Course Syllabus

Topics in Digital Media (CSCI 5550G)

Faisal Qureshi

http://faculty.uoit.ca/qureshi

Computer Science, Faculty of Science
University of Ontario Institute of Technology

Fall 2018

Description

This is an introductory graduate course in machine learning and computer vision. The course will focus on machine learning theory and methods for computer vision applications. The course is geared towards students who wish to develop a working knowledge of the recent advances in machine learning, and how these advances have lead to increasing powerful computer vision systems.

Machine learning deals with how to design computer programs that learn from "experience." Residing at the intersection of computer science and statistics, machine learning aims to extract useful information from data (often referred to as the *training data*) and leverages this information to create computer models capable of carrying out useful, non-trivial tasks, such as designing cars that can drive on their own, filters for blocking junk email, diagnostics tools for disease discovery, analyzing images for scene understanding, etc. By many accounts machine learning is the "greatest export" of computer science (and statistics) to other disciplines.

Computer vision deals with processing and analyzing digital images to extract useful properties about the real world. Computer vision, for example, can be used to extract 3D scene structure from a given set of photos, recognize people in images, identify actions in a video sequence, etc. Computer vision has also been used in specialized domains, such as medical imaging, say for analyzing CT scans or MRI photographs, satellite imaging, say for analyzing the health of a an ecosystem, etc. Computer vision has also found wide-spread use in entertainment and gaming industry.

Solving computer vision, it turns out, is a tough problem. Digital images after all are little more than a collection of pixels. Recent advances in machine learning, especially in deep learning, has opened up new avenues for computer vision research. The goal is simple: design algorithms and systems that will enable a computer to "learn to see" by "looking" at example pictures and videos. With this in mind, this course will explore machine learning approaches that have found wide-spread use in computer vision applications.

This course will mix lectures on a selection of topics with paper reading and discussion. The topics are selected to help you understand and implement the papers that you are asked to read, present, and discuss. The first 45 minutes of most classes will be devoted to lectures on one of the selected topics. The remain time will be used for paper presentation and discussion. The course will cover the following topics:

- image formation and camera models;
- optical flow;
- depth analysis;
- action recognition;
- convolution (filtering);
- regression;
- classification;
- clustering;
- dimensionality reduction; and
- neural networks and deep learning.

These topics provides a decent basis for understanding the papers that we plan to read and discuss in this course.

Pre-requisites

The course assumes that students are comfortable with statistics, basic linear algebra, and programming.

I recommend reading Part 1 of "Deep Learning" by I. Goodfellow, Y. Bengio and A. Courville to brush up on linear algebra and statistics. The book is available at here

We will be using Python for the programming part of this course. For Python, I recommend the Anaconda distribution, which comes pre-loaded for nearly all the packages that we will be using in this course. Of course you are welcome to use any variant/distribution of Python that suits you.

The paper also assumes that students are willing to read and comprehend large volumes of technical papers. Furthermore, that students have some experience with technical report writing.

Reading material

You will find the following computer vision books useful.

- *Computer Vision: Algorithms and Applications* by Richard Szleski
- *Computer Vision: Models, Learning, and Inference* by Simon J.D. Prince.

Following books are good resources for machine learning, especially deep learning

- *Neural Networks and Deep Learning: A Textbook* by Charu C. Aggarwal.
- *Deep Learning* by Ian Goodfellow and Yoshua Bengio.
- *Machine Learning: A Probabilistic Perspective* by Kevin P. Murphy and Francis Bach.
- *Understanding Machine Learning: From Theory to Algorithms* by Shai Shalev-Shwartz and Shai Ben-David
- *Pattern Recognition and Machine Learning* by Christopher M. Bishop.
- *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition* by Trevor Hastie and Robert Tibshirani.

These resources will not only help you understand the assigned papers. These resources may prove invaluable for your course projects.

Programming Resources

Here you'll find a number of tutorials showcasing Python use in machine learning. I strongly recommend that you become comfortable with the following four Python packages/environment:

- numpy;
- scipy;
- matplotlib; and
- jupyter notebook.

Instructor

Faisal Qureshi
Email: faisal.qureshi@uoit.net
Office: UA4032
Web: http://faculty.uoit.ca/qureshi

Lectures

Friday, 2:10-5:00 pm, ERC2056
First day of lectures, Friday, September 7, 2018.

Office hours

Friday, 1-2 pm in UA4032, or by appointment.

Grading

- Course project, 40% (**A student needs to get 60% marks in the project to successfully complete the course.**)
  - Proposal

    - Progress report
    - Presentations
    - Technical report
- Participation and interactions, 20%
    - Discussion, readings, QA, class exercises
- Paper presentation and leading discussion, 30%
- One pagers, 10%

Dates

- One pagers
    - Pager 1: Oct. 1, midnight
    - Pager 2: Oct. 15, midnight
    - Pager 3: Nov. 1, midnight
    - Pager 4: Nov. 15, midnight
- Project
    - proposal: Oct. 7, midnight
    - progress report: Nov. 20, midnight
    - final report: Dec. 15, midnight
    - project presentation: last two weeks of classes
- Presentations
    - Through out the term

Lectures

Attendance is mandatory. We will discuss topics in class that are not easily found in any single textbook.

Presentation

Each student will be assigned recent papers to read and present. The student will be responsible for leading the discussion for this paper. **Each student may be assigned to present multiple papers.**

Instructions for the presenter

- Duration, 40 minutes
- Create a slideshow
    - Easy to read
    - Avoid verbosity
    - Use figures, examples
    - Clear and easily understandable structure
    - Practice your talk before the lecture!
- Key questions
    - What does the paper do?
    - What are its limitations?
    - What are its strengths?
    - Is this paper reproducible?
    - How does paper support its key arguments?
    - What software does paper uses?
    - What datasets does paper uses?
    - How does this paper fit with the larger body of literature?

Instructions for the participants

- Read the paper before the lecture
- Be prepared to answer questions
- Be prepared to participate in the discussion
- Provide feedback to the presenter
  - Compliments, suggestions, criticism, thanks

Project

The course project is an independent exploration of a specific problem within the context of this course. A project can be implementation oriented—where a student implements a computer vision system—or application oriented—where a student attempts to solve a problem (of suitable difficulty) by applying machine learning techniques. The project topic will be selected in consultation with the instructor.

Project grade will depend on the ideas, how well you present them in the report, how well you position your work in the related literature, how thorough are your experiments and how thoughtful are your conclusions.

Course project is an individual effort.

Project proposal

- one page (12 pt)
- clear and concise problem statement
- discuss its relevance
- why is it an interesting problem to solve (level of difficulty)
- describe other related approaches
- sketch your approach
- list anticipated difficulties

Progress Report

- one page (12 pt)
- describe the problem you are working on (this should include any feedback that you've received on your project proposal)
- describe your approach in more detail
- summarize your accomplishments to date
- list next steps
- list any problems that you encountered, and how you solved otherwise
- identify any problems that you expect to encounter

Final in-class Presentation

- 15 minutes
- the problem description with a motivation
- a quick overview of related work
- the proposed solution
- a technical description of the solution
- encountered difficulties
- an evaluation
- future work and conclusion

Final Report

For your final project write-up you must use ACM SIG Proceedings Template (available at the ACM website). Project report is at most 12 pages long, plus extra pages for references. Your report must of "publishable quality," i.e., no typos, grammar error.

The final deadline for project report submission is 15th of December, midnight EST. *This is a firm deadline.* You will incur a penalty of 40% if you do not meet this deadline. These strict rules mimic conference submission process:

- a predefined format;
- limited amount of space to explain your ideas and contribution; and
- firm submission deadline.

One pager

A one pager is a summary of the paper (assigned reading for that week). A one pager should not be more than 1 page long (12 pt font). The summary should describe what the paper is doing, its strengths and weaknesses. It should also identify possible future directions for research. One pager is marked according to the following rubric:

- Not submitted, 0 marks;
- Submitted and is of satisfactory quality, 1 mark; and
- Submitted and is of exceptional quality (i.e., raises questions that go beyond the scope of the paper), 2 marks.

Tentative Schedule

- Week 1 - Introduction, image formation, stereo
- Week 2 - Linear regression
- Week 3 - Logistic regression, softmax
- Week 4 - Neural networks
- Week 5 - Image filtering
- Week 6 - Convolutional neural networks
- Week 7 - PCA
- Week 8 - Structure from motion
- Week 9 - Optical flow
- Week 10 - Action recognition
- Week 11 - Presentations
- Week 12 - Presentations

The list of assigned papers will be available after the first week of classes. Please check the course website for details.

Course work submission

Unless otherwise instructed, all course work should be submitted using Blackboard.

Remarking

It is extremely important that all work is fairly graded. Please submit a remark request by email within 5 days of receiving the grade. The email must contain the reasons for which you think the work should be remarked. *Please note that a remark may result in a lower grade.*

Late submission policy

The penalty for a late submission is 10% per day. An assignment or project will get a **zero** if submitted more than 48 hours after the submission deadline. A doctor's note will be needed to avoid late submission penalty.

Email traffic

The instructor and the TA will make every effort to respond to emails in a timely manner; however, it may take up to two working days to respond to an email. It simply means that emails sent right before a deadline may not be answered in time. Urgent emails may be sent to "faisal.qureshi@uoit.net" with the subject line "csci 5550g - fall 2018".

Discussions

Appropriate use of discussion groups include clarification of lecture material and assignments and other concerns and comments about the course that might of general interest to course participants. Please do not post assignment solutions to the discussion groups.

Collaboration

I encourage you to work together when discussing assignments/projects; however, it doesn't mean that you should share your written solutions or that you submit someone else's work as your own.

Course evaluation

It is important that every student participates in course evaluations. Course evaluations, which are completely anonymous, provide extremely useful feedback to the instructor and the TA, helping improve the course.

Important dates

UOIT academic calendar that lists important dates (and deadlines) is available at here.

Academic integrity

Assignments and tests must be strictly individual work. UOIT takes academic dishonesty very seriously. Please read and understand UOIT's policy on academic integrity available here

Accessibility

Students with disabilities may request to be considered for formal academic accommodation in accordance with the Ontario Human Rights Code. Students seeking accommodation must make their requests through the Centre for Students with Disabilities in a timely manner, and provide relevant and recent documentation to verify the effect of their disability and to allow the University to determine appropriate accommodations. More information about Student Accessibility Services (SAS) is available here.

Freedom of Information and Protection of Privacy Act

UOIT is governed by the Freedom of Information and Protection of Privacy Act ("FIPPA"). In addition to providing a mechanism for requesting records held by the university, this legislation also requires that UOIT not disclose the personal information of its students without their consent. FIPPA's definition of "personal information" includes, among other things, documents that contain both your name and your Banner ID. To ensure that your rights to privacy are protected, the I encourage you to use only your Banner ID on assignments or test papers being submitted for grading (the exception to this rule are midterm and final exams, since these are returned individually). This policy is intended to prevent the inadvertent disclosure of your information where graded papers are returned to groups of students at the same time. If you still wish to write both your name and your Banner ID on your tests and assignments, please be advised that UOIT will interpret this as an implied consent to the disclosure of your personal information in the normal course of returning graded materials to students. Please contact the UOIT Chief Privacy Officer at *accessandprivacy@uoit.ca* for more information.

Sexual Violence Policy

UOIT is committed to the prevention of sexual violence in all is forms. For any UOIT student who has experienced Sexual Violence, **UOIT can help**. UOIT will make accommodations to cater to the diverse backgrounds, cultures, and identities of students when dealing with individual cases.

If you think you have been subjected to or witnessed sexual violence:

- Reach out to a Support Worker, who are specially trained individuals authorized to receive confidential disclosures about incidents of sexual violence. Support Workers can offer help and resolutions options which can include safety plans, accommodations, mental health support, and more. To make an appointment with a Support Worker, call 905.721.3392 or email supportworker@uoit.ca
- Learn more about your options at: www.uoit.ca/sexualviolence