

1 Overview and Timeline

The project is intended to be a investigation of your choice in the topic areas relevant to the class (i.e. “structured prediction”, “learning with humans in the loop”, “learning representations”). The purpose of the project is to do research in an area of your interest in greater detail. It should be scoped to be original and publishable research. The project can take on many forms, including but not limited to:

- a. Projects that explore the application of machine learning ideas to an interesting “real-world” problem.
- b. Projects that involve a theoretical or empirical study of aspects of a learning method or model.
- c. Projects that extend existing learning algorithms, or that propose new learning algorithms.

Please propose a topic to us in your project proposal, and we will give you feedback on its feasibility. You will present the results of your project both in a short presentation in class, as well as in a final project report. More details are given below.

The important dates are:

- 9. February:** Form groups (no submission due)
- 16. February:** Project proposal due (via CMT as PDF)
- 7. April:** Status report due (via CMT as PDF)
- 24. April - 6. May:** Project presentations (in class)
- 12. May:** Final project report due (via CMT as PDF)
- 18. May:** Peer reviews of project reports due (via CMT)

2 Project Proposal

The project proposal should outline what you want to do in your project. We and your peers will give you feedback on your proposal to make sure the project you are proposing is feasible and appropriate. The proposal should contain the following sections:

- a. The Team. Names of the people working on the project.

- b. Motivation. Explain why this project is interesting and important.
- c. Statement of the Problem/Task. A statement of the problem, issue, or task that you're interested in studying. In particular, try to formulate the key questions (2 to 4 questions is probably a good number) that you will answer in the project.
- d. General Approach. A high-level description of the general approach you'll use to address the questions. This should include how you will evaluate and what evidence you are planning to gather (e.g. how you can answer the questions through experiments on data).
- e. Resources. A list of resources you have/need to conduct the project. This includes reading, software, datasets, etc. How are you planning to get these resources?
- f. Schedule. A schedule of work indicating the dates by which you plan to complete components of the project. Make sure the schedule is plausible.

The maximum length for the proposal is two pages in ICML style (including title, authors, and references). The ICML style is available at <http://www.icml2010.org/sub.html>.

3 Status Report

The status report contain a refined version of the proposal after incorporating instructor feedback and the peer feedback. Think of the feedback as advice which you can chose to follow, but in the end it is your decision how you do your research. In addition to the two-page proposal, the report should contain two pages describing your progress, any preliminary results, and as well as a summary of the changes that you made to the project proposal (i.e. total length 4 pages).

4 Project Presentation

The project presentation will be during the last two weeks of class. We will assign a time slot to each group. In the presentation, you should

- a. provide motivation for your project, explaining why it is important and interesting,
- b. explain your research questions,
- c. provide evidence,
- d. draw conclusion.

You can use the computer projector (e.g. via powerpoint) for your talk. Try to make the presentation interesting (e.g. by including a demo). More details on the logistics follow later.

5 Project Report

The final project should be submitted via CMT as a PDF file. There is no need to submit any other material like raw data, experiment logs, or source code.

A good model for writing the report is imagining that you are writing a conference paper. You have seen many good conference papers during the semester, so you should have an idea what a good paper looks like. The most important points to keep in mind for the report are

- to motivate your project,
- to discuss related literature,
- to make a convincing argument that supports your conclusions,
- to make sure that the reader understands what your project is about and how you came to your conclusions, and
- to make sure that credit is given to all software, literature, etc. that helped you in your work.

Also, keep in mind the criteria according to which the project will be evaluated (see section “Review Criteria and Grading”). The length limit for 8 pages (same as an ICML paper), and most research results can be fit into that length. Being concise is a good thing, but do not sacrifice clarity and completeness.

6 Peer Reviewing

You will receive and give peer feedback on all stages of the project, namely the proposal, the status report, the presentation, and the report. These reviews will be single blind, which means that the reviewer knows the author identities, but the author does not know the reviewer identities. Part of the code of academic integrity is to not review authors you feel conflicted with, and to not reveal to any of your peers which papers you reviewed.

When you write reviews, the goal is to be as objective and constructive as possible. Also, please be polite!

7 Review Criteria and Grading

The projects will be graded in the same spirit as research papers are assessed. Here is a list of things that we will be looking for:

- a. Originality (how novel is the research)
- b. Significance (are the questions you are asking interesting and important)

- c. Relevance (is the work relevant to the course)
- d. Quality of arguments (are claims supported, how convincing are the arguments you bring forward)
- e. Clarity (how clearly are goals and achievements presented)
- f. Connections to related work
- g. Scope/Size (in proportion to size of group)

Relative to each other, the proposal will account for 10% of the grade, the status report 10%, the presentation for 20%, and the report for 60%.

Feel free to come and talk to me during my office hour about the various aspects of your project so that we can make sure that you are on the right track.

Have fun, and good luck!