

Microsoft Azure AI Fundamentals (AI-900) 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

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. When training a predictive model for taxi fares, what feature is most relevant?**
 - A. The fare amount**
 - B. The trip distance of individual taxi journeys**
 - C. The number of passengers**
 - D. The driver's rating**
- 2. In Azure Cognitive Services, which service would be ideal for enhancing customer engagement through voice interaction?**
 - A. Face API**
 - B. Speech Service**
 - C. Computer Vision**
 - D. Text Analytics**
- 3. Is predicting a patient's likelihood of developing diabetes from medical history an example of anomaly detection?**
 - A. Yes**
 - B. No**
 - C. Only with extensive data analysis**
 - D. Only if it uses historical patient data**
- 4. Will the Face service be more accurate with more sample photos of each employee from different angles?**
 - A. Yes**
 - B. No**
 - C. Only for new employees**
 - D. Only if taken on different days**
- 5. To predict the animal population of an area, which type of Azure Machine Learning should you use?**
 - A. Classification**
 - B. Regression**
 - C. Clustering**
 - D. Anomaly detection**

6. Which method would you use to understand emotional tone in customer feedback?

- A. Key Phrase Extraction**
- B. Sentiment Analysis**
- C. Text Classification**
- D. Image Recognition**

7. Is "Quality Test" classified as a label in machine learning?

- A. Yes**
- B. No**
- C. Depends on context**
- D. Not defined**

8. Is it possible to use the Custom Vision service to analyze video files?

- A. Yes**
- B. No**
- C. Only under specific conditions**
- D. Only for short clips**

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

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

10. Organizing documents into groups based on textual similarities is an example of which machine learning concept?

- A. Classification**
- B. Regression**
- C. Clustering**
- D. Association**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. When training a predictive model for taxi fares, what feature is most relevant?

- A. The fare amount**
- B. The trip distance of individual taxi journeys**
- C. The number of passengers**
- D. The driver's rating**

When training a predictive model for taxi fares, the trip distance of individual taxi journeys is the most relevant feature because it directly influences the fare amount charged to passengers. In typical taxi fare structures, the fare is often calculated based on a combination of factors that include a base charge, a charge per mile or kilometer, and potentially time-based charges for waiting. Hence, the trip distance is a key determinant, as longer distances generally correlate with higher fares. The other features, while they may provide some context, do not have as strong a direct correlation to the fare. For instance, the fare amount itself is the target variable that you are trying to predict, so using it as a feature would be redundant. The number of passengers could influence the fare in some cases, but it is typically a smaller factor compared to distance. The driver's rating may impact customer satisfaction or choice of taxi, but it does not directly affect the fare calculation itself. Thus, trip distance emerges as the most pertinent feature for accurately predicting taxi fares.

2. In Azure Cognitive Services, which service would be ideal for enhancing customer engagement through voice interaction?

- A. Face API**
- B. Speech Service**
- C. Computer Vision**
- D. Text Analytics**

The Speech Service in Azure Cognitive Services is specifically designed to facilitate and enhance voice interactions, making it the ideal choice for improving customer engagement through spoken language. This service has capabilities for speech recognition, which converts spoken language into text, as well as text-to-speech, which can generate spoken language from text input. Additionally, it supports voice synthesis, enabling applications to communicate back to users in a more natural and engaging manner. The Speech Service also includes features such as speech translation and speaker recognition, which can further enrich the user experience and drive effective communication. By allowing customers to interact through voice, it creates a more intuitive and accessible means of engagement, making it especially useful in environments where typing or reading might be less practical. Other options like the Face API, Computer Vision, and Text Analytics focus on visual data processing and text analysis, respectively, and do not cater specifically to voice interaction. Thus, while they serve important roles in other aspects of AI-driven applications, they do not provide the same level of functionality dedicated to enhancing customer engagement through voice.

3. Is predicting a patient's likelihood of developing diabetes from medical history an example of anomaly detection?

- A. Yes**
- B. No**
- C. Only with extensive data analysis**
- D. Only if it uses historical patient data**

Predicting a patient's likelihood of developing diabetes based on their medical history is more accurately categorized as a form of predictive modeling or classification rather than anomaly detection. Anomaly detection typically focuses on identifying unusual patterns or outliers in data that do not conform to expected behavior, such as fraud detection or identifying network intrusions. In contrast, predicting diabetes involves analyzing patterns in historical medical data to assess risk factors and trends among a patient population. This prediction seeks to classify patients into risk categories rather than identify outliers. The objective here is to make informed predictions based on established correlations between patient data and diabetes outcomes, which aligns well with the concepts of supervised learning in machine learning, where models are trained to recognize these patterns and predict future occurrences. The other options, like extensive data analysis or using historical data, might apply to the process of prediction in a broader sense, but they do not fundamentally change the nature of the task, which is about assessing risks based on historical patterns rather than detecting anomalies.

4. Will the Face service be more accurate with more sample photos of each employee from different angles?

- A. Yes**
- B. No**
- C. Only for new employees**
- D. Only if taken on different days**

The Face service in Azure uses machine learning algorithms to recognize and analyze facial features from images. The accuracy of this service generally improves with a larger and more diverse dataset. When more sample photos of each employee are provided from different angles, the model can learn to identify the unique features of each person's face more effectively. Having various angles helps the system to account for possible variations in appearance due to head rotations, facial expressions, or lighting conditions. Essentially, the model benefits from more context and variability, which enhances its ability to generalize from the training data to real-world scenarios. This increased accuracy can lead to better recognition and less chance of confusion between different individuals. In this case, the correct choice illustrates the importance of data diversity and quantity in training machine learning models, particularly in services like facial recognition.

5. To predict the animal population of an area, which type of Azure Machine Learning should you use?

- A. Classification**
- B. Regression**
- C. Clustering**
- D. Anomaly detection**

To predict the animal population of an area, regression is the appropriate choice because this method is specifically designed to deal with continuous numerical outcomes. In this scenario, the goal is to forecast a specific quantity, which in this case is the number of animals in a given area. Regression techniques can handle various types of data relationships and can provide an estimate of population size based on factors such as environmental conditions, food availability, and previous population data. Regression models analyze historical data to identify patterns that can be used for prediction. For instance, if you have data on past animal populations alongside variables that may influence these populations, regression allows you to understand how these factors relate to the total number of animals, enabling you to produce numerical population estimates. The other options, such as classification, clustering, and anomaly detection, serve different analytical purposes. Classification is used for predicting categorical outcomes, clustering is about grouping data points based on similarity, and anomaly detection focuses on identifying outliers or abnormal patterns in the data. Since predicting a population involves estimating a numeric value rather than categorizing or grouping data, regression is the most suitable choice.

6. Which method would you use to understand emotional tone in customer feedback?

- A. Key Phrase Extraction**
- B. Sentiment Analysis**
- C. Text Classification**
- D. Image Recognition**

Sentiment analysis is the method specifically designed to determine the emotional tone behind a piece of text. This approach involves analyzing the language used in customer feedback to identify whether the sentiment conveyed is positive, negative, or neutral. By leveraging natural language processing techniques, sentiment analysis can help businesses gain insights into customer opinions and emotions, allowing them to respond appropriately to feedback. Key phrase extraction focuses on identifying significant terms or phrases within a text but does not provide insights into the emotional tone. Text classification allows for categorizing text into predefined groups but does not inherently assess the sentiment expressed. Image recognition pertains to analyzing visual content and is not applicable to text-based feedback. Thus, among the options presented, sentiment analysis is unequivocally the method most adept at understanding emotional tone in customer feedback.

7. Is "Quality Test" classified as a label in machine learning?

- A. Yes**
- B. No**
- C. Depends on context**
- D. Not defined**

In the context of machine learning, a label refers to the output or result that the algorithm predicts based on input data. Labels are used in supervised learning, where the model learns to map inputs to a specific output. "Quality Test" can indeed be classified as a label if it represents the desired classification or outcome in a predictive model. For instance, in a scenario where you're developing a model to predict product quality based on various features, "Quality Test" could denote the result of this prediction—such as pass or fail. If the context clearly identifies "Quality Test" as the predicted outcome derived from the input data, then it functions as a label within that machine learning framework. Labels categorize the training data and help in supervised learning tasks, making it crucial that they accurately represent what the model needs to predict. This understanding aligns with the broader function of labels in ML, making "Yes" the appropriate classification for "Quality Test" as a label.

8. Is it possible to use the Custom Vision service to analyze video files?

- A. Yes**
- B. No**
- C. Only under specific conditions**
- D. Only for short clips**

The rationale behind the ability of the Custom Vision service lies primarily in its focus on image classification and object detection rather than video analysis. Custom Vision is designed specifically to work with still images, where you can train models to recognize and classify images based on visual content. While one might consider that videos comprise a series of images, Custom Vision does not inherently support the capability to process video files directly. Thus, attempting to use it for analyzing video content would not align with the primary function of the service, leading to the conclusion that it cannot handle video analysis as a standard feature. In the realm of video analysis, other Azure services and tools are more appropriate, such as Azure Video Analyzer or Azure Media Services, which are built to handle video streams and provide analysis functions suited for moving images. This distinction is essential in understanding the limitations of the Custom Vision service concerning video files.

9. 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.

10. Organizing documents into groups based on textual similarities is an example of which machine learning concept?

- A. Classification**
- B. Regression**
- C. Clustering**
- D. Association**

Organizing documents into groups based on textual similarities is a clear example of clustering, which is a type of unsupervised machine learning. Clustering involves grouping data points that are similar to each other within the same cluster while keeping them distinct from those in other clusters. In this scenario, documents are analyzed to find patterns and similarities in their text, allowing the algorithm to organize them accordingly without having predefined categories or labels. This contrasts with classification, where the algorithm is trained on labeled data to predict a specific category for new instances. It also differs from regression, which focuses on predicting continuous values rather than grouping. Association refers to finding relationships between variables in large datasets, which does not apply to organizing documents based on textual similarities. Therefore, clustering is the most appropriate concept for this scenario.

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