

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

Practice Exam 5
Spring 2026

Exam is May 12, 2026, 8:00

Your Name: _____

Your NetID: _____

Instructions

- **SHOW YOUR WORK!** Correct numerical answers will not receive full credit unless you show sufficient derivation to convince us that you understand the answer.
- Have your ID ready; you will need to show it when you turn in your exam.
- 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 $x = 0.524979$.
- This will be a **CLOSED BOOK, CLOSED NOTES** exam. You are permitted to bring and use TWO 8.5x11 pages of notes, front and back, handwritten.
- No electronic devices (phones, tablets, calculators, computers etc.) are allowed.

Possibly Useful Formulas

$$P(X = x|Y = y)P(Y = y) = P(Y = y|X = x)P(X = x)$$

$$P(X = x) = \sum_y P(X = x, Y = y)$$

$$E[f(X, Y)] = \sum_{x, y} f(x, y)P(X = x, Y = y)$$

$$\text{Precision, Recall=Sensitivity, Specificity} = \frac{TP}{TP + FP}, \frac{TP}{TP + FN}, \frac{TN}{TN + FP}$$

$$\text{MPE=MAP: } f(x) = \arg \max (\log P(Y = y) + \log P(X = x|Y = y))$$

$$\text{Naive Bayes: } P(X = x|Y = y) \approx \prod_{i=1}^n P(W = w_i|Y = y)$$

$$\text{Laplace Smoothing w/OOV: } P(W = w_i) = \frac{k + \text{Count}(W = w_i)}{k + \sum_v (k + \text{Count}(W = v))}$$

$$\text{Fairness: } P(Y|A) = \frac{P(Y|f(X), A)P(f(X)|A)}{P(f(X)|Y, A)}$$

$$\text{Score: } v_{t+1}(j) = \max_i v_t(i) + \log a_{ij} + \log b_j(x_{t+1})$$

$$\text{Backpointer: } \psi_{t+1}(j) = \arg \max_i v_t(i) + \log a_{ij} + \log b_j(x_{t+1})$$

$$\text{Linear Regression: } \mathcal{L} = \frac{1}{n} \sum_{i=1}^n (\mathbf{w}^T \mathbf{x}_i - y_i)^2$$

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

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

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

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

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

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

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

$$R(x_d) = \frac{\frac{\partial \mathcal{L}}{\partial x_d} x_d}{\sum_{d'} \frac{\partial \mathcal{L}}{\partial x_{d'}} x_{d'}}$$

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

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

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

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

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

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

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

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

Q-Learning: $q(s_t, a_t) \leftarrow (1 - \eta) q(s_t, a_t) + \eta q_{\text{local}}(t)$

TD-Learning: $q_{\text{local}}(t) = r_t + \gamma \max_a q_t(s_{t+1}, a)$

SARSA: $q_{\text{local}}(t) = r_t + \gamma q_t(s_{t+1}, a_{t+1})$

Deep Q: $\mathcal{L}_{\text{critic}} = \frac{1}{2} (q_t(s_t, a_t) - q_{\text{local}}(t))^2$

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

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

Max: $v = \max \left(v, \max_c (\mathbb{E}[v(c)]) \right), \quad \alpha = \max \left(\alpha, \max_c (\mathbb{E}[v(c)]) \right), \beta = \text{Unchanged}$

Min: $v = \min \left(v, \min_c (\mathbb{E}[v(c)]) \right), \quad \alpha = \text{Unchanged}, \beta = \min \left(\beta, \min_c (\mathbb{E}[v(c)]) \right)$

Chance: $\mathbb{E}[v] = \sum_c P(c) v(c), \quad \alpha = \text{Unchanged}, \beta = \text{Unchanged}$

$$\begin{bmatrix} -1 \\ 1 \end{bmatrix}^T \mathbf{A} \boldsymbol{\beta} = 0, \quad \boldsymbol{\alpha}^T \mathbf{B} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = 0$$

$$a = a + \eta \frac{\partial E[r_A]}{\partial a}, \quad b = b + \eta \frac{\partial E[r_B]}{\partial b}$$


Question 1 (24 points)

A spam filter classifies emails as “Spam” or “Not Spam.” The prior probability of spam is 0.4.

- (a) (12 points) Suppose you pass a piece of e-mail through a pretty good spam filter. The probability the filter flags spam correctly is 0.9. The probability the filter incorrectly flags legitimate email as spam is 0.1. Your email is flagged as spam. Compute the posterior probability that the email is actually spam.
- (b) (12 points) Now suppose you don't have a good spam filter, so you must design one. A Naive Bayes classifier calls the email spam if $f(x) > g(x)$, where x is the set of words in the email, and where f and g are functions that you must design. You have a training corpus of 400 spam emails (200 containing the word “win,” 50 containing the word “probably”) and 600 non-spam emails (60 containing the word “win,” 120 containing the word “probably”). In terms of the Laplace smoothing parameter k , what are $f(x)$ and $g(x)$ for an e-mail that contains both “win” and “probably”?

Question 2 (24 points)

Consider the following hidden Markov model:



Suppose that Y_t and X_t are both binary for all t , with transition and observation probabilities abbreviated as $P(Y_t = j | Y_{t-1} = i) = a_{i,j}$ and $P(X_t = j | Y_t = i) = b_{i,j}$, and suppose it is known that $Y_0 = 0$.

(a) (12 points) In terms of the model parameters, what is $P(X_1 = 1 | Y_0 = 0, Y_2 = 1)$?

(b) (12 points) Let $v_t(j)$ and $\psi_t(j)$ be the Viterbi score and backpointer, respectively (the log probability of the best path ending in state j at time t , and the corresponding previous state). Suppose that $v_2(1) = -15$, $\psi_2(1) = 0$, and $X_2 = 1$. In terms of the model parameters, what is $v_1(0)$?

Question 3 (24 points)

The softmax function converts real-valued inputs, $-\infty < z_n < \infty$, into non-negative outputs that add up to one. Besides the standard softmax, there are many other functions that accomplish these goals, e.g.,

$$f_n(\mathbf{z}) = \frac{z_n^2}{\sum_m z_m^2} \quad (1)$$

(a) (12 points) For the function $f_n(\mathbf{z})$ in Eq. (1), find $\partial f_n / \partial z_s$ for both $s = n$ and $s \neq n$.

(b) (12 points) Equation (1) could be used to construct a self-attention layer in which a context vector \mathbf{c}_i is constructed from query, key, and value vectors \mathbf{q}_i , \mathbf{k}_t , and \mathbf{v}_t , all of which have dimension d :

$$\mathbf{c}_i = \sum_t \left(\frac{(\mathbf{q}_i^T \mathbf{k}_t / d)^2}{\sum_\tau (\mathbf{q}_i^T \mathbf{k}_\tau / d)^2} \right) \mathbf{v}_t \quad (2)$$

Express $\partial c_{i,p} / \partial q_{j,r}$ in terms of $\partial f_n / \partial z_s$ for appropriate values of s and n , where $c_{i,p}$ and $q_{j,r}$ are the p^{th} element of \mathbf{c}_i and the r^{th} element of \mathbf{q}_j , respectively.

Question 4 (24 points)

Suppose you are training a logistic regression network, $\sigma(z) = (1 + e^{-z})^{-1}$, with the following unusual loss function with respect to a desired output y :

$$\mathcal{L} = |\sigma(z) - y| \quad (3)$$

- (a) (12 points) What is $\partial\mathcal{L}/\partial z$? You may not leave $\sigma'(z)$, the derivative of $\sigma(z)$, in your answer; you must replace it with a simpler form, possibly including $\sigma(z)$.

- (b) (12 points) Suppose that z is computed by max-pooling a convolutional network with residual connections,

$$z = \max_i \max_j \left(b[i, j] + \sum_m \sum_n w[i - m, j - n] x[m, n] \right) \quad (4)$$

Use layer-wise relevance propagation to express, in terms of $\partial\mathcal{L}/\partial z$, the relevance of each filter weight $w[m, n]$ to the output $\sigma(z)$. Your answer should explicitly include $\partial\mathcal{L}/\partial z$ and may also explicitly include i^* and j^* defined by

$$(i^*, j^*) = \arg \max_i \arg \max_j \sum_m \sum_n w[i - m, j - n] x[m, n] \quad (5)$$

Question 5 (24 points)

Suppose the configuration space and workspace of a robot are related by a forward kinematic function: $\mathbf{w} = \phi(c)$, where c is an integer and \mathbf{w} is a real vector. Suppose that, from any configuration c , the robot can reach any configuration from $c - 2$ through $c + 2$ in one step. A path from start to goal is a sequence of states $\{c_0, \dots, c_T\}$ such that the robot spends exactly one second traveling from c_t to c_{t+1} at constant workspace velocity ($\frac{\partial \mathbf{w}}{\partial \tau}$ is constant for $t \leq \tau < t + 1$), c_0 is the starting state, $\mathbf{g} = \phi(c_T)$ is the goal in the workspace, and T is the (unknown) time required to reach the goal.

- (a) (12 points) Suppose you want to find a path that minimizes the mean-squared velocity of the robot,

$$\mathcal{L} = \int_0^T \left\| \frac{\partial \mathbf{w}}{\partial t} \right\|^2 dt, \quad (6)$$

You can find a path that minimizes \mathcal{L} by using Dijkstra's algorithm to search paths through configuration space. How should you define the distance $d(c_t, c_{t+1})$ so that Dijkstra's algorithm finds the desired path?

- (b) (12 points) Suppose that $\hat{h}(c) = \|\phi(c) - \mathbf{g}\|^2$, the squared Euclidean distance between $\phi(c)$ and \mathbf{g} . Suppose that you put c_0 on the frontier with $g(c_0) = 0$, then iterate the following loop: (1) From the frontier, pop the node with the lowest $g(c) + \hat{h}(c)$, (2) Add c to the explored set, (3) If $\phi(c) = \mathbf{g}$, end the search and report the path. (4) For each node c' that has $d(c, c') < \infty$ and for which c' is not yet in the explored set, add c' to the frontier with $g(c') = g(c) + d(c, c')$. Does this algorithm always find the shortest path? Why or why not?

Question 6 (24 points)

Consider an MDP in which the state variable is a real number, $s \in \mathfrak{R}$, and the reward in any state is equal to the value of the state, $r(s) = s$. Suppose there are two possible actions, $a \in \{-1, +1\}$, and suppose the state transition probabilities are

$$P(s'|s, a) = \begin{cases} p & s' = s + 2a \\ q & s' = s - 3a \\ 1 - p - q & s' = s \end{cases}$$


where the parameters p and q are assumed known in part (a), but assumed unknown in part (b).

- (a) (12 points) Suppose you start value iteration with initial utility estimates of $u_1(s) = s$ in iteration $t = 1$. Find $u_2(s)$ as a function of s , p , q , and γ .

- (b) (12 points) Suppose that you are training a policy network for which the probability of action $a = 1$ is $\pi_1(s) = \sigma(\mathbf{w}^T \begin{bmatrix} s \\ 1 \end{bmatrix})$, where $\mathbf{w}^T = [w_1, w_2]$ is a weight vector, $\sigma(\cdot)$ is the logistic sigmoid, and the probability of action $a = -1$ is $\pi_{-1}(s) = 1 - \pi_1(s)$. Suppose the REINFORCE algorithm starts with $\mathbf{w}^T = [0, 0]$, then observes the episode $(s_1, a_1, s_2, a_2) = (0, 1, 1, 1)$ and gains a reward of $r = 5$, which is better than the average reward $\mu = -1$. In terms of the learning rate η , what is the update $\Delta \mathbf{w}$?

Question 7 (24 points)

Consider the following game tree. First, Max performs action A or B. Second, Min performs action C or D. Third, a fair coin is flipped, and comes up either H or T. Finally, Max receives the number of points shown in the leaf node, and Min receives 9 minus that number of points.



- (a) (12 points) What is the value, to Max, of each of the nodes in the above tree? Write the value of each node inside the corresponding node. What actions should Max and Min perform in each node where they are allowed to act?
- (b) (12 points) Suppose you have a heuristic that can tell you which of two game states is more valuable to Max. Assuming the existence of such a heuristic, in what order should the game tree be evaluated by the alpha-beta algorithm so that it only needs to evaluate six of the eight leaf nodes?

Question 8 (32 points)

Alice and Bob are playing a simultaneous game in which each of them may either cooperate or defect. Their payoffs are shown in the matrix below.

		Bob	
		defects	cooperates
Alice	defects	1 -3	-9 -7
	cooperates	-9 -7	-1 3

- (a) (11 points) This game has a mixed equilibrium. What is it?
- (b) (10 points) How many Pareto-optimal outcomes does this game have, and what are they?
- (c) (11 points) It is possible to modify just one of the payoffs in this game (just one number in the chart above) so that Bob has a dominant strategy. What number can be modified, to what value, to give Bob a dominant strategy?

This page is scratch paper

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