

DS 4440: Project Requirements

Fall 2020

The culmination of this course is a final project which you will implement a non-trivial neural network. This project should be done alone, and ideally will build upon your mid-term topic survey. My hope is that this provides you an opportunity to take what we learn in this class and work on a project of interest to you, that you can then show off as part of your portfolio. Here are some example projects for inspiration (this is intended to be *illustrative*, absolutely not exhaustive!)

- Re-implement (from scratch!) a state-of-the-art model — or a stripped down version of it — that you reported on in your final project, and reproduce reported results.¹
- Analyze an existing model in a novel way, for example critically assessing datasets or metrics.
- Extend ‘minitorch’ (from earlier assignments) to implement a more advanced module or particular architecture.
- Work with our ‘evidence inference’ (<http://evidence-inference.ebm-nlp.com/>) dataset to try and improve the best models we’ve come up with: <https://arxiv.org/abs/2005.04177>.

You are expected to consult with me (and/or Dave, our TA) on the project to ensure appropriate and realistic scope; please reach out!

Proposal

You are to submit a one-page project proposal by 11/12, end of day. This will describe the following: (1) General project idea; (2) Dataset(s) to be used; (3) Metrics for evaluation; (4) Hypotheses concerning outcomes; (5) Anticipated difficulties/obstacles.

¹Note that if you selected a ‘task’ from paperswithcode.com you will need to be explicit in distinguishing your implementation from the reference open-source version; in particular, if the reimplementation itself is a major part of your project we will reference your implementation against the open-source variant — you must be able to explain yours! If your project is mostly about *extending* or *analyzing* an existing model, however, starting from an open-source implementation is completely fine. When in doubt, discuss with us.

Project Write-Up

By 12/11, end of day, you are to turn in a final write-up (as a PDF file, via Canvas) describing your project and accompanying results. This should roughly be structured as follows: (1) Abstract, (2) Introduction/motivation, (3) Experimental setup; (4) Results/discussion; (5) Conclusions/future work. There is no hard page minimum; the document should be sufficiently long to communicate your work.

Code

All code you write is to be submitted a single zipped repository. I strongly encourage you to also post your code to GitHub, but this is not strictly required. **A Jupyter notebook that produces the main analysis/results must be included.**

Presentation

You will present your projects to your peers (and me) on the last day of class. This will include a short talk (time limits TBD) outlining the project and findings, and then a walk through of your implementation/findings.

Grading

Projects will graded under the following general set of criteria.

- Project ambition/creativity (20%)
- Execution/implementation (30%)
- Final write-up (30%)
- In-class project presentation (20%)

Execution and implementation will be assessed both indirectly (via the artifacts produced and reported upon in the final report) and directly (via the code).