CS4100/5100 Final Exam

Out: 12/11, Due: 12/13, 5pm EST
Please make your answers **consise**!
Submit your exams via SVN (the same way you submitted your homeworks).

1. (Neural Network) Consider a single layer neural network that classifies an n-dimensional vector, $\mathbf{x} = (1, x_1, \ldots, x_n)^T \in R^n$ as a binary number. The network is parametrized by a weight vector, $\mathbf{w} = (w_0, w_1, \ldots, w_n)^T$. As discussed in class, it is standard to compute the output of the neural network as $y = round(g(\mathbf{w}^T\mathbf{x}))$, where g denotes the logistic function (*i.e.* the squashing function), and *round* rounds the output to either zero or one. Given training data, $(\mathbf{x}^1, y^1), \ldots, (\mathbf{x}^m, y^m)$, the loss function for the neural network is $\sum_{i=1}^{m}(y^i - g(\mathbf{w}^T\mathbf{x}^i))^2$.

a) (2 pts) Suppose we drop the logistic function from the neural network so that we now compute $y = round(\mathbf{w}^T\mathbf{x})$. What is the loss function now?

b)(3 pts) What is the update rule that would implement gradient descent for this error function?

c) (2 pts) Is there a way to compute $\mathbf{w}$ without using gradient descent (*i.e.* analytically)? Explain in words how it would work.

d) (3 pts) Derive the equation that would calculate $\mathbf{w}$ analytically. Show your work. (Hint: Suppose we want to calculate the vector $\mathbf{x}$ that minimizes $\|A\mathbf{x} - \mathbf{b}\|_2$ for a matrix A and a vector $\mathbf{b}$. The solution is: $\mathbf{x} = A^+\mathbf{b}$, where A^+ denotes the *pseudoinverse* of A. You may express your answer in terms of the pseudoinverse.)

2. (Clustering) Suppose we are given four scalars: $x_1 = 1, x_2 = 2, x_3 = 4, x_4 = 5$. We want to cluster these four scalars into two classes.

a) (k-means) (4 pts) Use k-means to cluster the four scalars. In the first iteration of k-means, we initialize the means of the two classes to $\mu_1 = 0, \mu_2 = 3$ and assign each point to the corresponding class as shown in the table below. Fill in the rest of the table for iterations 2 through 4.

Iteration	μ_1	μ_2	x_1 class	x_2 class	x_3 class	x_4 class
1	0	3	1	2	2	2
2						
3						
4						

b) (EM) (3 pts) We want to use EM to fit a mixture of two Gaussians to the four scalars. The model for a mixture of two Gaussians is:

$$P(x) = \pi_1 N(x|\mu_1, \Sigma_1) + (1 - \pi_1) N(x|\mu_2, \Sigma_2).$$

Write down the EM algorithm for fitting the above five parameters (π_1, μ_1, Σ_1, μ_2, Σ_2) to the four data points. How does this algorithm change if we fix $\pi_1 = 0.5$ so that the new model is:

$$P(x) = 0.5N(x|\mu_1, \Sigma_1) + 0.5N(x|\mu_2, \Sigma_2)?$$

c) (EM)(3 pts) We initialize the means of the two classes to $\mu_1 = 0, \mu_2 = 3$, the variances to $\sigma_1^2 = 1$ $\sigma_2^2 = 1$, and the weights to $\pi_1 = 0.5$. Fill in the rest of the table below. The weights, w_{ij}, in the right four columns denote the probability that data point i is generated by Gaussian j.

Iter	π_1	μ_1	μ_2	σ_1^2	σ_2^2	w_{11}	w_{21}	w_{31}	w_{41}	w_{12}	w_{22}	w_{32}	w_{42}
1	0	3											
2													
3													
4													

3. (Bayes Networks)

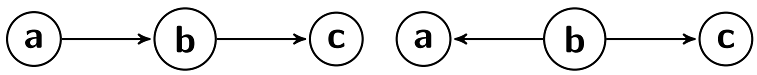

(a) (b)

a) (3 pts) Each of the two Bayes networks, (a) and (b), shown above can be used to define a joint distribution over the three variables by defining the appropriate conditional probability table associated with each arrow. Is there a joint distribution that can be represented using the network in (a) that cannot be represented with the network in (b)? Why or why not?

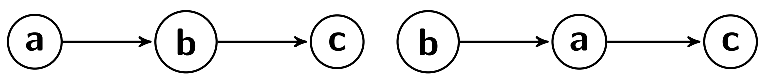

(a) (b)

b) (2 pts) Repeat the above question for the two networks show above. Is there a joint distribution that can be represented using the network in (a) that cannot be represented with the network in (b)? Why or why not?

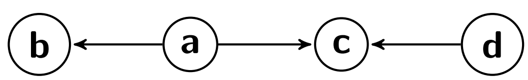

c) (3 pts) Write the joint distribution for the four variables above in terms of the conditional probability distributions associated with this Bayes network.

d) (2 pts) Consider four random variables, a, b, c, d. Draw six Bayes networks over these variables that encodes each of the following conditional independences.

- $a \not\perp c | b$
- $a \perp b | c$
- $a \not\perp b | \emptyset$
- $a \not\perp b | c, d$

4. (HMMs) Lance Armstrong is an important American cyclist who won the Tour de France each year between 1999 and 2005 (inclusive). However, it has recently been discovered that Armstrong used performance enhancing drugs (doping) during these years. Consider the following four years: 1999 ($time = 1$), 2000 ($time = 2$), 2001 ($time = 3$), and 2002 ($time = 4$). Since doping was very prevalent in 1999, suppose that there is a 50% probability that Armstrong doped in 1999 ($P(d_1) = 0.5$). Also, suppose that there is a 60% probability that Armstrong dopes in the following year given that he dopes in the current year ($P(d_{t+1}|d_t) = 0.6$). If Armstrong does not dope in the current year, then there is only a 10% chance he dopes in the following year ($P(d_{t+1}|\bar{d}_t) = 0.1$). Unfortunately, the appropriate cycling federations did not routinely apply accurate drug tests during the early 2000s. Therefore, we are uncertain about in exactly which years doping occurred. However, suppose that we know that Armstrong wins 60% of the time when he dopes and 30% of the time when he does not dope.

a) (2 pts) Formulate this as an HMM. Draw the corresponding Bayes network.

b) (3 pts) Use filtering to calculate the probability that Armstrong doped in each year between 1999 and 2002 given the prior years' information.

c) (4 pts) Use smoothing to calculate the probability that Armstrong doped in each year between 1999 and 2002 given all information.

d) (1 pts) Suppose that drug tests are administered in each of the four years. Update your HMM and update the corresponding Bayes network.

5. (RL) Consider the 9-state infinite horizon MDP below. All transitions are deterministic. All transitions have zero reward except for the transition from state d to a which has a reward of +5 and the transition from state i to a which as a reward of +6.

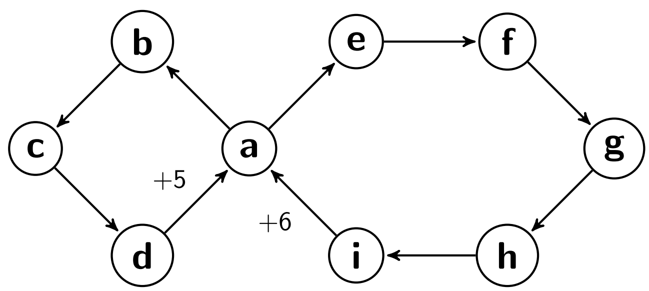

a) (2 pts) How many distinct policies are there in this MDP? What are the policies?

b) (4 pts) For each policy, calculate the value of state a. Show your work.

c) (2 pts) For what values of the discount factor, $\gamma \in [0, 1)$, does the optimal policy take the agent through state d? For what values of γ does the optimal policy take the agent through i?

6. (CSP) We will use backtracking with forward search to solve the 5-queens problem. Use the minimum remaining values (MRV) heuristic. Assume that one queen is placed in the upper left corner.

a) (2 pts) Formulate the 5-queens problem as a CSP. What are the variables and what are the constraints (be precise)?

b) (2 pts) After placing the first queen in the upper left corner, what are the possible values for the remaining variables?

c) (4 pts) Solve the rest of the 5-queens problem using forward search and MRV. Show the order in which the variables are filled and the remaining values for all variables at each step.

7. (Search) Consider the grid world below. Each empty square denotes a possible state where the agent may be located. At each step, the agent may move up, down, left, or right (U,D,L,R). Initially, the agent's state is completely unknown; in other words, the agent initially does not know where it is located. The shaded squares denote walls that the agent cannot cross through. If the agent executes a move that would take it into a wall or cause it to move outside the grid, then the agent's state remains unchanged. The objective is to reach the outlined square (the goal state) in the upper right corner.

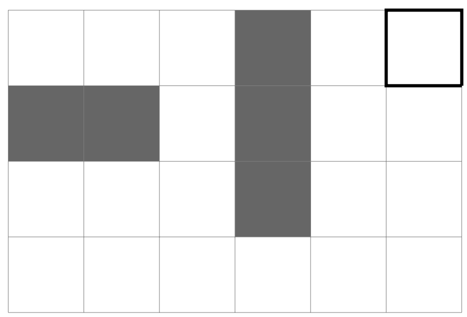

The objective of planning is to find a sequence of moves that cause the agent to reach the goal (even though the agent cannot sense its own location). For example, the sequence of actions, $(R, R, D, D, D, R, R, R, U, U, U)$, will cause the agent to reach the goal state regardless of regardless of where it starts.

a) (3 pts) Formulate this as a search problem over a finite search space. What are the states? How many states are there? What is the start state? What is the successor function (denote the state that is reached by executing action a from state x as $f(x, a)$). What is the goal test? (Hint: this is known as a conformant search problem.)

b) (2 pts) Give a non-trivial admissible heuristic for this problem.

c) (3 pts) Suppose that the agent's actions are stochastic such that the agent fails to move in the desired direction with probability ϵ. In this case, the state space becomes the space of probability distributions over the underlying states. Give an expression for $P(x|a_1, \ldots, a_t)$ in terms of $P(x|a_1, \ldots, a_{t-1})$.