EXAM 2 PRACTICE - ANSWER KEY

Practice Answer Key

CIS 350: AI for Cybersecurity - Smart Defense for the Digital Business

Colorado State University - Computer Information Systems

About this document

Answers and brief explanations for the Exam 2 practice questions. Use this to self-check AFTER attempting the practice exam. Understanding WHY an answer is right matters more than the letter.

Part A - Multiple Choice: answers

1:A 2:B 3:C 4:D 5:A
6:B 7:C 8:D 9:A 10:B
11:C 12:D 13:A 14:B 15:C
16:D 17:A 18:B 19:C 20:D

Part A - explanations

1. A - You have no labels
Unsupervised learning finds structure without any labels.

2. B - Many meaningless tiny groups
Too-large k fragments real groups into meaningless slivers.

3. C - Similar to each other
Clustering groups points that are similar (close in feature space).

4. D - Distance
k-means uses distance, so an unscaled big-range feature dominates.

5. A - Cutting the tree at a chosen height
You cut the dendrogram where it gives useful groups.

6. B - Isolation Forest
Isolation Forest (and LOF) detect outliers; k-means clusters.

7. C - Never-seen (novel) attacks
It learns 'normal' and flags deviations, so it catches novel activity.

8. D - Flag too many normal records (false alarms)
Contamination is the expected anomaly fraction; too high over-flags.

9. A - Scarce (few labels, lots of unlabeled data)
It shines when you have a few labels plus a large unlabeled pool.

10. B - Unlabeled
-1 is scikit-learn's 'unlabeled' marker.

11. C - How good an action's outcome was
Rewards are feedback on actions; the agent maximizes them over time.

12. D - Get stuck using a suboptimal habit
Without exploration it exploits a rut and misses better actions.

13. A - A labeled dataset of correct answers
RL learns from rewards by interacting, not from labeled data.

14. B - Reinforcement learning
RL learns adaptive, sequential strategies - fitting the defender-attacker game.

15. C - Finance staff with fake payment requests
BEC impersonates executives/vendors to trick finance into paying.

16. D - Multi-factor authentication (MFA)
MFA blocks takeover even when the password is known.

17. A - AI writes fluent, error-free, personalized lures
AI removes the grammar/spelling tells people once relied on.

18. B - Verifying via a known separate channel

Call back on a known number / use a code word before acting.

19. C - Signature-based detection

It mutates to slip past signature scanners; use behavior/anomaly detection.

20. D - Build better defenses

The goal is detection and defense - know the threat to stop it.

Part B - True / False: answers

21. FALSE

Unsupervised learning uses no labels.

22. TRUE

Look for the bend where more clusters stop helping.

23. FALSE

It learns 'normal' and flags the rest - no labels needed.

24. TRUE

Unlabeled data reveals structure the few labels can't.

25. TRUE

The policy maps situations to actions.

26. TRUE

Voice clones need only seconds of audio.

27. FALSE

Reuse makes it easier - one leak unlocks many accounts.

28. FALSE

MFA still blocks the login without the second factor.

29. TRUE

Each maps to a concrete countermeasure.

30. FALSE

It still leaves structural signals a classifier can catch.

Part C - Short Answer: model answers

31. In your own words, what is clustering, and how do you choose the number of clusters (k)?

Clustering groups similar records with no labels; choose k with the elbow method (the bend where more clusters stop reducing inertia much), cross-checked with a dendrogram.

32. What is an anomaly, and why is anomaly detection useful for novel attacks?

An anomaly is a rare record unlike the normal majority; anomaly detection learns 'normal' and flags deviations without labels, so it can catch brand-new (zero-day) attacks.

33. What is semi-supervised learning, and what is its main benefit?

Learning from a few labels plus lots of unlabeled data; the benefit is a much stronger model when labels are scarce (the usual security situation).

34. Explain exploration vs. exploitation in reinforcement learning.

Exploration = trying new actions to learn; exploitation = using known-good actions for reward. Good agents balance both; too little exploration gets stuck, too much never settles.

35. Name two ways attackers use AI and give one defense for each.

Any two: AI phishing -> ML filters + training; deepfake/voice fraud -> out-of-band verification; credential stuffing

-> MFA; adaptive malware -> behavior/anomaly detection.

Part D - Applied Scenarios: model answers

36. An online retailer (10 pts)

(a) Unsupervised learning / clustering.
(b) Anomaly detection (e.g., Isolation Forest) - learns normal, flags outliers, no labels.
(c) Semi-supervised learning - the 15 labels + the unlabeled pool.
(d) A flag means 'review,' not 'guilty' - false positives block real customers; verify first.

37. An IT-impersonation attack (10 pts)

(a) Social engineering / AI phishing (vishing), possibly a deepfake voice.
(b) AI produces fluent, personalized messages and can clone voices, removing the old tells.
(c) Awareness training: never share passwords; verify via a known channel; report it.
(d) MFA (so a shared password isn't enough) and/or verified-caller controls.

38. An RL auto-blocker (10 pts)

(a) A policy (which action to take in each state) from rewards/penalties by interacting.
(b) Speed and adaptivity - responds in seconds and adjusts as attacks change.
(c) A bad/manipulated policy can block legitimate users or be exploited.
(d) Keep a human in the loop for high-impact actions; train/test in a simulator first.