EXAM 2 PRACTICE - WEEKS 5-8

Practice Exam

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

Colorado State University - Computer Information Systems

About this document

This is a PRACTICE exam for Exam 2. The questions differ from the real exam but cover the same Weeks 5-8 material in the same format and length (100 points; 90 minutes): multiple choice, true/false, short answer, and applied scenarios. Try it closed-book and timed, then check the separate answer key. Exam 2 is taken online (asynchronous) in Canvas with LockDown Browser.

Total: 100 points, 90 minutes. Part A multiple choice (20 x 2 = 40), Part B true/false (10 x 1 = 10), Part C short answer (5 x 4 = 20), Part D applied scenarios (3 x 10 = 30). No heavy math, no coding.

Part A - Multiple Choice (20 x 2 = 40 pts)

1. The main advantage of unsupervised learning is that it works when:

A. You have no labels
B. You have a reward signal
C. You have only one row
D. You have a labeled test set

2. If you set k in k-means far too high, you tend to get:

A. One giant cluster
B. Many meaningless tiny groups
C. No clusters at all
D. A dendrogram

3. Two records placed in the same cluster are:

A. Always identical
B. From different classes
C. Similar to each other
D. Guaranteed anomalies

4. Feature scaling matters for k-means because the algorithm relies on:

A. Labels
B. Rewards
C. The alphabet
D. Distance

5. Hierarchical clustering lets you choose the number of clusters by:

A. Cutting the tree at a chosen height
B. Setting a reward
C. Adding labels
D. Deleting outliers

6. Which is an anomaly-detection method?

A. k-means
B. Isolation Forest
C. Logistic regression
D. A dendrogram

7. A key benefit of anomaly detection in security is that it can flag:

A. Only known malware
B. Nothing without labels
C. Never-seen (novel) attacks
D. Only phishing emails

8. If 'contamination' is set too high, the detector will:

A. Miss every anomaly
B. Never run
C. Ignore the features
D. Flag too many normal records (false alarms)

9. Semi-supervised learning is ideal when labels are:

A. Scarce (few labels, lots of unlabeled data)
B. Abundant for every row
C. Impossible to get at all
D. Replaced by rewards

10. Passing a label array with -1 entries to LabelSpreading means those rows are:

A. The test set
B. Unlabeled
C. Duplicates
D. Anomalies

11. In reinforcement learning, a reward tells the agent:

A. The correct label
B. The number of features
C. How good an action's outcome was
D. The cluster count

12. Too little exploration in reinforcement learning causes the agent to:

A. Try everything forever
B. Need labels
C. Delete its policy
D. Get stuck using a suboptimal habit

13. Compared with supervised learning, reinforcement learning does NOT need:

A. A labeled dataset of correct answers
B. Any environment
C. A reward signal
D. A policy

14. Which task best fits 'learn an adaptive defense against an adversary'?

A. Clustering
B. Reinforcement learning
C. A dendrogram
D. Feature scaling

15. Business Email Compromise (BEC) usually targets:

A. Printers
B. The weather service
C. Finance staff with fake payment requests
D. Game servers

16. The best single control against account takeover from leaked passwords is:

A. A longer username
B. Turning off logging
C. Reusing one password
D. Multi-factor authentication (MFA)

17. Why is 'spot the typo' weak advice for phishing now?

A. AI writes fluent, error-free, personalized lures
B. Typos are illegal
C. Email no longer exists
D. Filters removed all links

18. A voice-clone ('deepfake') phone scam is best stopped by:

A. Trusting the familiar voice
B. Verifying via a known separate channel
C. Acting fast because it's urgent
D. Replying with your password

19. Adaptive malware is dangerous mainly because it defeats:

A. Electricity
B. The keyboard
C. Signature-based detection
D. All humans equally

20. Studying how attackers use AI is valuable because it helps us:

A. Copy them for profit
B. Avoid computers
C. Win arguments
D. Build better defenses

Part B - True / False (10 x 1 = 10 pts)

21. Unsupervised learning needs a labeled training set. (T / F)
22. The elbow method helps you pick the number of clusters. (T / F)
23. Anomaly detection needs labeled attacks in order to work. (T / F)
24. Semi-supervised learning can beat supervised learning when labels are scarce. (T / F)
25. In reinforcement learning, the policy is what the agent learns. (T / F)
26. Deepfakes can imitate a real person's voice. (T / F)
27. Reusing the same password on many sites makes credential stuffing harder. (T / F)
28. MFA is useless once a password has leaked. (T / F)
29. Every attacker AI technique studied in Week 8 has a matching defense. (T / F)
30. AI-generated phishing cannot be detected by machine learning. (T / F)

Part C - Short Answer (5 x 4 = 20 pts)

31. In your own words, what is clustering, and how do you choose the number of clusters (k)?
32. What is an anomaly, and why is anomaly detection useful for novel attacks?

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

34. Explain exploration vs. exploitation in reinforcement learning.

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

Part D - Applied Scenarios (3 x 10 = 30 pts)

36. An online retailer (10 pts)

A retailer has millions of unlabeled transactions and wants to make sense of them.

(a) Which paradigm groups customers by behavior with no labels? (2 pts)
(b) How would you flag unusual (possibly fraudulent) transactions with no labels? (3 pts)
(c) They confirm 15 fraud cases. Which paradigm turns those few labels into a stronger classifier? (3 pts)
(d) Name one caution before auto-blocking flagged transactions. (2 pts)

37. An IT-impersonation attack (10 pts)

An employee gets an urgent text AND a phone call impersonating IT support, asking for their password.

(a) Name the attack type(s) involved. (2 pts)
(b) Why is it so convincing today? (2 pts)
(c) Give one PEOPLE defense. (3 pts)
(d) Give one TECHNOLOGY defense. (3 pts)

38. An RL auto-blocker (10 pts)

A security team is considering a reinforcement-learning agent that automatically blocks suspicious activity.

(a) What does the RL agent learn, and from what? (3 pts)
(b) Give one benefit of automating the response. (2 pts)
(c) Give one risk. (2 pts)
(d) What safeguard should apply to high-impact blocks? (3 pts)

End of practice exam. Check your answers with the separate answer key.