


HUAWEI CERTIFIED ICT ASSOCIATE – ARTIFICIAL INTELLIGENCE (HCIA-AI) PRACTICE EXAM (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://huaweicertifiedictassociate-ai.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	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 technique is commonly used in supervised learning?**
 - A. Using data without labels to train
 - B. Applying algorithms based solely on probabilities
 - C. Training a model with labeled datasets
 - D. Finding hidden structures in unlabelled data
- 2. What purpose does the Huawei HiAI Foundation serve regarding existing AI models?**
 - A. To create new models
 - B. To quickly convert and port them
 - C. To eliminate previous models
 - D. To enhance model efficiency
- 3. How do chatbots primarily improve customer support services?**
 - A. By handling complex inquiries with human-like responses
 - B. By providing quick responses to common queries
 - C. By replacing human agents with automated systems
 - D. By managing customer accounts and transactions
- 4. Define 'big data' in the context of AI.**
 - A. Small datasets easily handled by traditional applications
 - B. Large and complex datasets that traditional data processing applications cannot handle efficiently
 - C. Files stored in cloud platforms
 - D. Structured data stored in relational databases
- 5. A loss function reflects the error between a target output and an actual output in a neural network. Which of the following is a common loss function in deep learning?**
 - A. Means square loss function
 - B. Logarithmic loss function
 - C. Exponential loss function
 - D. Hinge loss function

- 6. Which component of AI is responsible for making decisions based on learned data?**
- A. Reasoning
 - B. Perception
 - C. Learning
 - D. Data Analysis
- 7. Which of the following best describes the execution capabilities of MindSpore?**
- A. It specializes only in non-AI tasks.
 - B. It can execute AI tasks on various hardware platforms.
 - C. It can only execute on high-end supercomputers.
 - D. It is limited to TensorFlow compatibility.
- 8. Which of the following are considered regularisation technologies?**
- A. Dropout
 - B. Momentum optimiser
 - C. L1 regularisation
 - D. L2 regularisation
- 9. Which steps belong to the TensorFlow development process?**
- A. Model testing and evaluation
 - B. Model restoration and invoking
 - C. Data analysis
 - D. Model visualization
- 10. What distinguishes a Generative Adversarial Network (GAN)?**
- A. It employs a single neural network model
 - B. It generates data through competition between two networks
 - C. It focuses solely on supervised learning
 - D. It is limited to classification tasks

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which technique is commonly used in supervised learning?

- A. Using data without labels to train
- B. Applying algorithms based solely on probabilities
- C. Training a model with labeled datasets**
- D. Finding hidden structures in unlabelled data

The technique commonly used in supervised learning involves training a model with labeled datasets. In supervised learning, the learning process operates under the premise that each training example is paired with an output label, allowing the model to learn the relationship between input data and the corresponding outputs. This is fundamental to developing models that can make predictions or classifications based on new, unseen data. The presence of labels provides a clear direction for the learning algorithm to adjust its parameters, optimizing its ability to generalize and make accurate predictions. In contrast, the other techniques mentioned are more aligned with different learning paradigms. Using data without labels to train refers to unsupervised learning, where the goal is to identify patterns without the guidance of labeled outcomes. Applying algorithms based solely on probabilities is more characteristic of probabilistic models but does not inherently imply the use of supervised techniques. Lastly, finding hidden structures in unlabelled data is another hallmark of unsupervised learning, focusing on extracting insights from datasets lacking explicit labels. Thus, training with labeled datasets is the defining feature of supervised learning, making it the correct choice.

2. What purpose does the Huawei HiAI Foundation serve regarding existing AI models?

- A. To create new models
- B. To quickly convert and port them**
- C. To eliminate previous models
- D. To enhance model efficiency

The Huawei HiAI Foundation's primary objective is to quickly convert and port existing AI models, enabling a seamless transition of these models to Huawei's AI platform. This process involves adaptively modifying existing models so that they can leverage the infrastructure and tools provided by Huawei, which may involve adjustments for performance optimization, compatibility, and functionality improvements. By focusing on the quick conversion and porting of existing models, the HiAI Foundation allows developers and organizations to efficiently utilize their pre-existing AI work without the need to start from scratch. This capability fosters innovation and speed in deploying AI applications, making it easier for developers to implement AI solutions on Huawei's devices and services. The other options do not accurately reflect the primary role of the HiAI Foundation; it does not focus on creating new models or eliminating previous ones, nor is its chief aim to enhance model efficiency independently, even though improved efficiency may result as a benefit of the conversion process.

3. How do chatbots primarily improve customer support services?

- A. By handling complex inquiries with human-like responses
- B. By providing quick responses to common queries**
- C. By replacing human agents with automated systems
- D. By managing customer accounts and transactions

Chatbots primarily enhance customer support services by providing quick responses to common queries. This efficiency allows businesses to address frequently asked questions promptly, improving customer satisfaction as users can receive immediate assistance without the need for human intervention. The ability of chatbots to instantly provide information on standard inquiries, such as business hours, product details, or troubleshooting steps, significantly reduces wait times and alleviates the pressure on human agents who can then focus on more complex issues that require personalized attention. In this context, while handling complex inquiries with human-like responses can be beneficial, it is not the primary function of most chatbots, which are typically designed to manage simpler tasks. Additionally, while the idea of replacing human agents with automated systems suggests a potential future direction of customer service, the current focus of chatbots is more on enhancing efficiency rather than complete replacement. Lastly, managing customer accounts and transactions, while a capability of certain systems, falls outside the typical scope of what most chatbots do in terms of general customer support.

4. Define 'big data' in the context of AI.

- A. Small datasets easily handled by traditional applications
- B. Large and complex datasets that traditional data processing applications cannot handle efficiently**
- C. Files stored in cloud platforms
- D. Structured data stored in relational databases

'Big data' refers to large and complex datasets that exceed the capabilities of traditional data processing applications to handle efficiently. In the context of artificial intelligence, big data is critical because it encompasses vast amounts of information generated from various sources, including social media, sensors, transactional data, and more. This enormous volume of data can be structured, semi-structured, or unstructured. The ability of AI systems to learn from and analyze these large datasets allows them to uncover patterns, trends, and insights that would not be apparent through conventional methods. For instance, machine learning models often rely on big data to improve their accuracy and predictive capabilities, making them more effective in real-world applications. In contrast, small datasets and structured data confined to relational databases do not represent the diverse and complex nature of big data. Additionally, big data often requires advanced frameworks and technologies, such as distributed computing and data lakes, to effectively store, process, and analyze the information.

5. A loss function reflects the error between a target output and an actual output in a neural network. Which of the following is a common loss function in deep learning?

A. Means square loss function

B. Logarithmic loss function

C. Exponential loss function

D. Hinge loss function

The mean squared loss function is commonly used in deep learning, especially in regression problems. This function calculates the average of the squares of the errors—that is, the difference between the predicted values produced by the model and the actual target values. The squaring of errors ensures that larger errors are penalized more heavily, which is beneficial for model training as it encourages the network to make more accurate predictions. The mean squared loss function is particularly favored when the relationship between the input and output can be modeled using continuous values, making it suitable for tasks such as predicting numerical values. This function provides a smooth gradient, which is vital for optimization algorithms like gradient descent, allowing the neural network to converge more effectively during training. While other loss functions like the logarithmic loss, exponential loss, and hinge loss each serve specific purposes in different contexts (such as classification tasks or handling specific types of errors), the mean squared loss function stands out as a foundational approach in deep learning for continuous output predictions.

6. Which component of AI is responsible for making decisions based on learned data?

A. Reasoning

B. Perception

C. Learning

D. Data Analysis

The component of AI responsible for making decisions based on learned data is reasoning. Reasoning involves the process of drawing conclusions and making decisions based on the patterns and insights gleaned from the data that the AI system has processed. This aspect of AI utilizes the rules and models developed during the learning phase to apply logical thought to situations, allowing the system to assess information and decide on the best course of action. While learning is critical for developing the model through which decisions can be made, it is specifically the reasoning aspect that allows the AI to apply this learned knowledge in practical situations. Perception primarily pertains to how AI systems interpret and understand the data they receive, and data analysis generally encompasses the techniques used to examine and organize data. However, none of these components focus solely on the decision-making aspect as reasoning does. Thus, reasoning stands out as the correct answer regarding the component of AI that leverages learned data to make decisions.

7. Which of the following best describes the execution capabilities of MindSpore?

- A. It specializes only in non-AI tasks.
- B. It can execute AI tasks on various hardware platforms.**
- C. It can only execute on high-end supercomputers.
- D. It is limited to TensorFlow compatibility.

MindSpore is a deep learning framework developed by Huawei that is designed to execute AI tasks efficiently across multiple hardware platforms, which is the hallmark of modern AI execution capabilities. This versatility is critical for developers and researchers who wish to leverage various types of hardware, including CPUs, GPUs, and even specialized hardware like Ascend chips, to optimize performance and resource utilization. By supporting various hardware setups, MindSpore allows for flexibility in deployment, making it accessible for different use cases, from edge computing to cloud environments. This broad compatibility is crucial for achieving scalability and efficiency in AI applications, as different hardware can be tailored to specific workloads, enhancing overall operational effectiveness. Other options, such as focusing solely on non-AI tasks, being limited to supercomputers, or restricting operations to compatibility with TensorFlow, do not accurately reflect the comprehensive nature of MindSpore's intended functionality. The framework is explicitly built to facilitate and enhance AI tasks across a wide range of platforms, underscoring the correct choice.

8. Which of the following are considered regularisation technologies?

- A. Dropout**
- B. Momentum optimiser
- C. L1 regularisation
- D. L2 regularisation

The correct choice encompasses several techniques aimed at preventing overfitting during the training of machine learning models. Regularization aims to impose a penalty on the model's complexity, helping it generalize better to unseen data. Dropout is a specific method used in neural networks where, during training, random neurons are "dropped out" or deactivated, which prevents the network from becoming overly reliant on any single neuron. This forces the model to learn more robust representations and helps in mitigating overfitting. By not allowing certain neurons to participate in every iteration, dropout introduces noise into the training process, leading the model to be more adaptable and better able to generalize. L1 and L2 regularization are also common regularization techniques that add a penalty to the loss function based on the magnitude of the weights. L1 regularization encourages sparsity in the weights, effectively shrinking some of them to zero, while L2 regularization penalizes the square of the weights' magnitudes, leading to smaller weights overall. Both contribute to reducing model complexity but differ in their approach. Momentum optimizer, on the other hand, is a technique used to accelerate the convergence of gradient descent algorithms rather than a regularization method. While it helps improve training speed and stability,

9. Which steps belong to the TensorFlow development process?

- A. Model testing and evaluation
- B. Model restoration and invoking**
- C. Data analysis
- D. Model visualization

The TensorFlow development process includes several key steps that are essential for building and deploying machine learning models effectively. Among these steps, model restoration and invoking play a crucial role. In the context of TensorFlow, model restoration refers to loading a previously trained model or checkpoint, which allows developers to resume training or make predictions without having to retrain the model from scratch. This step is critical for efficiency, particularly in scenarios where training on large datasets or complex models requires significant computational resources. Invoking the model involves using it to perform inference on new data, which is the ultimate goal of developing any machine learning model. After a model has been trained and potentially fine-tuned, it needs to be applied to real-world scenarios, and invoking facilitates this application. While model testing and evaluation, data analysis, and model visualization are important steps in the overall machine learning workflow, they do not specifically pertain to the direct processes of restoring and running a model after it has been created. Instead, they help in the preparation, assessment, and understanding of the models throughout their lifecycle. Thus, the focus on model restoration and invoking is critical as it ensures that the developed models can be efficiently retrained and applied to new data, marking essential elements of the TensorFlow development process

10. What distinguishes a Generative Adversarial Network (GAN)?

- A. It employs a single neural network model
- B. It generates data through competition between two networks**
- C. It focuses solely on supervised learning
- D. It is limited to classification tasks

A Generative Adversarial Network (GAN) is distinguished by its unique architecture that involves the competition between two neural networks: the generator and the discriminator. The generator's role is to create new data instances that resemble the training data, while the discriminator evaluates the authenticity of the generated data against real data. This adversarial process drives the generator to improve its output based on the feedback from the discriminator, leading to the generation of realistic data over time. This competition is what sets GANs apart from other neural network models, as it creates a dynamic where both networks are continually learning and adapting to each other's improvements. The generator aims to produce data that the discriminator cannot distinguish from real data, while the discriminator is tasked with improving its ability to tell real data from fake. This interplay is fundamental to the functionality of GANs and is responsible for their success in generating high-quality synthetic data in various domains, such as image and video generation. The other choices do not accurately describe the nature of GANs. A single neural network model does not apply, as there are two networks involved. GANs can operate in both supervised and unsupervised contexts, thus they do not focus solely on supervised learning. Furthermore, GANs are not limited to classification tasks;

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