

Asmt 6: Regression

Turn in through Canvas by 2:45pm:
Wednesday, April 1
100 points

Overview

In this assignment you will explore regression techniques on high-dimensional data.

You will use a few data sets for this assignment:

- <http://www.cs.utah.edu/~jeffp/teaching/cs5140/A6/X.csv>
- <http://www.cs.utah.edu/~jeffp/teaching/cs5140/A6/Y.csv>
- <http://www.cs.utah.edu/~jeffp/teaching/cs5140/A6/M.csv>
- <http://www.cs.utah.edu/~jeffp/teaching/cs5140/A6/W.csv>

For python, you can use the following approach to load the data:

```
X = np.loadtxt('X.csv', delimiter=',')
y = np.loadtxt('Y.csv', delimiter=',')
```

As usual, it is recommended that you use LaTeX for this assignment (or similar way to properly typeset math). If you do not, you may lose points if your assignment is difficult to read or hard to follow. Find a sample form in this directory: <http://www.cs.utah.edu/~jeffp/teaching/latex/>

1 Linear Regression & Cross-Validation (100 points)

We will find coefficients α to estimate $X\alpha \approx y$, using the provided datasets X and y . We will compare two approaches *least squares* and *ridge regression*. (e.g., in python as)

Least Squares: `Set alpha = LA.inv(X.T @ X) @ X.T @ y.T`

Ridge Regression: `Set alphas = LA.inv(X.T @ X + s*np.identity(50)) @ X.T @ y.T`

A (30 points): Solve for the coefficients α (or α s) using Least Squares and Ridge Regression with $s \in \{0.2, 0.4, 0.8, 1.0, 1.2, 1.4, 1.6\}$ (i.e. s will take on one of those 7 values each time you try, say obtaining $\alpha_{0.4}$ for $s = 0.4$). For each set of coefficients, report the error in the estimate \hat{y} of y as $\text{norm}(y - X\alpha, 2)$.

B (30 points): Create three row-subsets of X and Y

- $X_1 = X[:66, :]$ and $Y_1 = Y[:66]$
- $X_2 = X[33:, :]$ and $Y_2 = Y[33:]$
- $X_3 = \text{np.vstack}((X[:33, :], X[66:, :]))$ and $Y_3 = \text{np.vstack}((Y[:33], Y[66:]))$

Repeat the above procedure on these subsets and *cross-validate* the solution on the remainder of X and Y . Specifically, learn the coefficients α using, say, X_1 and Y_1 and then measure $\text{np.norm}(Y[66:] - X[66:, :] @ \alpha, 2)$.

C (15 points): Which approach works best (averaging the results from the three subsets): Least Squares, or for which value of s using Ridge Regression?

D (15 points): Use the same 3 test / train splits, taking their average errors, to estimate the average squared error on each predicted data point.

What is problematic about the above estimate, especially for the best performing parameter value s ?

E (10 points): Even circumventing the issue raised in part **D**, what *assumptions* about how the data set (X, Y) is generated are needed in an assessment based on cross-validation?

2 Bonus: Matching Pursuit (5 points)

Consider a linear equation $W = M * S$ where M is a measurement matrix filled with random values $\{-1, 0, +1\}$ (although now that they are there, they are no longer random), and W is the output of the sparse signal S when measured by M .

Use Matching Pursuit (as described in the book as **Algorithm 5.5.1**) to recover the non-zero entries from S . Record the order in which you find each entry and the residual vector after each step.