

# AWS Certified Machine Learning Specialty (MLS-C01) Practice Test (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 of the following is NOT a common metric used to evaluate machine learning models in classification tasks?**
  - A. Accuracy**
  - B. F1 Score**
  - C. Mean Squared Error**
  - D. Precision**
  
- 2. Which tool is designed for business analytics to create visualizations and perform ad hoc analysis?**
  - A. AWS Lambda**
  - B. Amazon QuickSight**
  - C. Amazon EC2**
  - D. Amazon EMR**
  
- 3. What is an ensemble method in machine learning?**
  - A. Combining multiple models to enhance prediction accuracy**
  - B. Training a single model on a complete dataset**
  - C. Using a single algorithm for all predictions**
  - D. Employing random selection of data points for training**
  
- 4. What AWS tool is used to simplify deployment and management of Docker applications?**
  - A. Amazon ECS**
  - B. Amazon EKS**
  - C. AWS Fargate**
  - D. Amazon Lightsail**
  
- 5. What key advantage does the Precision-Recall curve offer in classification tasks?**
  - A. It shows the computational efficiency of different models.**
  - B. It helps to visualize the cost associated with misclassifications.**
  - C. It provides a clear view of the trade-offs between precision and recall.**
  - D. It quantifies the error rate of model predictions.**

**6. Which service is best for handling high-availability compute infrastructure?**

- A. AWS Lambda**
- B. Amazon EC2**
- C. Amazon Kinesis Data Streams**
- D. AWS Batch**

**7. What transformation would be best used to stabilize variance and make the data more normal distribution-like?**

- A. Logarithmic transformation**
- B. Normalization**
- C. Standardization**
- D. Min-Max scaling**

**8. What feature engineering method transforms continuous features by centering their mean around zero and scaling to unit variance?**

- A. Standard Scaling**
- B. Min-Max Scaling**
- C. Robust Scaling**
- D. Feature Normalization**

**9. What service makes it easy to securely stream video from connected devices to AWS for analytics and processing?**

- A. Amazon Kinesis Video Streams**
- B. Amazon Elastic Transcoder**
- C. Amazon CloudFront**
- D. Amazon Rekognition Video**

**10. What is hyperparameter tuning in machine learning?**

- A. Testing different algorithms for model selection**
- B. Optimizing training parameters to enhance model performance**
- C. Cleaning data before model training**
- D. Creating new features for analysis**

## **Answers**

SAMPLE

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

SAMPLE

## **Explanations**

SAMPLE

**1. Which of the following is NOT a common metric used to evaluate machine learning models in classification tasks?**

- A. Accuracy**
- B. F1 Score**
- C. Mean Squared Error**
- D. Precision**

Mean Squared Error is not a common metric used for evaluating machine learning models in classification tasks; rather, it is primarily utilized in regression tasks. In classification, the goals are to predict discrete labels, such as classifying emails as 'spam' or 'not spam', while Mean Squared Error measures the average of the squares of the errors between predicted and actual values. This approach is more suitable for continuous numerical predictions. On the other hand, Accuracy, F1 Score, and Precision are crucial metrics for classification tasks. Accuracy indicates the proportion of correct predictions made by the model, providing a high-level overview of its performance. Precision measures the accuracy of positive predictions, which is especially important in cases where false positives carry a significant cost. The F1 Score is the harmonic mean of Precision and Recall, and it provides a single score that balances both, making it useful when dealing with imbalanced datasets. In summary, while the other metrics are integral to classification model evaluation, Mean Squared Error is tailored for regression analysis, thereby confirming that it is not commonly used in classification tasks.

**2. Which tool is designed for business analytics to create visualizations and perform ad hoc analysis?**

- A. AWS Lambda**
- B. Amazon QuickSight**
- C. Amazon EC2**
- D. Amazon EMR**

Amazon QuickSight is specifically designed for business analytics, enabling users to create visualizations and perform ad hoc analysis on data. It offers an interactive and user-friendly interface that allows both technical and non-technical users to derive insights from data without requiring extensive programming or data science skills. QuickSight supports a variety of data sources and has features like automatic data discovery, machine learning insights, and the ability to share dashboards and visualizations seamlessly within an organization. This makes it an excellent choice for organizations looking to empower their teams with data-driven decision-making capabilities. The other options do not align with the primary function of creating visualizations and conducting business analytics. AWS Lambda is a serverless compute service used for running code in response to events, Amazon EC2 provides scalable virtual servers for running applications, and Amazon EMR is focused on big data processing, utilizing frameworks like Apache Hadoop and Spark. While these services serve important roles within the AWS ecosystem, they do not specialize in the visualization and analytics capabilities that QuickSight offers.

### 3. What is an ensemble method in machine learning?

- A. Combining multiple models to enhance prediction accuracy**
- B. Training a single model on a complete dataset**
- C. Using a single algorithm for all predictions**
- D. Employing random selection of data points for training**

An ensemble method in machine learning refers to the technique of combining multiple models to enhance prediction accuracy. This approach leverages the strengths of individual models to improve overall performance. The underlying idea is based on the premise that different models might capture various aspects of the data, and when their predictions are aggregated (through methods like voting, averaging, or stacking), the combination often yields better results than any single model alone. Ensemble methods can reduce the risk of overfitting, as they smooth out the predictions by taking into account diverse perspectives from the various models involved. Popular ensemble techniques include bagging (e.g., Random Forest) and boosting (e.g., AdaBoost, Gradient Boosting), which intelligently combine models to achieve improved accuracy, robustness, and generalization on unseen data. The other options describe concepts that do not align with the definition of ensemble methods. Training a single model on a complete dataset relates to standard supervised learning, using a single algorithm focuses on a singular approach without leveraging diversity, and employing random selection of data points refers to techniques like bootstrapping but does not encapsulate the broader concept of model combination that defines ensemble methods.

### 4. What AWS tool is used to simplify deployment and management of Docker applications?

- A. Amazon ECS**
- B. Amazon EKS**
- C. AWS Fargate**
- D. Amazon Lightsail**

The correct answer is Amazon ECS. This service is specifically designed for running and managing Docker containers at scale. Amazon ECS (Elastic Container Service) allows you to easily deploy, manage, and scale containerized applications using Docker. It integrates seamlessly with other AWS services and offers features such as task definitions, which specify how Docker containers should run, and service management, which helps in maintaining the desired state of your application. While AWS Fargate is a technology that works with Amazon ECS and Amazon EKS to run containers without managing servers or clusters, it is primarily focused on the compute layer rather than being a comprehensive container management service. On the other hand, Amazon EKS (Elastic Kubernetes Service) is geared towards managing Kubernetes clusters, which can also run Docker containers but involves a steeper learning curve compared to ECS. Lastly, Amazon Lightsail is designed for simple cloud applications and does not primarily focus on container orchestration or management. Therefore, choosing Amazon ECS directly addresses the requirement of simplifying the deployment and management of Docker applications.

## 5. What key advantage does the Precision-Recall curve offer in classification tasks?

- A. It shows the computational efficiency of different models.
- B. It helps to visualize the cost associated with misclassifications.
- C. It provides a clear view of the trade-offs between precision and recall.**
- D. It quantifies the error rate of model predictions.

The key advantage of the Precision-Recall curve is its ability to provide a clear view of the trade-offs between precision and recall. Precision measures the accuracy of the positive predictions, indicating the proportion of true positive predictions out of all positive predictions made by the model. Recall, on the other hand, measures the model's ability to find all the relevant cases, reflecting the ratio of true positives to the total number of actual positives. In situations where the classes are imbalanced or where the cost of false positives and false negatives varies significantly, understanding this trade-off becomes crucial. The Precision-Recall curve plots precision against recall for different thresholds, allowing practitioners to select an appropriate balance based on the specific context of their problem, such as prioritizing fewer false positives over false negatives or vice versa. This visualization helps in making informed decisions about model performance, especially in binary classification tasks where the focus is on the positive class. The other options address different concepts that are not the primary purpose of the Precision-Recall curve, such as computational efficiency, cost visualization, or quantifying error rates, which do not directly relate to the strengths of this specific curve.

## 6. Which service is best for handling high-availability compute infrastructure?

- A. AWS Lambda
- B. Amazon EC2**
- C. Amazon Kinesis Data Streams
- D. AWS Batch

Amazon EC2 is designed for providing scalable and reliable compute resources that can be configured for high availability. It offers a variety of instance types suited for different workloads, enabling users to tailor their compute environment to meet specific application demands. One of the key features of EC2 is the ability to deploy instances across multiple Availability Zones (AZs) within a region. This multi-AZ deployment enhances fault tolerance and availability since it allows applications to remain operational even if one AZ becomes unavailable. Additionally, EC2 supports features such as Elastic Load Balancing and Auto Scaling, which further ensure that the infrastructure can handle varying loads while maintaining high availability. In contrast, while AWS Lambda is a great service for serverless computing, it is not specifically geared towards traditional high-availability setups in the way that EC2 is. Lambda manages the underlying infrastructure abstractly but may not provide the same level of tuning and control that is often required in high-availability scenarios. Amazon Kinesis Data Streams is focused on real-time data streaming rather than compute infrastructure. It enables the processing and analysis of streaming data but does not serve the same purpose as a general compute service like EC2. AWS Batch is designed for running batch jobs and can manage compute instances accordingly, but it is not specifically

**7. What transformation would be best used to stabilize variance and make the data more normal distribution-like?**

**A. Logarithmic transformation**

**B. Normalization**

**C. Standardization**

**D. Min-Max scaling**

The logarithmic transformation is a powerful technique used to stabilize variance and help make the data more closely resemble a normal distribution. This approach works particularly well for datasets with exponential growth patterns or when the data exhibits skewness, especially right skewness. By applying the logarithmic function, you compress the range of high values and expand the range of low values, which tends to reduce the influence of outliers and bring the overall data distribution closer to normality. In contrast, normalization, standardization, and min-max scaling are methods that do not specifically target variance stabilization or normality. Normalization typically rescales data to fit within a specific range, often between 0 and 1, without addressing the distribution's shape. Standardization (z-score normalization) shifts the data to have a mean of 0 and a standard deviation of 1 but does not inherently alter the overall distribution to resemble normality. Min-max scaling rescales features to a common range, which is useful for certain algorithms but does not directly help in achieving a normal distribution. Therefore, the logarithmic transformation stands out as the most effective approach for stabilizing variance and making the data distribution more normal-like.

**8. What feature engineering method transforms continuous features by centering their mean around zero and scaling to unit variance?**

**A. Standard Scaling**

**B. Min-Max Scaling**

**C. Robust Scaling**

**D. Feature Normalization**

The feature engineering method that transforms continuous features by centering their mean around zero and scaling to unit variance is known as Standard Scaling. This process involves calculating the z-score for each feature by subtracting the mean of the feature and then dividing by the standard deviation. The result is that the transformed data will have a mean of 0 and a standard deviation of 1, allowing for comparisons across features that may have different units and scales. Standard Scaling is particularly useful when you are dealing with algorithms sensitive to the scale of data, such as support vector machines and gradient descent-based algorithms, as they perform better when the features are centered and have similar variances. Min-Max Scaling, on the other hand, adjusts the features to a fixed range, typically between 0 and 1, making it different from Standard Scaling. Robust Scaling focuses on using the median and interquartile range to scale features, which helps with outliers but does not center the data around zero in the same way. Feature Normalization is a more general term that can include various techniques, but it is not specific to the mean and variance transformation that Standard Scaling provides.

## 9. What service makes it easy to securely stream video from connected devices to AWS for analytics and processing?

- A. Amazon Kinesis Video Streams**
- B. Amazon Elastic Transcoder**
- C. Amazon CloudFront**
- D. Amazon Rekognition Video**

Amazon Kinesis Video Streams is designed specifically for the purpose of securely streaming video from connected devices to the AWS cloud. This service allows developers to ingest, process, and analyze video data in real-time, making it an ideal solution for applications requiring video analytics and processing. With Kinesis Video Streams, it becomes seamless to create applications that can consume video streams, perform analytics, and manage the resulting data effectively. The service also supports various security features, ensuring that video data is transferred securely. This capability is crucial in use cases such as surveillance, media processing, and machine learning applications where video input is required. While Amazon Elastic Transcoder is a service for converting media files from one format to another, it does not handle real-time video streaming. Amazon CloudFront is a content delivery network that accelerates the delivery of static and dynamic web content, but it is not specifically tailored for video stream ingestion. Amazon Rekognition Video, on the other hand, is a service that provides analysis of video content for facial recognition and object detection, but it relies on input from other services or sources and is not focused on securely streaming video from devices.

## 10. What is hyperparameter tuning in machine learning?

- A. Testing different algorithms for model selection**
- B. Optimizing training parameters to enhance model performance**
- C. Cleaning data before model training**
- D. Creating new features for analysis**

Hyperparameter tuning refers specifically to the process of optimizing the hyperparameters of a machine learning model to improve its performance. Hyperparameters are the configurations that are set before the learning process begins and include settings such as the learning rate, the number of trees in a random forest, or the depth of a decision tree. Unlike parameters that are learned during training, hyperparameters are set by the practitioner and can significantly affect the model's ability to generalize on unseen data. By systematically exploring different combinations of hyperparameters, often using techniques such as grid search or random search, practitioners can discover the optimal settings that yield the highest performance on a validation dataset. This meticulous process enables the model to perform better by minimizing errors, improving accuracy, and enhancing its robustness across various scenarios. The other options reflect different aspects of the machine learning workflow. Testing different algorithms for model selection involves choosing among various machine learning models and is different from tuning hyperparameters within a specific model. Cleaning data is a crucial step before training begins, aimed at improving data quality, while feature engineering involves creating new input variables to better represent the underlying problem to the model. None of these processes relate directly to the optimization of hyperparameters.

# 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://awscertmachinelearningspecialty.examzify.com>**

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

**SAMPLE**