


CS451 - Assignment 6
A Gradient Descent Into Madness
Due: Sunday, October 27 by 11:59pm



For this assignment we're going to try out a few gradient descent variants (including the SVM). The amount of code you'll need to write for this will be pretty minimal.

1 High-level requirements

- Implement the gradient descent algorithm
- Your approach will need to support two loss functions: hinge loss and exponential loss
- Your approach will need to support three regularization type: none, L1 and L2
- You will provide a short writeup that includes:
 - an argument that your implementation is correct. This can be of the form of examining text snippets with justification and/or showing experimental results.
 - at least one experimental result

2 Implementation Requirements

To get you started I have included some starter code at:

<http://www.cs.middlebury.edu/~dkauchak/classes/cs451/assignments/assign6/assign6-starter.tar.gz>

This is basically the same starter code as assignment 5 except I have add the `KNNClassifier` (so that everything compiles file) and I have included a `GradientDescentClassifier`. This file is simply a copy of the `PerceptronClassifier` file with the addition of a few constants. This should serve as your starting point.

Implement a gradient descent classifier. Your implementation must:

- support a zero parameter constructor. By default you should use exponential loss, no regularization, $\lambda = 1$ and $\eta = 1$
- include a method `setLoss` that takes an int and selects the loss function to use
- include a method `setRegularization` that takes an int and selects the regularization method to use
- include a method `setLambda` that takes a double and sets the lambda to use
- include a method `setEta` that takes a double and sets the eta to use (Note, in practice we'd use some sort of schedule of etas, e.g. one for each iteration, however, to keep it simple we'll just have a constant learning rate.)
- support training and testing based on these parameters

3 Writeup

In addition to your code, include a writeup (in some reasonable file format) that includes two sections: **Algorithm Correctness** and **Experimentation**. Because of the small amount of coding required for this assignment, the writeup will be a non-trivial part of your grade, so make sure to devote sufficient time to this.

Algorithm Correctness

Include a short (1 page) justification of why your implementation of the hinge loss with L2 regularization is correct. This could include snippets of code, along with an explanation, and/or short experimental results, e.g. incremental output of a small problem. You will be graded based on how convincing and thorough your argument is.

Experimentation

Run one experiment that highlights something interesting about the algorithm, include the resulting data, and include a 2-3 sentence explanation/analysis. For example, you could investigate the optimal λ/η for one of the variants on one of the data sets, or you could investigate how the different loss/regularization approaches work (though pick some reasonable lambda). Plan on spending around an hour playing around with this. You will be graded based on the quality of your experimental setup and your presentation.

4 Extra Credit

For those who would like to experiment (and push themselves) a bit more (and of course, get a bit of extra credit) you can try out some of these extra credit options. If you try out these options, include an extra file called `extra.txt` that describes what extra credit you did.

- Allow for an η schedule rather than just a single constant. I'll let you be creative about the best way to do this, but please just create a new copy of your classifier code rather than trying to alter the existing code.
- Add additional regularization or loss functions. If you do this, add appropriate constants for selecting these.
- Do additional experimentation and include it in the write-up. The amount earned will be based on how much effort you put in and how creative the experiment is.

5 When You're Done

Make sure that your code compiles, that your files are named as specified and that you have followed the specifications exactly (i.e. method names, number of parameters, etc.).

Create a directory with your last name, followed by the assignment number, for example, for this assignment mine would be `kauchak6`. If you worked with a partner, put both last names.

Inside this directory, create a `code` directory and copy all of your code into this directory, maintaining the package structure.

Finally, also include your `writeup` file.

`tar` then `gzip` this folder and submit that file on the submission page on the course web page.

Commenting and code style

Your code should be commented appropriately (though you don't need to go overboard). The most important things:

- Your name (or names) and the assignment number should be at the top of each file
- Each class and method should have an appropriate JavaDoc
- If anything is complicated, it should include some comments.

There are many possible ways to approach this problem, which makes code style and comments very important here so that I can understand what you did. For this reason, you will lose points for poorly commented or poorly organized code.