

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 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**

- 2. Which paradigm learns by trial and error through reward signals to improve its policy?**
 - A. Reinforcement Learning**
 - B. Supervised Learning**
 - C. Symbolic AI**
 - D. Expert Systems**

- 3. Which type of neural networks is designed to handle sequential data more effectively than traditional neural networks?**
 - A. Convolutional Neural Networks**
 - B. Recurrent Neural Networks**
 - C. Transformers**
 - D. Generative Adversarial Network**

- 4. Which concerns relate to the handling and protection of personal data in AI applications?**
 - A. Data privacy issues**
 - B. Explainable AI (XAI)**
 - C. Generative AI**
 - D. Convolutional neural networks (CNNs)**

- 5. Which neural network type is designed to highlight patterns in input data through convolutions?**
 - A. Recurrent Neural Networks**
 - B. Convolutional Neural Networks**
 - C. Transformers**
 - D. Generative Adversarial Network**

- 6. Reacting to unforeseen events outside of the training environment describes which challenge?**
- A. Unpredictability in AI**
 - B. Data governance**
 - C. ISACA COBIT**
 - D. Lagging indicators**
- 7. Which component tries to generate fake data that the discriminator cannot distinguish from real data?**
- A. Generator in GAN**
 - B. Discriminator in GAN**
 - C. Training Loop**
 - D. Prompt Engineering**
- 8. Which concept is defined as balancing technological proficiency with ethical considerations, legal compliance, and strategic alignment?**
- A. Few-shot learning**
 - B. AI Governance**
 - C. ROI analysis in AI**
 - D. KPIs**
- 9. Which hardware is vital to the development and deployment of AI systems?**
- A. Central Processing Units (CPUs)**
 - B. Graphic Processing Units (GPUs)**
 - C. Random Access Memory (RAM)**
 - D. Hard Disk Drives (HDDs)**
- 10. Which term describes governance focusing on privacy, bias, responsibility, and autonomy in AI?**
- A. Ethical AI**
 - B. Federated learning**
 - C. Differential privacy**
 - D. Deepfakes**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. 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.

2. Which paradigm learns by trial and error through reward signals to improve its policy?

- A. Reinforcement Learning**
- B. Supervised Learning**
- C. Symbolic AI**
- D. Expert Systems**

Learning by trial and error with reward signals to improve behavior is reinforcement learning. In this approach, an agent interacts with an environment, chooses actions, receives feedback in the form of rewards, and updates its policy to maximize cumulative reward over time. The focus is on learning from feedback rather than from pre-labeled examples. This differs from supervised learning, which relies on labeled input-output pairs to learn a direct mapping; and from symbolic AI and expert systems, which depend on explicitly encoded rules and reasoning over a knowledge base. For instance, a robot learning to navigate a maze improves its policy by trying paths, receiving positive rewards for reaching the goal and negative signals for dead ends, and gradually choosing better routes.

3. Which type of neural networks is designed to handle sequential data more effectively than traditional neural networks?

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

Sequential data requires remembering what came before to make sense of what comes next. Recurrent neural networks incorporate feedback connections that pass the hidden state from one step to the next, creating a memory of prior inputs. This ability to carry information through time lets the network model temporal dependencies and context, which is essential for tasks like language, time series, and any sequence where order matters. Training them with backpropagation through time adjusts these recurrent connections so the network learns how past information influences current predictions. Convolutional networks excel at detecting patterns in spatial or local temporal structure but don't maintain a state across arbitrary-length sequences by default. Transformers use self-attention to relate all parts of the sequence, effectively modeling dependencies without recurrence, but that's a different approach to sequence handling that doesn't rely on the same internal memory loop. Generative adversarial networks focus on creating data distributions rather than modeling sequential dependencies over time. Since the question targets a design that inherently handles sequential data with memory across time, the recurrent neural network is the best fit.

4. Which concerns relate to the handling and protection of personal data in AI applications?

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

Handling and protecting personal data in AI applications centers on data privacy issues. This covers how data is collected, stored, used, and shared, and the safeguards needed to prevent misuse or exposure of individuals' information. It includes obtaining proper consent, minimizing data collected, retaining data only as long as necessary, enforcing access controls, encrypting data, and applying privacy-preserving techniques like anonymization or pseudonymization. It also involves governance and compliance with laws and regulations, addressing data subject rights, and considering risks such as data breaches or the potential for models to reveal sensitive information through privacy attacks. Explainable AI is about making AI decisions interpretable and understandable to humans, which aids accountability and trust but is not primarily about protecting personal data. Generative AI concerns the ability of models to create new content, which can raise different issues (like quality, misuse, or copyright) but not specifically the protection of personal data in the handling process. Convolutional neural networks are a type of model architecture used mainly for image processing, not a privacy protection concept.

5. Which neural network type is designed to highlight patterns in input data through convolutions?

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

Convolutional neural networks are built to detect patterns in input data by applying learnable kernels that slide over the input, creating feature maps that respond to specific patterns such as edges, textures, and shapes. This local pattern detection, combined with weight sharing across the input and hierarchical stacking of layers, lets the network capture spatial structures and invariances, which is especially powerful for images and other grid-like data. Recurrent networks focus on sequences and temporal dependencies, not on local pattern detection via convolutions. Transformers use self-attention to model relationships across the entire input rather than rely on convolutional pattern highlighting. Generative Adversarial Networks are a framework for training a generator and a discriminator to produce realistic data, not specifically defined by convolution-based pattern highlighting.

6. Reacting to unforeseen events outside of the training environment describes which challenge?

- A. Unpredictability in AI**
- B. Data governance**
- C. ISACA COBIT**
- D. Lagging indicators**

Reacting to unforeseen events outside of the training environment describes unpredictability in AI. When a model encounters inputs or scenarios it didn't see during training, its behavior can become unexpected or unsafe because it wasn't taught how to handle those cases. This challenge comes from distribution shifts and out-of-distribution data, where real-world situations differ from the data the model was built on. To cope, systems need awareness of uncertainty, mechanisms to detect anomalies, and safeguards like fallback paths or human review, so they don't rely on incorrect assumptions about what the model will do in new contexts. Data governance, COBIT, and lagging indicators address different concerns—data management rules, IT governance frameworks, and retrospective performance measurement, respectively—and don't directly describe the issue of handling novel, unseen situations.

7. Which component tries to generate fake data that the discriminator cannot distinguish from real data?

- A. Generator in GAN**
- B. Discriminator in GAN**
- C. Training Loop**
- D. Prompt Engineering**

In a Generative Adversarial Network, two networks compete: a generator and a discriminator. The generator creates synthetic data from random input, aiming to resemble real data. Its goal is to produce samples that the discriminator cannot reliably distinguish from real ones. The discriminator, on the other hand, evaluates inputs and tries to tell real data apart from generated data. The generator is the component that generates fake data to fool the discriminator, which is why it's the correct choice. The discriminator's role is the opposite—detect fake data—while the training loop coordinates the alternating training, and prompt engineering is unrelated to this GAN interaction.

8. Which concept is defined as balancing technological proficiency with ethical considerations, legal compliance, and strategic alignment?

- A. Few-shot learning**
- B. AI Governance**
- C. ROI analysis in AI**
- D. KPIs**

AI governance is the practice of balancing technological capability with ethical considerations, legal compliance, and strategic alignment. It establishes policies, standards, and processes to guide the development, deployment, and monitoring of AI systems so they are effective yet responsible. Key elements include accountability structures, risk management, data governance, transparency and explainability, and ongoing oversight by governance bodies. This approach helps ensure AI initiatives support business goals while respecting rights and complying with laws, reducing risks such as bias, privacy issues, and regulatory breaches. Techniques like few-shot learning, ROI analysis, and KPIs describe specific methods or metrics but do not define the overarching governance framework that coordinates technology with ethics, law, and strategy.

9. Which hardware is vital to the development and deployment of AI systems?

- A. Central Processing Units (CPUs)**
- B. Graphic Processing Units (GPUs)**
- C. Random Access Memory (RAM)**
- D. Hard Disk Drives (HDDs)**

AI development and deployment rely on heavy parallel computation. GPUs are built with thousands of cores that can perform many operations at once, which matches the way neural networks execute large numbers of matrix and vector calculations simultaneously. This parallelism provides far higher throughput for training and inference than a typical CPU, making it feasible to train large models and run real-time AI tasks. Modern GPUs also offer high memory bandwidth and specialized capabilities, like tensor cores, which speed up the precise math used in deep learning. Those features directly translate into shorter training times, the ability to handle bigger datasets, and faster, scalable deployment of AI systems. While a functioning system still needs CPU power, sufficient RAM, and storage, none of those match the parallel compute capabilities that GPUs bring to AI workloads. That's why GPUs are considered the most vital hardware for AI development and deployment.

10. Which term describes governance focusing on privacy, bias, responsibility, and autonomy in AI?

- A. Ethical AI**
- B. Federated learning**
- C. Differential privacy**
- D. Deepfakes**

Ethical AI describes governance that centers on privacy, bias, responsibility, and autonomy in AI systems. It encompasses principles like fairness and non-discrimination, privacy protection, transparency, accountability, and appropriate human oversight, ensuring that AI decisions respect individuals' privacy, avoid perpetuating or amplifying bias, and keep humans involved in important outcomes. Federated learning is a method for training models across multiple devices without centralizing data, aimed at privacy in the data process, not a governance term. Differential privacy is a technique that adds noise to protect individual data contributions, again a method rather than a governance framework. Deepfakes refer to realistically generated synthetic media and highlight policy concerns, but the term itself does not denote governance focused on privacy, bias, responsibility, and autonomy.

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