

# Ethics of Artificial Intelligence (AI) 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 item is listed as missing from n-gram models to resemble large language models?**
  - A. Transformer architecture**
  - B. Scale**
  - C. Post-training**
  - D. Data augmentation**
  
- 2. What is the key difference between act utilitarianism and rule utilitarianism?**
  - A. Act utilitarianism evaluates the consequences of each act; rule utilitarianism evaluates actions by conformity to rules with the best consequences.**
  - B. Act utilitarianism rejects any consideration of happiness.**
  - C. Rule utilitarianism is identical to Kantian ethics.**
  - D. Act utilitarianism focuses only on intentions.**
  
- 3. How does rule utilitarianism define a morally right act?**
  - A. An act is morally right if it produces the best consequences.**
  - B. An act is morally right if it conforms to a rule whose acceptance leads to the best consequences.**
  - C. Moral rules are irrelevant to morality.**
  - D. All rules are equally acceptable in all situations.**
  
- 4. In Chalmers' framing, what is a philosophical zombie (p-zombie)?**
  - A. An atom-for-atom identical to a human but without phenomenological experience**
  - B. An entity that is fully conscious**
  - C. A machine that can feel emotions**
  - D. A dream-like figure**
  
- 5. Compared to large language models, traditional n-gram models lack which capabilities?**
  - A. Contextual windowing.**
  - B. Ability to handle out-of-vocabulary words.**
  - C. A neural network-based predictor.**
  - D. Transformer-based architecture, large-scale training, and post-training capabilities.**

- 6. What concern about AI-generated art is raised in the material?**
- A. Theft of labor due to artists not being compensated**
  - B. Plagiarism as copying exact works**
  - C. Privacy violations in data used for training**
  - D. AI-generated art is always infringing**
- 7. Is a private conversation overheard by a friend in a public park a violation of privacy, according to this account?**
- A. No**
  - B. Yes**
  - C. Only if the overheard content is sensitive**
  - D. Only if you intended privacy**
- 8. What is a recommended strategy for technology deployment to address ethical concerns, as discussed under 'Slow Development'?**
- A. Slowing release pace to allow adjustment and safety review.**
  - B. Releasing updates as quickly as possible.**
  - C. Bypassing safety evaluations to accelerate benefits.**
  - D. Keeping releases secret until fully solved.**
- 9. What does RLHF stand for?**
- A. Reinforcement Learning from Human Feedback**
  - B. Randomized Language Handling Framework**
  - C. Recursive Latent Heuristic Functions**
  - D. Robust Linguistic Heuristic Formulation**
- 10. What is the primary aim of utilitarianism?**
- A. The greatest happiness for the greatest number.**
  - B. To maximize wealth and power.**
  - C. To follow rigid duties regardless of outcomes.**
  - D. To maximize personal freedom without regard to others.**

## Answers

SAMPLE

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

SAMPLE

## **Explanations**

SAMPLE

**1. Which item is listed as missing from n-gram models to resemble large language models?**

- A. Transformer architecture**
- B. Scale**
- C. Post-training**
- D. Data augmentation**

Post-training is the step that truly separates basic language models from large language models. N-gram models are simple probability estimators built from counting word sequences; they have no learning phase beyond those counts and no subsequent refinement stage. Large language models, on the other hand, undergo a two-stage process: a broad pretraining on massive text to learn general language patterns, followed by a post-training phase that tunes the model for specific behavior—such as instruction following, alignment with human preferences, and safety constraints—through fine-tuning and sometimes reinforcement learning from human feedback. This post-training stage is what endows LLMs with the ability to understand and follow complex prompts, adapt to different tasks, and produce coherent, context-aware responses.

**2. What is the key difference between act utilitarianism and rule utilitarianism?**

- A. Act utilitarianism evaluates the consequences of each act; rule utilitarianism evaluates actions by conformity to rules with the best consequences.**
- B. Act utilitarianism rejects any consideration of happiness.**
- C. Rule utilitarianism is identical to Kantian ethics.**
- D. Act utilitarianism focuses only on intentions.**

The main idea being tested is how the judging of right actions differs between evaluating each act versus evaluating the rules we follow. Act utilitarianism looks at the consequences of a specific action in a given moment and chooses the act that maximizes overall happiness in that situation. Rule utilitarianism, on the other hand, assesses actions by whether they conform to a rule that, if followed generally, would lead to the greatest happiness overall. This means a particular act that seems to produce good results might be rejected by rule utilitarianism if it would undermine a beneficial rule if followed widely. For example, lying to save a life could be justified by act utilitarianism because the outcomes are better in that instance, but rule utilitarianism would oppose lying because the rule against lying tends to produce better consequences when applied broadly, preserving trust and social cooperation. Both forms center happiness as the moral aim, but they differ in whether we evaluate the act itself or the rule behind the act.

### 3. How does rule utilitarianism define a morally right act?

- A. An act is morally right if it produces the best consequences.
- B. An act is morally right if it conforms to a rule whose acceptance leads to the best consequences.**
- C. Moral rules are irrelevant to morality.
- D. All rules are equally acceptable in all situations.

Rule utilitarianism defines a morally right act as one that follows a rule whose general adoption would maximize overall good. The key idea is not judging the single act by its immediate outcome, but judging the rule itself by the consequences of applying that rule broadly. If everyone followed a rule and that widespread practice led to the best consequences, then acting in line with that rule is morally right. This explains the shift from evaluating individual acts to evaluating the usefulness of rules as guides for action, aiming for predictable, fair, and cooperative behavior that promotes the greater good. This fits because the other views mix in different ideas: judging rightness by the act's own consequences is act utilitarianism; saying rules are irrelevant contradicts the very premise of rule-based ethics; and saying all rules are equally acceptable ignores choosing rules based on which ones maximize utility.

### 4. In Chalmers' framing, what is a philosophical zombie (p-zombie)?

- A. An atom-for-atom identical to a human but without phenomenological experience**
- B. An entity that is fully conscious
- C. A machine that can feel emotions
- D. A dream-like figure

A philosophical zombie, in Chalmers' framing, is something that is atom-for-atom identical to a human in every physical way but lacks subjective, phenomenological experience. The key idea is that all outward behavior, speech about feelings, and responses would be indistinguishable from a normal person, yet there would be no "what it's like" inner experience. Chalmers uses this to illustrate that conceiving a being with the same physical structure and functions but without conscious experience is possible, suggesting that conscious experience might not be entailed by physical facts alone. So why this option fits best: it captures the essence that the entity is physically indistinguishable from a human while lacking consciousness. The other descriptions describe beings that either are conscious, or are machines with emotions, or are purely fictional/dream-like figures, none of which match the specific claim about lacking subjective experience despite normal physical and behavioral properties.

**5. Compared to large language models, traditional n-gram models lack which capabilities?**

- A. Contextual windowing.**
- B. Ability to handle out-of-vocabulary words.**
- C. A neural network-based predictor.**
- D. Transformer-based architecture, large-scale training, and post-training capabilities.**

What sets large language models apart from traditional n-gram models is the combination of architecture, scale, and post-training adaptability they enable. Traditional n-gram models build probabilities from fixed-length sequences of words, relying on simple counts and smoothing. They operate within a limited context window and don't learn from data in the way neural models do, so they struggle to capture complex, long-range dependencies or adapt after deployment. Large language models use transformer-based architectures with attention, which lets them weigh information from far apart parts of the input and model dependencies across long sequences much more effectively. They're trained on massive datasets, giving them broad linguistic knowledge and nuanced pattern recognition that simple counts can't achieve. After the initial training, they can be fine-tuned, aligned with human preferences, or augmented with retrieval mechanisms, enabling post-training capabilities that adapt or improve the model over time. Therefore, the option mentioning transformer-based architecture, large-scale training, and post-training capabilities best captures what large language models offer that traditional n-gram models do not.

**6. What concern about AI-generated art is raised in the material?**

- A. Theft of labor due to artists not being compensated**
- B. Plagiarism as copying exact works**
- C. Privacy violations in data used for training**
- D. AI-generated art is always infringing**

The main idea here is that AI-generated art raises a concern about labor and fair compensation. When AI models are trained on vast collections of artwork created by humans, the question becomes whether and how artists should be compensated for the use of their labor in training these systems. The fear is that artists' work contributes to the capabilities of AI, but the creators aren't paid or credited in a way that reflects that contribution. That sense of "theft of labor" captures the ethical issue the material emphasizes: there's an economic and moral tension around who benefits from AI art and who pays the makers behind it. Other concerns—like whether a model copies an exact work and passes it off as original (plagiarism), or privacy issues related to data used for training—are related topics, but the material's focus is on the broader question of fair compensation and labor rights for artists. It's not claiming that AI-generated art is always infringing or that plagiarism is the sole problem, but rather that the central concern is the potential devaluation or nonpayment of human creative labor.

**7. Is a private conversation overheard by a friend in a public park a violation of privacy, according to this account?**

**A. No**

**B. Yes**

**C. Only if the overheard content is sensitive**

**D. Only if you intended privacy**

The situation tests how privacy expectations change with context. In a public park, conversations aren't clearly private, so there's a lower expectation of confidentiality. If a friend overhears a private chat in that setting, the account described treats it as not a privacy violation because the information was exposed in a public space and there's nothing shown about how the overheard content is used or disclosed beyond that moment. The idea is that privacy isn't just about whether something could be sensitive or whether you intended privacy; it's about where you are and what kind of expectation of privacy exists. So, the best answer is that there is no privacy violation in this account. The other options would imply additional conditions (sensitivity, or a difference in intent) that the account does not require for a violation.

**8. What is a recommended strategy for technology deployment to address ethical concerns, as discussed under 'Slow Development'?**

**A. Slowing release pace to allow adjustment and safety review.**

**B. Releasing updates as quickly as possible.**

**C. Bypassing safety evaluations to accelerate benefits.**

**D. Keeping releases secret until fully solved.**

Slow development means choosing a deployment pace that prioritizes safety, fairness, and accountability. Slowing the release gives time for thorough safety reviews, more extensive testing, and bias and privacy assessments, as well as opportunities to gather stakeholder input. This makes it possible to spot and fix potential harms before they affect a wide user base, and it supports stronger governance and oversight. When you scale up slowly, you can adjust designs and controls as you learn, reducing risk and building trust. Releasing updates quickly or bypassing safety checks would skip those protections, and keeping releases secret undermines accountability and the opportunity for external review and improvement.

## 9. What does RLHF stand for?

- A. Reinforcement Learning from Human Feedback**
- B. Randomized Language Handling Framework**
- C. Recursive Latent Heuristic Functions**
- D. Robust Linguistic Heuristic Formulation**

RLHF stands for Reinforcement Learning from Human Feedback. This approach uses human judgments to shape what the model should do, by feeding human preferences into the learning process. In practice, outputs are generated and humans provide feedback or comparisons, a reward model learns to predict those judgments, and the main model is fine-tuned with reinforcement learning to maximize that reward. This makes the model's behavior align more closely with what people want, improving usefulness and safety beyond what pure data-driven learning can achieve. The other phrases listed don't describe this well-established method and aren't recognized terms for aligning AI with human preferences.

## 10. What is the primary aim of utilitarianism?

- A. The greatest happiness for the greatest number.**
- B. To maximize wealth and power.**
- C. To follow rigid duties regardless of outcomes.**
- D. To maximize personal freedom without regard to others.**

Utilitarianism is a form of consequentialism that judges actions by their outcomes and aims to maximize overall happiness or welfare. The primary aim is to produce the greatest happiness for the greatest number, meaning we look at the net balance of good over bad for everyone affected and choose the action that increases total well-being. This view treats happiness (or welfare) as the key measure of value and requires impartial consideration of everyone's interests, not just our own. So a policy or action is good if, on balance, it leads to more happiness than unhappiness for the overall group, even if it sacrifices some individuals' interests. That's why this option is the best fit: it captures the goal of maximizing overall well-being for the most people. It's not focused on accumulating wealth or power, nor on following duties rigidly regardless of outcomes, nor on pursuing personal freedom without regard to others.

## 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](mailto:hello@examzify.com).**

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

**<https://ethicsofai.examzify.com>**

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

SAMPLE