# 6.867: Homework 1

Essentially all problems in machine learning involve some form of optimization, whether to fit a model to data (estimation) or to select a prediction based on a model (decision). In a few cases we can find the optimum analytically; in many others we do it numerically. In this assignment we will begin to explore numerical methods for some of the basic estimation problems that we have been learning about.

You will find a zip file with some useful code and data in the Resources section of the Piazza course page. You can do these assignments in any computational system that you are used to. We recommend Matlab or the Pylab/Numpy/Scipy/Matplotlib cluster of packages and we'll try to provide help for those two systems. If you use anything else, you're on your own...

You will be turning in a single, readable "paper" (a single PDF file) with your solutions. We will be emulating the process of submitting papers for publication to a conference. We will be using an actual conference review system (Easy Chair) to have these papers peer reviewed (by the other students). This means that your answers have to be readable and understandable to your peers and, where possible, interesting. Note that when explanations are called for, you will need to convince the reviewers that you understand what you're talking about. The course staff will serve as the Program Committee for the conference and make all final decisions. The details of this process will be posted on Piazza.

## Grading process and due dates

- This assignment may be done by pairs of students.
- You submit your paper via the Easy Chair site by 11:59 PM on Wednesday October 3.
- Each student who was an author or co-author on a submission will be assigned 3 papers to review.
- Reviews must be entered on the Easy Chair site by 11:59PM on Thursday, October 11.
- All reviews will be made visible shortly thereafter, and students will have a 2–3 day period in which to enter rebuttals of their reviews into EasyChair if they wish.

## Grading rubric

Your paper must be anonymous (no identifying information should appear in the PDF file). If

## Fall 2012

## MIT 6.867

# it is not, it will automatically receive a 20% deduction, and will be graded by a grumpy staff member.

paper must be no more than 6 pages long in a font no smaller than 10 point. It should include whatever tables, graphs, plots, etc., are necessary to demonstrate your work and conclusions. *It should not include code*.

Each of the four parts of the assignment will be graded on a scale from 0 to 5 (where 0 is failing and 5 is an A) on two aspects:

- **Content:** Did the solution answer the questions posed? Were the answers correct? Were the experiments well-designed or examples well chosen?
- **Clarity:** Were the results written up clearly? Were the plots labeled appropriate and described well? Did the plots support the points in the paper? Did the discussion in the paper illuminate the plots?

As a reviewer, you will be asked to provide a score for each section, and at at least two paragraphs of feedback, per review, explaining things that were done well and things that could have been improved upon.

Your overall score for this assignment will be:

- 80%: The average of all 8 scores on your assignemnt given by all three reviewers.
- 20%: A score for the quality of your reviews. This will be full credit, by default. But we will skim reviews and examine some carefully and may reduce this grade for review commentary that is sloppy or wrong.

The course staff will spot-check submissions and reviews, paying careful attention to cases where there were rebuttals. The staff will act as the program committee and determine a final score. Our overall goals in this process are:

- To motivate you to work seriously on the problems and learn something about the machine learning material in the process
- To engage you in thinking critically and learning from other students' solutions to the problems

We will arrange to give full credit to anyone who submits a serious and careful solution to the problems and who gives evidence of having read carefully the solutions they were assigned and who writes thoughtful reviews of them.

The following questions are the points that your paper should cover in order to receive full credit. Your presentation should roughly follow the order of these questions so that your reviewers can see what you're doing.

#### Fall 2012

## 1 Implement Gradient Descent

The simplest and most commonly referred-to method of optimization is gradient descent (see Bishop section 5.2.4 and Wikipedia). It's worth understanding its behavior on actual problems. There are many much more sophisticated optimization methods available in Matlab and Scipy; you should also benchmark one of those.

- Implement a basic gradient descent procedure to minimize scalar functions of a vector argument. Write it generically, so that you can easily specify the objective function and the function to compute the gradient. You should be able to specify the initial guess, the step size and the convergence criterion (a threshold such that the algorithm terminates when the objective values on two successive steps is below this value).
- 2. Test your gradient descent procedure on some functions whose optimal value you know, e.g. a quadratic bowl or the (negative of) Gaussian. Make sure that you try functions of more than one variable and that you try at least one convex function and one very non-convex function with multiple minima. Discss the effect of the choice of starting guess, the step size, and the convergence criterion on the resulting solution.
- 3. Write code to approximate the gradient of a function numerically at a given point using central differences (see the "Finite difference" article in Wikipedia). Verify its behavior on the functions you used in the question above by comparing the analytic and numerical gradients at various points.
- 4. Compare the behavior of your gradient descent procedure with one the more sophisticated optimizers available in Matlab (e.g. fminunc) or scipy.optimize (e.g. fmin\_bfgs). A good metric for comparison is the number of function evaluations required to reach convergence. You should instrument your code to keep track of function calls and call the Matlab or Scipy optimizers so that they print this.

## 2 Linear Basis Function Regression

Let's consider the linear combination of basis function class of regression models. We know how to get analytic solutions for the maximum likelihood weight vector (Bishop 3.15). We'll use this problem to "benchmark" the gradient descent solution.

We have provided you with a text file (curvefitting.txt) that has the 10 data points that were used to generate the plots in Bishop Figure 1.4. We have also given you code to read this data and some code illustrating how to generate plots, both in Python and Matlab.

1. Write a procedure for computing the maximum likelihood weight vector given (a) an array of 1-dimensional data points X, (b) a vector of Y values, and (c) the value of M, the maximum

## Fall 2012

order of a simple polynomial basis (as in Bishop)  $\phi_1(x) = x, \ldots, \phi_M(x) = x^M$ . Test your solution by replicating the plots in Bishop Figure 1.4 and the weight values in Table 1.1. You should be able to get very close agreement.

- 2. Now, write functions to compute the sum of squares error (SSE) (and its derivative) for a data set and a hypothesis, specified by the list of M polynomial basis functions and a weight vector. Verify your gradient using the numerical derivative code.
- 3. Use gradient descent on the SSE function to replicate the graphs in Bishop. Describe your experience with initial guesses, step sizes and convergence thresholds. Compare to using one of the more sophisticated optimizers. Explain your results in terms of the properties of the function being optimized and the properties of the algorithms.

## 3 Ridge Regression

- 1. Implement ridge regression both analytically and via gradient descent. Experiment with different values of  $\lambda$  on the simple data from Bishop's Figure 1.4, for various values of M. Describe your observations.
- 2. We have given you three additional data sets: there are two training data sets and one validation set. In general, we want to use the training data to optimize parameters and then use the performance of those parameters on the validation data to choose among models (e.g. values of M and values of λ), this is called "model selection". Show some values of M and λ, and show the effect on the fit in the test and validation data. Which values work best for each of the training data sets? Explain.

# 4 Generalizations

So far, we have been working with problems where the analytic solution is very easy and thus gradient descent might be used only for computational reasons, e.g. big data sets. Let's now consider some problems where analytic solutions don't exist.

1. We have seen that OLS regression solves the "least squares (LS)" fitting problem; it minimizes the sum of the squares of the prediction errors (the L<sub>2</sub> norm). Alternatively, we could minimize the sum of the absolute values of the errors (the L<sub>1</sub> norm). This is called "least absolute deviations (LAD)". Just as using least squares is equivalent to assuming that the outputs (Pr(Y | X)) have a Gaussian distribution, LAD is equivalent to assuming that the outputs have a Laplace distribution. Use gradient descent to find the weights that minimize LAD, using the quadratic regularization term as in ridge regression. Repeat your experiments from the ridge regression question using this criterion; illustrate and summarize your salient results.

## MIT 6.867

## Fall 2012

- 2. In ridge regression we used a quadratic regularizer (the sum of squares of the weights ( $L_2$  norm)). Bishop's equation (3.29) shows a generalization to other exponents. In particular, using q = 1 (the sum of the absolute values of the weights ( $L_1$  norm)) is called the "lasso". Use gradient descent to find the best value of the weights using the quadratic norm on errors and absolute value norm on the magnitude of the weights. Repeat your experiments from the ridge regression question using this new regularizer; illustrate and summarize your salient results.
- 3. Explain clearly the difference in the approaches of the previous two questions. When might you want least absolute deviations? When might you want the lasso? What, if anything, do these two approaches have to do with using absolute value loss L(p, a) = |p-a| in a prediction problem?