Project Report

Overview

In teams of up to 4 people, you will be working on a final project. This is the second part of the final project: the project report.

My hope is that, through this project, you carry out original research or take a deeper dive into the literature. The rubric is intended to give you a guideline, and my goal is for this report to be a manageable amount of work. If you are feeling overwhelmed by the rubric or this report, please talk to me. I'm here to help you find your way!

Details

Collaboration You may work in teams of up to 4 people. If you are having trouble finding a group of people to work with, please email the instructor. Only one report needs to be handed in per group.

Length The report should be 4 to 8 pages, not including appendices or bibliography. You are encouraged to keep the text short and informative. Illustrative figures are a great way to convey intuition, and we encourage you to think of creative ways to convey the main idea. You can include as many proofs, experiments, or additional details as you want in the appendix.

Format Please use 10pt, 11pt, or 12pt font with standard margins. You may use any format that you wish, but it should be readable. It should also have easily identifiable sections.

Tips Here are some tips on writing.

- Unless you have concise, informative proofs, long mathematical derivations are best put in the appendix. You can state the result in the main text, and defer the proof to the appendix.
- You are encouraged to attach your code.
- Axes on all plots should be labeled.
- Use a vector graphics format for your figures (e.g., pdf, eps, or svg) so that your figures don't look blurry.
- Try to use at most two font sizes.

Content The requirements of this project are intentionally flexible. I hope this gives you room to be creative. The following would all be acceptable projects:

- Inventing a new algorithm for either reinforcement learning or variational Bayesian inference.
- Deriving an interesting new connection between approximate Bayesian inference and reinforcement learning
- Performing a theoretical analysis of some poorly understood algorithm or approximation in reinforcement learning or variational Bayesian inference.

- A review that places the ideas of this course into a broader context, e.g., another field like physics or mathematics.

Marking Scheme

This marking scheme is worth 60 marks and is based heavily on Prof. Duvenaud's marking scheme and we refer you to his wonderful recommendations for more details.

- **Abstract** (**4 marks**) Summarize the main ideas of the project and its contributions.
  - Include the abstract at the top of the paper.
  - This should be clear and concise.
  - Do not describe every detail of the project.
  - Describe one or two main ideas.
- **Introduction** (**10 marks**) Describe the context of your paper and the key ideas.
  - The introduction is a critical part of any paper. Most people do not read past the introduction of a paper, so you should put effort into making the introduction approachable and clear.
  - Set the context: if you are trying to solve a certain problem in this project, describe why it is worth solving. If you are trying to prove a certain theorem, describe why it is worth proving.
  - You should also describe the main outline of your paper, including the key ideas and milestones.
- **Visualization** (**4 marks**) Visualize a main idea.
  - You should include at least one visualization of the main idea. The goal is to make the work you've done more accessible. You should create a new figure.
  - You may draw it by hand, as long as it's clear and readable.
  - We recommend OmniGraffle, which has a 14 day free trial.
  - We also recommend matplotlib in Python.
- **Formal description** (**10 marks**) Include a precise description of the main idea.
  - Your paper should be on a topic within the theme of the course, and your paper should involve in some way a model, loss function, problem domain, or mathematical conjecture. This part of the rubric asks that you describe that clearly and formally.
  - Include at least one of:
    - An algorithm box,
    - Equations describing your model,
    - A theorem or formally state conjecture,
    - A formal description of a problem domain.

  - You should also have a discussion that unpacks the intuition behind the formal description.

- **Related work** (**4 marks**) Explain exactly how your project relates to the literature.
  - If your project builds on previous work, clearly distinguish what they did from what your new contribution is.
  - It's OK if you do not find all related papers, but do your best to find a few closely related papers.
  - For closely related papers, include 1-2 sentence summaries.
  - Have a proper bibliography section and in-text citations.

- **Experiments or demonstration** (**10 marks**) Demonstrate a key idea.
  - Include at least one of:
    * If you are proving a theorem, include an illustrative example or counter-example.
    * If you are introducing a new algorithm, implement it for *a toy problem* and include a comparison to the simplest alternative algorithm for that toy problem.
    * If you are introducing a new machine learning model, train it on *a toy data set* and include a comparison to the simplest alternative model for that data.
    * If doing a review, include a table comparing the properties of the different approaches.
  - 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.
  - You are not expected to have your own source of compute for this part of the project. You can use Google Colab or UToronto Jupyter Hub. We do not expect you to run computationally intensive experiments for this part, but they should be carefully executed.
  - Toy data or toy problems are OK and even encouraged! 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.

- **Limitations** (**4 marks**) State the limitations of your approach.
  - Most new ideas have limitations and that is OK! In fact, it is expected.
  - Describe some settings in which you would 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.
  - If you are doing a review, explain the limitations of the ideas that you are reviewing.

- **Conclusions** (**2 marks**) Conclude your paper.
  - State the results that you obtained in your paper.
  - If you did experiments, repeat the main conclusions of your experiments.

  - Repeat the main ideas in your paper.

- **Novelty** (**10 marks**) Explore something new in your paper! This is a sliding scale.
  - (**0/10 marks**) If you did not include any new ideas, but you wrote a great review of the literature.
  - (**5/10 marks**) If you made some interesting innovations. Examples might include:
    * If you are proving a theorem, you used a somewhat uncommon proof technique or managed to extend some well-known result to a slightly new setting.
    * You introduced an algorithm that is an interesting, but small, twist on existing methods.
    * You used an existing model type in a new setting.
    * You mostly reviewed the literature, but you made some unexpected, detailed, and precise connections between apparently unrelated subfields.
  - (**10/10 marks**) You made very novel contributions. In this case, I would expect that your paper really contributes some kind of new insight.

- **Contributions** (**2 marks**) List the contributions of each team member.

Academic Integrity

Because this assignment will involve citing other people's work, it is important that cite properly. In general, you should follow U of T's Code of Behaviour on Academic Matters.