

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. Can Automated ML infer training data from the use case provided?**
 - A. Yes
 - B. No
 - C. Only under certain conditions
 - D. It depends on the algorithm
- 2. What does a validation set typically include?**
 - A. Data to train the model
 - B. Data used to test the model's predictions
 - C. A random subset of the entire dataset
 - D. A set of hyperparameter configurations
- 3. When creating an object detection model in the Custom Vision service, must you choose a class type?**
 - A. Yes
 - B. No
 - C. Only for images with multiple objects
 - D. Only for animal detection
- 4. What method is used to predict next month's store sales?**
 - A. Natural language processing
 - B. Machine learning (Regression)
 - C. Computer vision
 - D. Anomaly detection
- 5. What type of analysis helps in categorizing data into predefined groups?**
 - A. Clustering
 - B. Classification
 - C. Regression
 - D. Association
- 6. 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

- 7. What aspect of Azure ML designer enhances productivity for users?**
- A. The ability to create complex ML algorithms automatically
 - B. The drag-and-drop feature for module connectivity
 - C. The restriction of functions to prevent errors
 - D. The function of visualizing data only
- 8. To analyze images for your charity event, which tool should you use to ensure photos meet specific requirements?**
- A. Face service's Detect operation
 - B. Computer vision API
 - C. Text analytics service
 - D. Image recognition technology
- 9. Determining if an image contains a person is an application of which AI technology?**
- A. Augmented reality
 - B. Machine learning
 - C. Computer vision
 - D. Deep learning
- 10. What do we call the ability to analyze and ascertain whether individuals within an image are similar?**
- A. Grouping
 - B. Identification
 - C. Similarity
 - D. Classification

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Can Automated ML infer training data from the use case provided?

- A. Yes
- B. No**
- C. Only under certain conditions
- D. It depends on the algorithm

Automated Machine Learning (AutoML) tools are designed to streamline the machine learning process, making it more accessible for users without deep expertise in data science. However, these systems require a defined dataset to train models effectively. Automated ML does not infer or generate training data based solely on a given use case; instead, it relies on existing data provided by the user to learn patterns and make predictions. The use case may help guide the selection of algorithms and approaches, but it does not substitute for the necessity of having adequate and relevant training data. Therefore, stating that Automated ML cannot infer training data from the use case provided accurately reflects its operational parameters. A deeper understanding of what Automated ML offers highlights its reliance on existing datasets rather than generating new data from the use case, making 'No' the right selection.

2. What does a validation set typically include?

- A. Data to train the model
- B. Data used to test the model's predictions**
- C. A random subset of the entire dataset
- D. A set of hyperparameter configurations

A validation set is an essential component in the process of building and evaluating a machine learning model. Specifically, it is used to assess how well the trained model performs on unseen data during the training process. The validation set typically includes a distinct portion of the dataset that is not used during the training phase but is instead set aside to evaluate the model's performance after it has been trained. This assessment helps in tuning the model's parameters and prevents overfitting to the training data. By examining how well the model predicts outcomes on the validation set, practitioners can gain insights into its generalization capabilities and make necessary adjustments to improve performance before final testing. In contrast, the other choices do not accurately describe the primary purpose of a validation set. For instance, while the training set is specifically used to train the model, and a random subset refers to partitioning methods rather than a defined role in model evaluation, hyperparameter configurations are more related to model tuning rather than comprising the validation set itself. Thus, the role of the validation set as a tool for testing model predictions is pivotal in the model development lifecycle.

3. When creating an object detection model in the Custom Vision service, must you choose a class type?

- A. Yes
- B. No**
- C. Only for images with multiple objects
- D. Only for animal detection

In the context of using the Custom Vision service to create an object detection model, it is not mandatory to choose a class type when starting the project. You can set up an object detection model and train it without specifying predefined class types at the beginning. This flexibility allows for a more exploratory approach to developing models, where users can iteratively refine their training data and add classes as they improve their models over time. Choosing a class type can be beneficial for organization and clarity, especially when the project involves distinct categories of objects. However, it is not a strict requirement to start the model training process. Even with images containing multiple objects or focusing on specific categories like animals, one can begin the training without predetermined classes and later finalize or adjust them based on the data collected. This capability emphasizes the Custom Vision service's adaptability in accommodating various workflows and preferences, making it easier for developers to initiate their object detection tasks.

4. What method is used to predict next month's store sales?

- A. Natural language processing
- B. Machine learning (Regression)**
- C. Computer vision
- D. Anomaly detection

The method used to predict next month's store sales is machine learning, specifically regression techniques. Regression is a statistical approach utilized to understand the relationship between a dependent variable, in this case, store sales, and one or more independent variables, which could include factors such as historical sales data, promotions, seasonality, and economic indicators. By employing regression models, one can analyze past sales patterns and trends to make informed predictions about future sales figures. This is crucial for businesses to manage inventory, optimize staffing, and plan financial strategies. The regression approach in machine learning enables the model to learn from historical data, making it capable of forecasting future sales quantities with a certain level of accuracy. This predictive capability is essential for effective business operations and decision-making. Other methodologies like natural language processing, computer vision, and anomaly detection serve different purposes and are not directly applicable to the task of forecasting store sales based on historical data.

5. What type of analysis helps in categorizing data into predefined groups?

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

The analysis that assists in categorizing data into predefined groups is classification. This technique is used in supervised machine learning where the model learns from labeled training data. Each piece of training data is associated with a specific category or label, enabling the model to understand the characteristics that define each category. Classification is particularly beneficial for tasks like spam detection in emails, where emails are categorized as "spam" or "not spam," or in medical diagnosis, where patient data is classified into disease categories based on certain symptoms and test results. In contrast, clustering is an unsupervised technique that groups data points based on their similarities but does not rely on predefined categories. Regression focuses on predicting continuous values based on input features rather than categorizing them. Association analysis identifies relationships between variables or items in large datasets, like market basket analysis, without assigning elements to predefined groups. Therefore, classification is the most appropriate method for the stated task of categorizing data into established groups.

6. 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.

7. What aspect of Azure ML designer enhances productivity for users?

- A. The ability to create complex ML algorithms automatically
- B. The drag-and-drop feature for module connectivity**
- C. The restriction of functions to prevent errors
- D. The function of visualizing data only

The drag-and-drop feature for module connectivity in Azure ML designer significantly enhances productivity for users by allowing them to easily construct and design machine learning workflows without requiring extensive coding skills. This intuitive interface enables users to visually assemble different components of a machine learning pipeline, such as data input modules, data transformation, modeling algorithms, and evaluation metrics, simply by dragging and dropping these elements onto the canvas. This approach not only accelerates the design process but also makes it accessible to individuals who may not be deeply familiar with programming or data science concepts. Consequently, users can focus more on the experimentation and optimization of their models rather than spending excessive time on coding syntax or complex configuration. In contrast, while creating complex ML algorithms automatically sounds appealing, Azure ML designer mainly provides tools and modules that allow users to build algorithms using pre-defined components rather than generating them all automatically. The restriction of functions to prevent errors is beneficial, but it might pose limitations that can slow down the workflow. Likewise, visualizing data is an essential aspect but is only one part of a larger process, and not the primary feature that boosts overall productivity in the designer.

8. To analyze images for your charity event, which tool should you use to ensure photos meet specific requirements?

- A. Face service's Detect operation**
- B. Computer vision API
- C. Text analytics service
- D. Image recognition technology

The Face service's Detect operation is specifically designed to analyze faces within images, making it suitable for tasks such as identifying whether the conditions related to facial presence, orientation, and other attributes are met. This tool uses advanced algorithms to provide detailed information about the detected faces in images, which aligns well with the need to ensure that photos meet specific requirements, particularly if the requirements pertain to appearances or facial features. In contrast, the Computer Vision API generally focuses on broader image analysis, including the extraction of various features, object recognition, and understanding content beyond just faces. While it is a powerful tool for analyzing images, it may not specifically cater to the requirements related to facial analysis as efficiently as the Face service's Detect operation. The Text Analytics service specializes in processing and analyzing text data, extracting insights like sentiment and key phrases, and is not applicable for image analysis tasks. Image recognition technology is a broad term that encompasses various methods and tools for identifying and classifying objects within images, but without a specific implementation like the Face service's Detect operation, it would lack the targeted capabilities applicable for ensuring specific requirements intended for individual faces in photographs.

9. Determining if an image contains a person is an application of which AI technology?

- A. Augmented reality
- B. Machine learning
- C. Computer vision**
- D. Deep learning

Determining if an image contains a person falls under computer vision, which is a field of AI focused on enabling machines to interpret and understand visual information from the world. Computer vision involves using algorithms to process and analyze images, allowing systems to identify objects, classify them, and ultimately understand their context within the visual data. It is particularly relevant to tasks such as image recognition, object detection, and scene understanding. In this case, the task of recognizing a person in an image requires the application of techniques associated with computer vision. These techniques often integrate with machine learning models that have been trained to identify specific features associated with human figures, but it is the broader category of computer vision that encompasses the entire process of analyzing and interpreting the image. While machine learning is a foundational technology that enables systems to learn from data, it must be paired with computer vision techniques to specifically handle visual input. Deep learning, which is a subset of machine learning, can be employed for specific tasks within computer vision, such as using neural networks to improve the accuracy of image classification. Augmented reality, on the other hand, relates to the interaction between real-world and digital elements, which is not directly relevant to the task of detecting a person in an image. Thus, computer vision is

10. What do we call the ability to analyze and ascertain whether individuals within an image are similar?

- A. Grouping
- B. Identification
- C. Similarity**
- D. Classification

The ability to analyze and ascertain whether individuals within an image are similar is referred to as similarity. This concept focuses on measuring the likeness or resemblance between different objects—in this case, individuals within an image. In the context of image analysis and recognition, evaluating similarity is crucial for tasks such as face recognition, where the system determines how closely two faces match based on their features. This can include analyzing aspects like facial structure, skin texture, and other identifying characteristics. Grouping, identification, and classification, while related to the broader scope of image analysis and machine learning, do not specifically capture the essence of analyzing similarity. Grouping typically refers to organizing items based on shared characteristics, identification is about recognizing or tagging an individual or item, and classification involves categorizing data into predefined classes. While these processes may involve some element of similarity analysis, they address different objectives and methodologies compared to the direct assessment of how alike individuals are.

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