


CSE 546 Midterm Exam, Fall 2014

1. Personal info:
 - Name:
 - UW NetID:
 - Student ID:
2. There should be 14 numbered pages in this exam (including this cover sheet).
3. You can use any material you brought: any book, class notes, your print outs of class materials that are on the class website, including my annotated slides and relevant readings. You cannot use materials brought by other students.
4. If you need more room to work out your answer to a question, use the back of the page and clearly mark on the front of the page if we are to look at what's on the back.
5. Work efficiently. Some questions are easier, some more difficult. Be sure to give yourself time to answer all of the easy ones, and avoid getting bogged down in the more difficult ones before you have answered the easier ones.
6. You have 80 minutes.
7. Good luck!


Question	Topic	Max score	Score
1	Short Answer	24	
2	Decision Trees	16	
3	Logistic Regression	20	
4	Boosting	16	
5	Elastic Net	24	
Total		100	

1 [24 points] Short Answer


1. [6 points] On the 2D dataset below, draw the decision boundaries learned by the following algorithms (using the features x/y). **Be sure to mark which regions are labeled positive or negative, and assume that ties are broken arbitrarily.**



(a) Logistic regression ($\lambda = 0$)



(b) 1-NN



(c) 3-NN

2. [6 points] A random variable follows an *exponential* distribution with parameter λ ($\lambda > 0$) if it has the following density:

$$p(t) = \lambda e^{-\lambda t}, \quad t \in [0, \infty)$$

This distribution is often used to model waiting times between events. Imagine you are given i.i.d. data $T = (t_1, \dots, t_n)$ where each t_i is modeled as being drawn from an exponential distribution with parameter λ .

- (a) [3 points] Compute the log-probability of T given λ . (Turn products into sums when possible).

(b) [3 points] Solve for $\hat{\lambda}_{MLE}$.

3. [6 points]

A	B	C	Y
F	F	F	F
T	F	T	T
T	T	F	T
T	T	T	F

(a) [3 points] Using the dataset above, we want to build a decision tree which classifies Y as T/F given the binary variables A, B, C . Draw the tree that would be learned by the greedy algorithm with zero training error. You do not need to show any computation.

(b) [3 points] Is this tree optimal (i.e. does it get zero training error with minimal depth)? Explain in less than two sentences. If it is not optimal, draw the optimal tree as well.

4. [6 points] In class we used decision trees and ensemble methods for classification, but we can use them for regression as well (i.e. learning a function from features to real values). Let's imagine that our data has 3 binary features A, B, C , which take values 0/1, and we want to learn a function which counts the number of features which have value 1.

(a) [2 points] Draw the decision tree which represents this function. How many leaf nodes does it have?

(b) [2 points] Now represent this function as a sum of decision stumps (e.g. $\text{sgn}(A)$). How many terms do we need?

(c) [2 points] In the general case, imagine that we have d binary features, and we want to count the number of features with value 1. How many leaf nodes would a decision tree need to represent this function? If we used a sum of decision stumps, how many terms would be needed? No explanation is necessary.

2 [16 points] Decision Trees

We will use the dataset below to learn a decision tree which predicts if people pass machine learning (Yes or No), based on their previous GPA (High, Medium, or Low) and whether or not they studied.

GPA	Studied	Passed
L	F	F
L	T	T
M	F	F
M	T	T
H	F	T
H	T	T

For this problem, you can write your answers using \log_2 , but it may be helpful to note that $\log_2 3 \approx 1.6$.

1. [4 points] What is the entropy $H(\text{Passed})$?
2. [4 points] What is the entropy $H(\text{Passed} \mid \text{GPA})$?
3. [4 points] What is the entropy $H(\text{Passed} \mid \text{Studied})$?
4. [4 points] Draw the full decision tree that would be learned for this dataset. You do not need to show any calculations.

3 [20 points] Logistic Regression for Sparse Data

In many real-world scenarios our data has millions of dimensions, but a given example has only hundreds of non-zero features. For example, in document analysis with word counts for features, our dictionary may have millions of words, but a given document has only hundreds of unique words. In this question we will make l_2 regularized SGD efficient when our input data is sparse. Recall that in l_2 regularized logistic regression, we want to maximize the following objective (**in this problem we have excluded w_0 for simplicity**):

$$F(\mathbf{w}) = \frac{1}{N} \sum_{j=1}^N l(\mathbf{x}^{(j)}, y^{(j)}, \mathbf{w}) - \frac{\lambda}{2} \sum_{i=1}^d \mathbf{w}_i^2$$

where $l(\mathbf{x}^{(j)}, y^{(j)}, \mathbf{w})$ is the logistic objective function

$$l(\mathbf{x}^{(j)}, y^{(j)}, \mathbf{w}) = y^{(j)} \left(\sum_{i=1}^d \mathbf{w}_i \mathbf{x}_i^{(j)} \right) - \ln(1 + \exp(\sum_{i=1}^d \mathbf{w}_i \mathbf{x}_i^{(j)}))$$

and the remaining sum is our regularization penalty.

When we do stochastic gradient descent on point $(\mathbf{x}^{(j)}, y^{(j)})$, we are approximating the objective function as

$$F(\mathbf{w}) \approx l(\mathbf{x}^{(j)}, y^{(j)}, \mathbf{w}) - \frac{\lambda}{2} \sum_{i=1}^d \mathbf{w}_i^2$$

Definition of sparsity: Assume that our input data has d features, i.e. $\mathbf{x}^{(j)} \in \mathbb{R}^d$. In this problem, we will consider the scenario where $\mathbf{x}^{(j)}$ is sparse. Formally, let s be average number of nonzero elements in each example. We say the data is sparse when $s \ll d$. In the following questions, **your answer should take the sparsity of $\mathbf{x}^{(j)}$ into consideration when possible**. **Note:** When we use a sparse data structure, we can iterate over the non-zero elements in $O(s)$ time, whereas a dense data structure requires $O(d)$ time.

1. [2 points] Let us first consider the case when $\lambda = 0$. Write down the SGD update rule for \mathbf{w}_i when $\lambda = 0$, using step size η , given the example $(\mathbf{x}^{(j)}, y^{(j)})$.
2. [4 points] If we use a dense data structure, what is the average time complexity to update \mathbf{w}_i when $\lambda = 0$? What if we use a sparse data structure? Justify your answer in one or two sentences.

3. [2 points] Now let us consider the general case when $\lambda > 0$. Write down the SGD update rule for \mathbf{w}_i when $\lambda > 0$, using step size η , given the example $(\mathbf{x}^{(j)}, y^{(j)})$.

4. [2 points] If we use a dense data structure, what is the average time complexity to update \mathbf{w}_i when $\lambda > 0$?

5. [4 points] Let $\mathbf{w}_i^{(t)}$ be the weight vector after t -th update. Now imagine that we perform k SGD updates on \mathbf{w} using examples $(\mathbf{x}^{(t+1)}, y^{(t+1)}), \dots, (\mathbf{x}^{(t+k)}, y^{(t+k)})$, where $\mathbf{x}_i^{(j)} = 0$ for every example in the sequence. (i.e. the i -th feature is zero for all of the examples in the sequence). Express the new weight, $\mathbf{w}_i^{(t+k)}$ in terms of $\mathbf{w}_i^{(t)}$, k , η , and λ .

6. [6 points] Using your answer in the previous part, come up with an efficient algorithm for regularized SGD when we use a sparse data structure. What is the average time complexity per example? (Hint: when do you need to update \mathbf{w}_i ?)

4 [16 points] Boosting


Recall that Adaboost learns a classifier H using a weighted sum of weak learners h_t as follows

$$H(x) = \text{sgn} \left(\sum_{t=1}^T \alpha_t h_t(x) \right)$$


In this question we will use decision trees as our weak learners, which classify a point as $\{1, -1\}$ based on a sequence of threshold splits on its features (here x, y).

In the questions below, be sure to mark which regions are marked positive/negative, and assume that ties are broken arbitrarily.


1. [2 points] Assume that our weak learners are decision trees of depth 1 (i.e. decision stumps), which minimize the weighted training error. Using the dataset below, draw the decision boundary learned by h_1 .




2. [3 points] On the dataset below, circle the point(s) with the highest weights on the second iteration, and draw the decision boundary learned by h_2 .




3. [3 points] On the dataset below, draw the decision boundary of $H = \text{sgn}(\alpha_1 h_1 + \alpha_2 h_2)$. (Hint, you do not need to explicitly compute the α 's).




4. [2 points] Now assume that our weak learners are decision trees of maximum depth 2, which minimize the weighted training error. Using the dataset below, draw the decision boundary learned by h_1 .



5. [3 points] On the dataset below, circle the point(s) with the highest weights on the second iteration, and draw the decision boundary learned by h_2 .



6. [3 points] On the dataset below, draw the decision boundary of $H = \text{sgn}(\alpha_1 h_1 + \alpha_2 h_2)$.
(Hint, you do not need to explicitly compute the α 's).



5 [24 points] Elastic Net

In class we discussed two types of regularization, l_1 and l_2 . Both are useful, and sometimes it is helpful combine them, giving the objective function below (**in this problem we have excluded w_0 for simplicity**):

$$F(\mathbf{w}) = \frac{1}{2} \sum_{j=1}^n (y^{(j)} - \sum_{i=1}^d \mathbf{w}_i \mathbf{x}_i^{(j)})^2 + \alpha \sum_{i=1}^d |\mathbf{w}_i| + \frac{\lambda}{2} \sum_{i=1}^d \mathbf{w}_i^2 \quad (1)$$

Here $(\mathbf{x}^{(j)}, y^{(j)})$ is j -th example in the training data, \mathbf{w} is a d dimensional weight vector, λ is a regularization parameter for the l_2 norm of \mathbf{w} , and α is a regularization parameter for the l_1 norm of \mathbf{w} . This approach is called the Elastic Net, and you can see that it is a generalization of Ridge and Lasso regression: It reverts to Lasso when $\lambda = 0$, and it reverts to Ridge when $\alpha = 0$. In this question, we are going to derive the coordinate descent (CD) update rule for this objective.

Let g, h, c be real constants, and consider the function of x

$$f_1(x) = c + gx + \frac{1}{2}hx^2 \quad (h > 0) \quad (2)$$

1. [2 points] What is the x^* that minimizes $f_1(x)$? (i.e. calculate $x^* = \arg \min f_1(x)$)

Let α be an additional real constant, and consider another function of x

$$f_2(x) = c + gx + \frac{1}{2}hx^2 + \alpha|x| \quad (h > 0, \alpha > 0) \quad (3)$$

This is a piecewise function, composed of two quadratic functions:

$$f_2^-(x) = c + gx + \frac{1}{2}hx^2 - \alpha x$$

and

$$f_2^+(x) = c + gx + \frac{1}{2}hx^2 + \alpha x$$

Let $\tilde{x}^- = \arg \min_{x \in \mathbb{R}} f_2^-(x)$ and $\tilde{x}^+ = \arg \min_{x \in \mathbb{R}} f_2^+(x)$.
(**Note:** The argmin is taken over $(-\infty, +\infty)$ here)

2. [3 points] What are \tilde{x}^+ and \tilde{x}^- ? Show that $\tilde{x}^- \geq \tilde{x}^+$.

3. [6 points] Draw a picture of $f_2(x)$ in each of the three cases below.

(a) $\tilde{x}^- > 0, \tilde{x}^+ > 0$

(b) $\tilde{x}^- < 0, \tilde{x}^+ < 0$

(c) $\tilde{x}^- > 0, \tilde{x}^+ < 0$.

For each case, mark the minimum as either 0, \tilde{x}^- , or \tilde{x}^+ . (You do not need to draw perfect curves, just get the rough shape and the relative locations of the minima to the x -axis)

4. [4 points] Write x^* , the minimum of $f_2(x)$, as a piecewise function of g . (Hint: what is g in the cases above?)
5. [4 points] Now let's derive the update rule for \mathbf{w}_k . Fixing all parameters but \mathbf{w}_k , express the objective function in the form of Eq (3) (here \mathbf{w}_k is our x). You do not need to explicitly write constant factors, you can put them all together as c . What are g and h ?
6. [5 points] Putting it all together, write down the coordinate descent algorithm for the Elastic Net (**again excluding** w_0). You can write the update for \mathbf{w}_k in terms of the g and h you calculated in the previous question.