

Final Project

1. Logistics

For the course project, you will implement a research idea related to the course material. The purpose of the final project is to give you some experience working on a piece of original research and writing up your results in a paper-style format. You are expected to describe your research idea/application clearly in the project proposal, relate it to existing work, and document the project progress in the final report.

You **must form a group of two or three** to complete the project. Your report must clearly **list the contributions of each team member**. Once your group is formed, please sign up your group through Quercus. Instructions for sign-up can be found here: <https://qstudents.utoronto.ca/group-tool-the-student-side-of-things/>.

The project proposal is **due October 31st**. The final report is **due at 11:59 pm, December 14th**. The write-ups are to be submitted through Quercus.

2. Academic Integrity

Any work or ideas that are not your own must be appropriately acknowledged with a citation in the final report. This includes any code used for the experiments and text in the report itself. Plagiarism and other forms of academic misconduct are not tolerated in this course. The University of Toronto's Code of Behaviour on Academic Matters outlines the behaviours that constitute academic misconduct, the processes for addressing academic offenses, and the penalties that may be imposed. You are expected to be familiar with the contents of this document.

3. Large Language Model Policy

We will adopt a similar guideline to the International Conference on Machine Learning 2023 policy about the usage of Large Language Models (LLMs), such as ChatGPT, for the final report. The LLM policy **prohibits text produced entirely by LLMs (i.e., "generated")**. This does not prohibit authors from using LLMs for editing or polishing author-written text.

The LLM policy is largely predicated on the principle of being conservative with respect to guarding against potential issues of using LLMs, including plagiarism.

4. Writing Format

All submissions must be in PDF format. You may include algorithm blocks, tables, and figures. The write-ups should be prepared in the NeurIPS paper format: <https://nips.cc/Conferences/2021/PaperInformation/StyleFiles>. You may find online editors such as Overleaf helpful for writing the reports: <https://www.overleaf.com/latex/templates/neurips-2022/kxymzbpwsqx>.

Proposal

The project proposal is limited to two pages. It should roughly have the following sections:

- 1/4 page introduction
- 1/2 page related works
- 1/2 page method/algorithm
- 1/4 page abstract and references

The point of the proposal is mainly for feedback and formulating a plan for the final report. The proposal will **NOT** be graded. You will submit your proposal report through Quercus.

Final Report

You will expand your project proposal to include experiments and comprehensive method sections. You are expected to discuss the experimental results in detail and highlight any interesting findings. We recommend the final report to be **FOUR** pages plus the references. Appendix is allowed with no page limit, but the teaching staff reserves the right to judge the final project solely on the basis of the 4 pages of the main report. Any extra material is at the discretion of the reviewers and is not required. You will submit your final report through Quercus. You must also submit the code necessary to reproduce your experiments.

5. How to Choose a Project

The course projects should build on the course materials. You are encouraged to use neural networks as the function approximators for your method or application. There are two categories of projects to choose from:

Understanding and Analysis

For students who would like to have a more in-depth understanding of the course material, it is often a good idea to re-implement an existing method and re-evaluate the implementation against some standard benchmarks.

- Reproduce the experimental results from existing papers and perform sensitivity analysis on hyperparameters.
- Apply/extend existing algorithms to a new application/task/dataset.

If you choose to work on this category, you will need to implement and analyze the performance of **at least** two different deep learning algorithms/methods in a task domain, e.g., image recognition or natural language processing. You are asked to discuss the strengths and weaknesses of each approach backed by your experimental findings.

Doing a proper analysis for the existing methods is non-trivial. Here are two great examples of this type of study: <https://arxiv.org/pdf/1311.2901.pdf> and <https://arxiv.org/pdf/1506.02078.pdf>.

Exploratory Research

You may also choose to work on a novel research idea that may lead to a potential publication.

- Improve/fix an existing algorithm and evaluate the improvement on benchmark environments.
- Develop novel model architectures/algorithms to a new application/area/environment.

If you decide to work on a research idea, you will need to implement and compare the performance of your method against at least one existing approach in your problem.

For guidance on choosing a good research problem, see Bill Freeman's advice (<https://billf.mit.edu/sites/default/files/documents/cvprPapers.pdf>) and David Patterson's slides (<https://people.eecs.berkeley.edu/~pattnsn/talks/BadCareer.pdf>).

You are encouraged to choose a project related to your research. If you do so, your project proposal and final report must clearly outline how the project is connected to your research, what work had been completed prior to the course, and what contributions (if any) were made by individuals outside your project team. Please note, in this scenario, we will hold your work to a higher standard, expecting more depth, originality, and distinction in your approach and execution, given your prior familiarity with the subject.

6. Grading Scheme

You may receive full marks for the course project by choosing either of the two categories. There will be no advantage in terms of achieving a higher grade for choosing an exploratory project over an analysis one. The goal of the project is for you and your group to conduct original research.

The proposal and the final report will be graded according to the criteria used in top machine learning conference submissions. We will use the NeurIPS review criteria for this purpose:

- **Quality [35%]:** Is the report technically sound? Are claims well-supported by theoretical analysis or experimental results? Is this a complete piece of work, or merely a position report? Are the authors careful (and honest) about evaluating both the strengths and weaknesses of the work?

- **Clarity [25%]**: Is the report clearly written? Is it well-organized? (If not, feel free to make suggestions to improve the manuscript.) Does it adequately inform the reader? Are the figures/tables properly labeled? (A superbly written report provides enough information for the expert reader to reproduce its results.)
- **Originality [20%]**: Are the problems or approaches new? Is this a novel combination of familiar techniques? Is it clear how this work differs from previous contributions? Is related work adequately referenced? We recommend that you check the proceedings of recent NIPS conferences to make sure that each report is significantly different from papers in previous proceedings. Abstracts and links to many of the previous NeurIPS papers are available from <https://papers.nips.cc/>.
- **Significance [5%]**: Are the results important? Are other people (practitioners or researchers) likely to use these ideas or build on them? Does the report address a difficult problem in a better way than previous research? Does it advance the state of the art in a demonstrable way? Does it provide unique data, unique conclusions on existing data, or a unique theoretical or pragmatic approach?
- **Participation [15%]**: We will adopt a peer-review system through Quercus in which students will participate in reviewing the other classmates' reports. Reports will be made anonymously. We expect each student to review at least 2 reports. The participation score will be given based on the quality of the reviews by each student. You may find the following link helpful regarding how to write a good review with examples: <https://iclr.cc/Conferences/2022/ReviewerGuide>.

7. Lateness

There is **NO** late acceptance for the final project submissions, and any late reports will receive a **0%**. To avoid potential technical issues, we highly recommend submitting the final report well before the deadline, rather than at the last minute.