


MICROSOFT AZURE AI FUNDAMENTALS (AI-900) PRACTICE EXAM (SAMPLE)

STUDY GUIDE



**SAMPLE STUDY GUIDE
WITH LIMITED QUESTIONS**

**VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE**

<https://microsoftazureaifundamentals-ai-900.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. What is the main function of facial recognition technology?**
 - A. Extracting text from images
 - B. Identifying individuals in images
 - C. Locating objects within images
 - D. Classifying images into categories
- 2. When is a model considered to be overfitting?**
 - A. When it performs well on training data but poorly on validation data
 - B. When it has high accuracy on both training and validation data
 - C. When it uses less data than necessary
 - D. When it cannot learn any patterns from the data
- 3. What is the primary use of clustering in machine learning?**
 - A. To assign specific labels
 - B. To discover patterns
 - C. To reduce overfitting
 - D. To predict continuous values
- 4. What capability does object detection provide when analyzing images containing vehicles?**
 - A. Identify faces in images
 - B. Determine vehicle types
 - C. Locate vehicles in images
 - D. Recognize movie titles
- 5. What process would you use to separate humans from other objects in a photo?**
 - A. Classification
 - B. Detection
 - C. Verification
 - D. Segmentation

- 6. Which type of analysis determines if two images of a face belong to the same individual?**
- A. Verification
 - B. Classification
 - C. Detection
 - D. Segmentation
- 7. What is a critical principle for responsible AI concerning the handling of predictions when essential fields have unusual or missing values?**
- A. Fairness
 - B. Accountability
 - C. Reliability and safety
 - D. Bias mitigation
- 8. Which aspect of AI focuses on understanding the relationship between cause and effect in data?**
- A. Data representation
 - B. Model evaluation
 - C. Causal inference
 - D. Predictive modeling
- 9. Which method would you use to select the most relevant variables for a predictive model?**
- A. Feature engineering
 - B. Feature selection
 - C. Model optimization
 - D. Data augmentation
- 10. What is the primary goal of a classification model?**
- A. To predict continuous values
 - B. To group similar items
 - C. To assign labels to input data
 - D. To identify anomalies

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is the main function of facial recognition technology?

- A. Extracting text from images
- B. Identifying individuals in images**
- C. Locating objects within images
- D. Classifying images into categories

The main function of facial recognition technology is to identify individuals in images. This technology utilizes various algorithms to detect human faces in photos or video streams and then compares the features of those faces against a database to determine the identity of the person. It relies on biometric recognition to accurately differentiate between various individuals based on unique facial characteristics such as the distance between eyes, the shape of the jawline, and other distinct features. Facial recognition has applications in numerous areas, including security systems, law enforcement, personal device authentication, and social media tagging, among others. The ability to reliably identify individuals quickly and effectively is what sets facial recognition apart from other imaging technologies. In contrast, the other options listed serve different purposes. Extracting text from images relates to optical character recognition (OCR), which is not concerned with identifying individuals. Locating objects within images deals with object detection, a task focused on identifying and classifying specific objects rather than identifying people. Classifying images into categories involves image classification techniques that categorize an image based on its content without necessarily identifying specific individuals. Therefore, identifying individuals is the distinct function that defines facial recognition technology.

2. When is a model considered to be overfitting?

- A. When it performs well on training data but poorly on validation data**
- B. When it has high accuracy on both training and validation data
- C. When it uses less data than necessary
- D. When it cannot learn any patterns from the data

A model is considered to be overfitting when it performs exceptionally well on the training data but shows poor performance on validation data. This indicates that the model has learned the noise and details in the training data to the extent that it fails to generalize to new, unseen data. In essence, the model has memorized the training examples instead of learning the underlying patterns that would allow it to make accurate predictions on different datasets. Overfitting typically occurs when the model is too complex relative to the amount of training data available. This results in a model that captures noise and outliers as if they were valid patterns, leading to high accuracy on known data while failing to perform well on new data. Hence, a key indicator of overfitting is the disparity between training and validation performance, where the training accuracy is high, but validation accuracy is significantly lower. This is a critical concept in machine learning, as it underscores the importance of balancing model complexity with adequate training data to ensure the model can generalize effectively.

3. What is the primary use of clustering in machine learning?

- A. To assign specific labels
- B. To discover patterns**
- C. To reduce overfitting
- D. To predict continuous values

The primary use of clustering in machine learning is to discover patterns within a dataset. Clustering is an unsupervised learning technique that groups data points based on their similarities without needing pre-assigned labels. By identifying groups, or clusters, within the data, clustering algorithms help to reveal underlying structures, trends, and patterns that might not be apparent through other means. This is particularly useful in exploratory data analysis, marketing segmentation, or identifying anomalies in datasets, where the goal is to understand the distribution and relationships among the data rather than to predict specific outcomes or classify instances. Other alternatives focus on different tasks in machine learning. Assigning specific labels pertains to supervised learning methods like classification, while reducing overfitting generally involves strategies during model training to improve generalization. Predicting continuous values is a characteristic of regression tasks. Thus, clustering stands out primarily as a method for pattern discovery.

4. What capability does object detection provide when analyzing images containing vehicles?

- A. Identify faces in images
- B. Determine vehicle types
- C. Locate vehicles in images**
- D. Recognize movie titles

Object detection technology is specifically designed to identify and locate objects within images. When applied to images containing vehicles, object detection can accurately pinpoint where each vehicle is situated in the image, providing both its position and potentially bounding boxes around the vehicles. This capability is crucial in many applications, such as autonomous vehicles, traffic monitoring, and any scenario where understanding the spatial distribution of objects is important. While determining vehicle types might also be a function of image analysis, that aspect pertains more to classification rather than the fundamental task of locating objects. Identifying faces and recognizing movie titles do not relate to the core function of object detection in the context of vehicle recognition. Thus, the ability to locate vehicles is the primary and defining capability of object detection technology.

5. What process would you use to separate humans from other objects in a photo?

- A. Classification
- B. Detection**
- C. Verification
- D. Segmentation

The process for separating humans from other objects in a photo is best represented by detection. In the context of computer vision, detection involves identifying and localizing specific objects within an image, such as people, animals, vehicles, etc. This is typically done by drawing bounding boxes around the detected objects, which helps distinguish humans from other elements in the photo. Detection algorithms, such as YOLO (You Only Look Once) and SSD (Single Shot MultiBox Detector), are commonly used for this purpose. They analyze the entire image to identify the presence and position of objects in real time. While segmentation is also related to distinguishing objects within an image, it goes a step further by pixel-wise classification, basically outlining the exact shape of the objects. In contrast, verification usually deals with confirming the identity or authenticity of an object rather than distinguishing it from others in an image. Classification, on the other hand, assigns a label to whole images rather than distinguishing specific objects within them. Therefore, detection is the most appropriate choice for the task of identifying and separating humans from other objects in a photo.

6. Which type of analysis determines if two images of a face belong to the same individual?

- A. Verification**
- B. Classification
- C. Detection
- D. Segmentation

Verification is the correct choice as it specifically refers to the analysis process that confirms whether two images of a face belong to the same individual. In this context, face verification involves comparing the features of a person's face in one image to the features in another image, often using algorithms that measure similarity. This process typically results in a yes or no answer regarding the identity of the individual depicted in both images, making it ideal for applications such as biometric security systems or identity verification services. Classification, by contrast, involves categorizing an image into predefined classes or categories, but it does not determine identity between two instances. Detection pertains to identifying objects or faces within an image, rather than comparing images for identity purposes. Segmentation involves partitioning images into segments to simplify analysis, often targeting objects or regions within an image rather than directly comparing images. Thus, verification stands out as the best fit for this specific task involving facial comparison based on identity.

7. What is a critical principle for responsible AI concerning the handling of predictions when essential fields have unusual or missing values?

- A. Fairness
- B. Accountability
- C. Reliability and safety**
- D. Bias mitigation

The principle of reliability and safety is crucial for responsible AI, especially when dealing with predictions that involve unusual or missing values in essential fields. This principle emphasizes the need for AI systems to produce consistent, accurate, and dependable outcomes, particularly in critical scenarios where the integrity of the predictions could lead to significant consequences. In practice, when an AI model encounters unusual or missing values, it is vital for the system to have mechanisms in place to handle these situations appropriately. This could involve implementing fallback strategies, providing alerts to users, or ensuring that the model's reliability is maintained even when facing data irregularities. By prioritizing reliability and safety, organizations can enhance trust in their AI systems and ensure that decisions made based on AI predictions are sound and justified. Fairness focuses on eliminating bias and ensuring equitable outcomes for all users. Accountability refers to the responsibility of organizations and individuals for the decisions made by AI systems. Bias mitigation addresses the need to identify and reduce bias in AI models to promote fairness and equity. While all these principles are important for responsible AI, the specific context of handling predictions related to unusual or missing values aligns most closely with the principles of reliability and safety.

8. Which aspect of AI focuses on understanding the relationship between cause and effect in data?

- A. Data representation
- B. Model evaluation
- C. Causal inference**
- D. Predictive modeling

Causal inference is the aspect of AI that specifically focuses on understanding the relationship between cause and effect in data. This involves identifying whether a change in one variable leads to a change in another variable and establishing a clear cause-and-effect relationship. Causal inference seeks to analyze data in a way that goes beyond mere correlation, enabling practitioners to infer what might happen if a specific intervention is applied. In various applications, such as healthcare or economics, determining causality can be crucial for making informed decisions based on data. For instance, if a healthcare study suggests that a new treatment leads to better recovery rates, causal inference techniques would help confirm whether the treatment itself is responsible for this improvement or if other factors are at play. Other aspects mentioned, such as data representation, model evaluation, and predictive modeling, serve different purposes in AI. Data representation is concerned with how data is organized and encoded for processing, model evaluation involves assessing the performance of a predictive model, and predictive modeling focuses on using historical data to forecast outcomes. While these factors play important roles in machine learning and AI, they do not specifically address the fine nuances of causality in the way that causal inference does.

9. Which method would you use to select the most relevant variables for a predictive model?

- A. Feature engineering
- B. Feature selection**
- C. Model optimization
- D. Data augmentation

Feature selection is the process used to identify and select the most relevant variables for a predictive model. This method helps to improve the model's performance by eliminating irrelevant or redundant features that do not contribute significantly to the prediction. By focusing on the most impactful variables, the model can become more efficient, potentially leading to better accuracy and less overfitting. Feature selection techniques often assess the relationships between predictor variables and the target variable. This may include statistical tests, recursive feature elimination, or algorithms that inherently perform feature selection, such as tree-based models. The goal is to simplify the model without losing predictive power, making it easier to interpret and faster to run. On the other hand, feature engineering involves creating new features or modifying existing ones to improve model performance. This can be useful but does not specifically address the direct selection of relevant features. Model optimization is related to tuning hyperparameters and improving a model's performance but also does not focus on choosing the input variables. Lastly, data augmentation refers to techniques that increase the diversity of the training dataset by applying transformations, and it does not pertain to the selection of relevant features. Selecting the right features is crucial in the model-building process, as using irrelevant features can lead to a decrease in model performance, while the

10. What is the primary goal of a classification model?

- A. To predict continuous values
- B. To group similar items
- C. To assign labels to input data**
- D. To identify anomalies

The primary goal of a classification model is to assign labels to input data. In classification tasks, the model takes input features and then categorizes or classifies them into predefined classes or labels. This is useful in various applications such as email spam detection, where incoming messages are labeled as "spam" or "not spam," and image recognition, where images are classified into categories like "cat," "dog," or "car." Classification models operate based on training data that has features along with their corresponding labels, allowing them to learn patterns and relationships. Once trained, the model can then predict the label for new, unseen data points based on those learned patterns, thus achieving its primary goal. In this context, predicting continuous values relates to regression tasks rather than classification. Grouping similar items typically pertains to clustering, which involves unsupervised learning rather than assigning specific labels to known categories. Identifying anomalies is a distinct task generally handled by anomaly detection models, which focus on recognizing unusual patterns rather than classifying standard data points into defined categories.

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://microsoftazureaifundamentals-ai900.examzify.com>

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

SAMPLE