STAT 946
Deep Learning
Winter 2018

Instructor:
Ali Ghodsi
Dept. of Statistics & Actuarial Science
University of Waterloo
Office: M3 4208
E-mail: ali.ghodsi@uwaterloo.ca

Prerequisite: Some knowledge of machine learning, probability, calculus, linear algebra, and statistics.

Deep Learning Deep learning (also known as deep structured learning or hierarchical learning) attempts to learn representations of data with multiple levels of abstraction. Deep learning usually refers to a set of algorithms and computational models that are composed of multiple processing layers. These methods have significantly improved the state-of-the-art in many domains including, speech recognition, classification, pattern recognition, drug discovery, and genomics.

Tentative topics:

- Feedforward Deep Networks
- Optimization and regularization for Training Deep Models
- Convolutional Networks
- Sequence Modeling: Recurrent Nets
- Auto-Encoders
- Representation Learning
- Deep Generative Models (moment matching networks)
- Generative Adversarial Networks (GANs)
- Deep learning for Natural Language Processing
- Variational Autoencoders
- Attention model, Highway network, residual network, sequence-to-sequence models

Tentative Marking Scheme

Paper presentation 30%
Data Challenge 20%
Group Project 50%

Paper presentation:

- You need to choose a paper from NIPS 2017, ICLR 2018 or ICML 2018.
- Paper presentation starts **tentatively** late October.
- Attending all presentations is mandatory.

Marking Scheme for paper presentation::

- Presentation (10% of your final mark.)
- Primary contributor (10% of your final mark): You need to write a summary/critique of the paper that you would like to present on wikicoursenote.com by exactly one week before your presentation date.
- Secondary contributor (10% of your final mark): You must contribute to your peers summaries/critiques before their presentation dates.
    - You must contribute to AT LEAST 20 of your peers' papers.
    - Contributions can be technical or editorial, but AT LEAST 15 of them MUST be technical.

Project:

Final group project (presentation and reports up to 4 pages of PDF plus 1 page reference) are worth 50% of your final grade .

Your final project can be:

- A Kaggle completion. You may choose an *active* competition from *featured* or *research* categories. Kaggle Competitions in other categories (*in class, getting start* or *playground*) are not eligible for the final project.

- Develop a new algorithm. In this case, you will need to demonstrate (theoretically and/or empirically) why your technique is better (or worse) than other algorithms. (Note: A negative result does not lose marks, as long as you followed proper theoretical and/or experimental techniques).
- Application of Deep Learning to some domain. This could either be your own research problem, or you could try reproducing results of someone else's paper.

Note that you cannot borrow part of an existing thesis work, nor can you re-use a project from another course as your final project. Final project reports will be checked by Turnitin (Plagiarism detection software).

Final project presentation:

You will present your final project as a poster on Friday, December 14th. Poster presentations will take place on this date from 9:00 AM to 4:00 PM at M3 Atrium.

Acceptable forms for the poster is full poster or slides printed on A4 /letter size papers.

In addition to the poster, you will have to turn in a four-page single-spaced, single-column report plus one page reference. This is due both in hard copy and electronically (PDF) by Friday, December 14th.

Academic Honesty In paper critiques, projects and wikicoursenote, if you use ideas, plots, text and other intellectual property developed by someone else you have to cite the original source.

If you copy a sentence or a paragraph from work done by someone else, in addition of citing the original source you have to use quotation marks to identify the scope of the copied material.

Example: Plagiarism is an act of 'using ideas, plots, text and other intellectual property developed by someone else while claiming it is your original work.' [1]

Evidence of copying or plagiarism will cause a failing mark in the course. Note that final project reports and paper critiques will be checked by Turnitin (Plagiarism detection software).

Persons with Disabilities: The office for Persons with Disabilities (OPD), located in Needles Hall, Room 1132 collaborates with all academic departments to arrange appropriate accommodations for students with disabilities without compromising the academic integrity of the curriculum. If you require academic accommodations to lessen the impact of your disability, please register with OPD at the start of each academic term.

Communication

All communication should take place using the *Piazza* discussion board. Piazza is a good way to discuss and ask questions about the course materials, including assignments, in a public forum. It enables you to learn from the questions of others, and to avoid asking questions that have already been asked and answered. It also provides a forum for course personnel to make announcements and clarifications about assignments and other course-related topics. Students are expected to read Piazza on a regular basis.

Enrolling in Piazza

You will be sent an invitation to your UW email address. It will include a link to a web page where you may complete the enrollment process.

Piazza Guidelines

Here are some guidelines that you should keep in mind when posting items to Piazza:

1. Please remember that everything you post is public - everyone enrolled in this course will be reading it. As a result, in any posts you make, do not give away any details on how to do any of the assignments / data challenge . This could be construed as cheating, and you will be responsible as the poster. If you have questions about an assignment / data challenge that require you give specific details of your solution, you may still post to Piazza, but check *This is a private post - only visible to class instructors* . If the instructors and/or TAs feels that posting it to everyone is appropriate, they will do so.

2. Keep posts related to the course, concise, and topical. As students are all expected to read Piazza on a regular basis, try not to waste the time of readers.

3. Please be diligent about attempting to find the answer before you post a question. Piazza includes excellent search facilities – use them! Scan all of the questions that have already been asked. Better yet, read them along with the answers. You'll learn lots! Please do all you can to avoid duplicates.

4. Make it easy for other students to find your question – just in case they have the same question and want to see the answer.

    - Use a meaningful subject heading. "Help" and even "Help for A3Q2" is not very meaningful. "Clarify parameter order for A3Q2" is much better.
    - Tag your post with all the applicable tags. Start a tag by typing the hash character (#). A drop-down list of tags that are currently in use will appear. Use one of them, if applicable. If not, create a new one. However, any tag you create should be applicable to many posts not just yours.

5. Please don't post things to the group that provide no useful information to readers. Posts like "I have the same question as this one just posted", or "I agree with this comment" serve no useful purpose, and waste people's time.

6. Keep complaints about the course out of Piazza or mark them with the *This is a private post - only visible to class instructors* checkbox. If you have a concern about anything to do with the course, the best way to deal with it, and to get results, is to take it to the course instructor. Piazza is not a complaint forum.

Grades will be handled through *Learn*. Please log on frequently to Piazza and Learn. You are responsible for being aware of all material, information and email messages found on *Learn* and *Piazza* throughout the semester.

References

[1] Tec Encyclopedia. http://www.answers.com/topic/plagiarism