Department of Computer Science, University of Illinois at Chicago

Spring 2018

CS 594 — Advanced Machine Learning (CRN: 38551)

Course Syllabus

Although this course is listed as CS 594, it will count as a standard CS 500-level class, because it will become CS 512 pending the university's approval. In particular:

1. Taking this course in Spring 2018 will contribute to your course requirement in the same way as other standard 500-level courses (excluding 594, 590, 595-599).

2. For PhD students, it can be used to fulfill your WCP requirement.

Time: TR 5 – 6:15 PM

Classroom: **LH 210**

URL: via Blackboard / Piazza (https://piazza.com/uic/spring2018/cs594ml/home)

Instructor: Xinhua Zhang

Office: 1237 SEO

Phone: 312.413.2416

E-mail: zhangx@uic.edu (preferred)

Office hour: 4-5 PM Wednesday. Contact me by email if you would like to meet at other time.

TA: Vignesh Ganapathiraman (vganap2@uic.edu)

Office hour: 2-3 PM Tuesday **Office:** SEL-W 4029

Prerequisites

MATH 310/320, STAT 401, CS 251, CS 412; or consent of the instructor.
For graduate students, these prerequisites are only advisory.

A self-evaluation quiz (3%) will be given in class in the fourth lecture. For example:
Math: https://www.umiacs.umd.edu/~hal/courses/2013S_ML/math4ml.pdf
Stats: http://www.cs.princeton.edu/courses/archive/spring07/cos424/scribe_notes/0208.pdf

Course Goals:

- Students will be able to have an in-depth understanding of the principle and characteristics of advanced machine learning task settings (e.g., structured prediction, distributed optimization, deep learning for complex data).

- Students will be able to scale machine learning techniques to big datasets, by leveraging new structures in the data and new computational tools that emerge even after the completion of the course.
- Students will be able develop and analyze novel problem formulations and machine learning techniques that adapt to data analysis problems emerging in new applications.

Restrictions

Restricted to students in the following colleges/schools: Engineering or Graduate College. Capped to 35 students. Students who have taken CS 594 (Advanced Machine Learning) in Fall 2016 are NOT eligible for registration.

Credit Hours

4 graduate hours

Textbooks (Required)

[**BV**] Stephen Boyd and Lieven Vandenberghe. Convex Optimization. Cambridge University Press, 2004. PDF available at https://web.stanford.edu/~boyd/cvxbook

[**Murphy**] Kevin P. Murphy. Machine Learning: A Probabilistic Perspective. MIT Press, 2012. Free book available electronically via UIC library

[**GBC**] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, Deep Learning. MIT Press, 2016 Free book at: http://www.deeplearningbook.org Single PDF (with slight difference) for download (65 MB):

https://uofi.box.com/s/nv6idmcb5w9pyqay4966cxwdaaz547jy

[**Jordan**] Michael I. Jordan. Graphical Models. Unpublished lecture notes (Ch 2, 3, 4, 9, 10, 11)

Also used will be recent research papers and excerpts of relevant background material from available textbooks, and supplemental notes for specific topics.

Tentative Schedule (13 weeks)

1. (0.5 week) Introduction and Structured prediction – Structured SVMs
2. (3.5 weeks) Convex optimization (convex sets and function, convex formulation)
3. (1 week) Directed and undirected graphical models, Markov/Conditional Random Fields
4. (1.5 weeks) Graphical model inference: message passing and sampling (rejection, importance, MCMC)
5. (1 week) Unconstrained optimization, stochastic and parallel optimization

6. (1 week) Learning algorithms – Expectation maximization, contrastive divergence
7. (3.5 weeks) Deep learning: convolution networks, auto-encoder, deep generative models, understanding the properties and philosophy of non-convex optimization, and generalization performance
8. (1 week) Convex duality and advanced optimization algorithms (time permitting)

Tutorial on ACER cluster: 30 min on Jan 30

Quiz on math background: 50 min on Jan 25 (followed by 20-min introduction on AWS given by the TA)

Mid-term: 75 minutes
Project presentation: week 15

Teaching evaluation at mid-term (~week 8)

Grading and evaluation (tentative)

	%	date /deadline	Note
Convex optimization	0 %		Work on sample questions in preparation of the mid-term exam
Computing platforms	0%		Get familiar with AWS and ACER
Assignment 1	20 %	Fri Week 7	A project on conditional random fields (CRFs) with inference and learning
Mid-term	12 %	Week 9	Optimization and graphical model (one cheat sheet)
Assignment 2	20 %	Week 12	Continue the CRF project with stochastic optimization and parallel optimization using MPI on UIC's ACER cluster
Course Project	25 %	Week 16	Deep learning project (open) Oral presentation in week 15
Final exam	20 %		Cumulative (one cheat sheet)

Background quiz: 3%

Assignments and course project are in groups of up to 3 students. Midterm and final exams are individual.

Assignments 1 and 2 aim to develop students' ability to apply machine learning tools to different data analysis problems and perform thorough experimentation, benchmarking, and empirical analysis. Tasks will be specified in Assignments 1 and 2.

Mid-term and final exams will assess students' understanding of different machine learning tasks and techniques, especially their theoretical properties and underlying principles.

Course projects, including reports and presentations, aim to develop students' ability to create novel machine learning techniques for different real-data analysis problems by leveraging their specific structures and computational resources at hand.

The final exam will be comprehensive and cumulative.

1. The mid-term will be **75 minutes** with **closed book**. It will be given in the classroom during class time. Therefore, **no make-ups** will be given. Partial grading will be used.

   The mid-term covers everything from the first **eight** weeks (tentatively). Sample questions will be given.

2. In Assignment 1 and 2 you will build a classifier which recognizes "words" from images. This is a great opportunity to pick up practical experiences that are crucial for successfully applying machine learning to real world problems, and evaluating their performance with comparison to other methods.

3. Assignment 1 will set up the experiment background and the task will be: implement probabilistic inference algorithms such as message passing, train a CRF model using off-the-shelf solvers, compare with max-margin methods, and test the performance under transformations. Although you may use any language of your choice, we will only provide utility code in Matlab/Octave. Grading will be based on your results.

4. In Assignment 2, you will have a great opportunity to learn parallel optimization. We will use PETSc which is in C/C++ and uses MPI for message passing. All low-level constructs have been encapsulated into high level linear algebra commands and optimization solvers (PETSc package: http://www.mcs.anl.gov/petsc). You will run experiments on UIC's cluster: ACER. If your own machine gets multiple cores (highly likely) AND you are happy to install MPI and PETSc by yourself (caution: highly challenging for newbies, but we will provide *some* helpful guides), you may do experiments on your own machine.

   Assignment 2 will also involve stochastic optimization with approximate inference based on sampling.

5. The course project will be on deep learning, using the AWS platform. You will have the opportunity to design your own project. We plan to form 10 groups, and each group will have 15 minutes to present their projects orally in the two sessions on week 15 (12 min talk and 3 min Q&A). Only *initial* results will be needed for oral presentation, and the detailed results can be submitted in the final report in week 16.

All the five evaluations will be graded out of 100, and their weighted sum will be used to determine the final letter grade (A/B/…) **based on threshold**:

A: 75 - 100
B: 60 ~ 74
C: 40 ~ 59

D: 0 ~ 39

I reserve the right to **raise** your letter grade.

Communication, feedback, and discussion:

1. The web page on Blackboard will contain all materials relevant to the class, syllabus, assignments, lecture notes etc. You can also see you own grades.

2. For general announcements and notifications, I will send email to the whole class via Blackboard. Please check your email frequently, especially around deadlines (homework and exams).

3. Piazza will be used for Question and Answer. The system is highly catered to getting you help fast and efficiently from classmates and myself. Rather than emailing questions to me, I encourage you to post your *technical* questions on Piazza. If you have any problems or feedback for the Piazza developers, email team@piazza.com.

   Find our class page at: https://piazza.com/uic/spring2018/cs594ml/home

4. If you have any personal or non-technical questions, please send an email directly to the instructor (zhang@uic.edu) or the TA (if instructed to do so, e.g. dispute of grading).

General Policies on homework (project) and exams

1. **Late submissions**: Late submissions will not be accepted in any case, unless there is a **documented** personal emergency. Arrangements must be made with the instructor as soon as possible after the emergency arises, preferably well before the homework due date.

   Advice: If for whatever reason you don't manage to finish an assignment, hand in what you have. Partial credit will be given.

2. **Statute of limitations**: **Three weeks!** No grading questions or complaints — no matter how justified — will be listened to **three** weeks after the item in question has been returned.

3. **Group collaboration**: all members of each group should make nearly the same level of contribution to each project. So in a group of three, each member contributes 1/3 to Assignment 1, 1/3 to Assignment 2, and 1/3 to the course project. It is **not** allowed that one works on Assignment 1 alone, one on Assignment 2 alone, and one on the course project. All members should be responsible for the whole submission of the team, not only his/her own contributed part. You are supposed to understand the work of your teammates inside out, and be able to answer questions when asked. If one member plagiarized, then **all members** of the team will receive the **same** penalty.

Policy on Academic Integrity

Academic dishonesty will not be tolerated. Please see the CS department policy below on the topic; this policy specifies penalties for violations.

What is academic dishonesty? To hand in any work which is not 100% the student's creation, unless you are explicitly allowed to do so. Thus:

Exams: All work on all exams must be individually performed.

Projects: no team may give any other team any portion of their solutions or code, through any means. Students are not allowed to show each other any portions of code or homework, unless they are on the same team.

Important Note: almost every semester somebody is caught red-handed and as a consequence fails the class. Isn't it better to get a B or a C than an F?

CS department policy on academic dishonesty

The CS Department will not tolerate cheating by its students. The MINIMUM penalty for any student found cheating will be to receive an F for the course and to have the event recorded in a department and/or College record. The maximum penalty will be expulsion from the University. Cheating includes all the following, though this is not a complete list:

- Copying or any other form of getting or giving assistance from another student during any test, quiz, exam, midterm, etc.
- Plagiarism–turning in writing that is copied from some other source.
- Obtaining solutions to homework by posting to the Internet for assistance, purchasing assistance, obtaining copies of solutions manuals for instructors, and obtaining copies of previous year's homework solutions.
- Computer programs: Any time you look at another student's code, it is cheating. (Exception: If you are EXPLICITLY told that you may do so by the instructor.)

For computer programs, if for some reason we cannot determine who copied from whom, we may, at our discretion, give failing grades to both students.

It is the responsibility of all engineering and computer science professionals to safeguard their company's "trade secrets." An employee who allows trade secrets to be obtained by competitors will almost certainly be fired. So, YOU are responsible for making sure that your directories have permissions set so that only you can read your files, for being sure to log out at the end of working in the computer lab, etc.