Syllabus: Deep Learning for Vision (IMGS 684)
Fall 2018
Instructor: Prof. Christopher Kanan

Course 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 review neural networks and related theory in machine learning that is needed to understand how deep learning algorithms work. After gaining the prerequisite background knowledge, the class will review the latest algorithms that use deep learning to solve problems in computer vision and machine perception, and students will read and present recent papers on these systems. 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 produce new state-of-the-art deep learning algorithms. Students are expected to have programming experience and to be comfortable with probability, linear algebra and calculus. No prior background in machine learning or pattern recognition is required. Class 3, Credit 3 (F)

Prerequisites: Graduate Standing or permission of instructor

Class Location and Time: Mondays and Wednesdays, 10:00am - 11:15am, CAR-2155

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 required readings from other sources.

Instructor Contact:
Name: Dr. Christopher Kanan
Office Location: Building #76 (CAR), Room 3140
Office Hours: Wednesdays from 11:15am - 1:00pm (immediately following class)
No office hours on: 9/5, 9/26, 10/17, 10/24, 11/14, 11/21, 12/5
Shortened office hours on: 9/19 (45 mins)
Email Address: christopher.kanan@rit.edu

Teaching Assistants:
Name: Manoj Acharya
Office Location: CAR-2210 (kLab, location subject to change)
Office Hours: Wednesdays and Fridays 2pm-4pm
Email Address: ma7583@rit.edu

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

Homework:	34%
Basics Quiz:	10%
Peer Paper Reviews:	6%
Project:	25%
Class Lecture:	10%
Class Participation:	15%

Letter grades will be assigned as follows:

A : 93 or above	A- : 90-92	B+ : 86-89	B : 83-85	B- : 80-82
C+ : 76-79	C : 73-75	C- : 70-72	D : 60-70	F : under 60

Note that C- grades and below do not count toward the fulfillment of program requirements for a graduate degree at RIT. The professor may choose to "curve" the class by giving all students the same number of additional percentage 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 prepared in LaTeX and output to PDF format. If you don't already know LaTeX, this is an excellent opportunity to start using it. Many academic conferences and journals require LaTeX formatted submissions. Your solutions should include all diagrams, written explanations, code, and program output relevant to the problem.

Paper Presentations: Every student is required to lecture on a particular set of papers on a particular date. Sign up early if you want a particular topic.

1. Your group must meet with Prof. Kanan one week early to go over your slides (40% of your grade). Your presentation must be a coherent explanation of the topic area (one slide deck)
2. You should fill 70 minutes of content that reviews the papers, any background material, and puts them in context. You should not only explain what they did, but critically examine the paper to look for flaws.
3. You must give an interactive in-class quiz. More information will be provided in class.

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.

Peer Reviews: You will be assigned the task of reviewing three paper drafts produced by your classmates. You will be graded based on the efficacy and thoroughness of your reviews. Review criteria and guidelines will be distributed.

Project: You are required to complete a 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. Run your early ideas by Prof. Kanan via email or in office hours. *You may team with others, but this requires extremely good justification and pre-approval for Prof. Kanan*. Expectations will be greater for teams. 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 3-4 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. It should be typeset using LaTeX, and 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 format. It should be typeset using LaTeX, and submitted as a PDF (under 10MB). Read CVPR papers to get an idea for what the style and formatting should be.
4. **Project Presentation**: You will have to give a presentation on your project to the class.

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. Note that we will be unable to provide any programming help for projects.

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 RIT course may be prohibited by law.

Notes on Plagiarism: Plagiarism is a serious offense and is in violation of the RIT Student Academic Integrity Policy (http://www.rit.edu/academicaffairs/policiesmanual/d080). If you are unsure of what constitutes plagiarism in written documents, a good description can be found here: http://isites.harvard.edu/icb/icb.do?keyword=k70847&pageid=icb.page342054

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: RIT is 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 Disability Services Office. It is located in the Student Alumni Union, room 1150. If you receive accommodation approval, you must make me aware of this fact prior to the date that accommodations with be necessary.

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	8/27	Homework 0 Assigned	Introduction, Course Overview	Chapter 1 of GBC	Kanan
	8/29	Homework 1 Assigned	Machine Learning Basics: Simple Classifiers	Chapter 5 of GBC	Kanan
2	9/3	Homework 0 Due	No Class - Labor Day		
	9/5		Classifier Evaluation & Challenges	Chapter 5 of GBC	Kanan
3	9/10		Logistic Regression and Optimization	Chapter 5 of GBC	Kanan
	9/12		Training Neural Networks Part 1	LeCun, Bengio, & Hinton. (2015), *Nature*. Chapter 6 of GBC	Kanan
4	9/17	Homework 1 Due	Training Neural Networks Part 2	Chapter 8 of GBC	Kanan
	9/19		Training Neural Networks Part 3	Chapter 9 of GBC	Kanan
5	9/24	Homework 2 Assigned	Convolutional Neural Networks -Part 1	Krizhevsky et al. (2012) NIPS. Yosinski et al. (2014) NIPS Ioffe & Szegedy (2015) Batch Normalization He et al. (2016) CVPR.	Kanan
	9/26		Convolutional Neural Nets - Part 2	None	Kanan
6	10/1	Project Proposal Due	Machine Learning Basics Quiz Review	None	N/A
	10/3		Machine Learning Basics Quiz	Goodfellow et al. (2014)	Kanan
7	10/8	Proposal Feedback	No Class October Break		
	10/10	Homework 3 Assigned	Publishing in AI		Kanan
8	10/15	Revised Project Proposal Due	CNN Hardware & Architectures		Kanan
	10/17	Homework 2 Due	CNN Visualization & Limitations of Neural Networks		Kanan
9	10/22	Presentation sign up Due	Adversarial Attacks		Kanan
	10/24		Unsupervised Learning, Self-Supervised Learning, & Generative Methods		Kanan
10	10/29		Object Detection, Triplet Loss, Semantic/Instance Segmentation Recurrent Networks & BackProp Through Time	Arulkumaran et al. (2017)	Kanan
	10/31	Project Drafts Due (11/1) Mock Reviews Assigned	Reinforcement Learning - Introduction		Kanan
11	11/5		Lifelong Learning 1	Castro et al. (ECCV 2018) iCaRL Hou et al. (ECCV 2018) Shmelkov et al. (ICCV 2017)	
	11/7	Peer Reviews Due	Lifelong Learning 2	https://openreview.net/pdf?id=rJxF73R9tX Neal et al. (ECCV 2018) Van de Ven & Tolias (2018) Shah et al. (2018) Javed et al. (2018)	
12	11/12		Understanding Deep Learning	Jiang et al. (2018) Zhang et al. (2017) Ji & Telgarsky (2018) Bartlett et al. (2017) Hoffer et al. (2017) Collins et al. (2018)	
	11/14		CVPR Deadline - No Class		
13	11/19	Homework 3 Due	Vision Grounded Language	Ramakrishnan et al. (2018) Yi et al. (NIPS 2018) Yang et al. (HCOMP 2018) Noh et al. (2018)	
	11/21		No Class - Thanksgiving		
14	11/26		Visual Curiosity	Yang et al. (2018) Zhang et al. (2018) Teney & van den Hengel (2018) Uehara et al. (2018) Bonus, but not critical: Pathak et al. (2017)	
	11/28		Fairness and Bias in Machine Learning	Nature Article	

				Book Sutton et al. (2018) Doshi-Velez & Kim (2017)	
	12/12		Project Presentations		