

Instructions on Course Projects

Jian Tang

January 27, 2020

1 Goal

The goal of the course project is to apply deep learning techniques learned in class (it is fine if you use the techniques not introduced in class) to solve real-world problems or develop new deep learning techniques. You are expected to work in teams and learn to collaborate with your teammates. Each group should make a poster in the final class and participate in the poster session to present your results and communicate with other teams. A report should also be submitted at the end of the course.

2 Due Dates

Group Formation Deadline: 2019.1.30 (Students registering later than this deadline will be deducted 2 points)

Project Proposal Deadline: 2019.3.01

Poster Representation: 2019.04.09/(J1) 2019.4.14 (J2).

Final Report Deadline: 2018.4.25 (Note: 3 points will be deducted for reports handed in late than the deadline).

Handing in: Send the electronic version to jian.tang@hec.ca with title [MATH 80600A-J1/J2-Your Group ID] and also print a paper version and hand it in to the secretary office of Department of Decision Science.

3 Evaluation Criteria

The course project will take up 40 points. The evaluation criteria is composed of three parts: project proposals (5 points), poster presentation (10 points), and final report (25 points).

Project Proposal. The project proposal is a summary of your proposed research topic and study plan. It should include the background of the problem (context and motivation), problem definition, and a plan on how you want to study it. The project proposal should be at most 2 pages.

Poster. The suggested poster size will be updated later. Some examples of posters in conferences are available here: <https://postersession.ai/>

Final Report. The final report should give a comprehensive description of your projects. It should contain a section about the motivation and definition of your selected topics, a section summarizing the related work, a section on the techniques you used for solving the problem, an empirical section presenting your data sets and results with detailed analysis, and a conclusion section. The report should be at most 8 pages (not including references) using the NeurIPS format¹. The final report should be submitted in pdf format. Overall, it is evaluated according to the following perspectives:

- Originality of the problem studied (8 points): this part evaluates whether the studied problem is original or not.
- Related Work (2 points): the report should also summarize the relevant work in existing literature and point out the difference between previous work and current work.
- Experiment Design, Results and Analysis (10 points): this part measures the soundness of the experiments. A good experiment should be well designed, conducted, and presented with detailed analysis to convince the readers.
- Writing (5 points): the report should be easy to understand and contain few grammar/typo errors.

4 What could be a Good Project?

In general, you can think from the following directions to select your course project.

- The first type of project is a survey style on a specific topic. Taking the topic of text classification as an example, what are the most commonly used architectures and state-of-the-art methods for text classification? How are they different from each other? You can evaluate and compare all these methods on some data sets with a variety of settings. You should summarize the findings in your report.
- Another type is application style, which means you can apply deep learning techniques to new problems or data sets (especially those in your own domains). In this case, you should compare with classical methods in your selected problems and show why deep learning is better.
- The third type is much more difficult, which is developing new methods or architectures for a specific problem.

We will also provide some candidate topics in the course Website (it is not necessary to choose from the suggested list). In any case, if you have any problems on course projects, please reach out to the instructor or the teaching assistants for help.

¹<https://nips.cc/Conferences/2015/PaperInformation/StyleFiles>

5 Some Advice

Some advice (Modified from csc2515 at UofT):

- Be selective! Don't choose a project that has nothing to do with deep learning. Good problems are usually those with new data sets solved with deep learning techniques or new deep learning techniques.
- Be honest! You are not being marked on how good the results are. It doesn't matter if your method is worse than the ones you compare to provided you implemented it properly. What matters is that you try something sensible and clearly describe the problem, your method, what you did, and what the results were.
- Be modest! Don't pick a project that is way too hard. Usually, if you select the simplest thing you can think of to try, and do it carefully, it will take much longer than you think.
- Be careful! Don't do foolish things like test on your training data, set parameters by cheating, compare unfairly against other methods, include plots with unlabeled axes, use undefined symbols in equations, etc. Do sensible cross-checks like running your algorithms several times, leaving out small parts of your data, to make sure everything still works reasonably well. Make lots of pictures along the way.
- Learn! The point of the project is to give you a chance to "test drive" the process of apply deep learning techniques. Consider this an opportunity to learn how to write code to process and analyze data with data mining techniques, make nice figures, layout readable equations, describe your work concisely to a smart but uninitiated reader, etc.
- Have fun! If you pick something you think is cool, that will make getting it to work less painful and writing up your results less boring.