CS 6140: Machine Learning — Fall 2021— Paul Hand

Midterm 1 Study Guide and Practice Problems
Due: Never.

Names: [Put Your Name(s) Here]

This document contains practice problems for Midterm 1. The midterm will only have 5 problems. The midterm will cover material up through and including the bias-variance tradeoff, but not including ridge regression. Skills that may be helpful for successful performance on the midterm include:

1. Setting up and solving a linear regression problem with features that are nonlinear functions of the model's input.
2. Writing down the optimization problem for least squares linear regression using matrix-vector notation
3. Familiarity with matrix multiplication, in particular when multiplying by diagonal matrices
4. Evaluating the true positive rate, false positive rate, precision, and recall of a predictor for binary classification
5. Setting up a logistic regression problem and writing down the appropriate function that is being minimized
6. Computing the mean, expected value, and variance of uniform random variables
7. Explaining causes and remedies for overfitting and underfitting of ML models

Question 1.

Consider the following training data.

x_1	x_2	y
0	0	0
0	1	1.5
1	0	2
1	1	2.5

Suppose the data comes from a model $y = \theta_0 + \theta_1 x_1 + \theta_2 x_2$+noise for unknown constants $\theta_0, \theta_1, \theta_2$. Use least squares linear regression to find an estimate of $\theta_0, \theta_1, \theta_2$.

Response:

Question 2.

Consider the following training data:

x	y
1	3
2	1
3	0.5

Suppose the data comes from a model $y = cx^{\beta}$ + noise, for unknown constants c and β. Use least squares linear regression to find an estimate of c and β.

Response:

Question 3.

(a) Let $\theta^* \in \mathbb{R}^d$, and let $f(\theta) = \frac{1}{2}\|\theta - \theta^*\|^2$. Show that the Hessian of f is the identity matrix.

Response:

(b) Let $X \in \mathbb{R}^{n \times d}$ and $y \in \mathbb{R}^n$. For $\theta \in \mathbb{R}^d$, let $g(\theta) = \frac{1}{2}\|X\theta - y\|^2$. Show that the Hessian of g is $X^t X$.

Response:

Question 4.

Consider a binary classification problem whose features are in $\mathbb{R}^2$. Suppose the predictor learned by logistic regression is $\sigma(\theta_0 + \theta_1 x_1 + \theta_2 x_2)$, where $\theta_0 = 4, \theta_1 = -1, \theta_2 = 0$. Find and plot curve along which P(class 1) = 1/2 and the curve along which P(class 1) = 0.95.

Response:

Question 5.

Consider a 3-class classification problem. You have trained a predictor whose input is $x \in \mathbb{R}^2$ and whose output is $\text{softmax}(x_1 + x_2 - 1, 2x_1 + 3, x_2)$. Find and sketch the three regions in $\mathbb{R}^2$ that gets classified as class 1, 2, and 3.

Response:

Question 6.

Suppose $x \sim$Uniform$([-1,1])$ and $y = x + \varepsilon$, where $\varepsilon \sim$Uniform$([-\gamma,\gamma])$ for some $\gamma > 0$. Consider a predictor given by $f_\theta(x) = \theta_1 + \theta_2 x$, where $\theta \in \mathbb{R}^2$. Evaluate the risk of f_θ with respect to the square loss. Your answer should be a deterministic expression only depending on θ_1, θ_2, and γ.

Response:

Question 7.

You are training a logistic regression model and you notice that it does not perform well on test data.

- Could the poor performance be due to underfitting? Explain.
- Could the poor performance be due to overfitting? Explain.