

GARP RISK AND AI (RAI) PRACTICE EXAM (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://garprai.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 example is used to illustrate image generation from a text prompt among content modalities?**
 - A. GPT-3
 - B. DALL-E
 - C. Word2Vec
 - D. BERT

- 2. What do the model's coefficients measure in a multiple regression context?**
 - A. Dummy Variables
 - B. Multiple Regression Coefficients
 - C. Intercept
 - D. Non-Linear Terms

- 3. Which search strategy involves assuming the opponent minimizes your chances and uses the minimax principle?**
 - A. Adversarial Search
 - B. Combinational Explosion
 - C. Recursion
 - D. Brittleness

- 4. Linear regression is easily distorted by outliers, missing data, and measurement errors. Which term describes this data sensitivity?**
 - A. Data Sensitivity
 - B. Slope
 - C. Intercept
 - D. Dummy Variables

- 5. Which distance metric is defined as the straight-line distance in Euclidean space between two points?**
 - A. Manhattan distance
 - B. Chebyshev distance
 - C. Euclidean distance
 - D. Hamming distance

- 6. Which term denotes incorrectly flagging a negative case as positive?**
- A. False Positive (FP)
 - B. False Negative (FN)
 - C. Accuracy
 - D. Precision
- 7. Which metric summarizes the ROC curve into a single number that indicates separation quality?**
- A. Cleaning
 - B. Tokenization
 - C. Area Under the Curve
 - D. Area Under the Curve (AUC)
- 8. Making it easy for individuals to withdraw consent, access their data, or delete it.**
- A. Contextual Integrity
 - B. Control over Data
 - C. Right to be Forgotten
 - D. Explainability
- 9. What are the layers between the input and output called, which process information within a neural network?**
- A. Input Layer
 - B. Hidden Layers
 - C. Output Layer
 - D. Weights
- 10. What term refers to the connections between neurons that carry numbers and determine influence?**
- A. Input Layer
 - B. Hidden Layers
 - C. Weights
 - D. Output Layer

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which example is used to illustrate image generation from a text prompt among content modalities?

- A. GPT-3
- B. DALL-E**
- C. Word2Vec
- D. BERT

Generating an image from a text prompt shows how a model moves from language to a different content modality. DALL-E is built to take a descriptive text and produce a corresponding image, effectively bridging text and image domains. That makes it the example used to illustrate image generation from a text prompt. The other models focus on text in different ways: GPT-3 generates fluent prose, Word2Vec creates word embeddings to capture semantic relationships, and BERT models understand and contextualize text. So DALL-E uniquely demonstrates cross-modal generation from text to image.

2. What do the model's coefficients measure in a multiple regression context?

- A. Dummy Variables
- B. Multiple Regression Coefficients**
- C. Intercept
- D. Non-Linear Terms

In multiple regression, each coefficient represents the effect of one predictor on the outcome, holding all other predictors constant. It's the amount the predicted dependent variable changes for a one-unit increase in that predictor, with every other predictor kept fixed. This is a partial or conditional effect: it isolates the influence of that specific variable from the others in the model. Important complements: dummy-variable terms are used to encode categories, and their coefficients reflect differences from a reference group. The intercept is a separate term that estimates the expected outcome when all predictors are zero. Non-linear terms (like squares or interactions) introduce curvature or interactions, with their own coefficients describing those non-linear effects. The coefficients in a linear multiple regression specifically quantify the linear contribution of each predictor to the outcome, given the presence of the others.

3. Which search strategy involves assuming the opponent minimizes your chances and uses the minimax principle?

- A. Adversarial Search**
- B. Combinational Explosion
- C. Recursion
- D. Brittleness

Adversarial search models competitive environments where two players have opposing goals, such as in chess or checkers. It assumes the opponent will act to minimize your chances, and it uses the minimax principle to plan moves accordingly. In practice, you lay out a game tree of possible moves and responses. On your turn you pick the move that leads to the best outcome assuming the opponent will respond with the move that minimizes your eventual score. This step-by-step reasoning yields the minimax value for each position and guides the best move to choose. Alpha-beta pruning is a common enhancement that trims unnecessary branches to speed up the search without changing the result. The other terms aren't strategies for planning under an opponent's influence: combinational explosion describes the rapid growth of possibilities as a problem scales, recursion is a general technique for breaking problems into simpler ones, and brittleness refers to a system's lack of robustness.

4. Linear regression is easily distorted by outliers, missing data, and measurement errors. Which term describes this data sensitivity?

A. Data Sensitivity

B. Slope

C. Intercept

D. Dummy Variables

Data sensitivity describes how regression estimates react to imperfect data. Linear regression relies on clean data, and outliers, missing values, or measurement errors can pull the estimated line away from the true relationship, changing both the slope and intercept. So this term captures how fragile the model's results are to data quality. The slope and intercept are simply the estimated parameters of the line, not descriptors of how sensitive those estimates are to data issues. Dummy variables, meanwhile, are just a way to encode categorical predictors and don't address sensitivity to data quality. If you're concerned about data sensitivity, you'd look at robustness techniques and diagnostic measures that show how influential problematic observations are.

5. Which distance metric is defined as the straight-line distance in Euclidean space between two points?

A. Manhattan distance

B. Chebyshev distance

C. Euclidean distance

D. Hamming distance

Straight-line distance in ordinary Euclidean space is measured by the Euclidean distance. In two dimensions it follows the Pythagorean formula $\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$; in higher dimensions you take the square root of the sum of the squared differences across all coordinates. This represents the direct length of the line connecting the two points, regardless of any grid or axis-aligned paths. Other metrics describe distance in different ways: Manhattan distance sums the absolute differences along each axis, Chebyshev distance takes the largest single coordinate difference, and Hamming distance counts how many positions differ between equal-length vectors. So, the distance that embodies the straight-line concept is the Euclidean distance.

6. Which term denotes incorrectly flagging a negative case as positive?

A. False Positive (FP)

B. False Negative (FN)

C. Accuracy

D. Precision

In binary decision making, mislabeling a negative case as positive is a false positive. This error happens when the system or test indicates a positive result even though the true condition is negative, which can inflate the number of positives and affect metrics like the false positive rate and precision. In a confusion matrix, it's the instance where the predicted positive doesn't match the actual negative. The other terms don't describe this specific error: a false negative is when a positive case is incorrectly labeled negative; accuracy is the overall rate of correct predictions, not the nature of the misclassification; precision is the proportion of predicted positives that are truly positive, not the error type itself.

7. Which metric summarizes the ROC curve into a single number that indicates separation quality?

- A. Cleaning
- B. Tokenization
- C. Area Under the Curve
- D. Area Under the Curve (AUC)**

Discrimination quality in a ROC curve is captured by a single number called the area under the ROC curve, usually abbreviated as AUC. The ROC curve shows how the true positive rate grows as you trade off false positives by varying the threshold. By computing the area under that curve, you get a summary of how well the model separates positives from negatives across all possible thresholds. A higher AUC means stronger separation, with values between 0.5 (no better than random) and 1.0 (perfect separation). This threshold-independent measure makes it ideal for comparing models. The option that includes AUC is the correct one. The other terms, like cleaning or tokenization, are preprocessing steps and don't quantify separation quality.

8. Making it easy for individuals to withdraw consent, access their data, or delete it.

- A. Contextual Integrity
- B. Control over Data**
- C. Right to be Forgotten
- D. Explainability

The key idea here is giving individuals governance over their personal information—being able to withdraw consent, access their data, and have it deleted. This is best described by Control over Data, which covers managing consent, accessing data, and requesting deletion as part of one's overall data rights. Contextual Integrity deals with whether data flows align with social norms in a given context, not with user control. The Right to be Forgotten focuses mainly on deletion rather than the broader ability to control data throughout its lifecycle. Explainability is about making AI decisions understandable, not about data rights. So, the most comprehensive concept that fits the described ability is Control over Data.

9. What are the layers between the input and output called, which process information within a neural network?

- A. Input Layer
- B. Hidden Layers**
- C. Output Layer
- D. Weights

Hidden layers are the parts of a neural network that sit between the input and the output and do the processing. They transform the raw data as it flows through the network, applying linear combinations of inputs using weights and biases, then passing the result through nonlinear activation functions. This layering lets the model build increasingly abstract representations of the data, from simple patterns in early hidden layers to more complex features in deeper layers. The input layer simply provides the initial data, and the output layer produces the final prediction, so they are not doing the intermediate processing. Weights are the adjustable parameters inside the connections that the network learns during training; they aren't layers themselves.

10. What term refers to the connections between neurons that carry numbers and determine influence?

- A. Input Layer
- B. Hidden Layers
- C. Weights**
- D. Output Layer

In a neural network, the connections between neurons are defined by weights—the numerical values assigned to each connection. These weights determine how strongly a signal from one neuron influences the next one. A larger weight means a stronger influence, while a smaller weight means a weaker impact. During training, the network adjusts these weights to reflect how important each input is for the task, shaping the overall behavior and predictions. Think of it this way: when a neuron fires, the inputs it receives are multiplied by their respective weights before being summed and passed through an activation function. The weights are the numbers that carry influence through the network, encoding which inputs matter most for the final output. The other terms describe different parts of the network: the input layer simply accepts data, hidden layers perform computations, and the output layer produces the result.

SAMPLE

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://garprai.examzify.com>

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

SAMPLE