Deep Learning (CS 5787) - Syllabus
Spring 2021

Instructor & Lecturer: Prof. Christopher Kanan
Co-Instructor & Discussion Section Leader: 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 will study applications to problems in computer vision, 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	11:50am - 1:05pm	Discussion & Tutorials	Dr. Sun	Zoom
Tuesdays	7:15pm - 8:30pm	Theory & Lecture	Prof. Kanan	Zoom

Textbooks:

- "Deep Learning" (2016) by Ian Goodfellow, Yoshua Bengio, and Aaron Courville. A free online version is available here: http://www.deeplearningbook.org
- "Deep Learning with PyTorch" (2020) by Eli Stevens, Luca Antiga, and Thomas Viehmann. Available here: https://pytorch.org/assets/deep-learning/Deep-Learning-with-PyTorch.pdf
- 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 Zoom: TBD
Office Hours: TBD
Email Address: ck587@cornell.edu

Co-Instructor Contact:

Name: Dr. Jin Sun
Office Hour Zoom: Link
Office Hours: Fridays 4-5pm
Email Address: jinsun@cornell.edu

Teaching Assistants & Graders:

Role	Name	Office Zoom	Office Hours	Email
TA	Yin Li (yl3243)	Link	Mon. 4-6pm	yl3243@cornell.edu
TA	Andrew Bennett (awb222)	Link	Wed. 3-5pm	awb222@cornell.edu
TA	Sebastian Ament (sea79)	Link Pass: 542404	Wed. 9:30-10:30am Fri. 9:30-10:30am	sea79@cornell.edu
Grader	Irene Peradejordi (if76)	-	-	if76@cornell.edu
Grader	Rain Liang (xl778)	-	-	xl778@cornell.edu
Grader	Prashant Jain (pj263)	-	-	pj263@cornell.edu
Grader	Sachi Angle (sva22)	-	-	sva22@cornell.edu
Grader	Saloni Gandhi	-	-	sg2452@cornell.edu

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 reuse 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 will 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 / Discussion Topics	Main Reading	Presenter
1	2/8	No Class	No Class	N/A	N/A
	2/9	Homework 0 Assigned	Introduction; Course Overview	Chapter 1 of GBC	Kanan
2	2/15	Homework 1 Assigned	Demo of A Real World Learning Problem		Sun
	2/16	Homework 0 Due	Logistic Regression and Optimization	Chapter 5 of GBC	Kanan
3	2/22		Deep Learning Frameworks		Sun
	2/23		Training Neural Networks Backprop	LeCun, Bengio, & Hinton. (2015), *Nature*. Chapter 6 of GBC Ioffe & Szegedy (2015) Batch Norm	Kanan
4	3/1		Practical Tricks on Training Neural Networks		Sun
	3/2		More on Backprop Regularization: BatchNorm, etc. CNNs	Chapter 8 of GBC Chapter 9 of GBC	Kanan
5	3/8	Homework 1 Due	Auto-Differentiation		Sun
	3/9	Wellness Day - No Class	No Class	N/A	N/A
6	3/15		Data Collection		Sun
	3/16	Project Proposal Due	Hardware Acceleration Transfer Learning	Krizhevsky et al. (2012) NeurIPS He et al. (2016) CVPR Yosinski et al. (2014) NeurIPS	Kanan
7	3/22		Transfer Learning Tutorial		Sun
	3/23		CNN Architectures Network Visualization		Kanan
8	3/29		Transfer Learning Tutorial (Cont.)		Sun
	3/30	Homework 3 Released	Adversarial Attacks / Model Theft Adversarial Attack Defenses Bonus: Publishing in AI	https://blog.openai.com/adversarial-example-research/	Kanan
9	4/5		Neural Network Model Search		Sun
	4/6	Homework 2 Due	Recurrent Neural Networks Language + Vision Tasks		Kanan
10	4/12	TBD	Sequential Data Modeling Tutorial		
	4/13	Revised Project Proposal	Autoencoders		Kanan
11	4/19		Autoencoders tutorial		Sun
	4/20		Generative Adversarial Networks		Kanan
12	4/26	No Class - Wellness Day	No Class	N/A	N/A
	4/27		Deep Learning in Industry Getting a Job in AI		Kanan
13	5/3		Practical Tricks on Training GANs		Sun
	5/4	Homework 3 Due	Deep Reinforcement Learning		Kanan
14	5/10		Real-World Ready Tools / Evaluation Metrics		Sun
	5/11	Last Lecture	Neural Networks Limitations Neural Network Theory Deep Double descent Transformers for Language & Vision Wrap Up		Kanan
15	5/17		No Class		Sun
	5/18		No Class		N/A
16	5/21	Final Project Due			