

Due: Friday 8/9/2019 at 11:59pm (submit via Gradescope).

Policy: Can be solved in groups (acknowledge collaborators) but must be written up individually

Submission: Your submission should be a PDF that matches this template. Each page of the PDF should align with the corresponding page of the template (page 1 has name/collaborators, question 1 begins on page 2, etc.). **Do not reorder, split, combine, or add extra pages.** The intention is that you print out the template, write on the page in pen/pencil, and then scan or take pictures of the pages to make your submission. You may also fill out this template digitally (e.g. using a tablet.)

First name	
Last name	
SID	
Collaborators	

Q1. ML: Maximum Likelihood

Training Data

$(L = 1, F_1 = 1, F_2 = 1)$

$(L = 1, F_1 = 1, F_2 = 1)$

$(L = 0, F_1 = 1, F_2 = 1)$

$(L = 1, F_1 = 0, F_2 = 0)$

$(L = 0, F_1 = 0, F_2 = 0)$

$(L = 0, F_1 = 0, F_2 = 0)$

$(L = 0, F_1 = 0, F_2 = 1)$

You've decided to use a model-based approach to classification of text documents. Your goal is to build a classifier that can determine whether or not a document is about cats. You're taking a minimalist approach and you're only characterizing the input documents in terms of two binary features: F_1 and F_2 . Both of these features have domain $\{0, 1\}$. The thing you're trying to predict is the label, L , which is also binary valued. When $L = 1$, the document is about cats. When $L = 0$, the document is not.

The particular meaning of the two features F_1 and F_2 is not important for your current purposes. You are only trying to decide on a particular Bayes' net structure for your classifier. You've got your hands on some training data (shown above) and you're trying to figure out which of several potential Bayes' nets (also shown above) might yield a decent classifier when trained on that training data.

(a) Which of the Bayes' nets, once learned from the training data with maximum likelihood estimation, would assign non-zero probability to the following query: $P(L = 1|F_1 = 0, F_2 = 0)$? Fill in all that apply.

- (i) (ii) (iii) (iv)

(b) Which of the Bayes' nets, once learned from the training data with maximum likelihood estimation, would assign non-zero probability to the following query: $P(L = 1|F_1 = 0, F_2 = 1)$? Fill in all that apply.

- (i) (ii) (iii) (iv)

(c) Which of the Bayes' nets, once learned from the training data with Laplace smoothing using $k = 1$, would assign non-zero probability to the following query: $P(L = 1|F_1 = 0, F_2 = 1)$? Fill in all that apply.

- (i) (ii) (iii) (iv)

(d) What probability does Bayes' net (i), once learned from the training data with Laplace smoothing using $k = 1$, assign to the query $P(L = 1|F_1 = 0, F_2 = 1)$?

(e) As $k \rightarrow \infty$ (the constant used for Laplace smoothing), what does the probability that Bayes' net (i) assigns to the query $P(L = 1|F_1 = 0, F_2 = 1)$ converge to?

Q2. Spam Classification

The Naïve Bayes model has been famously used for classifying spam. We will use it in the “bag-of-words” model:

- Each email has binary label Y which takes values in $\{\text{spam, ham}\}$.
- Each word w of an email, no matter where in the email it occurs, is assumed to have probability $P(W = w | Y)$, where W takes on words in a pre-determined dictionary. Punctuation is ignored.
- Take an email with K words w_1, \dots, w_K . For instance: email “hi hi you” has $w_1 = \text{hi}, w_2 = \text{hi}, w_3 = \text{you}$. Its label is given by $\arg \max_y P(Y = y | w_1, \dots, w_K) = \arg \max_y P(Y = y) \prod_{i=1}^K P(W = w_i | Y = y)$.

- (a) You are in possession of a bag of words spam classifier trained on a large corpus of emails. Below is a table of some estimated word probabilities.

W	note	to	self	become	perfect
$P(W Y = \text{spam})$	1/6	1/8	1/4	1/4	1/8
$P(W Y = \text{ham})$	1/8	1/3	1/4	1/12	1/12

You are given a new email to classify, with only two words:

perfect note

Fill in the circles corresponding to all values of $P(Y = \text{spam})$ for which the bag of words with these word probabilities will give “spam” as the most likely label.

- 0 0.4 0.8
 0.2 0.6 1

- (b) You are given only three emails as a training set:

(Spam) dear sir, I write to you in hope of recovering my gold watch.

(Ham) hey, lunch at 12?

(Ham) fine, watch it tomorrow night.

Fill in the circles corresponding to values you would estimate for the given probabilities, if you were doing no smoothing.

- $P(W = \text{sir} | Y = \text{spam})$ 0 1/10 1/5 1/3 2/3 None of the above
 $P(W = \text{watch} | Y = \text{ham})$ 0 1/10 1/5 1/3 2/3 None of the above
 $P(W = \text{gauntlet} | Y = \text{ham})$ 0 1/10 1/5 1/3 2/3 None of the above
 $P(Y = \text{ham})$ 0 1/10 1/5 1/3 2/3 None of the above

- (c) You are training with the same emails as in the previous question, but now doing Laplace Smoothing with $k = 2$. There are V words in the dictionary. Write concise expressions for:

$P(W = \text{sir} | Y = \text{spam})$

$P(W = \text{watch} | Y = \text{ham})$

$P(Y = \text{ham})$

(d) On the held-out set, you see the following accuracies for different values of k :


k	0	1	2	10
accuracy	0.65	0.68	0.74	0.6

On the training set, the accuracy for $k = 0$ is 0.8. Fill in the circle of all plausible accuracies for $k = 10$ on the training set.

- 0.1 0.99
 0.7 None of the above

(e) **Becoming less naïve**

We are now going to improve the representational capacity of the model. Presence of word w_i will be modeled not by $P(W = w_i | Y)$, where it is only dependent on the label, but by $P(W = w_i | Y, W_{i-1})$, where it is also dependent on the previous word. The corresponding model for an email of only four words is given on the right.



(i) With a vocabulary consisting of V words, what is the *minimal* number of conditional **word** probabilities that need to be estimated for this model? The correct answer is among the choices.

- V V^2 2^V
 $2V$ $2V^2$ 2^{2V}

(ii) Select all expected effects of using the new model instead of the old one, if both are trained with a very large set of emails (equal number of spam and ham examples).

- The entropy of the posterior $P(Y|W)$ should on average be lower with the new model. (In other words, the model will tend to be more confident in its answers.)
 The accuracy on the **training** data should be higher with the new model.
 The accuracy on the **held-out** data should be higher with the new model.
 None of the above.

Q3. Potpourri

- (a) A single perceptron can compute the XOR function.
- True False
- (b) A perceptron is guaranteed to learn a separating decision boundary for a separable dataset within a finite number of training steps.
- True False
- (c) Given a linearly separable dataset, the perceptron algorithm is guaranteed to find a max-margin separating hyperplane.
- True False
- (d) You would like to train a neural network to classify digits. Your network takes as input an image and outputs probabilities for each of the 10 classes, 0-9. The network's prediction is the class that it assigns the highest probability to. From the following functions, select all that would be suitable loss functions to minimize using gradient descent:
- The square of the difference between the correct digit and the digit predicted by your network
- The probability of the correct digit under your network
- The negative log-probability of the correct digit under your network
- None of the above

Q4. Perceptron and Kernels

A kernel is a mapping $K(x, y)$ from pairs vectors in \mathbb{R}^d into the real numbers such that $K(x, y) = \Phi(x) \cdot \Phi(y)$ where Φ is a mapping from \mathbb{R}^d into \mathbb{R}^D where D is possibly different from d and even infinite. We say that a mapping $K(x, y)$ for which such Φ exists is a valid kernel.

(a) The following binary class data has two features, A and B .

Index	A	B	Class
1.	1	1	1
2.	0	3	-1
3.	1	-1	1
4.	3	0	-1
5.	-1	1	1
6.	0	-3	-1
7.	-1	-1	1
8.	-3	0	-1

(i) Select all true statements:

- This data is linearly separable.
- This data is linearly separable if we use a feature map $\phi((A, B)) = (A^2, B^2, 1)$.
- There exists a kernel such that this data is linearly separable.
- For all datasets in which no data point is labeled in more than one distinct way, there exists a kernel such that the data is linearly separable.
- For all datasets, there exists a kernel such that the data is linearly separable.
- For all valid kernels, there exists a dataset with at least one point from each class that is linearly separable under that kernel.
- None of the above.

We will be running both the primal (normal) binary (not multiclass) perceptron and dual binary perceptron algorithms on this dataset. We will initialize the weight vector w to $(1, 1)$ for the primal perceptron algorithm. Accordingly, we will initialize the α vector to $(1, 0, 0, 0, 0, 0, 0, 0)$ for the dual perceptron algorithm with the kernel $K(x, y) = x \cdot y$. Pass through the data using the indexing order provided. There is no bias term.

Write your answer in the box provided. Show your work outside of the boxes to have a chance at receiving partial credit.

(ii) What is the first misclassified point?

(iii) For the *primal* perceptron algorithm, what is the weight vector after the first weight update?

For your convenience, the data is duplicated on this page.

Index	A	B	Class
1.	1	1	1
2.	0	3	-1
3.	1	-1	1
4.	3	0	-1
5.	-1	1	1
6.	0	-3	-1
7.	-1	-1	1
8.	-3	0	-1

(iv) For the *dual* perceptron algorithm, what is the α vector after the first weight update?

(v) What is the second misclassified point?

(vi) For the *primal* perceptron algorithm, what is the weight vector after the second weight update?


(vii) For the *dual* perceptron algorithm, what is the α vector after the second weight update?

(b) Consider the following kernel function: $K(\mathbf{x}, \mathbf{y}) = (\mathbf{x} \cdot \mathbf{y})^2$ where $\mathbf{x}, \mathbf{y} \in \mathbb{R}^2$. Find a valid Φ map for this kernel. That is, find a vector-to-vector function ϕ such that $\phi(\mathbf{x}) \cdot \phi(\mathbf{y}) = K(\mathbf{x}, \mathbf{y}) = (\mathbf{x} \cdot \mathbf{y})^2$. Show work to have a chance at receiving partial credit. Any precise answer format is acceptable.

Q5. Clustering

In this question, we will do k -means clustering to cluster the points $A, B \dots F$ (indicated by \times 's in the figure on the right) into 2 clusters. The current cluster centers are P and Q (indicated by the \blacksquare in the diagram on the right). Recall that k -means requires a distance function. Given 2 points, $A = (A_1, A_2)$ and $B = (B_1, B_2)$, we use the following distance function $d(A, B)$ that you saw from class,

$$d(A, B) = (A_1 - B_1)^2 + (A_2 - B_2)^2$$



- (a) **Update assignment step:** Select all points that get assigned to the cluster with center at P :
- A B C D E F No point gets assigned to cluster P
- (b) **Update cluster center step:** What does cluster center P get updated to?


Changing the distance function: While k -means used Euclidean distance in class, we can extend it to other distance functions, where the assignment and update phases still iteratively minimize the total (non-Euclidian) distance. Here, consider the Manhattan distance:

$$d'(A, B) = |A_1 - B_1| + |A_2 - B_2|$$







We again start from the original locations for P and Q as shown in the figure, and do the update assignment step and the update cluster center step using Manhattan distance as the distance function:

- (c) **Update assignment step:** Select all points that get assigned to the cluster with center at P , under this new distance function $d'(A, B)$.
- A B C D E F No point gets assigned to cluster P
- (d) **Update cluster center step:** What does cluster center P get updated to, under this new distance function $d'(A, B)$?

Q6. Neural Networks: Representation



For each of the piecewise-linear functions below, mark all networks from the list above that can represent the function **exactly** on the range $x \in (-\infty, \infty)$. In the networks above, *relu* denotes the element-wise ReLU nonlinearity: $relu(z) = \max(0, z)$. The networks G_i use 1-dimensional layers, while the networks H_i have some 2-dimensional intermediate layers.











Q7. Decision Trees

You are given points from 2 classes, shown as '+'s and '.'s. For each of the following sets of points,

1. Draw the decision tree of depth at most 2 that can separate the given data completely, by filling in binary predicates (which only involve thresholding of a *single* variable) in the boxes for the decision trees below. If the data is already separated when you hit a box, simply write the class, and leave the sub-tree hanging from that box empty.
2. Draw the corresponding decision boundaries on the scatter plot, and write the class labels for each of the resulting bins somewhere inside the resulting bins.


If the data can not be separated completely by a depth 2 decision tree, simply cross out the tree template. We solve the first part as an example.



Q8. Deep Learning

(a) Perform forward propagation on the neural network below for $x = 1$ by filling in the values in the table. Note that (i), ..., (vii) are outputs after performing the appropriate operation as indicated in the node.

(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)




(b) Below is a neural network with weights a, b, c, d, e, f . The inputs are x_1 and x_2 . The first hidden layer computes $r_1 = \max(c \cdot x_1 + e \cdot x_2, 0)$ and $r_2 = \max(d \cdot x_1 + f \cdot x_2, 0)$. The second hidden layer computes $s_1 = \frac{1}{1 + \exp(-a \cdot r_1)}$ and $s_2 = \frac{1}{1 + \exp(-b \cdot r_2)}$. The output layer computes $y = s_1 + s_2$. Note that the weights a, b, c, d, e, f are indicated along the edges of the neural network here.

Suppose the network has inputs $x_1 = 1, x_2 = -1$.

The weight values are $a = 1, b = 1, c = 4, d = 1, e = 2, f = 2$.

Forward propagation then computes $r_1 = 2, r_2 = 0, s_1 = 0.9, s_2 = 0.5, y = 1.4$. Note: some values are rounded.




Using the values computed from forward propagation, use backpropagation to numerically calculate the following partial derivatives. Write your answers as a single number (not an expression). You do not need a calculator. Use scratch paper if needed.

Hint: For $g(z) = \frac{1}{1 + \exp(-z)}$, the derivative is $\frac{\partial g}{\partial z} = g(z)(1 - g(z))$.

$\frac{\partial y}{\partial a}$	$\frac{\partial y}{\partial b}$	$\frac{\partial y}{\partial c}$	$\frac{\partial y}{\partial d}$	$\frac{\partial y}{\partial e}$	$\frac{\partial y}{\partial f}$


(c) Below are two plots with horizontal axis x_1 and vertical axis x_2 containing data labelled \times and \bullet . For each plot, we wish to find a function $f(x_1, x_2)$ such that $f(x_1, x_2) \geq 0$ for all data labelled \times and $f(x_1, x_2) < 0$ for all data labelled \bullet .

Below each plot is the function $f(x_1, x_2)$ for that specific plot. Complete the expressions such that all the data is labelled correctly. If not possible, mark "No valid combination".



$$f(x_1, x_2) = \max(\underline{\text{(i)}} + \underline{\text{(ii)}}, \underline{\text{(iii)}} + \underline{\text{(iv)}}) + \underline{\text{(v)}}$$

- | | | | | | | |
|-------|-----------------------|----------------------|-----------------------|--------|-----------------------|---|
| (i) | <input type="radio"/> | x_1 | <input type="radio"/> | $-x_1$ | <input type="radio"/> | 0 |
| (ii) | <input type="radio"/> | x_2 | <input type="radio"/> | $-x_2$ | <input type="radio"/> | 0 |
| (iii) | <input type="radio"/> | x_1 | <input type="radio"/> | $-x_1$ | <input type="radio"/> | 0 |
| (iv) | <input type="radio"/> | x_2 | <input type="radio"/> | $-x_2$ | <input type="radio"/> | 0 |
| (v) | <input type="radio"/> | 1 | <input type="radio"/> | -1 | <input type="radio"/> | 0 |
| | <input type="radio"/> | No valid combination | | | | |



$$f(x_1, x_2) = \underline{\text{(vi)}} - \max(\underline{\text{(vii)}} + \underline{\text{(viii)}}, \underline{\text{(ix)}} + \underline{\text{(x)}})$$

- | | | | | | | |
|--------|-----------------------|----------------------|-----------------------|--------|-----------------------|---|
| (vi) | <input type="radio"/> | x_2 | <input type="radio"/> | $-x_2$ | <input type="radio"/> | 0 |
| (vii) | <input type="radio"/> | x_1 | <input type="radio"/> | $-x_1$ | <input type="radio"/> | 0 |
| (viii) | <input type="radio"/> | x_2 | <input type="radio"/> | $-x_2$ | <input type="radio"/> | 0 |
| (ix) | <input type="radio"/> | x_1 | <input type="radio"/> | $-x_1$ | <input type="radio"/> | 0 |
| (x) | <input type="radio"/> | x_2 | <input type="radio"/> | $-x_2$ | <input type="radio"/> | 0 |
| | <input type="radio"/> | No valid combination | | | | |