

Final Project

You are asked to carry out an original research project related to the course content. We envision several types of projects:

- Write a tutorial or review article.
 - This should not simply be a regurgitation of existing papers. There must be a clear value-added: you should clearly explain the relationships between the algorithms using a consistent notation, come up with illustrative examples, run experiments on toy problems which illustrate the differences between the algorithms.
- Try out an existing algorithm in a new setting.
- Invent a new algorithm.

Collaboration. You should form teams of 3–5 students. Your final report should list the contributions of each team member.

Due Dates

- **Project proposal:** Oct. 12
- **Presentations:** Nov. 24 and Dec. 1
- **Final report:** Dec. 10

Project Proposal

Each team will write a short (around 2 pages) research project proposal. It should include a description of a minimum viable project, some nice-to-haves if time allows, and a short review of related work. You don't have to do what your project proposal says — the point of the proposal is mainly to have a plan and to make it easy for us to give you feedback.

E-mail your proposal to `csc2541-submit@cs.toronto.edu`.

Presentation

Towards the end of the course everyone will present their project in a short, roughly 5 minute, presentation.

Final Report

At the end of the class you'll hand in a project report, in the format of a machine learning conference paper (e.g. NIPS¹). We recommend the report be about 8 pages plus references, but we do not enforce any minimum or maximum length.

E-mail your report to `csc2541-submit@cs.toronto.edu`. **You must also submit the code necessary to reproduce your experiments.**

¹<https://nips.cc/Conferences/2017/PaperInformation/StyleFiles>

Marking

85% of the marks will be given for meeting the requirements of the project and for the quality of the project proposal, presentation, and final report. (A mark of 85% corresponds to A-/A.) This includes:

1. **Abstract (5 points)** that summarizes the main idea of the project and its contributions.
 - Should be understandable to anyone in the course.
 - You don't need to say everything you did, just what the main idea was and one or two takeaways.
2. **Introduction (5 points)** that states the problem being solved and why we might want to solve it.
3. **Figure or diagram (10 points)** that shows the overall idea. The idea is to make your paper more accessible, especially to readers who are starting by skimming your paper.
 - For camera-ready diagrams, we recommend using Tikz, a LaTeX package. For Mac users, Omnigraffle is also convenient.
4. **Formal description (15 points)** of the model / algorithm / conjecture. Include at least one of:
 - An algorithm box.
 - Equations describing your model.
 - A theorem or formally stated conjecture.

Highlight how your approach differs from existing work.

5. **Related work section and bibliography (15 points).**
 - If your project builds on previous work, clearly distinguish what they did from what your new contribution is.
 - Also, include a short (2–3 sentence) summary of other closely related papers.
 - I realize you might not know about all related papers (or have time to carefully read all related papers), and that's OK for this project.
 - Using BibTeX is annoying at first, but Google Scholar can give you the BibTeX entries.
6. **Comparison or demonstration (20 points).** Include at least one of:
 - A proof of a theorem or conjecture, or an interesting counterexample.
 - An experimental comparison of the results of your method compared with a baseline. Qualitative evaluation is OK.
 - An experiment demonstrating a property that your model has that a baseline model does not. Experiments should also include a description of how you prepared your datasets, how you trained your model, and any tricks you used to get it to work.

- (for review/tutorial projects) An experiment that reveals interesting properties of or relationships between existing methods.

Toy data is OK! The point is to help the reader understand why or when we would want to use one approach over another, or to understand something better. Try to summarize the main takeaways. Negative results are fine, as long as you have an insightful and well supported explanation.

7. **Limitations (10 points)** of your approach.

- Describe some settings in which we'd expect your approach to perform poorly, or where all existing models fail.
- Give some examples of possible extensions, ways to address these limitations, or open problems.

8. **Conclusions (5 points)**

- State the results achieved in relation to the problem described in the introduction.
- Repeat the main takeaways from your paper.

9. **List of contributions.** You should list the contributions of each team member.

The other 15% will be given for going above and beyond. We are looking for some form of creativity, e.g. clever experiments that reveal an interesting phenomenon, tricks for circumventing obstacles, etc.

Standards will be higher if you are doing a project related to your research.

Rules about overlapping projects

- You are welcome to do a project related to your research.
 - Your project proposal and final report must **each** clearly explain the relationship to your research, what work was already done prior to the course, and what work (if any) was done by people not on the project team.
 - Our expectations will be higher in this case.
- If the project is closely related to one you are doing for another class:
 - You can't double-count work. Your project proposal and final report must **each** include a statement clearly explaining the relationship between the projects and which parts are meant to be counted for each class.
 - This statement must be approved (i.e. signed) in advance of the due date by the instructors of **both** classes.
 - You must include your project writeup for the other class along with your final report.