Deep Learning (CS 5787) - Syllabus
Spring 2022

Instructor & Lecturer: Prof. Christopher Kanan
Co-Instructor & Discussion Section Leader: Yongwhan Lim
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	5:55pm - 7:10pm	Discussion & Tutorials	Lim	Zoom
Wednesdays	5:55pm - 7:30pm	Theory & Lecture	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.

Slack Link: https://join.slack.com/t/slack-0iz4512/shared_invite/zt-1191t1era-qxtKD67n7moL1UJZ6rMcrw

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: Yongwhan Lim
Office Hour Zoom: Here (It might be easier to click on the link on Zoom tab from Canvas)
Office Hours: Mondays 1pm ET or by appointment!
Email Address: yongwhan.lim@cornell.edu

Teaching Assistants & Graders:

Role	Name	Office Zoom	Office Hours	Email
TA (FT)	Daniel Martin		MW 3pm-4pm	dm839@cornell.edu
TA (PT)	Kai Zhang		RF 11am-12pm	kz298@cornell.edu
TA (PT)	Shaden Shaar		TR 3:30pm-4:30pm	ss2753@cornell.edu
Grader	Grace Le	N/A	N/A	dl2228@cornell.edu
Grader	Haoran Yin	N/A	N/A	hy522@cornell.edu
Grader	Yiran Wang	N/A	N/A	yw2358@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. It takes time to implement, train, and evaluate neural networks. Depending on the hardware and implementation, some problems may take hours or more to train the network.

Your homework solutions must be typed and output to PDF format. Use LaTeX to write up your answers. Your solutions should include all diagrams, written explanations, code, and program outputs.

Homework Assignments and Topic (Tentative)	
Assignment #	**Topic**
0	LaTex, Linear Algebra, Prerequisite Skills
1	Optimization of Linear Networks
2	Transfer Learning
3	Training CNNs and RNNs
4	Generative Adversarial Networks

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 and other staff 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 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. 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 about eight pages long (not including references) and formatted in NeurIPS format. It should be submitted as a PDF (under 10MB).

Policy on Late Work: No credit will be given for the project report if it is turned in late. For the project proposal, a late submission will be treated as submitting a revised project proposal capping the score at 80%. For homework assignments, full points will be awarded only if the assignment is turned in at most one day late. Late homework assignments will be accepted up to 7 days late with a 20% penalty imposed, meaning the highest possible score will be 80% for any assignment that is 2-7 days late. No credit will be given for assignments turned more than 7 days late. An exception to this policy is that no assignments will be accepted after the project report deadline.

Programming Environment: For homework assignments, this course uses Python and PyTorch. 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	Presenter
1	1/24	No Class	No Class	N/A
	1/26	Homework 0 Assigned	Introduction; Course Overview	Kanan
2	1/31	Homework 1 Assigned	Demo of A Real World Learning Problem	Lim
	2/2	Homework 0 Due (Feb 3)	Review of Linear Algebra and Supervised Machine Learning	Kanan
3	2/7		Deep Learning Frameworks	Lim
	2/9		Optimization of Linear Neural Networks with Gradient Descent	Kanan
4	2/14		Practical Tricks on Training Neural Networks	Lim
	2/16		Optimization of Deep Neural Networks	Kanan
5	2/21		Transfer Learning Tutorial	Lim
	2/23	Homework 2 Assigned	Deep Neural Network Architecture Types (MLP, CNN, Transformer, RNN, GNN) & Transfer Learning	Kanan
6	2/28	February Break	No Class	N/A
	3/2	Homework 1 Due (Mar 4)	Convolutional Neural Networks & Hardware Acceleration	Kanan
7	3/7		Auto-Differentiation	Lim
	3/9		Convolutional Neural Networks & Hardware Acceleration	Kanan
8	3/14			Lim
	3/16		CNN Architectures Network Visualization	Kanan
9	3/21		Neural Network Model Search	Lim
	3/23		Adversarial Attacks / Model Theft Adversarial Attack Defenses	Kanan
10	3/28	Homework 2 Due (Mar 29)	Sequential Data Modeling Tutorial	Lim
	3/30	Project Proposal Due (Apr 1)	Recurrent Neural Networks Language + Vision Tasks	Kanan
11	4/4	Spring Break	No Class	N/A
	4/6	Spring Break	No Class	N/A
12	4/11		TBD	Lim
	4/13		Autoencoders	Kanan
13	4/18	Homework 3 Due (Apr 19)	Practical Tricks on Training GANs	Lim
	4/20		GANs	Kanan
14	4/25		Real-World Ready Tools / Evaluation Metrics	Lim
	4/27		Transformers	Kanan
15	5/2		Q&A	Lim
	5/4	Homework 4 Due (May 3)	Last Lecture	Kanan
16	5/9		No Class	Lim
	5/10	Final Project Due (May 10)	No Class	N/A
				N/A