Deep Learning (CS 5787) - Syllabus
Spring 2019

Instructor: Prof. Christopher Kanan
Co-Instructor: Dr. Jin Sun
3 Credit Hours

Catalog Description: Deep learning is an area of machine learning that has enabled enormous progress on long-standing problems in machine perception and reinforcement learning. This course will give you an understanding of neural networks, their limitations, and their current applications. Beyond reviewing state-of-the-art systems, students will be responsible for completing a project in deep learning. The course is designed to prepare students to evaluate and deploy state-of-the-art deep learning algorithms. Students are expected to understand the basics of machine learning, but no background in neural networks is expected. Students should have experience with python and should be comfortable with probability, linear algebra, and calculus.

Prerequisites: CS 5785 Applied Machine Learning or permission of instructor

Class Location and Times:

Mondays:	7:45pm - 9:00pm	[Discussion and Tutorials]	Dr. Sun	Bloomberg Center 131
Wednesdays:	6:05pm - 7:30pm	[Theory and Lecture]	Prof. Kanan	Bloomberg Center 131

Required Text: The main book for the class is “Deep Learning” (2016) by Ian Goodfellow, Yoshua Bengio, and Aaron Courville. An online version is available here: http://www.deeplearningbook.org There will be readings from other sources.

Instructor Contact:

Name: Prof. Christopher Kanan
Office Hour Location: Bloomberg Center 131
Office Hours: Wednesdays 7:30-8:30pm [Immediately after class]
Email Address: ck587@cornell.edu

Co-Instructor Contact:

Name: Dr. Jin Sun
Office Location: Bloomberg 364
Office Hours: Tuesdays: 2-3pm
Email Address: jinsun@cornell.edu

Teaching Assistants:

Name: Zhengqi Li
Office Location: Bloomberg 366
Office Hours: Wednesday 1-3pm
Email Address: zl548@cornell.edu

Name: Wenqi Xian
Office Location: Bloomberg 366
Office Hours: Fridays 3:30-4:30PM
Email Address: wx97@cornell.edu

Graders:

Name: Xiaohang Lu
Email Address: xl672@cornell.edu

Evaluation and Grading: The final course grade will be weighted as follows:

Homework:	50%
Basics Quiz:	10%
Project:	30%
Class Participation:	10%

We will follow standard Cornell grading guidelines to assign the percentage into a letter grade. The professor may choose to “curve” the class by giving all students the same number of additional points.

Homework: Your homework submissions must cite any references used (including articles, books, code, websites, and personal communications). All solutions must be written in your own words, and you must program the algorithms yourself. While there are not many homework assignments, they will be long and involved. You are responsible for starting on them early to ensure that you complete them by the deadline. If you start the day before, you will probably do poorly on the assignment.

Your homework solutions must be typed and output to PDF format. We encourage you to use LaTeX to write up your answers. Your solutions should include all diagrams, written explanations, code, and program output relevant to the problem.

Machine Learning Basics Quiz: This is an in-class quiz to ensure that each student understands the basics of neural networks and machine learning, before moving onto phase two of the course.

Team Project: You are required to complete a team project. Your project should be at the frontier of deep learning, but it does not necessarily need to move the frontier forward. You may use the programming language of your choice. Replicating results from a recent paper and comparing it to other works, would be a good project. An alternative is to build and rigorously evaluate a real-world application of deep learning. Run your early ideas by Prof. Kanan or Dr. Sun via email or in person. Unless you have good justification, each team should have 3-5 members. The schedule for the project is as follows:

1. **Project Proposal**: The project proposal should be clearly state what you plan to do. It should be four pages long (not including references). It should contain a list of three to six milestones and deadlines. You should list the questions the project will address and that will be discussed in the report. You should list what software you will be using or will build upon. Describe the datasets you will use and how how will you know if the project is successful. Describe the hypotheses you will test and the related work. The proposal should be a well organized document in continuous english, and it should not be merely an outline. You should be able to re-use much of the text for the final report. We encourage you to typeset it using LaTeX. It should be submitted as a PDF (under 10MB).
2. **Revised Project Proposal (optional)**: The revised proposal is an opportunity to improve your grade if you fail to do the project proposal effectively. You may submit a revised proposal that takes into account the comments received by the instructor and TA. The new grade will replace the original score, but the maximum score for the revised proposal is 80%.

3. **Project Report**: The project report will describe the project, i.e., what you did and the result. It should be six to eight pages long (not including references) and formatted in CVPR or NeurIPS format. It should be submitted as a PDF (under 10MB). Read CVPR or NeurIPS papers to get an idea for what the style and formatting should be.

Policy on Late Work: This is a graduate class. Late work will **not** be accepted. Assignments may involve a large time commitment, and you are unlikely to complete them by the deadline if you wait until the night before. I urge you to begin them immediately after they are assigned.

Programming Environment: You need to use Python for the homeworks. For the class project, you may use the programming environment of your choice, but most of our expertise in aiding you is in PyTorch.

Prior Course Materials: Unauthorized use of course materials from previous semesters (e.g., material you have received from others), is strictly prohibited.

New Course Materials: Course materials (slides, lectures, assignments, etc.) may not be re-distributed or posted elsewhere online. Redistribution of copyright protected material outside this course may be prohibited by law.

Notes on Plagiarism: Plagiarism is a serious offense and is in violation of the university policy. If you are unsure of what constitutes plagiarism in written documents, a good description can be found here: https://plagiarism.arts.cornell.edu/tutorial/index.cfm

Plagiarism does not just occur in written documents; it also occurs in code. Many of the algorithms we will code and problems we will solve have been solved by others who have posted code (in various programming languages) online. It is unacceptable (and it is considered plagiarism) to copy code developed by others and submit it as your own. (This includes code that is written by your fellow students!) Even making minor changes, such as changing variable names, function names, formatting, etc., is not enough to allow you to claim your submission as your own because the underlying structure of the code remains unchanged.

If you do consult any online sources of code, you must properly attribute the corresponding sections in your code to their original source, as you would add quotations, footnotes, or references in a written document. The consequences of plagiarism, whether in code or in written documents, are at the discretion of the instructor, and can be as severe as automatic failure of the course.

Academic Accommodations: We are committed to providing reasonable accommodations to students with disabilities. If you need accommodations such as special seating, note taking services, or extended time or a different environment due to a disability, please go to the Student Disability Services Office. If you receive accommodation approval, you must make me aware of this fact prior to the date that accommodations with be necessary.

Religious Observances: Cornell University is committed to supporting students who wish to practice their religious beliefs. Students are advised to discuss religious absences with their instructors well in advance of the religious holiday so that arrangements for making up work can be resolved before the absence.

Course Schedule: The following schedule lists dates for class topics. *The content in this schedule is tentative and subject to change.* It is your responsibility to attend class and to remain informed of any changes that may be announced.

Week	Date	Assignments	Class Topics	Main Reading	Presenter
1	1/21	No Class	No Class - Martin Luther King Jr. Day		N/A
	1/23		Introduction; Course Overview	Chapter 1 of GBC	Kanan
2	1/28		Example of A Simple Learning Problem		Sun
	1/30		Logistic Regression and Optimization	Chapter 5 of GBC	Kanan
3	2/4		Data Collection		Sun
	2/6		Training Neural Networks	LeCun, Bengio, & Hinton. (2015), *Nature*. Chapter 6 of GBC	Kanan
4	2/11	Homework 2 Assigned	Deep Learning Frameworks		Sun
	2/13		Backprop & CNNs	Chapter 8 of GBC Chapter 9 of GBC	Kanan
5	2/18	Homework 1 Due	Practical Tricks on Training Neural Networks 1		Sun
	2/20		CNN Architectures & Hardware Acceleration	Krizhevsky et al. (2012) NIPS. Yosinski et al. (2014) NIPS Ioffe & Szegedy (2015) Batch Normalization He et al. (2016) CVPR.	Kanan
6	2/25		Practical Tricks on Training Neural Networks 2		Sun
	2/27	Project Proposal Due	Adversarial Attacks What are Networks doing? Neural Networks Limitations		Kanan
7	3/4		Transfer Learning Tutorial		Sun
	3/6		Recurrent Neural Networks		Kanan
8	3/11		Neural Network Model Search		Sun
	3/13		Object Detection and Segmentation; UNe		Kanan
9	3/18	Homework 2 Due	How to Evaluate A Neural Network Model		Sun
	3/20	Revised Project Proposal	Reinforcement Learning Part 1		Kanan
10	3/25		Prelim / Midterm		Sun
	3/27		Reinforcement Learning Part 2		Kanan
11	4/1	No Class	Spring Break Holiday		N/A
	4/3	No Class	Spring Break Holiday		N/A
12	4/8		Train Your Own Object Detector (Local/Cloud)		Sun
	4/10		Generative Adversarial Networks		Kanan
13	4/15		Q&A		Sun
	4/17		Graph Networks		Kanan
14	4/22	Homework 3 Due	Sequential Data Modeling Tutorial		Sun
	4/24		Prof. Kanan Travelling - Guest Lecture		???
15	4/29		Q&A		Sun
	5/1		Topics I missed		Kanan
16	5/6	Final Project Due	Real-World Ready Tools		Sun
	5/8	No Class	No Class - Study Period		N/A
17	5/?				