CS281 Final Project Requirements

Objective

The objective of this final project is to explore new research in machine learning. The ideal outcome would be a paper that could be submitted to one of the top machine learning conferences, such as NIPS, ICML or AISTATS.

A strong paper along these lines is one that develops a new or improved algorithm (runs faster, scales better, makes better predictions) that learns to generalize from experience, broadly defined. Such a paper would demonstrate theoretical and/or empirical improvements over the state of the art. One type of paper along these lines would introduce a new probabilistic model that captures important characteristics of data that had previously been unexplored. Another type of paper might propose a new algorithm for performing inference and learning in existing models. A third type of paper might consider models and learning algorithms in important settings of constrained resources, such as with limited memory, real-time performance requirements, or streaming data.

Of course, such papers require innovative ideas about machine learning that may be difficult to come by in a single semester. It is helpful, therefore, to initially focus on a specific problem domain that you find important and exciting. Consider what the fundamental task is that needs to be solved and think about how it might map onto, e.g., regression or clustering. Catalogue the types of data that are available and consider how these might be exploited. What are the features that will help your algorithm make decisions or predictions? Don't be afraid to make assumptions that help establish an abstraction. Prefer abstractions and assumptions that may generalize beyond the immediate task at hand.

The next step is critical: define quantitative metrics for success on held-out test data. Classification accuracy? Area under the ROC curve? Predictive log probability? Rand index? If you are focused on resource constraints, then perhaps your metrics will be curves that measure prediction as a function of, e.g., memory usage or CPU time. If possible, consider several possible metrics so that it will be easy for the eventual readers of your paper to map your algorithm onto their priorities.

Once you have laid out the task, data and metrics, try to apply the dumbest, simplest possible algorithm first. *Do not* immediately try your new fancy idea. If you have a classification problem, try logistic regression first. Try things from the literature that seem like they would be applicable. Establish baselines for comparison that are honest attempts at doing well on the problem. After you've done this, you'll have a much better idea of what your algorithm is capable of. You may also learn that logistic regression is difficult to beat. What does that mean? Maybe you need to focus on extracting features rather than a fancy classifier. Applying simple things first will help you understand where the frontier of the problem lies and help you determine whether your abstraction actually provides the information that you require for the task.

In the end, it may be difficult to make a methodological contribution. If you have taken a problem-driven approach, however, then you have still done useful research by improving our understanding of how machine learning algorithms behave when applied to new problems. Although generally viewed as weaker papers than those that make methodological contributions, if timely and executed thoughtfully, they can be very impactful and widely read. For two strong examples of papers that are important, but fundamentally about empirical evaluation, I suggest looking at Jarrett et al. (2009) and Coates et al. (2011).

Collaboration

You may work on this project alone or with a partner, your choice. Larger groups will be considered with permission. With a partner, you will be able to (and indeed will be expected to) tackle a more difficult problem. Realize, however, that partners' grades will be extremely highly correlated as it will be impossible to disambiguate your contributions to the project.

Deliverables

There are four separate components of the final project. Each of these is due as a PDF file uploaded to the iSites dropbox at `http://isites.harvard.edu/icb/icb.do?keyword=k98807` by 23:59 of the day specified. The proposal and status report may be submitted up to a week late with a 50% penalty. There will be no extensions given for the poster and final report.

Project Proposal (25 October 2013) Write a two-page document that describes the plan for your project. This should clearly state what problem you are trying to solve. If you have developed a new model, explain what models this work will build on and how it resolves deficiencies. If it is a new algorithm for inference, explain the regimes for which you think it well be well-suited. If you are developing a new theoretical contribution, discuss the theorems you will prove. For problem-driven papers, discuss the data and the unique challenges that make this interesting. Identify relevant work and algorithms you intend to implement as baselines. This does not need to be a comprehensive document and I expect that it will be speculative. Your focus should be on identifying the questions you wish to answer about your data or your method and specifying clearly what "success" will mean. This proposal represents 10% of your total grade in the course.

Abstract and Status Report (22 November 2013) Write a one-page document that contains a paragraph that is a draft of the abstract for the final paper. Use the remainder of the page to describe the status of your project. What have you proved? What baselines have you established? Have there been unexpected results, good or bad? This page represents 5% of your overall course grade.

Poster Presentation (Tentatively 5 December 2013) We will have a class poster session, where you will present a conference-style poster. SEAS will pay for the poster printing (more details about this to come). You will also submit the poster as a PDF. The poster represents 5% of your overall course grade.

Final Report (11 December 2013) Using the NIPS conference paper format (available at `http://nips.cc`), write a paper of up to ten pages. This paper should have a typical conference style, with abstract, introduction, etc. You should clearly state what problem you are trying to solve, introduce and explain your approach, and review the relevant literature. It should explain in detail the experiments that were run, show their results and discuss conclusions that can be drawn. This report represents 20% of your overall grade.

References

Adam Coates, Honglak Lee, and Andrew Ng. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, 2011.

Kevin Jarrett, Koray Kavukcuoglu, Marc'Aurelio Ranzato, and Yann LeCun. What is the best multi-stage architecture for object recognition? In *Proc. International Conference on Computer Vision (ICCV'09)*. IEEE, 2009.