

# CPSC 340 Assignment 3 (due October 19)

## Linear Regression

### 0 Unofficial Course Evaluation

To help improve the course as we go along, or to suggest of how things could be done differently, please fill out the survey here:

<https://survey.ubc.ca/surveys/37-7d0090012c11ea5c07f0bca610f/cs340-asst3-winter16>

### 1 Vectors, Matrices, and Quadratic Functions

The first part of this question makes you review basic operations on vectors and matrices. If you are rusty on basic vector and matrix operations, see the notes on linear algebra on the course webpage. The second part of the question gives you practice taking the gradient of linear and quadratic functions, and the third part gives you practice finding the minimizer of quadratic functions.

#### 1.1 Basic Operations

Using the definitions below,

$$\alpha = 5, \quad x = \begin{bmatrix} 2 \\ 3 \end{bmatrix}, \quad y = \begin{bmatrix} 1 \\ 4 \end{bmatrix}, \quad z = \begin{bmatrix} 2 \\ 0 \\ 1 \end{bmatrix}, \quad A = \begin{bmatrix} 1 & 2 \\ 2 & 3 \\ 3 & 2 \end{bmatrix},$$

evaluate the following expressions (show your work, but you may use answers from previous parts to simplify calculations):

1.  $x^T x$ .
2.  $\|x\|^2$ .
3.  $x^T(x + \alpha y)$ .
4.  $Ax$
5.  $z^T Ax$
6.  $A^T A$ .

If  $\{\alpha, \beta\}$  are scalars,  $\{x, y, z\}$  are real-valued column-vectors of length  $d$ , and  $\{A, B, C\}$  are real-valued  $d$  by  $d$  matrices, state whether each of the below statements is true or false in general and give a short explanation.

7.  $yy^T y = \|y\|^2 y$ .

8.  $x^T A^T (Ay + Az) = x^T A^T Ay + z^T A^T Ax$ .
9.  $A^T (B + C) = BA^T + CA^T$ .
10.  $x^T (B + C) = Bx + Cx$ .
11.  $(A + BC)^T = A^T + C^T B^T$ .
12.  $(x - y)^T (x - y) = \|x\|^2 - x^T y + \|y\|^2$ .
13.  $(x - y)^T (x + y) = \|x\|^2 - \|y\|^2$ .

Hint: check the dimensions of the result, and remember that matrix multiplication is generally not commutative.

## 1.2 Converting to Matrix/Vector/Norm Notation

Using our standard supervised learning notation  $(X, y, w)$  express the following functions in terms of vectors, matrices, and norms (there should be no summations or maximums).

1.  $\sum_{i=1}^n |w^T x_i - y_i|$ .
2.  $\max_{i \in \{1, 2, \dots, n\}} |w^T x_i - y_i| + \frac{\lambda}{2} \sum_{j=1}^n w_j^2$ .
3.  $\sum_{i=1}^n z_i (w^T x_i - y_i)^2 + \lambda \sum_{j=1}^d |w_j|$ .

You can use  $Z$  to denote a diagonal matrix that has the values  $z_i$  along the diagonal.

## 1.3 Minimizing Quadratic Functions as Linear Systems

Write finding a minimizer  $w$  of the functions below as a system of linear equations (using vector/matrix notation and simplifying as much as possible). Note that all the functions below are convex so finding a  $w$  with  $\nabla f(w) = 0$  is sufficient to minimize the functions (but show your work in getting to this point).

1.  $f(w) = \frac{1}{2} \|w - v\|^2$ .
2.  $f(w) = \frac{1}{2} \|w\|^2 + w^T X^T y$ .
3.  $f(w) = \frac{1}{2} \|Xw - y\|^2 + \frac{1}{2} w^T \Lambda w$ .
4.  $f(w) = \frac{1}{2} \sum_{i=1}^n z_i (w^T x_i - y_i)^2$ .

Above we assume that  $v$  is a  $d$  by 1 vector, and  $\Lambda$  is a  $d$  by  $d$  diagonal matrix with positive entries along the diagonal.

Hint: Once you convert to vector/matrix notation, you can use the results from class to quickly compute these quantities term-wise. As a sanity check for your derivation, make sure that your results have the right dimensions.

## 2 Linear Regression and Nonlinear Bases

In class we discuss fitting a linear regression model by minimizing the squared error. This classic model is the simplest version of many of the more complicated models we will discuss in the course. However, it typically performs very poorly in practice. One of the reasons it performs poorly is that it assumes that the target  $y_i$  is a linear function of the features  $x_i$  with an intercept of zero. This drawback can be addressed by adding a bias variable and using nonlinear bases (although nonlinear bases may increase to over-fitting).


In this question, you will start with a data set where least squares performs poorly. You will then explore how adding a bias variable and using nonlinear (polynomial) bases can drastically improve the performance. You will also explore how the complexity of a basis affects both the training error and the test error. In the final part of the question, it will be up to you to design a basis with better performance than polynomial bases. If you are not familiar with Matlab, to get you started please see the notes on Matlab commands on the course webpage.

## 2.1 Adding a Bias Variable

If you run the script *example\_basis*, it will:

1. Load a one-dimensional regression dataset.
2. Fit a least-squares linear regression model.
3. Report the training error.
4. Report the test error (on a dataset not used for training).
5. Draw a figure showing the training data and what the linear model looks like.

Unfortunately, this is an awful model of the data. The average squared training error on the data set is over 28000 (as is the test error), and the figure produced by the demo confirms that the predictions are usually nowhere near the training data:



The y-intercept of this data is clearly not zero (it looks like it's closer to 200), so we should expect to improve performance by adding a *bias* variable, so that our model is

$$y_i = w^T x_i + w_0.$$

instead of

$$y_i = w^T x_i.$$

Write a new function, *leastSquaresBias*, that has the same input/model/predict format as the *leastSquares* function, but that adds a *bias* variable  $w_0$ . Hand in your new function, the updated plot, and the updated training/test error.

Hint: recall that adding a bias  $w_0$  is equivalent to adding a column of ones to the matrix  $X$ . Don't forget that you need to do the same transformation in the *predict* function.

## 2.2 Polynomial Basis

Adding a bias variable improves the prediction substantially, but the model is still problematic because the target seems to be a *non-linear* function of the input. Write a new function, `leastSquaresBasis(x,y,p)`, that takes a data vector  $x$  (i.e., assuming we only have one feature) and the polynomial order  $p$ . The function should perform a least squares fit based on a matrix  $Z$  where each of its rows contains the values  $(x_i)^j$  for  $j = 0$  up to  $p$ . E.g., `leastSquaresBasis(x,y,3)` should form the matrix

$$Z = \begin{bmatrix} 1 & x_1 & (x_1)^2 & (x_1)^3 \\ 1 & x_2 & (x_2)^2 & (x_2)^3 \\ \vdots & & & \\ 1 & x_n & (x_n)^2 & (x_n)^3 \end{bmatrix},$$

and fit a least squares model based on it. [Hand in the new function, and report the training and test error for  \$p = 0\$  through  \$p = 10\$ . Explain the effect of  \$p\$  on the training error and on the test error.](#)

Hints: To keep the code simple and reduce the chance of having errors, you may want to write a new function `polyBasis` that you can use for transforming both the training and testing data.

## 2.3 Manual Search for Optimal Basis

Polynomials are a flexible class of functions, but there is structure in this data that is not well-modelled by polynomials. Try to find a nonlinear basis that gives the best performance on this dataset in terms of test error. [Report the basis that you use and the training/test score that you achieve.](#)

Hint: the data seems to have periodic behaviour, and it's possible to obtain training and test errors below 60.

# 3 Non-Parametric Bases and Cross-Validation

Unfortunately, in practice we often don't know what basis to use. However, if we have enough data then we can make up for this by using a basis that is flexible enough to model any reasonable function. These may perform poorly if we don't have much data, but can perform almost as well as the optimal basis as the size of the dataset grows. In this question you will explore using Gaussian radial basis functions (RBFs), which have this property. These RBFs depend on a parameter  $\sigma$ , which (like  $p$  in the polynomial basis) can be chosen using a validation set. In this question, you will also see how cross-validation allows you to tune parameters of the model on a larger dataset than a strict training/validation split would allow.

## 3.1 Proper Training and Validation Sets

If you run the demo `example_rbf`, it will load a dataset and split the training examples into a "train" and a "validation" set. It will then search for the best value of  $\sigma$  for the RBF basis (it also uses regularization since  $Z^T Z$  tends to be very close to singular). Once it has the "best" value of  $\sigma$ , it re-trains on the entire dataset and reports the training error on the full training set as well as the error on the test set.

Unfortunately, there is a problem with this function. Because of this problem, the RBF basis doesn't perform much better than a linear model. [What is the problem with this training/validation/testing procedure? How can you fix this problem?](#)

Hint: you may want to plot the models that are considered during the search for  $\sigma$ .

### 3.2 Cross-Validation

Using the standard solution to the above problem, a strange behaviour appears: if you run the script more than once it might choose different values of  $\sigma$ . It rarely performs too badly, but it's clear that the randomization has an effect on the value of  $\sigma$  that we choose. This variability would be reduced if we had a larger “train” and “validation” set, and one way to simulate this is with *cross-validation*. [Modify the training/validation procedure to use 10-fold cross-validation to select  \$\sigma\$ , and hand in your code. What value of  \$\sigma\$  does this procedure typically select?](#)

### 3.3 Cost of Non-Parametric Bases

When dealing with larger datasets, an important issue is the dependence of the computational cost on the number of training examples  $n$  and the number of features  $d$ . [What is the cost in big-O notation of training the model on  \$n\$  training examples with  \$d\$  features under \(a\) the linear basis and under \(b\) Gaussian RBFs \(for a fixed  \$\sigma\$ \)? What is the cost of classifying  \$t\$  new examples under each of these two bases? When are RBFs cheaper to train? When are RBFs cheaper to test?](#)

### 3.4 Non-Parametric Bases with Uneven Data


There are reasons why this dataset is particularly well-suited to Gaussian RBFs:

1. The period of the oscillations stays constant.
2. We have evenly sampled the training data across its domain.

If either of these assumptions are violated, the performance with our Gaussian RBFs might be much worse. [Consider a scenario where either 1 or 2 above is violated, and describe a way that you could address this problem.](#)

## 4 Robust Regression and Gradient Descent

The script *example\_outliers* loads a one-dimensional regression dataset that has a non-trivial number of ‘outlier’ data points. These points do not fit the general trend of the rest of the data, and pull the least squares model away from the main downward trend that most data points exhibit:



## 4.1 Weighted Least Squares in One Dimension

One of the most common variations on least squares is *weighted* least squares. In this formulation, we have a weight  $z_i$  for every training example. To fit the model, we minimize the weighted squared error,

$$f(w) = \frac{1}{2} \sum_{i=1}^n z_i (w^T x_i - y_i)^2.$$

In this formulation, the model focuses on making the error small for examples  $i$  where  $z_i$  is high. Similarly, if  $z_i$  is low then the model allows a larger error.

Write a model function, `weightedLeastSquares(X,y,z)`, that implements this model (note that Q1.3.4 asks you to show how this formulation can be solved as a linear system). Apply this model to the data containing outliers, setting  $z = 1$  for the first 400 data points and  $z = 0.1$  for the last 100 data points (which are the outliers). [Hand in your function and the updated plot.](#)

## 4.2 Smooth Approximation to the L1-Norm

Unfortunately, we typically do not know the identities of the outliers. In situations where we suspect that there are outliers, but we do not know which examples are outliers, it makes sense to use a loss function that is more robust to outliers. In class, we discussed using the sum of absolute values objective,

$$f(w) = \sum_{i=1}^n |w^T x_i - y_i|.$$

This is less sensitive to outliers than least squares, but it is non-differentiable and harder to optimize. Nevertheless, there are various smooth approximations to the absolute value function that are easy to optimize. One possible approximation is to use the log-sum-exp approximation

$$|r| \approx \log(\exp(r) + \exp(-r)).$$

Using this approximation, we obtain an objective of the form

$$f(w) = \sum_{i=1}^n \log(\exp(w^T x_i - y_i) + \exp(y_i - w^T x_i)).$$

which is smooth but less sensitive to outliers than the squared error. [Derive the gradient  \$\nabla f\$  of this function with respect to  \$w\$ . You should show your work but you do not have to express the final result in matrix notation.](#)

### 4.3 Robust Regression

The function *example\_gradient* is the same as *example\_outlier*, except that it fits the least squares model using a *gradient descent* method. One advantage of this strategy is that it only costs  $O(nd)$  for an iteration of the gradient method, which is faster than forming  $X^T X$  which costs  $O(nd^2)$ . Of course, we need to know the *number* of gradient iterations in order to precisely compare these two strategies, but for now we will assume that the number of gradient iterations is typically often reasonable.

The usual input to a gradient method is a function that, given  $w$ , returns  $f(w)$  and  $\nabla f(w)$ . See *funObj* in the *leastSquaresGradient* function for an example. Note that *leastSquaresGradient* also has a numerical check that the gradient code is approximately correct, since implementing gradients is often error-prone.

A second advantage of gradient-based strategies is that they are able to solve problems that do not have closed-form solutions, such as the formulation from the previous section. The function *robustRegression* has most of the implementation of a gradient-based strategy for fitting the robust regression model under the log-sum-exp approximation. The only part missing is the function and gradient calculation inside the *funObj* code. [Modify this function to implement the objective function and gradient based on the smooth approximation to the absolute value function \(from the previous section\). Hand in your code, as well as the plot obtained using this robust regression approach.](#)