Course Project

The course project will be individual. No group projects. However, I encourage you to discuss your papers and project with other students. The project aims to introduce you to the AI/ML research landscape and help you dive deeper into a topic of your choice (as the lecture content is necessarily narrow). The project aims to emulate the research paper publishing process including literature review, paper writing, implementation, presentation, and peer reviews.

Peer Reviews of Project

Throughout the semester, you will be required to turn in project checkpoints that will be reviewed by your peers (similar to the peer review process for academic publishing). Like all conferences, this process will be double blind: reviewers will not know the identity of authors and vice versa. To support this, like all conferences, **you should NOT put your name on your submissions or on your reviews.** Also, like conferences, reviews will be confidential. The only person who will be privy to the reviews will be the reviewer, the instructor, and the author. For the final term paper, each student will be required to read other student term papers and prepare conference-style reviews, primarily indicating clarity and the quality of the implementation effort.

Project Checkpoints

For each checkpoint and the final term paper, you will be required to write the document in LaTeX meeting the typesetting conventions of ICML (see ICML 2022 author instructions for LaTeX template and instructions). For references, use BibTex (*.bib file) as in the ICML template. I strongly suggest using https://www.overleaf.com/ for writing your paper since it includes all necessary software for doing LaTeX.

1. Draft review and critique of selected papers – You will be required to submit a draft of your review and critique of the three papers (2-3 pages, i.e., at least 2 full pages).
2. Preliminary implementation results – You will be required to submit a draft describing your experimental motivation, experimental setup, preliminary implementation efforts, and your planned next steps (2-3 pages of content). You may have failed to reproduce the results or have run into significant problems. This is okay at this stage, but you will need to write about your efforts.

Research Paper Selection (at least 3)

Students will be required to select and read *at least* three (but could be more) related recent conference or journal research papers in the fields of AI, computer vision, natural language processing, or machine learning. Specifically, the papers must have been:

1. Published in 2018, 2019, 2020, 2021, or 2022.
2. Published in one of the following venues (*=preferred venues):
    a. Machine Learning (NeurIPS*, ICML*, ICLR*, AISTATS, UAI, JMLR)
    b. Artificial Intelligence (broader) (AAAI*, IJCAI)
    c. Computer Vision (CVPR*, ICCV, ECCV)
    d. Natural Language Processing (ACL*, NAACL, EMNLP)
3. Contain material that can be implemented. (at least one paper)

I may be willing to accept papers even if they do not fit the criteria above, but they must be approved by me one week before the due date. Note that arXiv papers are not allowed unless they have been published already in the venues above. When you cite a paper, **make sure to cite the final publication venue (e.g., ICML) and NOT arXiv**.

For a journal paper, your citations should contain (at least) the paper title, authors, journal, volume, year, pages and pdf URL. For a conference paper, your citations should contain (at least) the paper title, authors, conference, year, and pdf URL. Below is an example of BibTeX entry for a conference.

```
@inproceedings{inouye2018deep,
 title = {Deep Density Destructors},
 author = {Inouye, David I. and Ravikumar, Pradeep},
 booktitle = {International Conference on Machine Learning (ICML)},
 year = {2018},
 url = {http://proceedings.mlr.press/v80/inouye18a/inouye18a.pdf}
}
```

Implementation

You will be required to do one of the following:

1. Reimplement method in paper: Reimplement and evaluate the method from at least one paper. If an implementation of the paper is already available (e.g. from the author's website), you **must state this in your report** and compare your implementation to the existing implementation, both in terms of code and performance. Ideally, you should replicate the experiments presented in the paper, but I will not require this.
2. Rerun experiments AND extend methods in paper: If code already exists to reproduce experiments, you should first rerun some of the original experiments. Then you must propose, implement, and evaluate a significant extension of the paper. For this option, you can build off of any existing implementation, but your implementation must extend or alter the original method in a significant way.

Your part of the implementation for any of the above must be nontrivial. A good guideline is that your implementation should be at least four pages of code. Ultimately, I will determine whether or not the implementation meets the non-triviality requirement. If you have questions about your implementation, please contact me to discuss.

*The implementation must be in **Python**. Given that state-of-the-art methods require many GPU hours to train, I would suggest getting your implementation to work on simple benchmark datasets first (e.g., MNIST, FashionMNIST, CIFAR10). If you have time, evaluate on more complex datasets.*

Term Paper

You will be required to write a six page paper (specifically, 5.5-6.5 pages of content excluding references) in LaTeX meeting the typesetting conventions of ICML (see link above). I strongly suggest using Overleaf.com for LaTeX. Approximately two to three pages of this paper should be a substantive critique of the three (or more) papers that you have read. And approximately three to four pages of this paper should be a description of your

implementation, evaluation and discussion of the method from at least one of the papers. (Note: **You should not put your name on the term paper to accommodate for peer reviews.** See "Peer Reviews" section above.)

5 Minute Spotlight Video

To facilitate the large class size, you will be required to create a 5-minute video presentation of your project. These videos will be made available to everyone in the class. The video should be between 4-5 minutes but not longer than 5 minutes. This "spotlight" presentation should cover:

1. Motivation and problem definition (including prior work if appropriate).
2. A description of your implementation.
3. A description and discussion of your evaluation.
4. One slide to highlight the take-home messages.
5. (Optional) Future directions.

You should only have between 5-10 slides to fit within the 5 minutes. One minute per slide is usually reasonable. You can see real examples of video abstracts from top machine learning conference at https://nips.cc/Conferences/2018/Schedule?type=Poster (look for "3-min video" links). I do not expect your video to be at the quality level of these videos, but it hopefully gives you some ideas. The video can be in any format you want including animations, slides with narration, video of you presenting, or combination of the above. However, the focus should be on **clarity** rather than *fanciness*.

Presentations

(Tentative, final details TBD) For presentations, we will likely use Zoom breakout rooms for simultaneous live presentation sessions to accommodate the large class. You will be required to sign up as a presenter and a discussant. The discussant will watch another student's presentation beforehand and prepare discussion questions to ask the presenter live.