other architectures have strengths—convolutional networks focus on local patterns, transformers use attention to relate many positions in a sequence, and GANs aren't built for sequential dependencies—but they don't naturally maintain step-to-step context in the same flowing, memory-based way.

6. What term describes media that show real or imagined individuals, often used for fake news?

- A. Deepfakes**
- B. Synthetic media
- C. Media manipulation
- D. Impersonation videos

Deepfakes are AI-generated media that depict real or imagined people in ways they never actually appeared, often swapping faces or synthesizing appearance and voice to create convincing but false content. This realism makes them a common vehicle for fake news and misinformation, because viewers may be fooled into believing the depicted person said or did something they did not. Synthetic media is a broader category that includes deepfakes as a subset. The broader term covers all artificially created media, not just those that impersonate specific individuals. Media manipulation and impersonation videos are less precise labels in this context: manipulation is a wide-ranging term for altering media, and impersonation videos isn't a standard, widely used term for this phenomenon.

7. Which regulation right to access, correct, erase data, and object to processing in the EU?

- A. GDPR (EU)**
- B. CCPA 2020
- C. Personal Information Protection Law (PIPL - China)
- D. EU AI Act of 2024

The rights to access, correct, erase, and object to processing are provided by the GDPR. This regulation gives individuals in the EU a set of data subject rights that empower them to control their personal information. Access allows you to see what data is held and how it's used; rectification lets you correct inaccuracies; erasure (the right to be forgotten) enables deletion under certain conditions; and the right to object lets you stop or limit processing, including for marketing or profiling. The GDPR applies to personal data across the EU and to organizations outside the EU under specific circumstances, making it the standard framework for these rights in Europe. In contrast, the California-based CCPA addresses similar ideas but in a US context, PIPL governs data in China, and the EU AI Act focuses on regulating AI systems rather than general data rights.

8. Which learning paradigm is used to detect novel patterns in data without relying on labeled examples?

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

Detecting novel patterns in data without labeled examples is the realm of unsupervised learning. It focuses on letting the data speak for itself, uncovering structure, groupings, and relationships without guidance from predefined targets. This makes it ideal for identifying new or unexpected patterns that weren't annotated beforehand. Techniques in this paradigm include clustering to group similar items, dimensionality reduction to simplify data while preserving its structure, and anomaly detection to spot unusual observations. In contrast, reinforcement learning relies on feedback from interacting with an environment, probabilistic reasoning deals with uncertainty and probabilistic models, and symbolic AI uses explicit rules and logic. Because the goal here is to discover structure and patterns from unlabeled data, unsupervised learning is the best fit.

9. If a metric specifies the outcomes to be achieved, it is best described as a KGI rather than a KPI. Which option reflects this idea?

- A. KPI
- B. KGI**
- C. ROI analysis
- D. Process metric

The idea being tested is the difference between outcome-focused indicators and performance indicators. A metric that specifies the end result to be achieved is describing a goal or objective, not just how well a process performs. That makes it a Key Goal Indicator, which is used to measure whether strategic goals and desired outcomes are reached. A KPI, by contrast, measures performance of activities or processes that contribute to those goals—like throughput, cycle time, or defect rate. ROI analysis is a financial evaluation of return, not a direct measure of goal achievement. A process metric tracks how a process performs, not the ultimate outcome. So when the metric centers on the end result to be achieved, it aligns with a Key Goal Indicator.

10. Which set of elements are essential components of an AI governance policy?

- A. Data handling, privacy, fairness, accountability, risk management.
- B. Data handling, privacy, fairness, accountability, risk management, roles, and approval workflows.**
- C. Roles and approval workflows only.
- D. Data handling, privacy, fairness, accountability, risk management.

A solid AI governance policy covers both what needs to be controlled in the AI lifecycle and how that control is organized. Data handling sets the rules for how data is collected, stored, processed, cleaned, and retained, emphasizing quality and provenance. Privacy ensures personal information is protected through controls like minimization, consent, and compliance with laws and regulations. Fairness focuses on preventing bias and ensuring equitable outcomes across users and scenarios. Accountability makes sure there is a clear owner for decisions and actions, with traceable records of what was done and by whom. Risk management identifies potential harms and uncertainties from AI systems and puts in place measures to reduce or monitor those risks. Including roles and approval workflows is essential because governance isn't just about what should be done; it's about who is authorized to do it and how changes get reviewed and approved. Roles define responsibilities (for data handling, model development, testing, deployment, monitoring), while approval workflows introduce checks and balances before critical actions—such as training on new data, deploying a model, or making changes to governance controls—are executed. This creates accountability, prevents unauthorized alterations, and provides audit trails for oversight. The other options omit one or more of these governance elements. Without roles and approval workflows, you lose the structural oversight that ensures the other components are consistently applied in practice.

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