

Midterm exam CS 189/289, Fall 2015

- You have **80 minutes** for the exam.
- **Total 100 points:**
 1. **True/False:** 36 points (18 questions, 2 points each).
 2. **Multiple-choice questions:** 24 points (8 questions, 3 points each).
 3. **Three descriptive** questions worth 10, 15, 15 points.
- The exam is closed book, closed notes except your one-page crib sheet.
- No calculators or electronic items.
- For true/false questions, fill in the True/False bubble.
- **WARNING** For multiple-choice questions, fill in the bubbles for **ALL CORRECT CHOICES (in some cases, there may be more than one). NO PARTIAL CREDIT: all correct answers must be checked and no incorrect answers should be checked.**

First name	
Last name	
SID	
First and last name of student to your left	
First and last name of student to your right	

For staff only	
T/F	/36
Multiple choice	/24
Problem I	/15
Problem II	/15
Problem III	/10
Total	/100

Notation:

\mathbf{X} : the training data matrix of dimension (N, d) , of N rows representing samples and d columns representing features.

\mathbf{x} : an input data vector of dimension $(1, d)$ of components $x_i, i=1:d$.

\mathbf{x}^k : a training example of dimension $(1, d)$ is a row of \mathbf{X} , $k=1:N$.

\mathbf{w} : weight vector of a linear model of dimension $(1, d)$ such that

$$f(\mathbf{x}) = \mathbf{w} \mathbf{x}^T = \mathbf{x} \mathbf{w}^T = \sum_{i=1:d} w_i x_i$$

\mathbf{y} : target vector of dimension $(N, 1)$ of components y^k .

α : weight vector of dimension $(N, 1)$ of kernel method $f(\mathbf{x}) = \sum_{k=1:N} \alpha_k k(\mathbf{x}, \mathbf{x}^k)$

$k(\mathbf{u}, \mathbf{v})$: a kernel function (a similarity measure between two samples \mathbf{u} and \mathbf{v}).

True/False (36 points):

1. Stochastic gradient descent performs less computation per update than batch gradient descent.

TRUE

FALSE

2. A function is convex if its Hessian is negative semidefinite.

TRUE

FALSE

3. If $N < d$, the solution to $\mathbf{X}\mathbf{w}^T = \mathbf{y}$ is unique.

TRUE

FALSE

4. A support vector machine computes $P(y|\mathbf{x})$.

TRUE

FALSE

5. Adding a ridge to $X^T X$ guarantees that it is invertible.
TRUE FALSE
6. Grid search is less prone to being trapped in a local minimum than other heuristic search methods.
TRUE FALSE
7. The bootstrap method involves sampling without replacement.
TRUE FALSE
8. A non linearly-separable training set in a given feature space can always be made linearly-separable in another space.
TRUE FALSE
9. Using the kernel trick, one can get non-linear decision boundaries using algorithms designed originally for linear models.
TRUE FALSE
10. Logistic regression cannot be kernelized.
TRUE FALSE
11. Ridge regression, weight decay, and Gaussian processes use the same regularizer: $\|\mathbf{w}\|^2$.
TRUE FALSE
12. Hebb's rule computes the centroid method solution, if the target values are $+1/N_1$ and $-1/N_0$ (N_1 and N_0 are the number of examples of each class)
TRUE FALSE

13. Any kernel method can be thought of as a parametric method in a possibly infinite dimensional space.

TRUE

FALSE

14. Nearest neighbors is a parametric method.

TRUE

FALSE

15. A symmetric matrix is positive semidefinite if all its eigenvalues are positive or zero.

TRUE

FALSE

16. Zero correlation between any two random variables implies that the two random variables are independent.

TRUE

FALSE

17. The Linear Discriminant Analysis (LDA) classifier computes the direction maximizing the ratio of between-class variance over within-class variance.

TRUE

FALSE

18. If we repeat an experiment twice and get p-values p_1 and p_2 , the minimum of the two p-values is the p-value of the overall experiment.


TRUE

FALSE

Multiple choice questions (24 points)

1. You trained a binary classifier model which gives very high accuracy on the training data, but much lower accuracy on validation data. The following may be true:
 - This is an instance of overfitting.
 - This is an instance of underfitting.
 - The training was not well regularized.
 - The training and testing examples are sampled from different distributions.
2. Okham in the 14th century is credited to have stated that one should “shave off unnecessary parameters of a model”. Which of the following implement that principle:
 - Regularization.
 - Maximum likelihood estimation.
 - Shrinkage.
 - Empirical risk minimization.
 - Feature selection.
3. Good practices to avoid overfitting include:
 - Using a two part cost function which includes a regularizer to penalize model complexity.
 - Using a good optimizer to minimize error on training data.
 - Building a structure of nested subsets of models and train learning machines in each subset, starting from the inner subset, and stopping when the cross-validation error starts increasing.
 - Discarding 50% of randomly chosen samples.

4. Wrapper methods are hyper-parameter selection methods that:
- Should be used whenever possible because they are computationally efficient.
 - Should be avoided unless there are no other options because they are always prone to overfitting.
 - Are useful mainly when the learning machines are “black boxes”.
 - Should be avoided altogether.
5. Three different classifiers are trained on the same data. Their decision boundaries are shown below. Which of the following statements are true?



- The leftmost classifier has high robustness, poor fit.
 - The leftmost classifier has poor robustness, high fit.
 - The rightmost classifier has poor robustness, high fit.
 - The rightmost classifier has high robustness, poor fit.
6. What are support vectors:
- The examples farthest from the decision boundary.
 - The only examples necessary to compute $f(\mathbf{x})$ in an SVM.
 - The class centroids.
 - All the examples that have a non-zero weight α_k in a SVM.

7. Which of the following does not converge to a solution if the training samples are not linearly separable?

- Linear Logistic Regression.
- Linear Soft margin SVM.
- Linear hard-margin SVM.
- The centroid method.
- Parzen windows.

8. The number of test examples needed to get statistically significant results should be:

- Larger if the error rate is larger.
- Larger if the error rate is smaller.
- It does not matter.

Three descriptive problems

Problem I: Gradient descent (15 points).

Given N training data points $\{(\mathbf{x}^k, y^k)\}$, $k=1:N$, \mathbf{x}^k in \mathbb{R}^d , and labels in y^k in $\{-1,1\}$, we seek a linear discriminant function $f(\mathbf{x}) = \mathbf{w} \cdot \mathbf{x}$ optimizing the loss function $L(z) = e^{-z}$, for $z=y f(\mathbf{x})$.

Question I.1 (3 points) Is $L(z)$ a large margin loss function? Justify your answer (a graphical justification may be useful).

Question 1.2 (4 points) Derive the stochastic gradient descent update $\mathbf{w} \leftarrow \mathbf{w} + \Delta \mathbf{w}$ for $L(z)$, where $\Delta \mathbf{w}$ is the difference between two consecutive values of \mathbf{w} :

Question 1.3 (3 points) We call $R_{\text{emp}}(\mathbf{w}) = \sum_{k=1:N} L(z^k)$, where $z^k = y^k f(\mathbf{x}^k)$, the “empirical risk”. Derive the batch gradient update for the empirical risk:

Question 1.4 (3 points) Suppose you also want to include a penalty term $\lambda \|\mathbf{w}\|^2$ to the risk functional that you wish to minimize. Derive the batch gradient update for the regularized risk $R_{\text{reg}}(\mathbf{w}) = R_{\text{emp}}(\mathbf{w}) + \lambda \|\mathbf{w}\|^2$:

Question 1.5 (2 points) How do you estimate λ (answer in at most 3 words)?

Problem II. Classification concept review (15 points).

Question II.1. **Centroid method.** Now consider a 2-class classification problem in a 2-dimensional feature space $x=[x_1, x_2]$ with target variable $y=\pm 1$. The training data comprises 7 samples as shown in Figure 1 (4 black diamonds for the positive class and 3 white diamonds for the negative class). The 7 samples are also numbered for your reference.




Figure 1: Data for Problem II.1 Centroid method question.

Question II.1.A (2 points): Draw on Figure 1 the **centroids of the two classes** (mark them with a circled “+” for the positive class and a circled “-” for the negative class). **Join the centroids with a thick dashed line.** Draw the **decision boundary** of the centroid method with a **thick solid line.**

Question II.1.B (1 point) What is the **training error rate**?

Question II.1.C (2 points) Is there any sample such that upon its removal, the decision boundary changes in a manner that the removed sample goes to the other side (Answer “yes” or “no”)?

Question II.1.D (2 points) What is the **leave-one-out error rate**?

Question II. 2: **Support Vector Machine (SVM)**. Consider again the same training data as in Question II.1, replicated in Figure 2, for your convenience. The “maximum margin classifier” (also called linear “hard margin” SVM) is a classifier that leaves the largest possible margin on either side of the decision boundary. The samples lying on the margin are called support vectors.




Figure 1: Data for Problem II.2 SVM method question.

Question II.2.A (2 points): Draw on Figure 2 the **decision** boundary obtained by the **linear hard margin SVM** method with a **thick solid line**. Draw the **margins** on either side with **thinner dashed lines**. **Circle the support vectors**.

Question II.2.B (1 point) What is the **training error rate**?

Question II.2.C (1 point) The removal of which sample will change the decision boundary?

Question II.2.D (2 points) What is the leave-one-out error rate?

Question II.2.E (1 point) A method is more robust if the difference between training error and leave-one-out error is smaller. **Which of the two methods (centroid or SVM) is more robust?**

Question II.3.F (1 point) A method has a better fit if it has a smaller training error. **Which of the two methods has the best fit?**

Problem III. Newton-Raphson for least-square regression (10 points)

[Hard problem, attempt only if you have time.]

In this problem, we will derive an optimization algorithm which we did not study in class, called the **Newton-Raphson** algorithm. The algorithm makes updates in a manner that often allows reaching the solution faster than regular gradient descent.

Suppose we start with an initial value of a $(1, d)$ vector \mathbf{w} that we call $\mathbf{w}^{(0)}$. We know that the first order Taylor approximation of $\nabla_{\mathbf{w}}R(\mathbf{w}^{(1)})$, at the point $\mathbf{w}^{(0)}$ is:

$$\nabla_{\mathbf{w}}R(\mathbf{w}^{(1)}) \approx \nabla_{\mathbf{w}}R(\mathbf{w}^{(0)}) + (\mathbf{w}^{(1)} - \mathbf{w}^{(0)}) \nabla_{\mathbf{w}}^2R(\mathbf{w}^{(0)})$$

Question III.1 (3 points). We want to minimize $R(\mathbf{w}^{(1)})$ using this approximation of $\nabla_{\mathbf{w}}R(\mathbf{w}^{(1)})$. Find the update equation for the value of $\mathbf{w}^{(1)}$. This is called the Newton-Raphson update. Notes: This is not a trick question, you just have to solve for $\mathbf{w}^{(1)}$ after equaling $\nabla_{\mathbf{w}}R(\mathbf{w}^{(1)})$ to 0. You can assume that the (d, d) Hessian matrix $\nabla_{\mathbf{w}}^2R(\mathbf{w}^{(0)})$ is invertible.

Question III.2 (4 points). Consider now the linear regression problem: We are given a training data matrix X of dim (N, d) and a target vector \mathbf{y} of dim $(N, 1)$ and want to find a weight vector \mathbf{w} of dim $(1, d)$ such that $f(\mathbf{x}) = \mathbf{x} \mathbf{w}^T$ approximates \mathbf{y} best, in the least square sense. The risk functional is: $R(\mathbf{w}) = (\mathbf{X}\mathbf{w}^T - \mathbf{y})^T (\mathbf{X}\mathbf{w}^T - \mathbf{y})$. We will assume that we are in the “regression case” $N > d$ and that the Hessian is invertible. Find the Newton-Raphson update for $\mathbf{w}^{(1)}$.

Question III.3 (3 points). Recall the solution to the problem we found in class using the normal equations or the solution found by solving for $\nabla_{\mathbf{w}} R(\mathbf{w}) = 0$ directly. Compare with the solution obtained in question (2). How many iterations of the Newton-Raphson update do we need to perform for linear regression?