COMPSCI 271 - Final Project

Due date: Dec. 6

1 Project Scope

For this project you will evaluate a machine learning technique theoretically, or some benchmark or real world data. Acceptable projects include the following:

1. Proving a new theorem about an existing learning method (ambitious)
2. Proposing a modification to an existing learning method and evaluating this modification against the standard method on several benchmark data sets.
3. Developing a hypothesis about how different learning methods compare and evaluating this hypothesis using these learning methods on many different benchmark data sets.
4. Proposing an application of machine learning to an interesting problem that you care about, selecting a method, doing your best to make it work, and discussing the results.

This is only a partial list. Other ideas are also possible if you clear them with me first.

2 Project Proposal

You should turn in a proposal for your project to me (email is fine) *no later than October 16*. This should be a less than one page description of what you plan to do. You are, of course, free to turn in a proposal earlier.

I will give feedback on your proposal within a few days of submission.

3 Project Milestone

No later than November 15, you should turn in an approximately page description of the progress you have made since your proposal. This is your opportunity to document your initial efforts and identify areas where you think you might need help.

This will not be graded, but if you miss this opportunity I will be less sympathetic about last minute problems that arise. (NB: They almost always do.)

4 Project Write Up

Turn in a document describing your efforts and a link to any code you've produced. As a rough guideline, you should aim for a writeup of 10-15 single column, double spaced pages. However, I won't be counting pages and this is only a very rough guideline. You shouldn't pad your writeup with worthless text to reach this size. If you can describe what you've done clearly and concicsely in much less space than this, that's great.

I'll expect your writeup to address (at least) the following issues:

1. What problem does your project address? If you have chosen an application area, please remember that I may not be an expert in the application you have chosen, so be sure to describe the application area clearly.
2. What methods did you use to address the problem?
3. What is the reason you picked the methods you picked? Can you justify theoretically or empirically that this was the best choice?
4. How did the machine learning techniques that you applied perform?
5. How did you validate your results, i.e., what were the training and test sets?
6. What difficulties did you encounter and how did you try to overcome them?
7. What would be the next step if you were to extend this project?
8. What did you learn from this?

Remember to include a complete bibliography.