

# CS221 Midterm

CS221  
November 18, 2014

Name: \_\_\_\_\_  
by writing my name I agree to abide by the honor code

SUNet ID: \_\_\_\_\_

**Read all of the following information before starting the exam:**

- This test has 3 problems and is worth 150 points total. It is your responsibility to make sure that you have all of the pages.
- Keep your answers precise and concise. Show all work, clearly and in order, or else points will be deducted, even if your final answer is correct.
- Don't spend too much time on one problem. Read through all the problems carefully and do the easy ones first. Try to understand the problems intuitively; it really helps to draw a picture.
- Good luck!

Problem	Part	Max Score	Score
1	a	10	
	b	10	
	c	10	
	d	10	
	e	10	
2	a	10	
	b	10	
	c	10	
	d	10	
	e	10	
3	a	10	
	b	10	
	c	10	
	d	10	
	e	10	

Total Score:  +  +  =

## 1. Enchaining Realm (50 points)

This problem is about machine learning.

### a. (10 points)

Suppose we want to predict a real-valued output  $y \in \mathbb{R}$  given an input  $x = (x_1, x_2) \in \mathbb{R}^2$ , which is represented by a feature vector  $\phi(x) = (x_1, |x_1 - x_2|)$ .

Consider the following training set of  $(x, y)$  pairs:

$$\mathcal{D}_{\text{train}} = \{((1, 2), 2), ((1, 1), 1), ((2, 1), 3)\}. \quad (1)$$

We use a modified squared loss function, which penalizes overshooting twice as much as undershooting:

$$\text{Loss}(x, y, \mathbf{w}) = \begin{cases} \frac{1}{2}(\mathbf{w} \cdot \phi(x) - y)^2 & \text{if } \mathbf{w} \cdot \phi(x) < y \\ (\mathbf{w} \cdot \phi(x) - y)^2 & \text{otherwise} \end{cases} \quad (2)$$

Using a fixed learning rate of  $\eta = 1$ , apply the stochastic gradient descent algorithm on this training set starting from  $\mathbf{w} = [0, 0]$  after looping through each example  $(x, y)$  in order and performing the following update:

$$\mathbf{w} \leftarrow \mathbf{w} - \eta \nabla_{\mathbf{w}} \text{Loss}(x, y, \mathbf{w}). \quad (3)$$

For each example in the training set, calculate the loss on that example and update the weight vector  $\mathbf{w}$  to fill in the table below:

	$x$	$\phi(x)$	$\text{Loss}(x, y, \mathbf{w})$	$\nabla_{\mathbf{w}} \text{Loss}(x, y, \mathbf{w})$	weights $\mathbf{w}$
Initialization	n/a	n/a	n/a	n/a	$[0, 0]$
After example 1	(1,2)				
After example 2	(1,1)				
After example 3	(2,1)				

**b. (10 points)**

Consider the following set of 6 points in the plane:

$$\{A = (0, 0), B = (1, 0), C = (0, 3), D = (1, 3), E = (5, 7), F = (8, 12)\} \quad (4)$$

You'd like to partition the points into two clusters, where each cluster is represented by a centroid  $\mu_k$  that is *constrained to lie on the diagonal line*; formally,  $\mu_k = [c_k, c_k]$  for  $k \in \{1, 2\}$ .

Recall that the reconstruction loss is the sum of squared distances from each point to the centroid it is assigned to. Modify the K-means algorithm to minimize this loss while respecting the diagonal constraint.

Suppose we have  $K = 2$  clusters which are initialized with  $c_1 = 3$  and  $c_2 = 11$ . Break ties, if any, by assigning a point to the centroid farther away from the origin. Run two iterations of K-means, filling out the values below:

<b>Iteration</b>	<b>Task</b>	<b>Value</b>
1.	Assignment $z_i = 1$ :	(i) _____
1.	Assignment $z_i = 2$ :	(ii) _____
1.	New $c_1$ :	(iii) _____
1.	New $c_2$ :	(iv) _____
1.	Reconstruction loss after updating centroids:	(v) _____
2.	Assignment $z_i = 1$ :	(vi) _____
2.	Assignment $z_i = 2$ :	(vii) _____
2.	New $c_1$ :	(viii) _____
2.	New $c_2$ :	(ix) _____
2.	Reconstruction loss after updating centroids:	(x) _____

**c.** (10 points)

Given two images  $x_1$  and  $x_2$ , our goal is to predict whether they are of the same person ( $y = 1$ ) or not ( $y = -1$ ).

We build the following model: We have a feature extractor that maps an image  $x$  to a feature vector  $\phi(x) \in \mathbb{R}^d$ . For each  $j = 1, \dots, K$ , define  $h_j(x) = \mathbf{v}_j \cdot \phi(x)$ . Intuitively, each parameter  $\mathbf{v}_j \in \mathbb{R}^d$  corresponds to a direction;  $h_j(x)$  corresponds to the projection of  $\phi(x)$  along that direction.

Given a weight vector  $\mathbf{w} = (w_1, \dots, w_K)$ , the classification score on an input  $(x_1, x_2)$  is defined as:

$$s(x_1, x_2) = \sum_{j=1}^K w_j h_j(x_1) h_j(x_2). \quad (5)$$

As an example, for  $K = d = 2$ , if the parameters were  $\mathbf{v}_1 = (1, 0)$ ,  $\mathbf{v}_2 = (0, 1)$ ,  $\mathbf{w} = (3, 4)$ , then the classification score would be  $3\phi(x_1)_1\phi(x_2)_1 + 4\phi(x_1)_2\phi(x_2)_2$ .

Prove that there exists a new feature extractor  $A$  mapping  $(x_1, x_2)$  to a  $d^2$ -dimensional feature vector  $A(x_1, x_2) \in \mathbb{R}^{d^2}$ , such that for any scoring function  $s$  (defined by  $\mathbf{v}_1, \dots, \mathbf{v}_K, \mathbf{w}$ ), there exists a new weight vector  $\mathbf{u} \in \mathbb{R}^{d^2}$  with  $s(x_1, x_2) = \mathbf{u} \cdot A(x_1, x_2)$ . In other words, the scoring functions in (5) can be represented by linear functions with appropriate features. You must define  $\mathbf{u}$  explicitly.

**d.** (10 points)

Suppose we are performing classification where the input points are of the form  $(x_1, x_2) \in \mathbb{R}^2$ . We can choose any subset of the following set of features:

$$\mathcal{F} = \left\{ x_1^2, x_2^2, x_1x_2, x_1, x_2, \frac{1}{x_1}, \frac{1}{x_2}, 1, \mathbf{1}[x_1 \geq 0], \mathbf{1}[x_2 \geq 0] \right\} \quad (6)$$

For each subset of features  $F \subseteq \mathcal{F}$ , let  $D(F)$  be the set of all decision boundaries corresponding to linear classifiers that use features  $F$ .

For each of the following sets of decision boundaries  $E$ , provide the minimal  $F$  such that  $D(F) \supseteq E$ . If no such  $F$  exists, write ‘none’.

- $E$  is all lines:

---

(7)

- $E$  is all circles centered at the origin:

---

(8)

- $E$  is all circles:

---

(9)

- $E$  is all axis-aligned rectangles:

---

(10)

- $E$  is all axis-aligned rectangles whose lower-right corner is at  $(0, 0)$ :

---

(11)

e. (10 points)

For each of the following problems, circle all correct options. There may be multiple correct options for a particular question.

- Hyperparameters should be tuned to minimize error on the
  - Training set
  - Development set
  - Test set
- The majority algorithm (which outputs the most common label based on the training data) for classification has
  - High bias
  - High variance
  - Low bias
  - Low variance
- Suppose our learning algorithm for a particular task has very low training error but large test error. Which of the following changes to our algorithm could fix this?
  - Increasing the number of features
  - Decreasing the number of features
  - Increasing the number of SGD iterations
  - Decreasing the number of SGD iterations
  - Adding more training examples
  - Adding a regularization penalty

## 2. States (50 points)

Sabina has just moved to a new town, which is represented as a grid of locations (see below). She needs to visit various shops  $S_1, \dots, S_k$ . From a location on the grid, Sabina can move to the location that is immediately north, south, east, or west, but certain locations have been blocked off and she cannot enter them. It takes one unit of time to move between adjacent locations. Here is an example layout of the town:

	(2,5)	(3,5)	(4,5)	
(1,4)	<b>S<sub>1</sub></b> (2,4)	(3,4)	<b>S<sub>2</sub></b> (4,4)	(5,4)
(1,3)	(2,3)		(4,3)	(5,3)
	(2,2)	(3,2)	(4,2)	<b>S<sub>3</sub></b> (5,2)
<b>House</b> (1,1)	(2,1)	<b>S<sub>4</sub></b> (3,1)	(4,1)	(5,1)

Sabina lives at (1, 1), and no location contains more than one building (Sabina's house or a shop).

### a. (10 points)

Sabina wants to start at her house, visit the shops  $S_1, \dots, S_k$  **in any order**, and then return to her house as quickly as possible. We will construct a search problem to find the fastest route for Sabina. Each state is modeled as a tuple  $s = (x, y, A)$ , where  $(x, y)$  is Sabina's current position, and  $A$  is some auxiliary information that you need to choose. If an action is invalid from a given state, set its cost to infinity. Let  $V$  be the set of valid (non-blocked) locations; use this to define your search problem. You may assume that the locations of the  $k$  shops are known. You must choose a minimal representation of  $A$  and solve this problem for general  $k$ . Be precise!

- Describe  $A$ : \_\_\_\_\_
- $s_{\text{start}} =$  \_\_\_\_\_
- $\text{Actions}((x, y, A)) = \{\text{N, S, E, W}\}$

•  $\text{Succ}((x, y, A), a) =$  \_\_\_\_\_

\_\_\_\_\_

•  $\text{Cost}((x, y, A), a) =$  \_\_\_\_\_

•  $\text{IsGoal}((x, y, A)) =$  \_\_\_\_\_

Now we will tweak the problem. Sabina must visit the shops  $S_1, S_2, \dots, S_k$  **in that order**. She is allowed to step in the location of a shop without visiting it. Assuming that each state is still  $(x, y, A)$ , how would you modify  $A$  to solve the constrained version of this problem? Again, make  $A$  as minimal as possible and be precise.



**b.** (10 points)

Sabina is now again allowed to visit the shops **in any order**. But she is impatient and doesn't want to wait around for your search algorithm to finish running. In response, you will use the A\* algorithm, but you need a heuristic. For each pair of shops  $(S_i, S_j)$  where  $i \neq j$  and  $1 \leq i, j \leq k$ , define a **consistent** heuristic  $h_{i,j}$  that approximates the time it takes to ensure that shops  $S_i$  and  $S_j$  are visited and then return home. Computing  $h_{i,j}(s)$  should take  $O(1)$  time.

Now, let each  $h_{i,j}$  be an arbitrary consistent heuristic. Let  $h(s)$  be the heuristic that takes a weighted sum of these heuristics; that is,  $h(s) = \sum_{i \neq j} \alpha_{i,j} h_{i,j}(s)$ , for some numbers  $\alpha_{i,j} \geq 0$ . For which values of  $\alpha_{i,j}$  is the heuristic  $h$  **guaranteed** to be consistent?



**d.** (10 points)

Sabina wants to go from her house (located at 1) to the gym (located at  $n$ ). At each location  $s$ , she can either (i) deterministically walk forward to the next location  $s + 1$  (takes 1 unit of time) or (ii) wait for the bus. The bus comes with probability  $\epsilon$ , in which case, she will reach the gym in  $1 + \alpha(n - s)$  units of time, where  $\alpha$  is some parameter. If the bus doesn't come, well, she stays put, and that takes 1 unit of time.



We have formalized the problem as an MDP for you:

- State:  $s \in \{0, 1, \dots, n\}$  is Sabina's location

- Actions( $s$ ) = {Walk, Bus}

- $\text{Reward}(s, \text{Walk}, s') = \begin{cases} -1 & \text{if } s' = s + 1 \\ -\infty & \text{otherwise} \end{cases}$

- $\text{Reward}(s, \text{Bus}, s') = \begin{cases} -1 - \alpha(n - s) & \text{if } s' = n \\ -1 & \text{if } s' = s \\ -\infty & \text{otherwise} \end{cases}$

- $T(s, \text{Walk}, s') = \begin{cases} 1 & \text{if } s' = s + 1 \\ 0 & \text{otherwise} \end{cases}$

- $T(s, \text{Bus}, s') = \begin{cases} \epsilon & \text{if } s' = n \\ 1 - \epsilon & \text{if } s' = s \\ 0 & \text{otherwise} \end{cases}$

- $\text{IsEnd}(s) = \mathbf{1}[s = n]$

Compute a closed form expression for the value of the “always walk” policy and the “always wait for the bus” policy (using some or all of the variables  $\epsilon, \alpha, n$ ).

- $V_{\text{Walk}}(s) =$  \_\_\_\_\_

- $V_{\text{Bus}}(s) =$  \_\_\_\_\_

- For what values of  $\epsilon$  (as a function of  $\alpha$  and  $n$ ) is it advantageous to walk rather than take the bus?

e. (10 points)

Not surprisingly, buses operate strangely in this town, and we will now assume instead that Sabina doesn't know the reward function nor the transition probabilities. She decides to use reinforcement learning to find out. She starts by going around town using the two different modes of transportation:

$s_0$	$a_1$	$r_1$	$s_1$	$a_2$	$r_2$	$s_2$	$a_3$	$r_3$	$s_3$	$a_4$	$r_4$	$s_4$	$a_5$	$r_5$	$s_5$
1	Bus	-1	1	Bus	-1	1	Bus	3	3	Walk	1	4	Walk	1	5

Run the Q-learning algorithm once over the given data to compute an estimate of the optimal Q-value  $Q_{\text{opt}}(s, a)$ . Process the episodes from left to right. Assume all Q-values are initialized to zero, and use a learning rate of  $\eta = 0.5$  and a discount of  $\gamma = 1$ .

•  $\hat{Q}(1, \text{Walk}) =$  \_\_\_\_\_

•  $\hat{Q}(1, \text{Bus}) =$  \_\_\_\_\_

•  $\hat{Q}(3, \text{Walk}) =$  \_\_\_\_\_

•  $\hat{Q}(3, \text{Bus}) =$  \_\_\_\_\_

•  $\hat{Q}(4, \text{Walk}) =$  \_\_\_\_\_

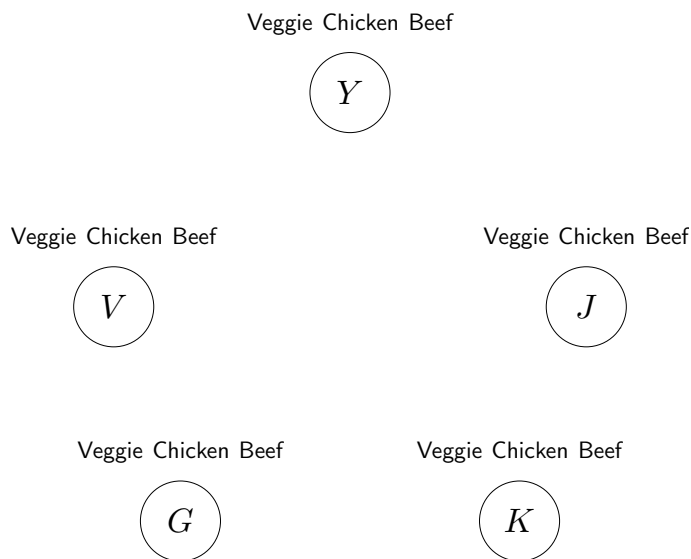
•  $\hat{Q}(4, \text{Bus}) =$  \_\_\_\_\_

### 3. Variables and Logic (50 points)

It's Friday night, and you and your friends go out to dinner in anticipation of the Big Game the next day (go Card!). Since you've just been working on your final project for CS221, you are seeing CSPs and Bayesian networks everywhere!

a. (10 points)

You and your friends (Veronica, Jarvis, Gabriela, Kanti) sit around a table like this:



There are three dishes on the menu: the vegetarian deep dish pizza, the chicken quesadilla, and the beef cheeseburger. Each person will order exactly one dish.

But what started out as a simple dinner has quickly turned into a logistical nightmare because of all the constraints you and your friends impose upon yourselves:

1. Each person must order something different than the people sitting immediately next to him/her.
2. You ( $Y$ ) are vegetarian.
3. If Veronica ( $V$ ) orders beef, then Jarvis ( $J$ ) will order veggie.
4. Kanti ( $K$ ) and Jarvis ( $J$ ) cannot both get non-chicken dishes.

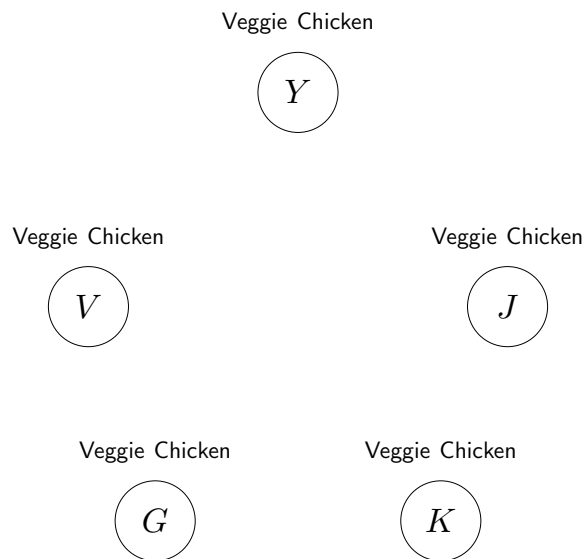
Draw the potentials for the above constraints and write the propositional formula above each potential (e.g.,  $[Y = \text{Veggie}]$ ). Then for each pair of variables, enforce arc consistency in both directions, crossing out the appropriate values from the domains.

b. (10 points)

Your server comes by your table and tells you that they are out of beef today, so you all decide to rework your constraints. Now they are:

1. There is a preference for people sitting next to each other to order different dishes.  
Formally, we have 5 potentials:  $f(Y, J) = \mathbf{1}[Y \neq J] + 1$ ,  $f(J, K) = \mathbf{1}[J \neq K] + 1$ , etc.
2. You ( $Y$ ) are vegetarian.

With the 2 constraints above, what is a maximum weight assignment and what is its weight? If there are many assignments with the same maximum weight, give any one. For convenience, the updated table is given below.



**c. (10 points)**

Now that you and your friends have your food, you are watching the highlights from yesterday's football game. As you watch a particular four-down sequence, you have an epiphany: football is a glorified Bayesian network.

On each play  $t \in \{1, 2, 3, \dots\}$ , the quarterback will either (1) run the ball, (2) throw a short pass, or (3) throw a long pass. Let  $X_t \in \{1, 2, 3\}$  denote the quarterback's move, and  $E_t$  be the number of yards advanced as a result. Assume that  $X_1$  is chosen uniformly at random from  $\{1, 2, 3\}$ .

Before the game starts, the team decides on one of the two strategies ( $S \in \{A, B\}$ ), which is chosen uniformly at random and is held constant throughout the game: if  $S = A$ , then the quarterback uses  $\alpha = \frac{1}{2}$ ; if  $S = B$ , then the quarterback uses  $\alpha = \frac{1}{4}$ .

On each  $t$ , the quarterback performs the same move as the last time ( $X_t = X_{t-1}$ ) with probability  $\alpha$ , and one of the other two each with probability  $\frac{1-\alpha}{2}$ . When a play  $x \in \{1, 2, 3\}$  is made, the ball is advanced  $5x$  yards with probability  $2^{-x}$  and 0 yards with probability  $1 - 2^{-x}$ .

- Draw the Bayesian network for corresponding to the football model described above, using  $t \in \{1, 2, 3, 4\}$ . Recall that the variables are  $S, X_1, \dots, X_4, E_1, \dots, E_4$ .

- Draw the factor graph associated with the Bayesian network corresponding to this football model. You are not required to write the potentials for each factor.



**d.** (10 points)

You saw that the first play was a long pass ( $X_1 = 3$ ) which resulted in advancing the ball  $E_1 = 15$  yards (this happened with probability  $\frac{1}{8}$ ). Then while you were busy constructing your Bayesian network, you missed what happened in the second time step, but you found out that only five yards were gained on the third play ( $E_3 = 5$ ). You do not remember the yard line the ball was previously at, so you have no knowledge about what type of play was run on the second down or what its result was. What is the probability that the team is playing strategy  $A$ ? In other words, compute:

$$\mathbb{P}(S = A \mid X_1 = 3, E_1 = 15, E_3 = 5).$$

e. (10 points)

You realize that the football model you made up was pretty lousy, and it's better to learn the model from data. You will use the same Bayesian network as before, but now just learn the weights. In this new model, any play can gain either 0, 5, or 10 yards. Teams still play with one of the two strategies: strategy  $S = A$  and strategy  $S = B$ , although they may not have the same interpretation they had before. As more highlights flash across the screen, the sportscaster announces the strategy and you write the plays down furiously on your napkin:

- Strategy A: (2, 0), (1, 0)
- Strategy B: (1, 5), (1, 0), (1, 10), (3, 0)
- Strategy A: (3, 5), (2, 0), (1, 10), (1, 0), (3, 10)

This means that the first sequence was with strategy A, which involved two plays: (short pass, 0 yardage) and (run, 0 yardage). The next sequence involved strategy B, etc. Each sequence is independent.

Use Laplace smoothing with  $\lambda = 1$ .

$x_t$	$e_t$	$p(e_t   x_t)$
1	0	
1	5	
1	10	
2	0	
2	5	
2	10	
3	0	
3	5	
3	10	

$s$	$x_{t-1}$	$x_t$	$p(x_t   x_{t-1}, s)$
A	1	1	
A	1	2	
A	1	3	
A	2	1	
A	2	2	
A	2	3	
A	3	1	
A	3	2	
A	3	3	
B	1	1	
B	1	2	
B	1	3	
B	2	1	
B	2	2	
B	2	3	
B	3	1	
B	3	2	
B	3	3	