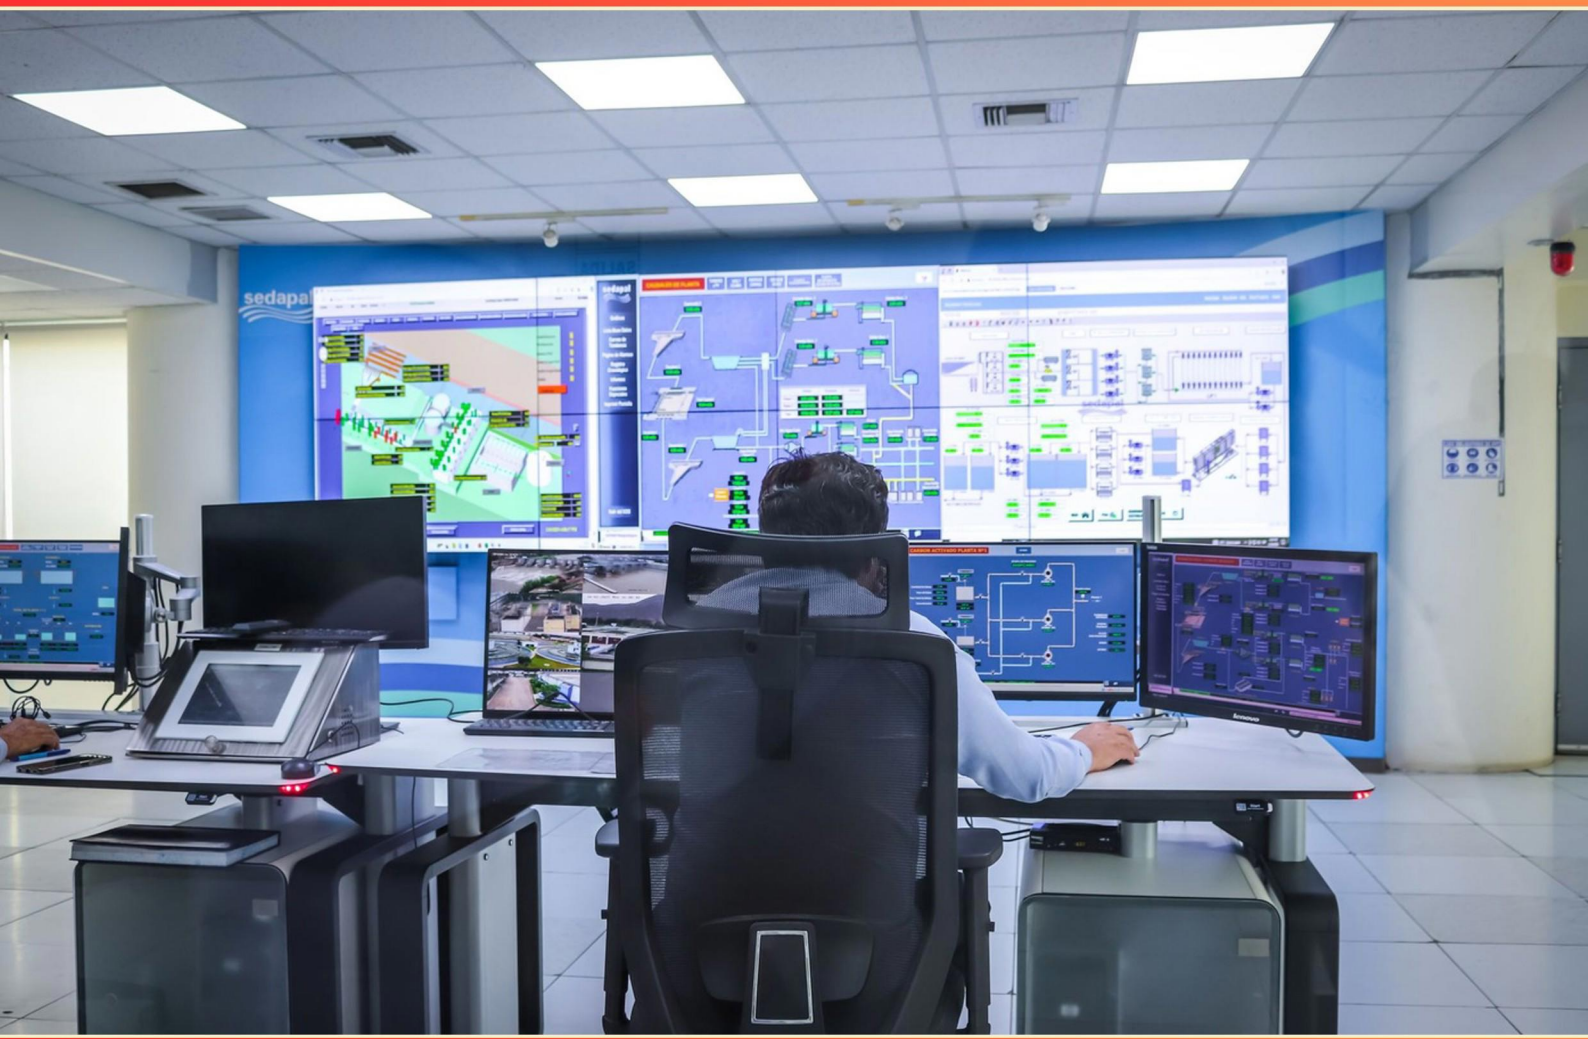
MICROSOFT AZURE AI FUNDAMENTALS (AI-900) PRACTICE EXAM (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR THE FULL VERSION STUDY GUIDE

https://microsoftazureaifundamentals-ai900.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.

SAMPLE

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!

SAMPLE

Questions

SAMPLE

1. **When developing an AI system for self-driving cars, which principle should be applied to ensure consistent operation during unexpected circumstances?**

   A. Transparency

   B. Reliability and safety

   C. Inclusiveness

   D. Accountability

2. **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

3. **Is forecasting housing prices based on historical data a valid example of anomaly detection?**

   A. Yes

   B. No

   C. Only in certain cases

   D. It depends on the data source

4. **In Azure ML designer, where must you deploy a model to create a real-time inference pipeline?**

   A. Azure Functions

   B. Azure Kubernetes Service (AKS)

   C. Azure Blob Storage

   D. Azure Virtual Machines

5. **In a classification model, are labels required to be numeric?**

   A. Yes

   B. No

   C. Only for binary classification

   D. Only for multiclass classification

SAMPLE

6. What does the privacy and security principle emphasize regarding personal data?

A. It must be anonymized

B. It must be visible only to approved users

C. It must be shared with third parties

D. It must be collected for analysis

7. Does a validation set verify that all training data was used to train the model?

A. Yes

B. No

C. Only in supervised learning

D. Depends on the machine learning algorithm

8. What does Automated ML enable users to specify for predictions?

A. A dataset for training only

B. A dataset and label to predict

C. Only the labels to predict

D. Predefined features and parameters

9. What type of AI workload assists in identifying negative sentiments from customer chats?

A. Predictive Analytics

B. Natural Language Processing

C. Supervised Learning

D. Image Processing

10. Which principle dictates that AI systems must not discriminate based on gender or race?

A. Inclusiveness

B. Fairness

C. Privacy and security

D. Reliability and safety

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. When developing an AI system for self-driving cars, which principle should be applied to ensure consistent operation during unexpected circumstances?

A. Transparency

B. Reliability and safety

C. Inclusiveness

D. Accountability

In the context of developing AI systems for self-driving cars, applying the principle of reliability and safety is crucial to ensure that the vehicle can react appropriately to unexpected circumstances. This principle focuses on creating a system that is dependable under various conditions, including those that are unforeseen or potentially hazardous. Reliability ensures that the AI system performs consistently, regardless of variations in the environment, such as adverse weather conditions, sudden obstacles, or erratic behavior from other road users. Safety is paramount, as self-driving cars must not only operate correctly under normal conditions but also manage risks effectively to protect passengers, pedestrians, and other road users. This principle involves rigorous testing, validation, and adherence to safety standards throughout the development process. It ensures that the AI can make accurate decisions even in situations it has not previously encountered, thereby instilling confidence in its operation. While the other principles, such as transparency, inclusiveness, and accountability, are also important in the broader context of ethical AI and its implementation, they are not as directly focused on the critical operational requirements necessary for self-driving technology to function effectively and safely in unpredictable scenarios. Reliability and safety take precedence in this specific use case, making it the most relevant principle for this question.

2. 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.

SAMPLE

3. Is forecasting housing prices based on historical data a valid example of anomaly detection?

A. Yes

B. No

C. Only in certain cases

D. It depends on the data source

Forecasting housing prices based on historical data is primarily focused on predicting future trends and patterns rather than identifying irregular or outlier data points within the existing dataset. Anomaly detection typically aims to identify data points that significantly differ from the majority of data, indicating unusual behavior that requires further investigation. In contrast, forecasting analyzes historical trends to generate predictions, without specifically targeting any anomalies present in the dataset. This distinction makes the assertion that forecasting housing prices is an example of anomaly detection inaccurate. Hence, the most relevant answer is that it is not a valid example of anomaly detection.

4. In Azure ML designer, where must you deploy a model to create a real-time inference pipeline?

A. Azure Functions

B. Azure Kubernetes Service (AKS)

C. Azure Blob Storage

D. Azure Virtual Machines

To create a real-time inference pipeline in Azure ML designer, deploying the model on Azure Kubernetes Service (AKS) is the optimal choice. AKS is a managed Kubernetes service that simplifies the deployment, management, and scaling of containerized applications, including machine learning models. When using AKS, models can be deployed as containerized applications that provide RESTful endpoints for real-time inference. This capability is crucial for scenarios where low latency and high availability are required, allowing applications to send data to the model and receive predictions almost instantly. AKS also supports auto-scaling, enabling resources to be adjusted dynamically based on demand, which is essential for handling varying workloads effectively. In contrast, while Azure Functions can offer serverless capabilities for running code in response to events, they are not specifically tailored for the deployment of complex machine learning models that require robust scaling and orchestration. Azure Blob Storage serves primarily as storage for data and models, but does not provide the necessary infrastructure for real-time inference. Azure Virtual Machines allow for hosting applications, but they lack the container orchestration benefits provided by AKS, making them less efficient for deploying machine learning models in a scalable, responsive manner.

SAMPLE

5. In a classification model, are labels required to be numeric?

A. Yes

B. No

C. Only for binary classification

D. Only for multiclass classification

In a classification model, labels do not need to be numeric as they can also be categorical. Classification models are designed to sort data into predefined classes or categories, and these labels can be represented as strings or any other categorical types, not just numbers. For example, in a model trained to classify types of animals, the labels could be "cat," "dog," and "bird." These categorical labels provide meaningful identification of classes to the model. During preprocessing, if numerical representations are necessary for algorithms or libraries being used—especially those that expect numeric input—these categorical labels can be encoded into numeric formats. However, this encoding is a step of data preprocessing, not a requirement of the classification model itself. Thus, the core function of classification is to allocate instances to classes based on their features, and this can effectively be done with non-numeric labels too.

6. What does the privacy and security principle emphasize regarding personal data?

A. It must be anonymized

B. It must be visible only to approved users

C. It must be shared with third parties

D. It must be collected for analysis

The privacy and security principle emphasizes that personal data must be protected and visible only to approved users. This means that organizations handling personal data should implement measures to ensure that only authorized individuals have access to sensitive information. This principle aligns with various privacy regulations and frameworks, which advocate for safeguarding individuals' personal information against unauthorized access and potential misuse. By restricting visibility to approved users, organizations can help maintain trust and ensure compliance with legal and ethical standards surrounding data protection. Other options, while they touch on aspects of data handling, do not capture the essence of the privacy and security principle as effectively. Anonymization and sharing with third parties may not always be applicable or suitable depending on the context and regulatory requirements. Collecting data for analysis is a common practice, but it does not inherently address privacy and security concerns regarding personal data access.

SAMPLE

7. Does a validation set verify that all training data was used to train the model?

A. Yes

B. No

C. Only in supervised learning

D. Depends on the machine learning algorithm

A validation set serves a critical role in assessing the performance of a machine learning model, but it does not verify that all training data was used to train the model. Instead, the primary function of a validation set is to evaluate how well the model can generalize to unseen data. During the training process, the model learns from the training data, and the validation set allows for tuning the model's parameters and making adjustments to improve performance. When a model is trained, it utilizes the training dataset. The validation dataset, which is separate from the training dataset, is used to test the model's predictions and overall performance. This separation ensures that the evaluation of the model is done independently of the data it was trained on. If the validation data were part of the training data, it would lead to overfitting, where the model performs well on the training data but poorly on new, unseen data. Thus, while a validation set is an essential tool in the model evaluation process, it does not serve to confirm the usage of all training data. It specifically focuses on validating the model's performance based on a different set of data.

8. What does Automated ML enable users to specify for predictions?

A. A dataset for training only

B. A dataset and label to predict

C. Only the labels to predict

D. Predefined features and parameters

Automated ML in Azure enables users to specify both a dataset and a label to predict, which is fundamental for supervised machine learning tasks. When you select a dataset, you provide the raw input data that will be used for training the model. The label, on the other hand, represents the target variable or outcome that you want the model to predict based on the input features in your dataset. This combination allows Automated ML to understand what to learn from the data and what predictions to make, streamlining the process of model creation and enabling non-experts to utilize machine learning effectively. In contrast, specifying only a dataset for training would limit the model's ability to understand what it is being trained to predict. Providing only the labels to predict without the corresponding dataset would not furnish the model with the necessary context or input data for making accurate predictions. Finally, focusing solely on predefined features and parameters overlooks the flexibility of including the entire dataset and target label for more robust model training. Thus, the option of specifying both a dataset and a label to predict is crucial for the effectiveness of Automated ML processes.

SAMPLE

9. What type of AI workload assists in identifying negative sentiments from customer chats?

A. Predictive Analytics

B. Natural Language Processing

C. Supervised Learning

D. Image Processing

The identification of negative sentiments from customer chats falls under the domain of Natural Language Processing (NLP). NLP is a branch of artificial intelligence that focuses on the interaction between computers and humans through natural language. It involves the analysis, understanding, and generation of human language in a valuable way. When assessing customer chats for sentiment, NLP algorithms are employed to analyze the text, understand the context, and determine the sentiment expressed, whether it's positive, negative, or neutral. This process often involves various techniques like sentiment analysis, which uses machine learning models trained on labeled data to recognize patterns associated with different sentiments. In contrast, predictive analytics involves forecasting future outcomes based on historical data and patterns rather than analyzing the sentiment of textual information. Supervised learning is a type of machine learning where models are trained on labeled datasets, which can indeed incorporate sentiment analysis, but it doesn't specifically denote the application to chat analysis. Image processing is focused on manipulating and analyzing visual data, which is not relevant to identifying sentiments in text. Therefore, Natural Language Processing is the most appropriate choice for this type of workload.

10. Which principle dictates that AI systems must not discriminate based on gender or race?

A. Inclusiveness

B. Fairness

C. Privacy and security

D. Reliability and safety

The principle that dictates that AI systems must not discriminate based on gender or race is fairness. Fairness in AI ensures that the decisions made by AI systems do not favor one group over another, particularly marginalized or historically disadvantaged groups. It emphasizes the need for equitable treatment and the importance of mitigating bias in algorithms and data that could lead to discriminatory outcomes. To achieve fairness, AI systems must be designed and trained with diverse datasets that accurately represent the population they serve, taking care to eliminate potential biases that may arise from skewed data representations. This consideration is vital in applications ranging from hiring algorithms to loan approvals, where biased decisions could impact individuals' lives significantly. While inclusiveness is related to fairness, focusing on representation and involvement, it does not specifically address the necessity of non-discrimination as clearly as fairness does. The other principles, such as privacy and security, and reliability and safety, are important in their own rights but do not specifically pertain to discrimination issues within AI systems.

SAMPLE

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