

Presentation and Code Notebook - v2

As a reminder, the course syllabus contains the full submission policies that cover every assignment, including this one.

Overview

Either by yourself or with a partner, you will be presenting a recent paper from the academic literature and preparing a code notebook to exemplify a key idea in that paper. The goal of this assignment is to help you practice the following skills:

- Reading academic papers, including the ability to fill in any gaps in your knowledge that are necessary for understanding someone else's work;
- distilling an academic paper into its key idea and reproducing some of its empirical findings;
- and presenting a paper's contribution in a short academic talk that provides a broader context.

"Fresh" academic papers are usually more challenging to engage with than textbooks. Even very influential papers can be poorly written and challenging to reproduce. So these skills will hopefully help throughout your research career. Because we want to cover a reasonable number of papers each week, if students end up dropping the course, we may require that students present alone.

Deliverables

The following must be submitted through MarkUs on the due date.

1. Presentation slides in `.pdf` format.
2. Python code notebook in `.ipynb` format.
3. If you used large language models, then you must submit all chat transcripts directly related to the preparation of the assignment in `.txt` format.

Timing and Due Date

- In the first two weeks of the course, you will work with me to pick a week in which to present. You will present on the date that we mutually agree.
- Two weeks before your presentation, meet with me at office hours to confirm the paper that you're presenting. This will also give us a chance to discuss the context behind the paper.
- A few days before your presentation, you will meet with one of the teaching staff to practice your presentation and get feedback.
- On the day of your presentation, **your slides and code notebook are due at the start of the class.**

Presentation

Your team will be preparing a 20-minute academic presentation on a paper from the literature. We will pick a paper for you to present from the a list of three papers that are related to the core reading of that week (see the course website).

The goal of your presentation is to *explain the 1-2 key ideas or findings* in this paper and to describe their context in the literature. You do not need to cover all aspects of the paper in your presentation. It is usually good to spend most of your presentation setting up the background necessary to understand the main ideas of the paper. If there's more than a couple key findings, then you should pick the 2 that you think are the most relevant to the course. I can help you do this.

We will be spending the first two weeks on motivation, definitions, and basic tools that will help you approach the paper that you are presenting. Still, we cannot guarantee that these two weeks will provide you with *all* of the necessary facts to fully understand any of the papers in the recent references. Figuring out what you need to know is an aspect of this assignment. We are happy to provide pointers, if you come to us with specific questions.

Tips

- Cite the paper and the paper's authors on the first slide.
- You will need to build a bit of a background in order to understand the key contribution of your paper.
 - Read / skim some of the other papers in the same stack for your presentation week.
 - Follow the citation graph, use Google Scholar.
 - Your presentation should be able to comment on where your paper sits in the literature.
- I find the title slide of presentations to be the hardest. It's not always clear what to say, and I often start rambling. So, even if I don't script my whole presentation, I find it useful to script the first few sentences of my presentation.
- Ask for help!

Code Notebook

Your team will be preparing an original Python Colab notebook to demonstrate a key idea or reproduce an experimental result from the paper that you are presenting. The purpose of this is to help everyone understand the paper better.

Your main goal with the notebook is to *demonstrate one key idea* in the paper. It is not expected that you reproduce all of the empirical findings of the paper, and you are not expected to use an unreasonable amount of computational resources. The notebook should be self-contained and fully runnable on Google Colab (<https://colab.research.google.com>), *i.e.*, it should not contain bugs and not require that we manually sort out dependencies or install obscure packages ourselves.

The requirement for this section is somewhat vague, because it will depend heavily on the specific paper that you are presenting. The following are suggestions:

- If the paper already has a toy experiment, you can reimplement that and use it to demonstrate the properties of the algorithm presented in the paper.
- If the paper does not have a toy experiment, you could design a novel toy setting that highlights the strengths (or reveals the weaknesses) of the methods proposed in the paper.
- If your paper is particularly challenging to design experiments for, you could also design an animation that provides intuition regarding the key ideas in the paper.

It is not a requirement that this contains significant original work or insights, but you can get an extra mark if it does.

Marking Scheme

This marking scheme is worth 25 marks.

Presentation

- **(3 mark(s))** Provides the necessary background to understand the key idea.
- **(3 mark(s))** Explains the key idea clearly at the beginning of the presentation and at the end.
- **(3 mark(s))** Outlines the scope and limitations of the approach.
- **(3 mark(s))** Presents a visual demonstration of one of the ideas in the paper. This can be a run-through of the code notebook.
- **(1 mark(s))** Does not run over time. You can decide the exact balance of each section, but you should aim for 20 minutes and definitely not exceed 25 minutes.
- **(1 mark(s))** Present to one of the course staff a few days before your presentation to get feedback.
- **(1 mark(s))** Cites the paper on the first slide.

Code notebook

- **(3 mark(s))** Cites the paper and explains in a text cell which keys ideas are being demonstrated in the code notebook.
- **(6 mark(s))** Produces plots or visualizations when run that accurately represent the key ideas as well as relevant baselines.
- **(1 mark(s))** Contains an original insight or design, and explicitly points out which aspect is original.

FAQ

If the code notebook needs to be fully self-contained, how can I use external sources of compute or memory to run larger experiments? You can cache results from experiments that are not fully runnable on Google Colab and load those experimental results into the notebook. The spirit of the “fully runnable” requirement is to avoid submissions that fail to run due to a bug or an obscure dependency that is not automatically installed by the notebook.

Academic Integrity

Because this assignment involves presenting someone else’s work, it is important that all of the material that you produce adequately cites the original paper. In general, you should follow U of T’s Code of Behaviour on Academic Matters. In particular, please make sure that the first slide of your presentation clearly cites the original authors (as well as the presenters) and that your notebook cites the original paper.