


HUAWEI CERTIFIED ICT ASSOCIATE – ARTIFICIAL INTELLIGENCE (HCIA-AI) PRACTICE EXAM (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	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. What is the missing phase in the three phases of AI: computing intelligence, _____, cognitive intelligence?
 - A. Enhanced intelligence
 - B. Behavioral intelligence
 - C. Perceptual intelligence
 - D. Emotional intelligence
2. In the code `SVR(kernel='rbf', C=1e3, gamma=0.1)`, does it mean we use linear kernel functions?
 - A. True
 - B. False
 - C. Conditional true
 - D. Not applicable
3. What is reinforcement learning primarily focused on?
 - A. Minimizing training data
 - B. Maximizing data preprocessing
 - C. Learning through rewards and penalties
 - D. Creating large neural networks
4. What is the primary function of autoencoders in unsupervised learning?
 - A. To classify data into labeled categories
 - B. To learn efficient representations for data
 - C. To predict future outcomes based on historical data
 - D. To generate synthetic data samples
5. What does the ROC curve illustrate in machine learning?
 - A. The relationship between data preprocessing and feature selection
 - B. The trade-off between true positive rate and false positive rate for a binary classifier
 - C. The performance of regression models against benchmarks
 - D. The distribution of data points across different classes

- 6. Which concept should be prioritized to achieve true AI capabilities?**
- A. Traditional computing methods
 - B. Data-driven machine learning approaches
 - C. Simple algorithmic processes
 - D. Static programming techniques
- 7. What is the goal of using decision trees in supervised learning?**
- A. To minimize error in predictions
 - B. To visualize data relationships
 - C. To create random outputs
 - D. To eliminate the need for data
- 8. Which of the following intelligences did Howard Gardner propose?**
- A. Mathematical, Linguistic, Spatial, Musical
 - B. Space, Logic, Language, Music
 - C. Emotional, Practical, Traditional, Creative
 - D. Visual, Analytical, Abstract, Theoretical
- 9. Which technology innovation is not employed by MindSpore to enhance AI applications?**
- A. New way of cooperation
 - B. New programming paradigm
 - C. New execution paradigm
 - D. New programming language
- 10. What does the term "loss function" signify in model training?**
- A. It measures the model's runtime
 - B. It evaluates the model's performance
 - C. It signifies the complexity of the model
 - D. It quantifies the model's data inputs

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is the missing phase in the three phases of AI: computing intelligence, _____, cognitive intelligence?

- A. Enhanced intelligence
- B. Behavioral intelligence
- C. Perceptual intelligence**
- D. Emotional intelligence

The correct answer, which is perceptual intelligence, fits seamlessly within the three phases of AI: computing intelligence, perceptual intelligence, and cognitive intelligence. Perceptual intelligence refers to the ability of AI systems to interpret and understand sensory data, such as images, sounds, and other forms of input that resemble human perception. In this phase, AI focuses on sensory input processing and the extraction of meaningful information from that data. This is an essential step in AI development, as it allows machines to perceive their environment in a way that mimics human sensory capabilities, effectively leading to more complex decision-making processes. The concept of perceptual intelligence encompasses technologies like computer vision, speech recognition, and other modalities that enable machines to interact and understand their surroundings. This foundational phase establishes the groundwork for the subsequent cognitive intelligence phase, where the AI applies learned knowledge and reasoning skills to make decisions or solve problems. While enhanced intelligence, behavioral intelligence, and emotional intelligence are important concepts in the broader context of artificial intelligence, they do not represent a specific connectivity as a 'missing' phase in the sequence presented. Enhanced intelligence generally describes the improvement in AI capabilities, behavioral intelligence offers insights into actions taken based on learned behaviors, and emotional intelligence relates to understanding and responding to human emotions.

2. In the code `SVR(kernel='rbf', C=1e3, gamma=0.1)`, does it mean we use linear kernel functions?

- A. True
- B. False**
- C. Conditional true
- D. Not applicable

The use of the 'rbf' kernel in the provided Support Vector Regression (SVR) code indicates that a Radial Basis Function kernel is being employed rather than a linear kernel function. The 'rbf' kernel is nonlinear and is commonly used for problems where the relationship between the features and the target variable is not linearly separable. Setting the kernel to 'linear' would involve a different specification within the code. In contrast, the 'rbf' kernel function allows the model to handle more complex patterns in the data. The parameters 'C' and 'gamma' in the SVR function play significant roles in controlling the regularization and the influence of each training example, but they do not change the fact that a linear kernel is not being used in this scenario. Therefore, stating that the code does mean we use linear kernel functions is incorrect; the presence of 'rbf' clearly designates it as a nonlinear kernel. This reinforces the choice that the assertion is false.

3. What is reinforcement learning primarily focused on?

- A. Minimizing training data
- B. Maximizing data preprocessing
- C. Learning through rewards and penalties**
- D. Creating large neural networks

Reinforcement learning is a type of machine learning that emphasizes learning through interactions with an environment. In this framework, an agent learns to make decisions by taking actions and receiving feedback in the form of rewards or penalties. The goal is to maximize the cumulative reward over time, which drives the agent to adjust its behavior based on past experiences. This approach is distinct as it involves sequential decision-making where the consequences of actions unfold over time, and the agent must balance exploration (trying new actions to discover their effects) with exploitation (choosing actions known to yield high rewards). The focus on rewards and penalties is what differentiates reinforcement learning from other learning paradigms. The other options do not align with the primary focus of reinforcement learning. Minimizing training data, maximizing data preprocessing, and creating large neural networks are more associated with different aspects of machine learning and neural network design, rather than with the core principles that guide reinforcement learning.

4. What is the primary function of autoencoders in unsupervised learning?

- A. To classify data into labeled categories
- B. To learn efficient representations for data**
- C. To predict future outcomes based on historical data
- D. To generate synthetic data samples

The primary function of autoencoders in unsupervised learning is to learn efficient representations for data. Autoencoders are a type of neural network designed to learn a compressed encoding of input data. They work by passing the data through an encoder that reduces the dimensionality, followed by a decoder that attempts to reconstruct the original data from the encoded representation. This process allows the autoencoder to capture the essential features of the data while discarding noise and redundancy. Through this mechanism, autoencoders can uncover the underlying structure of the data, making them useful for tasks like dimensionality reduction, feature extraction, and anomaly detection. By learning to represent the data in a lower-dimensional space, they can facilitate the understanding and analysis of complex datasets without requiring labeled information. Other choices pertain to functions that are not the primary focus of autoencoders. Classifying data into labeled categories involves supervised learning techniques, while predicting future outcomes is typically associated with forecasting models. Generating synthetic data samples aligns more with generative models rather than the main function of autoencoders, which centers on representation learning.

5. What does the ROC curve illustrate in machine learning?

- A. The relationship between data preprocessing and feature selection
- B. The trade-off between true positive rate and false positive rate for a binary classifier**
- C. The performance of regression models against benchmarks
- D. The distribution of data points across different classes

The ROC curve, or Receiver Operating Characteristic curve, is a crucial concept in evaluating the performance of binary classifiers. It illustrates the trade-off between the true positive rate (sensitivity) and the false positive rate (1-specificity) at various threshold settings. By plotting these rates on a graph, the ROC curve allows practitioners to visualize how changing the threshold for classifying a positive instance affects both the ability to correctly identify positive instances and the likelihood of incorrectly labeling negative instances as positive. The area under the ROC curve (AUC) provides a single scalar value that summarizes the overall performance of the model across all classification thresholds. A model with a higher AUC indicates better performance in distinguishing between the positive and negative classes. The other options do not accurately represent the purpose of the ROC curve. The relationship between data preprocessing and feature selection pertains to data preparation techniques, while the performance of regression models is evaluated through different metrics, such as R-squared or mean squared error. Lastly, the distribution of data points across classes relates more to visualization techniques like scatter plots rather than the ROC curve itself, which focuses specifically on binary classification outcomes and their trade-offs.

6. Which concept should be prioritized to achieve true AI capabilities?

- A. Traditional computing methods
- B. Data-driven machine learning approaches**
- C. Simple algorithmic processes
- D. Static programming techniques

To achieve true AI capabilities, prioritizing data-driven machine learning approaches is crucial as it emphasizes the importance of data in training models that can learn and adapt from patterns within that data. Unlike traditional computing methods, which rely heavily on pre-defined rules and logic, machine learning utilizes algorithms that can automatically improve their performance as they are exposed to more data. This adaptability enables AI systems to handle complex tasks and make predictions or decisions based on new, unseen input. Data-driven approaches foster the ability to uncover insights and relationships within large datasets, allowing for greater generalization and accuracy in various applications such as natural language processing, image recognition, and predictive analytics. The effectiveness of AI relies on its ability to analyze vast amounts of data, derive knowledge, and continuously improve through learning. In contrast, relying on simple algorithmic processes or static programming techniques would limit the potential of AI because these methods do not leverage the power of data learning, making them less effective for dynamic and complex scenarios that modern AI applications encounter.

7. What is the goal of using decision trees in supervised learning?

A. To minimize error in predictions

B. To visualize data relationships

C. To create random outputs

D. To eliminate the need for data

The primary goal of using decision trees in supervised learning is to minimize error in predictions. Decision trees are predictive models that use a tree-like structure to represent decisions and their possible consequences. Each node of the tree represents a feature (or attribute), and each branch represents a decision rule, leading to the final output which is the prediction outcome. In the context of supervised learning, decision trees aim to accurately classify data points or predict outcomes based on the patterns learned from the training data. The model is trained on labeled data, meaning it learns from input-output pairs, adjusting its structure to minimize the difference between the predicted outcomes and the actual outcomes. This is often referred to as reducing prediction error. Minimizing error in predictions is crucial in supervised learning, as the effectiveness of a model is typically assessed based on how well it performs on unseen data (validation/test sets) after being trained. Decision trees achieve this by recursively splitting the data based on the feature values that result in the highest information gain, effectively creating rules that capture the underlying relationships in the data. Other potential objectives, such as visualizing data relationships, are secondary outcomes of decision trees but not their primary purpose in a supervised learning context.

8. Which of the following intelligences did Howard Gardner propose?

A. Mathematical, Linguistic, Spatial, Musical

B. Space, Logic, Language, Music

C. Emotional, Practical, Traditional, Creative

D. Visual, Analytical, Abstract, Theoretical

Howard Gardner proposed the Theory of Multiple Intelligences, which identifies various types of intelligence that individuals may possess. Among these intelligences, the categories highlighted include linguistic intelligence, logical-mathematical intelligence, spatial intelligence, musical intelligence, and others. The correct choice reflects a specific interpretation of Gardner's framework that includes categories such as space (spatial), logic (logical-mathematical), language (linguistic), and music (musical). This grouping aligns with Gardner's emphasis on recognizing and valuing different ways in which people learn, think, and interact with the world. Each type of intelligence represents a distinct way of processing information and allows individuals to excel in varied fields. The other choices do not encapsulate Gardner's original theory as accurately, mixing elements that do not directly correspond to his established categories. For example, while emotional intelligence has been popularized, it does not directly correspond with Gardner's initial classifications. Similarly, terms like practical and traditional do not feature explicitly in the framework he developed, which primarily focuses on distinct forms of cognitive ability rather than broader or more subjective categories.

9. Which technology innovation is not employed by MindSpore to enhance AI applications?

- A. New way of cooperation
- B. New programming paradigm
- C. New execution paradigm
- D. New programming language**

MindSpore is an open-source deep learning framework developed by Huawei, and it utilizes several key innovations to enhance AI applications. One of the foundational aspects of MindSpore is its design to support a new way of cooperation, allowing for more efficient collaboration between various components and systems within AI applications. Additionally, MindSpore introduces a new programming paradigm that emphasizes the separation of model definition from execution, which facilitates easier debugging and optimization for AI development. Alongside this, it employs a new execution paradigm that optimizes performance by adapting to various hardware environments, thus ensuring that applications can run efficiently on different devices. While MindSpore does support various programming languages and can interact with them, it does not specifically offer a new programming language as part of its core framework. Instead, it builds upon existing languages and tools to enhance AI application development. This focus on leveraging established programming languages, rather than creating an entirely new one, differentiates it from other frameworks that might introduce novel languages.

10. What does the term "loss function" signify in model training?

- A. It measures the model's runtime
- B. It evaluates the model's performance**
- C. It signifies the complexity of the model
- D. It quantifies the model's data inputs

The term "loss function" is a critical concept in machine learning and model training. It serves as a quantitative measure of how well a model's predictions align with the actual outcomes. Essentially, the loss function evaluates the model's performance by computing the difference between the predicted values generated by the model and the true values from the training data. This assessment is vital because it provides feedback to the model about how far off its predictions are, guiding the optimization process in adjusting the model parameters to improve accuracy. A well-defined loss function helps in training the model efficiently, as it directs the learning algorithm on how to minimize errors during predictions. By continuously optimizing this function, the model learns to make better predictions over time, enhancing its performance on unseen data.

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://huaweicertifiedictassociate-ai.examzify.com>

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

SAMPLE