Deep Learning (CS 5787) - Syllabus
Spring 2020

Instructor & Lecturer: Prof. Christopher Kanan
Discussion Section: Dr. Jin Sun
3 Credit Hours

Catalog Description: Students will learn deep neural network fundamentals, including, but not limited to, feed-forward neural networks, convolutional neural networks, network architecture, optimization methods, practical issues, hardware concerns, recurrent neural networks, dataset acquisition, dataset bias, adversarial examples, current limitations of deep learning, and visualization techniques. We still study applications to problems in computer vision and to a lesser extent natural language processing and reinforcement learning. The course is designed to prepare students to evaluate and deploy state-of-the-art deep learning algorithms.

Course Frequency: Offered in Spring

Prerequisites: CS 5785 Applied Machine Learning or permission of instructor

Preparation Summary: 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.

Class Location and Times:

Mondays	12:30pm - 1:45pm	Discussion & Tutorials	Dr. Sun	Bloomberg Auditorium (131)
Wednesdays	6:30pm - 8:00pm	Theory & Lecture	Prof. Kanan	Bloomberg Auditorium (131)

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

Student Outcomes:
1. Demonstrate understanding of deep neural network fundamentals.
2. Demonstrate the ability to characterize deep neural network performance.
3. Gain experience deploying deep learning models or in conducting deep learning R&D.

Instructor Contact:

Name: Prof. Christopher Kanan
Office Hour Location: Bloomberg Center Auditorium
Office Hours: Immediately after class on Wednesdays until 9:30pm
Email Address: ck587@cornell.edu

Co-Instructor Contact:

Name: Dr. Jin Sun
Office Hour Location: Bloomberg 312
Office Hours: Tuesdays 2-3pm
Email Address: jinsun@cornell.edu

Teaching Assistants & Graders:

Role	Name	Office	Office Hours	Email
TA	Qianqian Wang	Bloomberg 377	Wed 1-3pm	qw246@cornell.edu
TA	Wen-Ding Li	Slack Call	Fri. 1-3pm	wl678@cornell.edu
TA	Zhengqi Li	Bloomberg 316	Thurs. 2-4pm	zl548@cornell.edu
Grader	Siyu Yao	-	-	-
Grader	Xiran Sun	-	-	-
Grader	Bowen Zhang	-	-	-
Grader	William Krasnoff	-	-	-
Grader	Ishan Virk	-	-	-
Grader	Konstantinos Kallidromitis	-	-	-

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

Homework: 70%
Project: 30%

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 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.

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 5 members. The schedule for the project is as follows:

1. **Project Proposal**: The project proposal should 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 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: Given the accelerated pace of Cornell Tech's degrees and that many of you are travelling for job interviews, we will permit some "slip days" for homework assignments in which no penalty will be incurred. After the slip day period, **homework assignments will no longer be accepted**. That said, we expect you to be mature and to not procrastinate. Homework assignments may involve a large time commitment, and you are unlikely to complete them by the deadline if you wait until the night before, if for no other reason than some take a few hours to train the models. We urge you to start assignments immediately after they are posted online. No slip days will be given for the project proposal or final project report and you should avoid using slip days unless absolutely unavoidable.

Programming Environment: You need to use Python for the homework assignments. The deep learning framework to be used is PyTorch -- assignments cannot be done in Keras, Tensorflow, or some other framework. For the class project, you may use the programming language and framework of your choice, but most of our expertise is in PyTorch.

Academic Integrity: Each student in this course is expected to abide by the Cornell University Code of Academic Integrity. Any work submitted by a student in this course for academic credit will be the student's own work. The policy can be found on the university's website here: https://theuniversityfaculty.cornell.edu/academic-integrity/.

You are encouraged to study together and to discuss information and concepts covered in lecture and the sections with other students. You can give "consulting" help to or receive "consulting" help from such students. However, this permissible cooperation should never involve one student having possession of a copy of all or part of the work done by someone else, in the form of an email, an email attachment file, a diskette, or a hard copy. Should copying occur, both the student who copied work from another student and the student who gave material to be copied will both automatically receive a zero for the assignment. Penalty for violation of this Code can also be extended to include failure of the course and University disciplinary action.

During examinations, you must do your own work. Talking or discussion is not permitted during the examinations, nor may you compare papers, copy from others, or collaborate in any way. Any collaborative behavior during the examinations will result in failure of the exam, and may lead to failure of the course and University disciplinary action.

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 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	1/20	No Class	No Class - Martin Luther King Jr. Day		N/A
	1/22	Homework 0 Assigned	Introduction; Course Overview	Chapter 1 of GBC	Kanan
2	1/27	Homework 1 Assigned	Demo of A Real World Learning Problem		Sun
	1/29		Logistic Regression and Optimization	Chapter 5 of GBC	Kanan
3	2/3	Homework 0 Due	Deep Learning Frameworks		Sun
	2/5		Training Neural Networks	LeCun, Bengio, & Hinton. (2015), *Nature*. Chapter 6 of GBC	Kanan
4	2/10		Practical Tricks on Training Neural Networks		Sun
	2/12		Backprop & CNNs	Chapter 8 of GBC Chapter 9 of GBC	Kanan
5	2/17	Homework 1 Due	Auto-Differentiation		Sun
	2/19		More CNNs Regularization Hardware Acceleration Transfer Learning	Krizhevsky et al. (2012) NIPS. Yosinski et al. (2014) NIPS Ioffe & Szegedy (2015) Batch Normalization He et al. (2016) CVPR.	Kanan
6	2/24	No Class	February Break		Sun
	2/26	Project Proposal Due	CNN Architectures Network Visualization Neural Networks Limitations		Kanan
7	3/2		Transfer Learning Tutorial		Sun
	3/4		Neural Network Theory Deep Double descent Adversarial Attacks Model Theft	https://blog.openai.com/adversarial-example-research/	Kanan
8	3/9		Data Collection		Sun
	3/11		Prelim Preparation		Kanan
9	3/16		Sequential Data Modeling Tutorial		Sun
	3/18	Homework 2 Due	No Class		Kanan
10	3/23		Prelim / Midterm Cancelled Due to COVID-19		Sun
	3/25	Revised Project Proposal Homework 3 Released	Adversarial Attack Defenses Recurrent Neural Networks		Kanan
11	3/30	No Class	Spring Break Holiday		N/A
	4/1	No Class	Spring Break Holiday		N/A
12	4/6		Practical Tricks on Training GANs		Sun
	4/8		Language + Vision Tasks Autoencoders Generative Adversarial Networks		Kanan
13	4/13		Neural Network Model Search		Sun
	4/15		Reinforcement Learning Part 1		Kanan
14	4/20	Homework 3 Due	Evaluation Metrics		Sun
	4/22		Reinforcement Learning - Part 2		Kanan
15	4/27		Real-World Ready Tools		Sun
	4/29	Last Lecture	Self-Supervised Learning & BERT Multi-Label & Multi-Task Learning Semantic Segmentation Theory vs. Practice: Deep Learning in the Wild		Kanan

16	5/4	Final Project Due Homework 4 Due	TBD		Sun
	5/6	No Class	No Class - Study Period		N/A