

# AI Engineering Degree Practice Exam (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

|                                    |           |
|------------------------------------|-----------|
| <b>Copyright</b> .....             | <b>1</b>  |
| <b>Table of Contents</b> .....     | <b>2</b>  |
| <b>Introduction</b> .....          | <b>3</b>  |
| <b>How to Use This Guide</b> ..... | <b>4</b>  |
| <b>Questions</b> .....             | <b>5</b>  |
| <b>Answers</b> .....               | <b>8</b>  |
| <b>Explanations</b> .....          | <b>10</b> |
| <b>Next Steps</b> .....            | <b>16</b> |

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 statement about decision trees is NOT true?**
  - A. They can easily handle both numerical and categorical data.**
  - B. They often require feature scaling to improve performance.**
  - C. They work by dividing the dataset into subsets based on feature values.**
  - D. They can overfit the training data if not properly constrained.**
- 2. What is a benefit of using marginalization in the context of conditional probabilities?**
  - A. It simplifies complex conditional dependencies**
  - B. It allows exploration of all conditional outcomes**
  - C. It avoids the reliance on Bayesian priors**
  - D. It increases dimensionality of the analysis**
- 3. What is a key advantage of using a GPU in training models?**
  - A. Increased memory capacity**
  - B. Lower energy consumption**
  - C. Parallel processing capabilities**
  - D. Faster data retrieval**
- 4. What is the role of the softmax function in machine learning?**
  - A. To convert scores to probabilities**
  - B. To calculate mean values**
  - C. To generate random samples**
  - D. To compress data dimensions**
- 5. What is a common characteristic of the k-means clustering algorithm?**
  - A. It identifies overlapping clusters.**
  - B. It requires prior knowledge of the data labels.**
  - C. It relies on strictly defined cluster boundaries.**
  - D. It can adjust the clusters dynamically.**

**6. What is meant by "data preprocessing" in machine learning?**

- A. Cleaning and organizing raw data before model training**
- B. The process of collecting data from various sources**
- C. Analyzing data to determine its usefulness**
- D. Testing data for quality and consistency**

**7. Which is NOT a characteristic of supervised learning?**

- A. Requires labeled data**
- B. Used for clustering**
- C. Typical for prediction tasks**
- D. Evaluation based on accuracy metrics**

**8. What is a convolutional neural network (CNN) primarily used for?**

- A. Natural language processing tasks**
- B. Image processing tasks such as recognition and classification**
- C. Predicting time-series data**
- D. Developing recommendation systems**

**9. Which of the following is an application of regression?**

- A. Predicting the outcome of a football match**
- B. Forecasting rainfall amount for the next day**
- C. Classifying emails as spam or not**
- D. Identifying fraudulent transactions**

**10. Which of the following is an example of a classification problem?**

- A. To predict the average temperature tomorrow**
- B. To predict the category a customer belongs to**
- C. To summarize sales figures for the month**
- D. To reduce dimensionality of a dataset**

## **Answers**

SAMPLE

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

SAMPLE

## **Explanations**

SAMPLE

## 1. Which statement about decision trees is NOT true?

- A. They can easily handle both numerical and categorical data.
- B. They often require feature scaling to improve performance.**
- C. They work by dividing the dataset into subsets based on feature values.
- D. They can overfit the training data if not properly constrained.

The statement that decision trees often require feature scaling to improve performance is not true. Decision trees inherently work by splitting the data based on the feature values at each node, evaluating one feature at a time for the best split. This means that they can naturally handle data without the need for feature scaling, such as normalization or standardization, which is often essential for algorithms that compute distances, like k-nearest neighbors or gradient descent-based methods. Because decision trees make decisions based on the order and value of features rather than their specific scales, they maintain effective performance without transformed ranges. Thus, the requirement for feature scaling is not applicable to decision trees, setting this statement apart as incorrect in the context of their functionality.

## 2. What is a benefit of using marginalization in the context of conditional probabilities?

- A. It simplifies complex conditional dependencies**
- B. It allows exploration of all conditional outcomes
- C. It avoids the reliance on Bayesian priors
- D. It increases dimensionality of the analysis

Using marginalization in the context of conditional probabilities indeed simplifies complex conditional dependencies. Marginalization involves integrating or summing out certain variables in a probabilistic model to focus on the relationships of interest. This process allows analysts to reduce complexity by eliminating unnecessary variables, making it easier to analyze the conditional probabilities of the remaining variables. By simplifying these complex relationships, marginalization makes it possible to derive cleaner and more interpretable results from a model. This is particularly useful in scenarios with high-dimensional data or intricate dependencies, where directly calculating conditional probabilities would be computationally expensive or analytically intractable. The other options do not correctly capture the primary benefit of marginalization. While exploring all outcomes (as mentioned in one option) is related to understanding probabilities, it does not specifically highlight the clarity gained through simplification. Moreover, marginalization does not inherently avoid reliance on Bayesian priors or increase dimensionality; instead, it often helps to manage such issues by providing a way to effectively manage the dimensions involved.

### 3. What is a key advantage of using a GPU in training models?

- A. Increased memory capacity**
- B. Lower energy consumption**
- C. Parallel processing capabilities**
- D. Faster data retrieval**

A key advantage of using a GPU in training models is its parallel processing capabilities. Unlike traditional CPUs, which typically have a limited number of cores optimized for sequential processing tasks, GPUs are designed with thousands of smaller cores that can handle numerous operations simultaneously. This architecture makes GPUs exceptionally well-suited for tasks that involve large-scale matrix and tensor calculations, which are prevalent in machine learning and deep learning models. When training these models, especially neural networks, the computations required can be very intensive and involve processing large datasets. The parallel nature of GPUs allows for these calculations to be distributed across many cores at once, drastically reducing the time needed to train a model compared to using a CPU that processes tasks sequentially. As a result, tasks such as forward and backward propagation in neural networks can be performed much more efficiently, leading to faster convergence and model development. The other options, while they touch on aspects that may be relevant in different contexts, do not capture the fundamental reason why GPUs are favored in training machine learning models.

Increased memory capacity is not specifically associated with GPUs, as both GPUs and CPUs can have varying memory configurations. Lower energy consumption is generally not a distinguishing feature of GPUs compared to CPUs; in fact, GPUs can consume a significant amount of power depending on

### 4. What is the role of the softmax function in machine learning?

- A. To convert scores to probabilities**
- B. To calculate mean values**
- C. To generate random samples**
- D. To compress data dimensions**

The softmax function plays a critical role in machine learning, particularly in multi-class classification problems. Its primary purpose is to convert raw scores, also known as logits, into a probability distribution over multiple classes. This conversion enables the model to provide a clear interpretation of the outputs as probabilities that sum to one, making it easier to understand and compare the likelihood of different class outcomes. Mathematically, the softmax function takes a vector of scores and exponentiates each score, normalizing these values by dividing by the sum of the exponentiated scores. This results in each output value being in the range of 0 to 1, which is indicative of a probability. Consequently, the class with the highest score after applying the softmax function can be interpreted as the predicted class, as it corresponds to the highest probability. The mathematical formulation of the softmax function is particularly effective because it emphasizes larger scores while diminishing the impact of smaller ones. This property is essential in many applications such as neural networks for classification tasks, enabling models to make confident and interpretable predictions based on input data. The remaining options do not align with the specific functionality of the softmax function. For example, calculating mean values is related to averaging and not indicative of generating

## 5. What is a common characteristic of the k-means clustering algorithm?

- A. It identifies overlapping clusters.**
- B. It requires prior knowledge of the data labels.**
- C. It relies on strictly defined cluster boundaries.**
- D. It can adjust the clusters dynamically.**

K-means clustering is characterized by its reliance on strictly defined cluster boundaries, which is foundational to the algorithm's functionality. In k-means, the process begins by partitioning data into a predetermined number of clusters, denoted as "k." Each cluster is represented by its centroid, which is the mean location of all data points assigned to that cluster. During the algorithm's iterative process, each data point is assigned to the nearest centroid based on a distance metric, typically Euclidean distance. This assignment is what establishes the strict boundaries of clusters, as each point can belong to only one cluster based on proximity. The centroids are then recalculated based on the new assignments, and the process repeats until the clusters stabilize. This clear demarcation of clusters underlines the algorithm's effectiveness in forming distinct groups in structured datasets but also highlights its limitation when dealing with datasets with overlapping clusters. Unlike some other clustering methods that account for more fluid and probabilistic boundaries, k-means creates well-defined, non-overlapping clusters, emphasizing its characteristic focus on strict boundary enforcement.

## 6. What is meant by "data preprocessing" in machine learning?

- A. Cleaning and organizing raw data before model training**
- B. The process of collecting data from various sources**
- C. Analyzing data to determine its usefulness**
- D. Testing data for quality and consistency**

Data preprocessing in machine learning refers to the steps taken to clean and organize raw data before training a model. This is a crucial phase in the machine learning pipeline as it significantly impacts the performance of the model. During preprocessing, various tasks are typically carried out, such as removing duplicates, handling missing values, normalizing or standardizing data, and encoding categorical variables. The goal is to convert raw data into a format that is suitable for modeling, ensuring that the model trains effectively on relevant, high-quality data. While other processes like collecting data and testing for quality do form parts of the broader data handling lifecycle, they do not specifically address the transforming and preparing of data to make it ready for model training, which is the essence of data preprocessing. Hence, the focus on organizing and cleaning raw data distinguishes this concept as a foundational element in developing machine learning systems.

## 7. Which is NOT a characteristic of supervised learning?

- A. Requires labeled data
- B. Used for clustering**
- C. Typical for prediction tasks
- D. Evaluation based on accuracy metrics

In supervised learning, the model is trained using labeled data, meaning that each training example comes with a corresponding output label. This characteristic allows the model to learn the mapping from inputs to outputs, making it suitable for prediction tasks, where the goal is to forecast outcomes based on new, unseen input data. Evaluation in supervised learning typically involves assessing the model's accuracy, meaning it measures how well the model's predictions match the actual labels in the test data. Clustering, on the other hand, is an unsupervised learning technique where the goal is to group similar data points together without any labeled outputs. Since there is no designated output in clustering, this method is not characteristic of supervised learning. Therefore, indicating that supervised learning is used for clustering correctly identifies a difference between the two paradigms.

## 8. What is a convolutional neural network (CNN) primarily used for?

- A. Natural language processing tasks
- B. Image processing tasks such as recognition and classification**
- C. Predicting time-series data
- D. Developing recommendation systems

A convolutional neural network (CNN) is specifically designed to process data that has a grid-like topology, making it particularly effective for tasks involving image data. The architecture of a CNN incorporates convolutional layers, which apply various filters to the input image, allowing the network to automatically learn spatial hierarchies of features. This hierarchical feature extraction is vital for identifying patterns such as edges, textures, and more complex structures within images. CNNs excel in image recognition and classification tasks by leveraging their ability to recognize visual patterns through layers of abstraction. For example, in an image classification scenario, the network can take a raw pixel input and progressively transform it through multiple layers to output a classification label, such as identifying whether the image contains a cat or a dog. While convolutional neural networks can still be applied in broader contexts — such as video processing or even in certain natural language processing applications by transforming text into a suitable format — their primary and most successful application remains in the realm of image processing.

## 9. Which of the following is an application of regression?

- A. Predicting the outcome of a football match
- B. Forecasting rainfall amount for the next day**
- C. Classifying emails as spam or not
- D. Identifying fraudulent transactions

Regression is a statistical method used for predicting and modeling the relationship between a dependent variable and one or more independent variables. It is particularly effective when the task involves estimating a continuous outcome based on various input features. In the context of forecasting rainfall amounts, regression models can analyze historical weather data, such as temperature and humidity levels, to produce precise predictions about future rainfall. This continuous nature of the output (amount of rainfall) makes regression a suitable technique for such an application. Other choices involve different types of analyses. Predicting the outcome of a football match (first choice) typically involves categorical outcomes (win, loss, draw) and is better suited to classification methods rather than regression. Classifying emails as spam or not (third choice) also involves a binary classification approach, where the aim is to categorize each email rather than predict a continuous outcome. Identifying fraudulent transactions (fourth choice) also falls under classification, as the goal is to label transactions as either fraudulent or legitimate based on certain parameters. Thus, forecasting rainfall aligns with the nature of regression, making it the most appropriate application among the provided options.

## 10. Which of the following is an example of a classification problem?

- A. To predict the average temperature tomorrow
- B. To predict the category a customer belongs to**
- C. To summarize sales figures for the month
- D. To reduce dimensionality of a dataset

A classification problem involves assigning a label or category to an input based on its features. In this case, predicting the category a customer belongs to is a clear example of classification. This involves analyzing customer data and determining which predefined category, such as "high-value," "medium-value," or "low-value" customer, the individual fits into. The process typically includes training a model on labeled data and then using that model to predict the category for new, unseen data. In contrast, predicting the average temperature for tomorrow is typically a regression problem, as it involves estimating a continuous numerical value rather than a category. Summarizing sales figures for the month also does not involve classification; it's more about aggregating data to produce insights. Finally, dimensionality reduction refers to techniques used to reduce the number of input variables in a dataset, which is not applicable to classification tasks. Thus, the selection of predicting a customer's category appropriately identifies the essence of a classification problem.

# 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://ai-engineering.examzify.com>**

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

**SAMPLE**