



Fall 2018 Deep Learning: Syllabus and Schedule

Course Description:

This course is an introduction to deep learning, a branch of machine learning concerned with the development and application of modern neural networks. Deep learning algorithms extract layered high-level representations of data in a way that maximizes performance on a given task. For example, asked to recognize faces, a deep neural network may learn to represent image pixels first with edges, followed by larger shapes, then parts of the face like eyes and ears, and, finally, individual face identities. Deep learning is behind many recent advances in AI, including Siri's speech recognition, Facebook's tag suggestions and self-driving cars.

We will cover a range of topics from basic neural networks, convolutional and recurrent network structures, deep unsupervised and reinforcement learning, and applications to problem domains like speech recognition and computer vision. Prerequisites: a strong mathematical background in calculus, linear algebra, and probability & statistics (students will be required to pass a math prerequisites test), as well as programming in Python and C/C++. There will be assignments and a final project.

Time/Location: Mon/Wed 2:30-4:15pm in room CAS 313

Sections: EC500 K1 / CS591 K1

Instructor:

Brian Kulis, bkulis@bu.edu; office hours: Tuesday 11:0-12:30pm in PHO 441

Teaching Assistants:

Ali Siahkamari, siaa@bu.edu, office hours: M/W 4:30-6:30pm, CS Undergrad Lab (730 Comm Ave #302)

Xide Xia, xidexia@bu.edu, office hours: M/W 4:30-6:30pm, CS Undergrad Lab

Weichao Zhou, zwc662@bu.edu

Mehrnoosh Sarmashghi, msmshgi@bu.edu

Blackboard: registered students can access via <https://learn.bu.edu>

Course Pre-requisites

This is an upper-level undergraduate/graduate course. All students should have the following skills:

- Calculus, Linear Algebra
- Probability & Statistics
- Ability to code in Python
- Background in machine learning

In addition, students must complete and pass the Pre-Quiz on prerequisite math knowledge – see schedule below. Students who cannot pass the Pre-Quiz must drop the class.

Syllabus and Schedule	Lectures	Assignments	Reading	Videos
---------------------------------------	--------------------------	-----------------------------	-------------------------	------------------------

Schedule*

	Topic (Instructor)	Details	Homework
Wed Sep 5	1. Course overview	What is deep learning? DL successes; syllabus & course logistics; what is on the pre-quiz?	
Mon Sep	Math Prerequisite Quiz	there will be no make-up quiz	

10			
Wed Sep 12	2. Math review I	Gradient descent, logistic regression. Reading: Goodfellow Ch5.9-5.10	return pre-quiz, hw1 out
Mon Sep 17	3. Math review II	Probability, continuous and discrete distributions; maximum likelihood. Reading: Goodfellow Ch5.1-5.6	
Wed Sep 19	4. Intro to neural networks	cost functions, hypotheses and tasks; training data; maximum likelihood based cost, cross entropy, MSE cost; feed-forward networks; perceptron; neuroscience inspiration; Reading: Goodfellow Ch6.2	
Mon Sep 24	5. SCC/TensorFlow Overview (Katia Oleinik)	How to use the SCC cluster; introduction to Tensorflow. Please bring your laptop to class, this will be an interactive tutorial.	SCC Info
Wed Sep 26	6. Learning in neural networks (Kate Saenko)	output vs hidden layers; linear vs nonlinear networks; Reading: Goodfellow Ch6.1-6.3	ps1 due 11:59pm Ps2 out
Mon Oct 1	7. Backpropagation	learning via gradient descent; recursive chain rule (backpropagation); if time: bias-variance tradeoff, regularization; output units: linear, softmax; hidden units: tanh, RELU; Reading: backprop notes , Goodfellow Ch6.5	
Wed Oct 3	8. Deep learning strategies I	GPU training, regularization, etc; project proposals	
Tue Oct 9	9. Deep learning strategies II	Optimization algorithms, dropout, batch normalization	
Wed Oct 10	10. CNNs	Convolutional neural networks; Reading: Goodfellow Ch9.1-9.3	Ps2 due 11:59pm Ps3 out
Mon Oct 15	11. Unsupervised deep learning I	Autoencoders	Project proposal due
Wed Oct 17	12. Unsupervised deep learning II	Generative Adversarial Networks	
Mon Oct 22	13. RNNs I	recurrent neural networks; sequence modeling; backpropagation through time; vanishing/exploding gradient problem; gradient clipping, long-short term memory (LSTM)	
Wed Oct 24	14. RNNs II	more intuition about RNNs, LSTMs; toy addition problem; language modeling; bi-directional RNN	Ps3 due 11:59pm Ps4 out
Mon Oct 29	15. Deep Belief Nets I	Probabilistic modeling	
Wed Oct 31	16. Variational Methods	Variational Autoencoders	
Mon Nov 5	17. No class	Brian out of town	
Wed Nov 7	18. Deep Belief Nets II; Attention and Memory	Applications of Deep Belief Nets and related models; Neural Turing Machines Reading: Neural Turing Machines paper; Neural Machine Translation by Jointly Learning to Align and Translate paper; http://distill.pub/2016/augmented-rnns/ (optional)	

Mon Nov 12	19. Deep Reinforcement Learning I	Overview of RL	
Wed Nov 14	20. Deep reinforcement learning II	Policy Gradient	Ps4 due 11:59pm
Mon Nov 19	21. Deep Reinforcement Learning III	Actor-critic, Q-learning	Progress report due in class Template
Fri Nov 23			PS5 out
Mon Nov 26	22. Image/Video Captioning, Autonomous Driving		
Wed Nov 28	23. Other NLP applications	Parsing, recursive neural networks	
Mon Dec 3	24. Speech and Audio Applications	ResNet and WaveNet	
Wed Dec 5	No class	NIPS Conference	
Fri Dec 7			Ps5 due 11:59pm
Mon Dec 10	Project presentations I		project presentation (due 12:00pm on day of presentation)
Wed Dec 12	Project Presentations II		
Fri Dec 14			Project report due at 5:00pm template

*schedule is tentative and is subject to change.

Textbook

The required textbook for the course is

- Ian Goodfellow, Yoshua Bengio, Aaron Courville. [Deep Learning](#).

Other recommended supplemental textbooks on general machine learning:

- Duda, R.O., Hart, P.E., and Stork, D.G. [Pattern Classification](#). Wiley-Interscience. 2nd Edition. 2001.
- Theodoridis, S. and Koutroubas, K. [Pattern Recognition, Edition 4](#). Academic Press, 2008.
- Russell, S. and Norvig, N. [Artificial Intelligence: A Modern Approach](#). Prentice Hall Series in Artificial Intelligence. 2003.
- Bishop, C. M. [Neural Networks for Pattern Recognition](#). Oxford University Press. 1995.
- Hastie, T., Tibshirani, R. and Friedman, J. [The Elements of Statistical Learning](#). Springer. 2001.
- Koller, D. and Friedman, N. [Probabilistic Graphical Models](#). MIT Press. 2009.

Recommended online courses

- <http://cs231n.stanford.edu/> CS231n: Convolutional Neural Networks for Visual Recognition

- <http://web.stanford.edu/class/cs224n/> CS224n: Natural Language Processing with Deep Learning
- <http://rll.berkeley.edu/deeprcourse/> CS 294: Deep Reinforcement Learning
- <http://distill.pub/> Very nice explanations of some DL concepts

Deliverables/Graded Work

There will be six homework assignments, each consisting of written and/or coding problems, and a final project. The project will be done in teams of 3-4 students and will have several deliverables including a proposal, progress update(s), final report and a final in-class presentation. The course grade consists of the following:

- Math prerequisite quiz 5%
- Homeworks, best 4 of 5 45%
- Project (including all components) 45%
- Class participation 5%

Software/Hardware

Programming assignments and projects will be developed in the Python programming language. We will also use the Tensorflow deep learning library for some homeworks and for the project. Students are expected to use the [Shared Computing Cluster \(SCC\)](#) and/or their own machines to complete work that does not require a GPU. For the projects, we will provide GPU resources.

If you do not already have a CS account and would like one, you should stop by the CS undergraduate lab (EMA 302) and activate one. This process takes only a few minutes, and can be done at any time during the lab's operating hours: <<http://www.bu.edu/cs/resources/laboratories/undergraduate-lab/>>

Late Policy

Late work will incur the following penalties

- Final project report and presentation: 20% off per day up to 2 days
- Homework 20% off per day, up to 3 days

Academic Honesty Policy

The instructors take academic honesty very seriously. Cheating, plagiarism and other misconduct may be subject to grading penalties up to failing the course. Students enrolled in the course are responsible for familiarizing themselves with the detailed BU policy, available [here](#). In particular, plagiarism is defined as follows and applies to all written materials and software, including material found online. Collaboration on homework is allowed, but should be acknowledged and you should always come up with your own solution rather than copying (which is defined as plagiarism):

Plagiarism: Representing the work of another as one's own. Plagiarism includes but is not limited to the following: copying the answers of another student on an examination, copying or restating the work or ideas of another person or persons in any oral or written work (printed or electronic) without citing the appropriate source, and collaborating with someone else in an academic endeavor without acknowledging his or her contribution. Plagiarism can consist of acts of commission-appropriating the words or ideas of another-or omission failing to acknowledge/document/credit the source or creator of words or ideas (see below for a detailed definition of plagiarism). It also includes colluding with someone else in an academic endeavor without acknowledging his or her contribution, using audio or video footage that comes from another source (including work done by another student) without permission and acknowledgement of that source.

Religious Observance

Students are permitted to be absent from class, including classes involving examinations, labs, excursions, and other special events, for purposes of religious observance. In-class, take-home and lab assignments, and other work shall

be made up in consultation with the student's instructors. More details on BU's religious observance policy are available [here](#).