Syllabus for CSC 578: Neural Networks and Deep Learning, Fall 2017

Overview

This course focuses on the algorithms, implementation, and application of neural networks for learning about data. Students will learn how neural networks represent data and learn in supervised and unsupervised contexts with applications to language processing, classification, and regression problems. Topics include learning algorithms, and optimization methods, deep learning methods for deriving deep representations from surface features, recursive networks, Boltzmann machines and convolutional networks.

Learning Outcomes

By the end of this course you should be able to:

- Describe the mathematical foundations underlying neural networks
- Use matrix programming techniques to implement and modify a learning algorithm
- Implement deep learning techniques for dealing with structure in language and image processing applications
- Select an appropriate neural network / deep learning approach for a given task

Prerequisites

(CSC 412 AND CSC 478) OR (CSC 403 AND IS 467)

Textbooks

- **NNDL**: Neural Networks and Deep Learning, by Michael Nielsen. Available for free online.
- **DLB**: Deep Learning Book, by Goodfellow, Bengio, and Courville. MIT Press. Also available for free online, or bound from your favorite bookseller.

Attendance

Class attendance is strongly encouraged, but not mandatory. However, if you are absent from class you are responsible for understanding the material and for finding out about any announcements made in that class.

Class Plan

The following class plan is tentative and subject to change as the course progresses.

- **Class 1:** (9/11) Course overview. Intro to NNs
- **Class 2:** (9/18) How Neural Nets work. Matlab lab (second half of class in Daley 505)
- **Class 3:** (9/25) Backpropagation
- **Class 4:** (10/2) Improving Neural Networks I
- **Class 5:** (10/9) Improving Neural Networks II and Why are Deep Networks Hard to Train?
- **Class 6:** (10/16) Deep Learning
- **Class 7:** (10/23) Sequence Modeling

- **Class 8:** (10/30) Natural Language Processing, TensorFlow
- **Class 9:** (11/6) AutoEncoders, Recursive AutoEncoders
- **Class 10:** (11/13) Final project workshop in Daley 505
- **Class 11:** (11/20) Project presentations

Instructor Information

Email	peterh@cdm.depaul.edu
Home Page	http://reed.cs.depaul.edu/peterh/
Phone	312-362-5736
Office Hours	Monday, Wednesday 3:30-5:00pm (except 9/13, 10/4, 11/1) or by arrangement
Address	CDM Center 717 DePaul University 243 South Wabash Avenue Chicago, IL 60604-2301 USA

Assessment

Your final grade will be based on:

- Weekly online "warm-ups" from readings: 10 points total
- Weekly participation: 15 points total
- 3 Homework Projects: 15 points each
- Final Project: 20 points
- Final Project Presentation: 5 points
- Meaty comments on 5 other final project presentations: 5 points

The grading scale will be:

Points	Grade
93.3	A
90	A-
86.6	B+
83.3	B
80	B-
76.6	C+
73.3	C
70	C-
66.6	D+
60	D
< 60	F

Weekly Warm-ups

I use a "Just In-Time Teaching" methodology which is intended to help students by aligning what's done in and outside of class. So each week (after the first), you will be given some "warm-up" questions or exercises to help you come to grips with what you've read, and to help me see what you're getting and what you're not. (That way the class time can be tailored to your needs.) The

warm-ups will generally consist of around 6 questions, and must be completed before noon (i.e. by 11:59 AM) on the day of the class. (This applies to both in-class and online students.) The lowest score will be dropped.

Weekly Participation

The participation aspect of the class is designed to keep you involved with the class on a weekly basis, and to give you a chance to practice new concepts learned in class. It can also be a forum for asking follow-up questions. There will be two types of participation points available. You will need both to get full participation points.

Regular participation points (10 of the 15 points) can be earned by answering the questions that I post in the weekly online discussion forums. (Note: In-class students can get partial participation credit for participating in in-class discussion, but should still participate in the online forums.) To encourage you to keep up with the class, each weekly discussion forum will only be open for 10 days starting from the relevant class date.

Bonus participation points (5/15 points) can be earned by being a "good citizen" in the class's learning community. This can come through such postings as:

- Asking an important question that didn't come up in class (explain why it is important)
- Answering someone else's question
- Suggestions for related topics or tutorials or datasets that can expand on what we looked at in class. There is an explosion of information available now on these topics, and there's no way that the class can cover it all. This is your opportunity to help others find something useful. **Note:** You won't get bonus points by just posting a link to something. You've got to also describe what it is, and why it is important / interesting.

To earn bonus participation points, you should make an additional post to the relevant weekly discussion forum, and clearly label what type of bonus you think it is.

Projects

The projects for this course will involve applying different machine learning techniques to different problems. Some of the projects will require programming in Matlab (or Octave) though we will spend very little time discussing how the language works. So if you are not confident with your ability to learn a new programming language largely on your own, this is not the class for you.

The final project will be on a topic of your choice (with instructor approval of the topic) and include a presentation to the class that describes your project at the class meeting during finals week. This will be described in more detail later.

Late submission policy

Unless otherwise stated, written assignments are due via D2L at the time and date posted on the course homepage. You are expected to complete all of the written assignments by the deadline. Late homework submissions will be accepted via D2L with the following penalty:

If assignment is turned in…	**Penalty will be…**
within 3 days of due date	10% of the total points for each day it is late
3 days or more after due date	will NOT be accepted

On Plagiarism

You are encouraged to discuss all homeworks and projects with your classmates. You are, however, required to complete the assignments **on your own**. In particular, this means that you are not allowed to "cut and paste" text from anywhere else, *or to paraphrase* someone else's work, unless: 1) it is a *very small* part of your submission, 2) the copied text is clearly indicated (i.e. surrounded by quotation marks), and 3) the source is clearly identified (with citation and full reference information).

All assignments will be submitted to "Turn it in" for automatic plagiarism testing. This system is very good at finding things that have been copied, so just don't do it.

School policies on instructor evaluation, email, plagiarism and incompletes.