AI PROMPT ENGINEERING AND KEY CONCEPTS IN MACHINE LEARNING AND NLP PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR THE FULL VERSION STUDY GUIDE

https://aipromptengrmachinelearningnlp.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 15

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. What is dataset leakage in evaluation sets, and why is it problematic?

A. Leakage occurs when evaluation data never overlaps with training data.

B. Leakage refers to data leakage prevention measures.

C. Leakage is when the model's predictions are leaked to the user.

D. Leakage occurs when evaluation data overlaps with training data, inflating scores.

2. Which method uses a small number of examples to guide the model's responses?

A. Few-shot

B. Zero-shot

C. Cognitive verifier pattern

D. Chain-of-thought (CoT)

3. What is the graphical representation of data and information used to communicate insights?

A. Data visualization

B. Data

C. Preprocessing

D. Tool creation

4. Explain token budget management in chat prompts and strategies to maximize information within it.

A. Plan tokens for system, user, and assistant; use concise prompts, clear formatting, and summarization; reuse context with shorter history; consider compression and retrieval.

B. Always include the full conversation history in every message.

C. Discard older content to save tokens.

D. Use lengthy system messages to reserve space.

5. Which term is defined as 'a parameter in the LLM that influences the level of randomness or creativity in generated outputs'?

A. Specificity

B. Temperature

C. False Positive

D. True Negative

SAMPLE

6. Which statement best describes zero-shot prompting in language models?

A. Zero-shot uses no examples to guide the model.

B. Zero-shot uses a single example to demonstrate the desired format.

C. Zero-shot relies on chain-of-thought reasoning.

D. Zero-shot requires fine-tuning on task data.

7. Which statement correctly contrasts instruction-tuning and prompt-tuning?

A. Instruction-tuning trains the model on many tasks described by natural-language instructions; prompt-tuning trains soft prompt vectors prepended to inputs while keeping the base model fixed.

B. Instruction-tuning updates the model weights; prompt-tuning updates the training data.

C. Both involve changing the base model weights; none involve prompts.

D. Instruction-tuning applies only to supervised datasets; prompt-tuning applies only to unsupervised data.

8. Which set of prompting methods is typically described as including few-shot and zero-shot prompting?

A. Advanced Prompting Technique

B. Self-Consistency

C. Generated Knowledge Prompting

D. Iterative Prompting

9. What is model drift and its implications for prompt maintenance?

A. Data Drift Is The Only Issue; Prompts Don't Need Maintenance

B. Hardware Failures Are The Main Concern For Prompts

C. Model Updates Change Behavior, Prompt Maintenance Required

D. Models Never Drift After Deployment

10. Which term describes the presence of unfair or prejudiced outcomes in AI systems due to biases in data, algorithms, or design?

A. Bias

B. Ethics

C. Generated Knowledge Prompting

D. Iterative Prompting

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is dataset leakage in evaluation sets, and why is it problematic?

A. Leakage occurs when evaluation data never overlaps with training data.

B. Leakage refers to data leakage prevention measures.

C. Leakage is when the model's predictions are leaked to the user.

D. Leakage occurs when evaluation data overlaps with training data, inflating scores.

Dataset leakage in evaluation sets happens when information from the evaluation data ends up available during training. If some evaluation examples are present in the training data, or if preprocessing and feature engineering are computed using the full dataset (including the evaluation portion), the model can learn from data it should be tested on. This makes the model appear to perform better than it would on truly unseen data because it has effectively memorized or exploited information from the test set. Why this is a problem is that evaluation is meant to measure how well the model generalizes to new, unseen data. Leakage creates an inflated performance estimate, which can mislead you into choosing models that won't generalize in real use. It also masks the true difficulty of the task and undermines trust in the evaluation results. To prevent leakage, keep a strict train/validation/test separation, ensure preprocessing and feature engineering are fitted only on the training data, and use time- or event-appropriate splits for sequential data. This helps ensure the evaluation reflects genuine generalization.

2. Which method uses a small number of examples to guide the model's responses?

A. Few-shot

B. Zero-shot

C. Cognitive verifier pattern

D. Chain-of-thought (CoT)

Few-shot prompting uses a small number of examples included in the prompt to steer the model's outputs. By showing a handful of input-output pairs, you demonstrate the exact pattern, format, or style you want the model to follow. The model treats those examples as demonstrations and generalizes from them to handle new inputs in the same way. This is a form of in-context learning—no changes to the model's weights are needed, just the prompt you provide. This differs from zero-shot, which asks the model to perform the task based only on the instruction, with no illustrative examples. Chain-of-thought prompts aim to elicit a step-by-step reasoning process in the model's output, focusing on internal reasoning rather than simply guiding the final answer through examples. The cognitive verifier pattern isn't a standard, widely used term for this particular mechanism of guiding behavior with a few examples, so it doesn't capture the same idea as few-shot prompting.

3. What is the graphical representation of data and information used to communicate insights?

A. Data visualization

B. Data

C. Preprocessing

D. Tool creation

Data visualization is the graphical representation of data and information used to communicate insights. By turning numbers into visuals like charts, graphs, and dashboards, it makes patterns, trends, and relationships easy to spot, helping people understand complex data quickly and make informed decisions. The other terms refer to different ideas: data themselves are raw facts, preprocessing is cleaning and shaping that data before analysis, and tool creation is about building software or systems. So, presenting insights through visuals is exactly what data visualization is.

SAMPLE

4. Explain token budget management in chat prompts and strategies to maximize information within it.

A. Plan tokens for system, user, and assistant; use concise prompts, clear formatting, and summarization; reuse context with shorter history; consider compression and retrieval.

B. Always include the full conversation history in every message.

C. Discard older content to save tokens.

D. Use lengthy system messages to reserve space.

Token budget management is about working within a limited number of tokens for each chat turn, balancing tokens across the system instructions, the user's message, and the model's reply. The goal is to convey as much useful information as possible without hitting the limit, so the model can respond accurately and completely. To maximize information within the budget, plan how many tokens to allocate to the system, user, and assistant portions. Keep prompts concise and free of unnecessary filler, and use clear formatting so the model can quickly parse structure and intent. Summarize long past interactions instead of reprinting them in full, preserving essential context while saving tokens. Reuse context with shorter history by maintaining a compact, meaningful thread and rewriting earlier content into brief summaries that capture the key points. Consider compression and retrieval: store relevant knowledge externally or in compressed representations and fetch only the pieces needed for the current query, rather than carrying everything in the prompt. This approach is effective because it preserves the model's ability to follow instructions and maintain continuity, while freeing tokens to provide a complete, high-quality response. By avoiding verbose repetition and unnecessary system messages, you prevent token waste. In contrast, including the full conversation history uses token capacity that could be spent on a better answer, discarding older content bluntly can sever important context, and lengthy system messages waste tokens that could be used for substantive content.

5. Which term is defined as 'a parameter in the LLM that influences the level of randomness or creativity in generated outputs'?

A. Specificity

B. Temperature

C. False Positive

D. True Negative

Temperature controls randomness in the outputs of a language model. It works by shaping the probability distribution the model uses to pick the next word. When you increase the temperature, the distribution becomes flatter, so the model is more willing to select less likely words, which boosts creativity and variety but can reduce coherence. Lowering the temperature makes the distribution peakier, so the model sticks to high-probability words and produces more predictable, focused responses. Specificity isn't a standard term for this kind of behavior in language models, and false positive or true negative are terms from binary classification and evaluation, not generation settings. So the parameter described—the one that modulates randomness or creativity—is temperature.

SAMPLE

6. Which statement best describes zero-shot prompting in language models?

A. Zero-shot uses no examples to guide the model.

B. Zero-shot uses a single example to demonstrate the desired format.

C. Zero-shot relies on chain-of-thought reasoning.

D. Zero-shot requires fine-tuning on task data.

Zero-shot prompting tests the model's ability to perform a task using only the task description or instruction, with no demonstration examples given in the prompt. It relies on the model's pre-trained knowledge to generalize from that instruction alone. Providing no examples to guide the model is the defining feature, which is exactly what zero-shot means. Using a single example would be demonstrating the task format, which is not zero-shot—it's a one-shot (or few-shot) setup. Chain-of-thought reasoning describes including a step-by-step line of thought, which can be used in various prompting styles and isn't what defines zero-shot. Fine-tuning on task data changes the model's parameters through training; zero-shot inference uses the fixed model without task-specific training, so it isn't about training either.

7. Which statement correctly contrasts instruction-tuning and prompt-tuning?

A. Instruction-tuning trains the model on many tasks described by natural-language instructions; prompt-tuning trains soft prompt vectors prepended to inputs while keeping the base model fixed.

B. Instruction-tuning updates the model weights; prompt-tuning updates the training data.

C. Both involve changing the base model weights; none involve prompts.

D. Instruction-tuning applies only to supervised datasets; prompt-tuning applies only to unsupervised data.

The main idea is to distinguish how each method shapes the model's behavior. Instruction-tuning broadens the model's ability to follow natural-language instructions by updating the model's own parameters through training on many tasks described by instructions. This means the internal weights are adjusted so the model learns a general capability to perform a variety of tasks when given instructions. Prompt-tuning, on the other hand, keeps the base model fixed and learns a lightweight set of prompt parameters (soft prompts) that are prepended to inputs. These prompts steer the model's outputs without changing the underlying weights, making the adaptation more about conditioning rather than retraining the model. So the statement that instruction-tuning updates the model weights while prompt-tuning updates the prompts is the correct contrast. It's also incorrect to say prompt-tuning changes training data or that both methods modify the base weights, and it's inaccurate to claim one method is restricted to supervised data while the other only uses unsupervised data.

SAMPLE

8. Which set of prompting methods is typically described as including few-shot and zero-shot prompting?

A. Advanced Prompting Technique

B. Self-Consistency

C. Generated Knowledge Prompting

D. Iterative Prompting

Few-shot and zero-shot prompting show how we guide a model's responses purely through prompt design. In zero-shot prompting, you give the task with no examples, relying on the model's ability to generalize from instructions alone. In few-shot prompting, you provide a small set of example pairs to illustrate the task and pattern, helping the model infer what to do. Both are grouped under advanced prompting techniques because they rely on crafting the input to steer behavior rather than altering the model itself. Other methods focus on different angles: self-consistency uses multiple reasoning paths to improve reliability, iterative prompting refines the prompt or the interaction over turns, and generated knowledge prompting centers on bringing in or creating factual content within the prompt. So the set that typically includes few-shot and zero-shot prompting is advanced prompting techniques.

9. What is model drift and its implications for prompt maintenance?

A. Data Drift Is The Only Issue; Prompts Don't Need Maintenance

B. Hardware Failures Are The Main Concern For Prompts

C. Model Updates Change Behavior, Prompt Maintenance Required

D. Models Never Drift After Deployment

Model drift means the model's behavior can change over time, especially after updates, retraining, or new data being incorporated. This matters for prompts because the same prompt that guided the model to produce a desired result can start behaving differently once a new model version is deployed or the training data shifts. To keep outputs aligned with goals, prompts need maintenance: you adjust wording, add system messages or constraints, and sometimes revise the prompt template so the intended intent still leads to the correct behavior with the updated model. This also often involves versioning prompts, running tests, and monitoring outputs to catch regressions early. Hardware issues aren't the main factor here, and while data drift matters, the focus in this context is that changes to the model itself drive changes in how prompts behave, so ongoing prompt maintenance is necessary.

10. Which term describes the presence of unfair or prejudiced outcomes in AI systems due to biases in data, algorithms, or design?

A. Bias

B. Ethics

C. Generated Knowledge Prompting

D. Iterative Prompting

Bias in AI refers to the presence of unfair or prejudiced outcomes that arise when data, algorithms, or design choices reflect or amplify societal prejudices. If training data underrepresents a group, the model may perform worse for that group, leading to discriminatory predictions. Similarly, optimization that focuses on overall accuracy can hide harms to minority groups, producing biased results. Design decisions, such as which features to include or how thresholds are set, can also embed these biases into the system. This is why the term that describes this phenomenon is bias. Ethics deals with what ought to be, not the description of unfair outcomes, and the other two options relate to prompting methods rather than fairness or bias.

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://aipromptengrmachinelearningnlp.examzify.com

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

SAMPLE