Syllabus: Deep Learning for Vision (IMGS 789)
Fall 2016
Instructor: Prof. Christopher Kanan
Teaching Assistant: Tyler Hayes

Course Description: Deep learning is an area of machine learning that has enabled enormous progress on long-standing problems in computer vision. 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. Students are expected to have programming experience and to be comfortable with 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: Tuesdays and Thursdays, 12:30pm - 1:45pm, CAR-1275

Required Text: The main book for the class is "Deep Learning" (2016) by Ian Goodfellow, Yoshua Bengio, and Aaron Courville. The print version of the book has not yet been released, but an online version is available here: http://www.deeplearningbook.org
There will be many required readings from other sources.

Instructor Contact:

Name: Dr. Christopher Kanan
Office Location: Building #76 (CAR), Room 3140
Office Hours: By appointment
Email Address: christopher.kanan@rit.edu

Teaching Assistant:

Name: Tyler Hayes
Office Location: CAR-1150
Office Hours: Mondays from 11am-1pm; Fridays from 1pm-3pm
Email Address: tlh6792@rit.edu

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

Homework:	30%
Basics Quiz:	10%
Project:	30%
Class Presentation:	15%
Class Participation:	5%
Paper Quizzes:	10%

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. You must meet with the TA one week early to go over your slides (40% of your grade)
2. You must come up with three questions about the readings one week early
3. You should fill 60 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.

MyCourses Paper Quizzes: These are short 1-3 question quizzes before some of the paper readings to ensure you got the gist of the paper.

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.

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. The schedule for the project is as follows:

1. **Project Proposal**: The project proposal will clearly state what you plan to do. 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.
2. **Revised Project Proposal**: The revised proposal should take into account the comments received by the instructor and TA. It should be two to four pages long (not including references) and formatted in NIPS or CVPR format. It should be typeset using LaTeX, and submitted as a PDF (under 10MB).
3. **Project Report**: The project report will describe the project, i.e., what you did and the result. It should be four to eight pages long (not including references) and formatted in NIPS or CVPR format. It should be typeset using LaTeX, and submitted as a PDF (under 10MB). Read NIPS and CVPR papers to get an idea for what the style and formatting should be.

Policy on Late Work: Late work will **not** be accepted. You will have 3 to 5 weeks to do each homework assignment. 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 may use MATLAB or 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.

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	8/23		Introduction, Course Overview	Chapter 1 of GBC	Kanan
	8/25		Machine Learning Basics: Simple Classifiers	Chapter 5 of GBC	Kanan
2	8/30		Classifier Evaluation & Challenges	Chapter 5 of GBC	Kanan
	9/1		Logistic Regression and Optimization	Chapter 5 of GBC	Kanan
3	9/6	Homework 1 Assigned	Training Neural Networks Part 1	LeCun, Bengio, & Hinton. (2015), *Nature*. Chapter 6 of GBC	Kanan
	9/7		Training Neural Networks Part 2	Chapter 8 of GBC	Kanan
4	9/13		Convolutional Neural Networks	Chapter 9 of GBC	Kanan
	9/15		Convolutional Neural Networks - Practical Issues, Finetuning, & Video	Krizhevsky et al. (2012) NIPS. Yosinski et al. (2014) NIPS.	Kanan
5	9/20		Dropout & Batch Normalization	Srivastava et al. (2014) JMLR. Ioffe & Szegedy (2015) Batch Normalization	Kanan
	9/22	Project Proposal Due	Machine Learning Basics Quiz Review	None	Kanan
6	9/27	Homework 2 Assigned	Machine Learning Basics Quiz	None	N/A
	9/29	Homework 1 Due (9/30)	Residual Networks - Deeper is Better	He et al. (2016) CVPR. Veit, Wilber, & Belongie (2016) arXiv.	Kanan
7	10/4		CNN Visualization	Yosinski et al. (2015) Mahendran & Vedaldi (2016)	Yousefhussien
	10/6		Fooling CNNs Robustness to Adversarial Examples	Nguyen, Yosinski & Clune (2014) NIPS. Szegedy et al. (2013) arXiv. Goodfellow, Shlens, & Szegedy (2015).	T. Kleynhans Sanketh
8	10/11		No Class - RIT Policy		
	10/13	Revised Project Proposal Due	Object Detection with CNNs	Redmon et al. (2016) YOLO. CVPR. Ren et al. (2015) Faster R-CNN.	A. Chandrashekar Mahshad Hezaveh
9	10/18		Semantic Segmentation	Long et al. (2015) CVPR.	Kemker
	10/20	Homework 2 Due	Recurrent Networks & BackProp Through Time	Chapter 10 of GBC	Prof. Ptucha
10	10/25	Homework 3 Assigned	LSTMs	Karpathy et al. (2015) arXiv. Lipton et al. (2015) arXiv.	Dillon Graham
	10/27		Auto-encoders & Unsupervised Methods	Vincent et al. (2010) JMLR. Chapter 14 of GBC.	Starynska
11	11/1		Image Captioning & Attention DenseCap	Xu et al. (2016) arXiv. Johnson et al. (2016) CVPR.	D. Sharma A. Shringarpure
	11/3		VQA & Attention	Xu & Saenko (2016) ECCV. Fukui et al. (2016) arXiv.	S. Ravi Raagav
12	11/8		Neural Module Networks	Andreas et al. (2016) CVPR.	Krishnamurthy
	11/10		CVPR Deadline - No Class	None	N/A
13	11/15		CVPR Deadline - No Class	None	N/A
	11/17		Neural Turing Machines	Graves, Wayne, & Danihelka (2014) arXiv.	S. Langroudi
14	11/22		Neural Memory Networks	Weston, Chopra, & Bordes (2014) arXiv.	K. Kafle
	11/24		Thanksgiving - No Class	None	N/A
15	11/29		Generative Systems	Goodfellow et al. (2014) NIPS. Nguyen et al. (2016) NIPS.	K. Wilson
	12/1	Homework 3 Due	Reinforcement Learning Overview	Reviews on RL and Q-Learning Mnih et al. (2014) NIPS.	R. Sanchez A. Casson
16	12/6		Deep Reinforcement Learning Part 1	Mnih et al. (2013) NIPS. Minh et al. (2015) Nature.	Sulabh Kumra Aneesh Rangnekar
	12/8	Final Project Report Due	Deep Reinforcement Learning Part 2	Silver et al. (2016) Nature. Hasselt et al. (2015) AAAI.	M. Ali