CS260: Machine Learning Theory
Final Project Guidelines
Fall 2011

1 Project Overview

All students enrolled in CS260 are required to complete a class project. This project is intended to provide you with a chance to explore recent research on a particular subtopic of machine learning theory in more depth on your own. The project is fairly open-ended. You may choose to explore any topic that excites you, as long as your project touches on the theoretical foundations of learning.

Projects must be completed in groups of **two or three** students. There are **no exceptions** to this rule. Some of the papers you read will involve ideas and techniques that are new to you, and it can be extremely useful to discuss them with your partners. Plus, the research process is more fun when you get to share your ideas and discoveries with friends! The same criteria will be used to evaluate groups of two and groups of three; in other words, you will not be expected to do more if there are three members in your group as opposed to two. Collaboration is strongly encouraged!

2 Project Types

There are three primary styles of projects that I would recommend for most students: a literature synthesis project, an empirical project, and a theoretical research project. These are described below. However, you are welcome to pursue projects that don't fit neatly into one of these categories. If you have such a project in mind, make an appointment to meet with me to discuss whether or not it sounds appropriate for the class.

Please keep in mind when selecting your project that this is a class on learning *theory*. Your project **must involve a theoretical component**. Simply implementing a machine learning algorithm (whether or not it is novel) and measuring its accuracy on data is **not** a valid project for this class. However, this can certainly be a valid *component* of a bigger project.

- **Literature Synthesis:** If you are interested in learning more about current research on learning theory, but not ready to jump into a research project just yet, you may choose to prepare a literature synthesis (i.e., survey paper) for your final project. This synthesis may be on any subtopic of learning theory, such as learning bounds for domain adaptation, privacy-preserving machine learning, or machine learning for finance. These and other suggestions are listed on the course website, along with some links to relevant papers to get you started. Your survey should not just summarize existing work. You should strive to provide insight about this work, such as connections between different papers, assumptions that you found unrealistic or unmotivated, or interesting open problems that could be studied. Your grade will be based on both the clarity of your explanations and the insights that you provide. It should be clear that you have given the topic some real thought, and are not simply restating results that appear in the papers you survey.
- **Empirical Project with a Comparison to the Theory:** If you have an idea for an exciting domain in which machine learning techniques could be applied, you may wish to do an empirical project. For a

typical empirical project, you should plan to implement a selection of the algorithms covered in class (e.g., the Perceptron, Weighted Majority, Adaboost, or SVMs) or minor variations of these algorithms, and compare their performance on a real-world data set of your choice. (Be sure to separate your data into training and test sets to get an accurate measure of generalization error. You might want to take a look at some examples of experimental results sections from machine learning conference papers to get a feel for other conventions.)

In addition to empirical results, your report should include a detailed technical discussion of the reasons *why* one algorithm may have outperformed others, making reference to theoretical results discussed in class or in the literature. For example, you may choose to compare the performance of each algorithm to the performance that would be predicted by the theory. If the theoretical predictions are inaccurate, discuss why this might be the case. For example, were there assumptions that we had to make for the analysis that were violated by your data set?

It can sometimes be enlightening to generate one or more "synthetic" data sets that satisfy nice properties (e.g., by choosing data points and labels at random from a particular distribution) and compare the performance of the algorithms on these data sets too. This will allow you to discuss what properties of the data lead to better or worse performance for various algorithms.

If you like, you could additionally try implementing algorithms that we did not discuss in class, or algorithms of your own. Just remember that your paper **must** have a theoretical component.

- **Theoretical Research Project:** Finally, you have the option of attempting to generate new theoretical results of your own. There are (at least) two good options here. First, you could choose a known open problem related to models that we have examined in class or other models that have been studied in the learning theory community and write about your attempts to solve this problem. It's ok if you don't end up solving the problem, as long as you make a reasonable attempt, clearly explain what you tried, and explain why what you tried didn't work.

  Alternately, you could design your own model of learning, for example by extending one of the models that we discussed in class. If you choose this option, you should provide a clear motivation for your new model and explain your design choices. (What natural learning problem does the model represent? What assumptions are you making and why? What phenomena does your model capture that existing models do not?) You should also attempt to provide some preliminary results for your model, such as classes that are learnable or algorithms that yield provably low error. Again, it is ok if you make a reasonable effort but are unable to prove the results that you hoped for. In this case you should explain why the techniques you tried failed, and what this says about the model.

  If you choose to do a theoretical research project, you must do enough of a literature search to be reasonably certain that your ideas are novel. I will expect to see a brief section in your report covering related work.

  Theoretical research projects can be really hard! You are certainly not expected to prove groundbreaking theorems in a couple of weeks. Always start by trying to prove something simple, and go on from there. Make an appointment to meet with me if you get stuck.

  Sometimes high quality research projects started for a class can lead to conference paper submissions or even thesis topics. There are quite a few potentially relevant conferences with deadlines in the spring, including COLT, ICML, and UAI. The deadline for NIPS is typically in early June. (The websites for these conferences are also great places to start looking for project ideas!)

3 Important Dates and Milestones

There are four deliverables associated with the class project. Each of these is described below. All submissions are one per group.

- **Project Proposals:** Proposals should be submitted by email (`jenn@cs.ucla.edu`) before the start of class (2pm) on Monday, October 31. Submit your proposal in plain text, right in the body of the email. Be sure to copy all group members on the email so that I can send feedback to everyone.

  Your project proposal should contain a preliminary project title, a list of team members, and a brief (approximately 3 paragraphs) summary of what you plan to do. It should also include a preliminary list of the primary references that you plan to use. (Hint: Claiming that you don't need any outside references is probably not a good idea!)

  Remember that if you propose a literature synthesis, you will be expected to provide some new insight into the area that you will survey. If you propose an empirical project, you should describe in your proposal how you will connect your empirical results to theory.

  The proposal deadline is purposely set late so that you will have lots of time to consider different options, but you are free to submit your proposal earlier if you're ready to get started and want some feedback. This could be especially useful if you plan to propose a research project that will take more time to complete.

- **Progress Reports:** A brief progress report should be submitted by email (`jenn@cs.ucla.edu`) before the start of class (2pm) on Wednesday, November 16. This is roughly the midpoint between the date that you submit your initial proposal and the date that the final reports are due. You may submit your progress report either in plain text in your email or as a pdf attachment. Please limit it to one page.

  Your progress report should include a **brief** description of the work that you have already completed, and a summary of the next steps that you plan to take. The purpose of this report is to help you make sure that you are on track for finishing your project in time. By this point, you should have read some relevant papers and made some initial progress, and should have a clear idea of what to do next.

- **Project Presentations:** Each group will be responsible for a short in-class presentation describing the idea behind your project as well as your results. All presentations will take place on Monday, November 28 and Wednesday, November 30. Each group will be given approximately 10 minutes for their presentation, including time for questions. The exact amount of time will depend on the number of groups.

  Project presentations will be scored by both me and your peers in the class. Scores will be based on the clarity of the presentation, the structure of the presentation (we will talk about this more in class), the technical merit of the project, and the relevance of the project to the course. All students enrolled in CS260 are required to attend the presentations and participate in the peer scoring.

  You are not required to use slides for your presentation, but if you would like to, please email a pdf version of your slides to me by 10am on the day of your presentation.

- **Final Report:** The primary component of your project is the final report, which will be due at 5pm on Monday, December 5. You are welcome to submit your report earlier than this, but **no reports will be accepted after the deadline**. Final reports must be submitted as a hard copy in person. You can

either drop off your report with me in 4532H Boelter Hall, or if I am not around, leave it with Edna Todd in 4532N. Please make sure you give your report directly to either me or Edna; **do not** just leave it under a door.

Final reports must be no more than **six pages** including any figures and references, must be in 11 point or larger font, and must use single column format. As long as you follow these requirements, you are free to use any format that you like. If you received help from anyone outside your group (including your advisor, another faculty member, or other students in the class), you must acknowledge their contributions appropriately in the report.

Your grade will be based in part on the clarity of your report, so please make sure your final report is written clearly! You may wish to show a draft of your report to a student from a different project group in order to get feedback about the presentation of your results. This is strongly encouraged! The ability to communicate your results clearly is a very important part of the research process.

Your report will also be evaluated based on the technical quality of your work. This means that techniques you use should be reasonable, stated results should be accurate, proofs should be correct and complete, and any gaps in your arguments should be noted and explained.

Keep the following general criteria in mind not only when writing your report, but when picking your project topic:

- Literature synthesis projects should provide a thorough overview of the area being surveyed, but will also be evaluated based on the novelty and technical merit of the insights that you provide. Again, these insights could include things like connections between different papers, assumptions that you found unrealistic or unmotivated, or interesting open problems that could be studied. You should plan to devote at least two pages of your report to your own insights and critiques of the literature.
- Empirical projects will be evaluated largely on the connections that you draw between your empirical results and the theory. Do the theoretical results correctly predict algorithm performance on your data? If not, explain why. You should plan to devote at least two pages of your report to this comparison.
- Theoretical research projects will be evaluated primarily on the reasonableness of your model, definitions, techniques, and results. You should explicitly motivate all of these aspects of your paper. Again, it's ok if you don't end up solving a problem completely, as long as you make a reasonable attempt, clearly explain what you tried, and explain why what you tried didn't work.