

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
CS440/ECE448 Artificial Intelligence

Exam 3
Spring 2026

April 6, 2026

Your Name: _____

Your NetID: _____

Instructions

- Please write your name on the top of every page.
- Have your ID ready; you will need to show it when you turn in your exam.
- This will be a **CLOSED BOOK, CLOSED NOTES** exam. You are permitted to bring and use only one 8.5x11 page of notes, front and back, handwritten.
- No electronic devices (phones, tablets, calculators, computers etc.) are allowed.
- Make sure that your answer includes only the variables that it should include, but **DO NOT** simplify explicit numerical expressions, including expressions involving standard functions (exp, log, sin, cos) and expressions involving the sum, product, maximum, or argmax of a complete list of alternatives. For example, the answer $x = \max\left(\sin(0.3), \frac{1}{1+\exp(-0.1)}\right)$ is MUCH preferred (much easier for us to grade) than the answer $x = 0.524979$.

Possibly Useful Formulas

$$\sigma(z) = \frac{1}{1 + e^{-z}}, \quad \sigma'(z) = \sigma(z)(1 - \sigma(z))$$

$$f = g(\mathbf{w}^T \mathbf{x}), \quad \frac{\partial f}{\partial \mathbf{w}} = g'(\mathbf{w}^T \mathbf{x}) \mathbf{x}$$

$$\text{soft max}(\mathbf{z}) = \frac{\exp(\mathbf{z})}{\sum_i \exp(z_i)}, \quad \frac{\partial \text{soft max}_k(\mathbf{z})}{\partial \mathbf{z}} = \text{soft max}_k(\mathbf{z}) (\mathbb{1}_k - \text{soft max}(\mathbf{z}))$$

$$f[k] = \sum_i w[i]x[k-i], \quad \frac{\partial f[k]}{\partial w[i]} = x[k-i]$$

$$\mathbf{C} = \text{soft max} \left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}} \right) \mathbf{V}$$

$$S: \{\mathcal{V}_P, \mathcal{V}_Q\} \rightarrow \{\mathcal{V}_Q, \mathcal{C}\} \text{ s.t. } S(P) = S(Q) = U$$

$$\frac{x'}{x} = \frac{y'}{y} = \frac{-f}{z}$$

$$R_{c,d} = \frac{\frac{\partial z_c}{\partial x_d} x_d}{\sum_{d'} \frac{\partial z_c}{\partial x_{d'}} x_{d'}} R_c$$

$$\mathcal{C}_{\text{obs}} = \{\mathbf{c} : \exists \mathbf{b} : \phi(\mathbf{b}, \mathbf{c}) \in \mathcal{W}_{\text{obs}}\}$$

$$\text{Admissible: } \hat{h}(n) \leq h(n)$$

$$\text{Consistent: } \hat{h}(n) - \hat{h}(m) \leq d(n, m)$$

$$\text{Bellman: } u(s) = r(s) + \gamma \max_a \sum_{s'} P(s'|s, a) u(s')$$

$$\text{Policy Eval: } u_i(s) = r(s) + \gamma \sum_{s'} P(s'|s, \pi_i(s)) u_i(s')$$

$$\text{Policy Update: } \pi_{i+1}(s) = \arg \max_a \sum_{s'} P(s'|s, a) u_i(s')$$

$$\text{Value Iteration: } u_{i+1}(s) = r(s) + \gamma \max_a \sum_{s'} P(s'|s, a) u_i(s')$$

$$\text{Model-based: } P(s_{t+1}|s_t, a_t) = \frac{N(s_t, a_t, s_{t+1}) + k}{\sum_{s'} (N(s_t, a_t, s') + k)}$$

$$\text{TD-Learning: } q(s_t, a_t) \leftarrow (1 - \eta)q(s_t, a_t) + \eta \left(r_t + \gamma \max_a q_t(s_{t+1}, a) \right)$$

$$\text{SARSA: } q(s_t, a_t) \leftarrow (1 - \eta)q(s_t, a_t) + \eta (r_t + \gamma q_t(s_{t+1}, a_{t+1}))$$

$$\text{Actor-Critic: } \mathcal{L}_{\text{actor}} = - \sum_a \pi_a(s) q(s, a)$$

$$\text{REINFORCE: } \Delta \mathbf{W} = \eta (r - \mu) \sum_t \frac{\partial \log \pi_{a_t}(s_t)}{\partial \mathbf{W}}$$

Question 1 (28 points)

Consider a neural network with the following architecture. The input is a scalar x , and there is a weight vector $\mathbf{w} = [w_1, \dots, w_n]^T$ and a scalar weight u . The scalar outputs f , g and h are defined as

$$\begin{aligned}f &= \mathbf{w}^T \mathbf{x}, \\g &= \cos(f), \\h &= ug\end{aligned}$$

where $\cos(\cdot)$ denotes cosine.

- (a) (14 points) Suppose that you know $\frac{\partial \mathcal{L}}{\partial h}$. In terms of $\frac{\partial \mathcal{L}}{\partial h}$ and any of the variables f , g , h , and/or \mathbf{w} , find $\frac{\partial \mathcal{L}}{\partial g}$.


- (b) (14 points) Suppose that you know $\frac{\partial \mathcal{L}}{\partial g}$. In terms of $\frac{\partial \mathcal{L}}{\partial g}$, and in terms of f , g , h , and/or \mathbf{w} , find $\frac{\partial \mathcal{L}}{\partial \mathbf{x}}$.

Question 2 (15 points)

Consider a one-dimensional convolutional neural network with input image $x[i, j]$, weights $w[i]$, and output image $f[i, j]$ related as:

$$f[k, l] = \sum_i x[i, l] w[k - i]$$

In terms of the elements of x , w , and/or f , what is $\frac{\partial f[k, l]}{\partial x[i, j]}$?

Question 3 (28 points)

- (a) (14 points) The search graph above starts at node A, and ends at node G. Create a table showing, in the first column, the node that is expanded by depth-first search at each step of the search process (starting with A, ending with G), and in the second column, the set of nodes that are in the frontier after the node in the first column has been expanded. Optionally, you may list a priority next to each node in the frontier if you wish. Note that, since DFS is not complete, you might need to put a node on the frontier many times with many different priorities. Ties are broken in alphabetical order.

- (b) (14 points) Suppose you're given the following partial heuristic for an A* search: $\hat{h}(C) = 2$, $\hat{h}(D) = 2$, $\hat{h}(G) = 0$. You must now choose the heuristics $\hat{h}(A)$ and $\hat{h}(B)$. What is the definition of an **admissible** heuristic? What are the largest values of $\hat{h}(A)$ and $\hat{h}(B)$ that would result in an admissible heuristic?

Question 4 (30 points)

Consider an MDP with two states, $s \in \{0, 1\}$, and two actions, $a \in \{0, 1\}$. The states have rewards $r(0) = 3$ and $r(1) = 7$, and the transition probabilities are:

s, a	$P(s' = 0 s, a)$	$P(s' = 1 s, a)$
0, 0	0.8	0.2
0, 1	0.3	0.7
1, 0	0.4	0.6
1, 1	0.1	0.9

- (a) (15 points) Suppose you start value iteration with initial utility estimates of $u_1(0) = 3$ and $u_1(1) = 7$. Use value iteration, with $\gamma = \frac{1}{2}$, to find an updated utility estimate, $u_2(0)$, for state $s = 0$.

- (b) (15 points) Explain in words why reinforcement learning must use a policy that balances the competing demands of **exploration** and **exploitation**.

This page is scratch paper

This page is scratch paper, you may tear it off if you wish