

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

STUDY GUIDE



Prepare with structure. Pass with confidence



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. You are designing an AI system that empowers everyone, including people who have hearing, visual, and other impairments. This is an example of which Microsoft guiding principle for responsible AI?**
 - A. Inclusiveness
 - B. Transparency
 - C. Accountability
 - D. Fairness
- 2. What guiding principle would indicate the need for improvements in an AI system based on user feedback?**
 - A. Inclusiveness
 - B. Reliability and safety
 - C. Transparency
 - D. Fairness
- 3. A medical research project uses machine learning for early detection of different brain haemorrhage types in images. This is an example of what type of machine learning?**
 - A. Regression
 - B. Classification
 - C. Clustering
 - D. Association
- 4. Which term describes the calculated probability of a correct image classification?**
 - A. Precision
 - B. Recall
 - C. Accuracy
 - D. F1 Score
- 5. Can the Custom Vision service be utilized to detect objects within an image?**
 - A. Yes
 - B. No
 - C. Only in predefined categories
 - D. Only with a third-party library

- 6. 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
- 7. What type of machine learning involves predicting a continuous output?**
- A. Regression
 - B. Classification
 - C. Clustering
 - D. Dimensionality reduction
- 8. What technology should be used to extract text from handwritten documents?**
- A. Text Analytics
 - B. Optical Character Recognition (OCR)
 - C. Speech Recognition
 - D. Language Understanding
- 9. What type of AI is used to adapt and improve responses based on customer interactions in chatbots?**
- A. Unsupervised Learning
 - B. Reinforcement Learning
 - C. Natural Language Processing
 - D. Generative Models
- 10. What is an example of knowledge mining in an AI context?**
- A. Predicting user behavior
 - B. Automated chatbot answering refund questions
 - C. Classifying images using neural networks
 - D. Analyzing stock prices using regression

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. You are designing an AI system that empowers everyone, including people who have hearing, visual, and other impairments. This is an example of which Microsoft guiding principle for responsible AI?

- A. Inclusiveness
- B. Transparency
- C. Accountability
- D. Fairness

The choice of inclusiveness aligns perfectly with the goal of designing an AI system that empowers individuals with varying abilities, including those with hearing, visual, and other impairments. Inclusiveness is a guiding principle that emphasizes the importance of ensuring that AI technologies are accessible and beneficial to all members of society, regardless of their physical, cognitive, or emotional capabilities. This principle advocates for creating systems that consider diverse user needs and experiences, ultimately fostering a more equitable environment. By focusing on inclusiveness, the AI system not only addresses the barriers faced by people with disabilities but also emphasizes the value of diverse perspectives and participation in the development and deployment of AI technologies. This commitment ensures that the benefits of AI are available to everyone and helps in building trust and promoting responsible innovation within the community. The other principles, while essential in the context of responsible AI, do not specifically capture the essence of creating systems that cater to individuals with different abilities as inclusiveness does.

2. What guiding principle would indicate the need for improvements in an AI system based on user feedback?

- A. Inclusiveness
- B. Reliability and safety
- C. Transparency
- D. Fairness

The principle of reliability and safety is crucial for indicating the need for improvements in an AI system based on user feedback. This principle emphasizes that an AI system should be dependable and operate safely in various conditions. When users provide feedback that highlights issues such as errors, unpredictability, or safety concerns, this feedback signals that the system may not be meeting the reliability and safety expectations. Improving reliability ensures that the AI system performs its intended function consistently while prioritizing user safety. By focusing on reliability and safety, developers can address any shortcomings highlighted by users, leading to a more robust and trustworthy AI product. Thus, the guiding principle of reliability and safety directly relates to how user feedback can drive enhancements in the AI system's performance and user experience.

3. A medical research project uses machine learning for early detection of different brain haemorrhage types in images. This is an example of what type of machine learning?

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

This example illustrates the use of classification in machine learning. In this context, the goal is to detect and categorize different types of brain hemorrhages based on image data. Classification tasks involve training a model on labeled datasets, where the labels correspond to specific categories or classes—in this case, various types of brain hemorrhages. Once the model is trained, it can analyze new, unseen images and predict which category they belong to. In medical imaging, this approach is particularly valuable because it allows for swift and accurate diagnosis, thereby supporting timely interventions. The classification process deals with distinct outcomes, making it essential in applications like medical diagnosis, where the distinction among various conditions is crucial for effective treatment and intervention strategies.

4. Which term describes the calculated probability of a correct image classification?

- A. Precision
- B. Recall
- C. Accuracy**
- D. F1 Score

The term that describes the calculated probability of a correct image classification is accuracy. Accuracy measures the proportion of true results (both true positives and true negatives) among the total number of cases examined. In the context of image classification, it indicates how often the model correctly identifies the images out of all the images it processed. For instance, if a model classifies 80 out of 100 images correctly, its accuracy would be 80%. This metric provides a straightforward way to understand the effectiveness of the model in making correct predictions, which is particularly useful in scenarios where the classes are balanced. Precision, recall, and F1 score are also important metrics used to evaluate model performance, but they serve different purposes. Precision focuses on the correctness of positive predictions, recall measures the model's ability to identify all relevant instances, and the F1 score is the harmonic mean of precision and recall, providing a balance between the two. However, none of these terms directly reflect the overall probability of correct classification across both classes, making accuracy the clear choice for this specific inquiry.

5. Can the Custom Vision service be utilized to detect objects within an image?

- A. Yes**
- B. No
- C. Only in predefined categories
- D. Only with a third-party library

The Custom Vision service is designed to enable users to build and customize machine learning models tailored to their specific image classification and object detection needs. The primary purpose of this service is to train models using labeled images that users provide, enabling the detection of objects within images. This means that users can upload images, tag the specific objects of interest, and then train a model that learns to recognize those objects in new, unseen images. One of the key features of the Custom Vision service is its flexibility, allowing users to create custom tags and categories for their specific use cases. This is particularly valuable in environments where predefined categories may not suffice, as users can define their own categories based on their unique requirements. This makes the capability to detect objects within an image not only feasible but a core function of the Custom Vision service. It empowers businesses and developers to leverage machine learning for applications such as retail inventory management, quality control in manufacturing, and many other scenarios that require object detection.

6. 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.

7. What type of machine learning involves predicting a continuous output?

- A. Regression
- B. Classification
- C. Clustering
- D. Dimensionality reduction

Regression is the type of machine learning that focuses on predicting a continuous output. In regression tasks, the goal is to model the relationship between one or more independent variables (also known as features or predictors) and a dependent variable (the output). For instance, predicting the price of a house based on its characteristics such as size, location, and age falls under regression because the price is a continuous numerical value. This method utilizes various algorithms to understand how the dependent variable changes with changes in the independent variables, enabling predictions for new, unseen data points. Common regression techniques include linear regression, polynomial regression, and more advanced models like regression trees and neural networks designed for regression tasks. In contrast, classification involves predicting discrete labels or categories, clustering focuses on grouping similar data points without predefined labels, and dimensionality reduction aims to simplify datasets by reducing the number of features while preserving essential information. These concepts differentiate them from regression, which is specifically about forecasting continuous outputs.

8. What technology should be used to extract text from handwritten documents?

- A. Text Analytics
- B. Optical Character Recognition (OCR)
- C. Speech Recognition
- D. Language Understanding

Optical Character Recognition (OCR) is specifically designed to convert different types of documents, including handwritten notes, into machine-readable text. This technology uses pattern recognition and machine learning techniques to identify characters in scanned images of documents. The ability to process handwriting is particularly useful in contexts such as digitizing historical documents or automating data entry from handwritten forms. Text Analytics is focused more on extracting insights and understanding content from structured text rather than converting images of handwritten documents into text. Speech Recognition, on the other hand, deals with converting spoken language into text through voice input rather than written formats. Language Understanding is concerned with comprehending the intent behind text in natural language processing applications and does not perform character recognition tasks like OCR does. Thus, OCR stands out as the most appropriate technology for extracting text from handwritten documents.

9. What type of AI is used to adapt and improve responses based on customer interactions in chatbots?

- A. Unsupervised Learning
- B. Reinforcement Learning**
- C. Natural Language Processing
- D. Generative Models

The type of AI that is utilized to adapt and improve responses based on customer interactions in chatbots is reinforcement learning. This approach allows systems to learn from their own actions and experiences in an environment by receiving feedback in the form of rewards or penalties. In the context of chatbots, reinforcement learning enables the system to evaluate its conversational responses and understand which interactions lead to successful outcomes (positive feedback) or unsatisfactory interactions (negative feedback). As a result, the chatbot can continuously improve its performance and make more informed decisions in future conversations, leading to better user experiences. The other options, while related to AI, serve different purposes. Unsupervised learning focuses on finding patterns or groupings in data without labeled outputs, which is not specifically aimed at adapting responses over time. Natural Language Processing is essential for enabling chatbots to understand and generate human language but does not inherently include the learning mechanism to adapt responses based on interactions. Generative models, while capable of creating new content, do not primarily focus on improving responses based on feedback from interactions in the way that reinforcement learning does.

10. What is an example of knowledge mining in an AI context?

- A. Predicting user behavior
- B. Automated chatbot answering refund questions**
- C. Classifying images using neural networks
- D. Analyzing stock prices using regression

In an AI context, knowledge mining refers to the process of extracting valuable insights and patterns from large amounts of unstructured or semi-structured data. The ability to utilize natural language processing (NLP) to understand and respond to user inquiries is a key aspect of knowledge mining. Automated chatbots that can answer refund questions exemplify knowledge mining, as they not only retrieve specific information from a database but also interpret questions and generate appropriate responses. In this scenario, the chatbot leverages previous interactions and knowledge bases to provide answers, making it a practical application of knowledge mining—transforming unstructured queries into structured, actionable insights for users. The other options, while they represent important AI capabilities, do not focus primarily on the extraction of knowledge from unstructured data. Predicting user behavior involves machine learning models that analyze patterns in structured data. Classifying images utilizes neural networks for pattern recognition in image data, and analyzing stock prices typically involves using statistical methods on structured datasets. Hence, they do not encapsulate the broader concept of knowledge mining in the same way as the chatbot scenario does.

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