Midterm, CS-GY 6923

Sample Questions

Show all of your work to receive full (and partial) credit.

Always, Sometimes, Never

Indicate whether each of the following statements is **always** true, **sometimes** true, or **never** true. Provide a one or two short justification or example to explain your choice.

1. For random events $p(x \mid y) < p(x, y)$.

ALWAYS SOMETIMES NEVER

2. Consider a loss function L. If $\nabla L(\beta) = 0$, then β is a minimizer of L.

ALWAYS SOMETIMES NEVER

3. The empirical risk of a model is lower than the population risk.

ALWAYS SOMETIMES NEVER

4. The linear classifier found by logistic regression minimizes error rate (0-1 loss) on the training data.

ALWAYS SOMETIMES NEVER

5. Suppose we learn a linear classifier. I.e., we learn parameters β and classify an input vector x in class 1 if $\mathbb{1}[\langle x, b\rangle > \lambda]$. Increasing λ Increases the recall of our classifier.

ALWAYS SOMETIMES NEVER

Short Answer

5. You are trying to develop a machine learning algorithm for classifying data $\mathbf{x}_1, \ldots, \mathbf{x}_n \in \mathbb{R}^d$ into catagories $1, \ldots, q$. You have decided to use linear classification for the problem.

   (a) You know you can find a good linear classifier for *binary* classification (dividing into $q = 2$ classes) using logistic regression. You are considering using either the **one-vs-all** or **one-vs-one** approach to adapting this approach to the multiclass problem. In a few sort sentences describe why you might use one over the other.

   (b) Your coworker suggests the following alternative approach: let's try to learn a parameter vector $\boldsymbol{\beta} \in \mathbb{R}^d$ and classify using the following model:

$$f_{\beta}(\mathbf{x}) = \begin{cases} 1 \text{ if} & \langle \boldsymbol{\beta}, \mathbf{x} \rangle \leq 1 \\ 2 \text{ if} & 2 < \langle \boldsymbol{\beta}, \mathbf{x} \rangle \leq 3 \\ 3 \text{ if} & 3 < \langle \boldsymbol{\beta}, \mathbf{x} \rangle \leq 4 \\ \vdots & \\ q-1 \text{ if} & q-2 < \langle \boldsymbol{\beta}, \mathbf{x} \rangle \leq q-1 \\ q \text{ if} & q-1 < \langle \boldsymbol{\beta}, \mathbf{x} \rangle \end{cases} \tag{1}$$

(c) Describe **one potential issue** and **one potental benefit** of your coworker's method over the approaches mentioned in (a). There is no one "right" answer here.

(d) For the two datasets D_1 and D_2 below, indicate which of the three approaches (**one-vs-one**, **one-vs-all**, or your **coworkers approach**) would lead to an accurate solution to the multiclass classification problem. No explanation is required, but having one might help you earn partial credit.

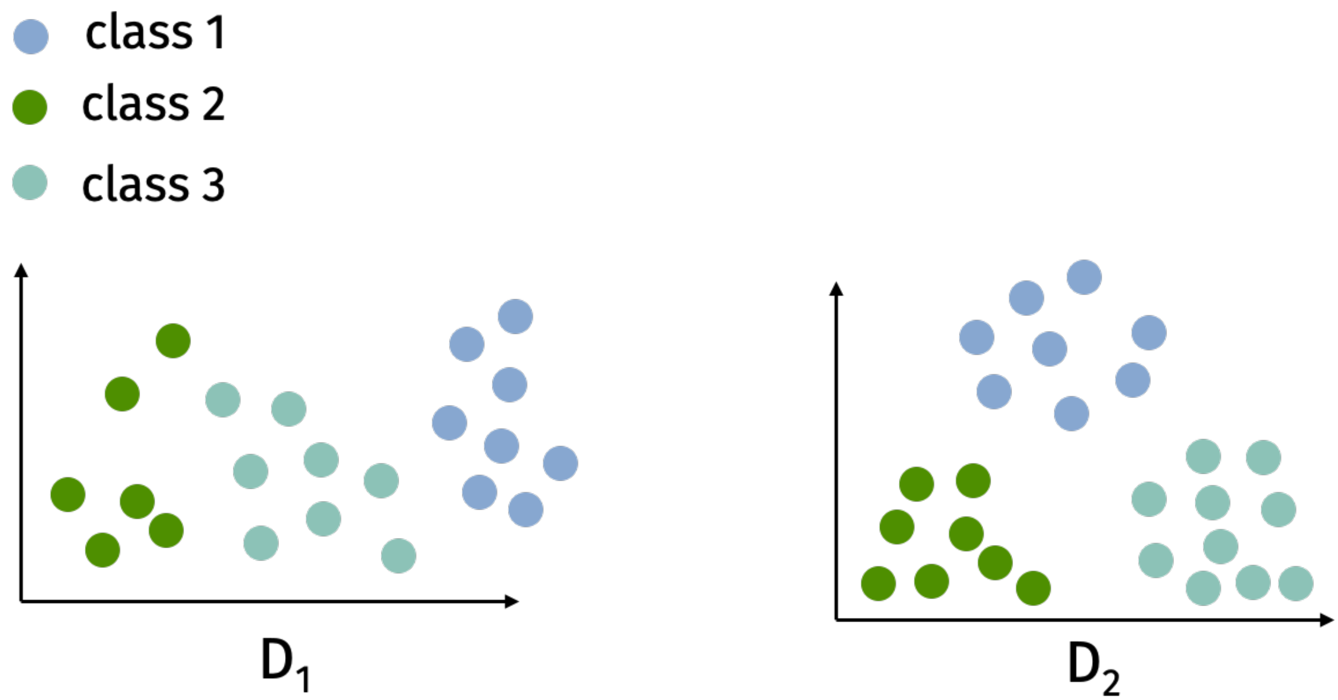

6. We are given data with just one predictor variable and one target: $(x_1, y_1), \ldots, (x_n, y_n)$, with the goal of fitting a degree two polynomial model using unregularized multiple linear regression with data transformation. The goal is to find the best coefficients $\beta_0, \beta_1, \beta_2$ for predicting y as $\beta_0 + \beta_1 x + \beta_2 x^2$.

Consider the following three transformed data matrices:

$$X_1 = \begin{bmatrix} 1 & x_1 & x_1^2 \\ \vdots & \vdots & \vdots \\ 1 & x_n & x_n^2 \end{bmatrix}, X_2 = \begin{bmatrix} 1 & x_1^2 - x_1 & x_1^2 \\ \vdots & \vdots & \vdots \\ 1 & x_n^2 - x_n & x_n^2 \end{bmatrix}, \text{ and } X_3 = \begin{bmatrix} 1 & 2x_1^2 - x_1 & 2x_1 - 4x_1^2 \\ \vdots & \vdots & \vdots \\ 1 & 2x_n^2 - x_n & 2x_n - 4x_n^2 \end{bmatrix}$$

Which of the above matrices can be used to solve this problem? In other words, if we train a multiple linear regression problem with X_i can we obtain an optimal degree two polynomial fit for $y_1, \ldots, y_n$. Justify your answer in words, or with equations.

7. Write each of the following models as transformed linear models. That is, find a parameter vector $\boldsymbol{\beta}$ in terms of the given parameters a_i and data transformation $\phi(\mathbf{x})$ such that $y = \langle \boldsymbol{\beta}, \phi(\mathbf{x}) \rangle$. Also, show how to recover the original parameters a_i from the parameters β_j:

   (a) **Example:** $y = a_1 x_1^2 + a_2 \log(a_3 x_2)$.

   Solution: Notice that $y = a_1 x_1^2 + a_2 \log(x_2) + a_2 \log(a_3)$. Let $\phi([x_1, x_2]) = [x_1^2, \log(x_2), 1]$. Set $a_1 = \beta_1$, $a_2 = \beta_2$, $a_3 = e^{\beta_3/a_2}$.

   (b) $y = \begin{cases} a_1 + a_2 x & \text{if } x < 1 \\ a_3 + a_4 x & \text{if } x \geq 1 \end{cases}$

   (c) $y = (1 + a_1 x_1) e^{-x_2 + a_2}$.

   (d) $y = (a_1 x_1 + a_2 x_2) e^{-x_1 - x_2}$.