


AAISM DOMAIN 2 PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://aaismdomain2.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 NAS component focuses on training and validating candidate models to compare and select the best one with cost efficiency in mind?**
 - A. Performance estimation strategy
 - B. Search strategy
 - C. Training procedure
 - D. Model selection
- 2. Which term describes a neural network architecture where information moves forward in a single pass and is commonly used for image classification?**
 - A. Input Layer
 - B. Association rules
 - C. Feedforward neural network (FNN)
 - D. Hidden Layer
- 3. Which term describes risk where data used to train or validate the model can be manipulated to embed backdoors or biases?**
 - A. Improper output handling
 - B. Data and model poisoning
 - C. Unbounded consumption
 - D. System prompt leakage
- 4. A system that manipulates human behavior in a way that could cause severe harm is categorized as which risk?**
 - A. Low/No Risk AI System
 - B. High-Risk AI Systems
 - C. Limited Risk AI Systems
 - D. Unacceptable Risk AI Systems
- 5. Which method predicts the probability of an event based on independent variables and is commonly used for binary classification?**
 - A. Linear Regression
 - B. K-Means Clustering
 - C. Logistic Regression
 - D. Support Vector Machines

- 6. Which technique reduces high-dimensional data to a lower-dimensional representation while preserving essential information?**
- A. Regression
 - B. Association rules
 - C. Input Layer
 - D. Dimensionality reduction
- 7. ML algorithms use which learning methods?**
- A. Supervised learning only
 - B. Supervised or Unsupervised learning
 - C. Reinforcement learning only
 - D. Unsupervised learning only
- 8. Which practice involves evaluating the AI solution in a real-world-like environment prior to full deployment?**
- A. Scenario analysis
 - B. Piloting
 - C. Bias checks
 - D. Data privacy impact assessment
- 9. Which topic transforms high-dimensional data into a lower-dimensional space to simplify the data while retaining its essential features?**
- A. Neural networks (NNs)
 - B. Input Layer
 - C. Dimensionality reduction
 - D. Feedforward neural network (FNN)
- 10. Data subjects can gain access to data that organizations have about them, including the ability to receive a copy of that information or direct access to the organization's database. Which right does this describe?**
- A. Right to access
 - B. Right to be informed
 - C. Right to rectify
 - D. Right to be forgotten

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which NAS component focuses on training and validating candidate models to compare and select the best one with cost efficiency in mind?

A. Performance estimation strategy

B. Search strategy

C. Training procedure

D. Model selection

The idea being tested is how NAS assesses many candidate architectures quickly and cost-effectively. That role belongs to the performance estimation strategy. It's about training and validating each candidate in a lightweight way to get a comparable measure of performance without fully training every model to convergence. By using proxies, partial training, weight sharing, early stopping, or learned surrogates, you can estimate how good an architecture will be while controlling compute and time, which is essential when you're comparing many options. The other components serve different purposes. The search strategy determines how you explore the space of architectures to propose candidates. The training procedure defines how a given architecture is trained. The model selection step happens after evaluation, deciding which architecture to deploy based on the estimated performance.

2. Which term describes a neural network architecture where information moves forward in a single pass and is commonly used for image classification?

A. Input Layer

B. Association rules

C. Feedforward neural network (FNN)

D. Hidden Layer

In a feedforward neural network, information moves in one direction—from the input through any hidden layers to the output—with no cycles or loops. This single-pass flow means each input is processed through a fixed sequence of computations to produce a result, which is ideal for tasks like image classification where you map image features to a class probability in one forward pass. For image classification, this structure is powerful because the network can learn hierarchical representations: early layers detect simple features, later layers combine them into more complex concepts, and the final layer assigns a class. While modern image classifiers often use convolutional structures to handle image data efficiently, those networks remain feedforward in that data flows forward through layers without feedback. The other terms don't describe this forward-only architecture. The input layer is just the starting point, not the entire forward-pass design. A hidden layer is one component within many networks, not the overall architectural property. Association rules are a data-mining concept unrelated to neural network architectures.

3. Which term describes risk where data used to train or validate the model can be manipulated to embed backdoors or biases?

- A. Improper output handling
- B. Data and model poisoning**
- C. Unbounded consumption
- D. System prompt leakage

Data poisoning is the risk where the data used to train or validate a model is tampered with to embed backdoors or biases. By injecting carefully crafted examples into the training set, an attacker can make the model learn hidden behaviors that only activate under specific triggers, effectively creating a backdoor. Poisoned validation data can also skew evaluation results and encourage biased or unsafe model behavior once deployed, because the model has learned from misleading signals. Other terms describe different issues: improper output handling concerns how outputs are managed or sanitized, unbounded consumption relates to resource or data ingestion limits, and system prompt leakage involves hidden instructions or prompts leaking into the model's behavior.

4. A system that manipulates human behavior in a way that could cause severe harm is categorized as which risk?

- A. Low/No Risk AI System
- B. High-Risk AI Systems
- C. Limited Risk AI Systems
- D. Unacceptable Risk AI Systems**

When evaluating AI risk classes, a system that manipulates human behavior in a way that could cause severe harm is considered unacceptable risk. This category covers scenarios where the harm to people's safety or autonomy is so serious that deployment is not justified, even with safeguards. Unacceptable risk means the system should not be allowed to operate. This differs from high-risk systems, which may be deployed only under strict controls and risk management; limited risk and low/no risk describe progressively less severe or easier-to-manage threats with lighter or no required safeguards. In the described case, the potential for severe harm from behavioral manipulation clearly places it in the unacceptable risk category.

5. Which method predicts the probability of an event based on independent variables and is commonly used for binary classification?

- A. Linear Regression
- B. K-Means Clustering
- C. Logistic Regression**
- D. Support Vector Machines

This question is about modeling the probability of a binary event using a method that yields probability estimates based on several input features. The best fit is logistic regression. It takes a linear combination of the predictors, $z = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots$, and maps it through a sigmoid function to produce a probability: $P(Y=1|X) = 1 / (1 + \exp(-z))$. The coefficients are typically estimated by maximum likelihood, which aligns the model's predicted probabilities with the observed outcomes. Logistic regression is ideal for binary classification because its output is always between 0 and 1, making it easy to interpret as a probability. It works with multiple predictors (continuous or categorical, once encoded), and the coefficients tell you how each predictor changes the odds of the event. You can then choose a decision threshold (commonly 0.5, but adjustable) to assign class labels. In contrast, linear regression predicts a continuous outcome and can give values outside the 0–1 range, which isn't meaningful for probability. K-means clustering is an unsupervised method for grouping data, not for predicting event probabilities. Support vector machines focus on creating a decision boundary; while probabilistic outputs can be obtained with additional calibration, SVMs aren't inherently modeling probabilities in the straightforward way logistic regression does.

6. Which technique reduces high-dimensional data to a lower-dimensional representation while preserving essential information?

- A. Regression
- B. Association rules
- C. Input Layer
- D. Dimensionality reduction**

Dimensionality reduction is about mapping data from a high-dimensional space to a lower-dimensional one while keeping the important structure and information intact. It works by identifying the most informative directions or representations—often the axes that capture the most variance or preserve meaningful relationships—so the essence of the data isn't lost even as you reduce its size. This helps with visualization, reduces noise and overfitting, and makes computation more efficient. Methods range from linear approaches like PCA to nonlinear ones like t-SNE or autoencoders, each aiming to keep the core patterns and structure of the data. Regression, on the other hand, focuses on predicting a target variable and doesn't inherently compress the data into a smaller, information-rich representation. Association rules uncover patterns of co-occurring items, not a reduced feature space. The input layer is simply a stage in a neural network that receives data; it isn't a technique by itself for reducing dimensionality.

7. ML algorithms use which learning methods?

- A. Supervised learning only
- B. Supervised or Unsupervised learning**
- C. Reinforcement learning only
- D. Unsupervised learning only

Machine learning methods are typically grouped into supervised and unsupervised learning. Supervised learning uses labeled data to learn a mapping from inputs to outputs, which is why it's used for tasks like classification and regression. Unsupervised learning works with unlabeled data to uncover structure, such as clusters or latent patterns, and includes techniques like clustering and dimensionality reduction. Some approaches blend these ideas or adapt to different data setups, but reinforcement learning is a separate paradigm where an agent learns by interacting with an environment and receiving rewards, rather than fitting to labeled targets or discovering structure in unlabeled data. So, the statement that machine learning uses supervised or unsupervised learning captures the main ways many ML problems are approached.

8. Which practice involves evaluating the AI solution in a real-world-like environment prior to full deployment?

- A. Scenario analysis
- B. Piloting**
- C. Bias checks
- D. Data privacy impact assessment

Testing an AI solution in a setting that mirrors real use before rolling it out widely is piloting. This approach puts the AI into a limited, real-world-like environment with actual users, data, and workflows, so you can observe how it performs in realistic conditions, not just in a lab or mock setup. You'll see how it handles real inputs, integrates with current systems, and what issues surface once it's exposed to everyday use. This hands-on exposure helps you measure real-world performance, catch operational or usability problems, and understand what support, monitoring, or adjustments are needed before a full deployment. Other practices focus on different aims: imagining how the system might respond to hypothetical scenarios, checking for fairness and bias, or assessing privacy and data protection implications, rather than validating behavior in a live-like operational setting.

9. Which topic transforms high-dimensional data into a lower-dimensional space to simplify the data while retaining its essential features?

- A. Neural networks (NNs)
- B. Input Layer
- C. Dimensionality reduction**
- D. Feedforward neural network (FNN)

Dimensionality reduction is the process of transforming high-dimensional data into a lower-dimensional space while keeping the most important information intact. By reducing the number of features, it simplifies the data, helps with visualization, and speeds up processing, all without discarding the essential structure. Classic PCA projects data onto principal components that capture the most variance; other methods like t-SNE or UMAP handle nonlinear relationships; autoencoders use neural networks to learn compact representations. This focus on lowering dimensionality while preserving core features is why this topic fits the description. In contrast, neural networks, the input layer, or a standard feedforward network describe models or components that do not inherently reduce dimensionality unless paired with a dedicated reduction mechanism.

10. Data subjects can gain access to data that organizations have about them, including the ability to receive a copy of that information or direct access to the organization's database. Which right does this describe?

A. Right to access

B. Right to be informed

C. Right to rectify

D. Right to be forgotten

The main concept here is the right of individuals to access their personal data held by an organization. This right lets data subjects confirm whether their data is being processed, obtain a copy of the data, and view details about how and why it's being used, who it's shared with, and how long it will be stored. The scenario describes being able to receive a copy or directly access the organization's database, which is exactly what this right guarantees. It's distinct from being informed (providing initial notices about data processing), rectifying (correcting inaccuracies), or being forgotten (erasing data).

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://aaismdomain2.examzify.com>

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

SAMPLE