Department of Computer Science, University of Illinois at Chicago Spring 2022

CS 512 — Advanced Machine Learning

Course Syllabus

Week 1-2: Zoom for Sync session and office hour (both the instructor and the TA).

Instructor:	Prof. Xinhua Zhang (zhangx@uic.edu)
Teaching Assistant:	Omid Memarrast (omemar2@uic.edu)
Time:	MW 4:30 – 5:45 PM
Venue:	SES 130 (online for the first two weeks)
Instructor's Office hour	3-4 PM, Thursday (or by appointment) North End, Level 3, Richard Daley Library (also on **Zoom**)
TA's Office hour	12-1 PM, Tuesday SELW 4211 (also on **Zoom)**

Course links (bookmark whenever appropriate):

Piazza	https://piazza.com/class/kxlyrpi2tjo6ls You can register using your UIC netid with no need of instructor's approval.
Master Schedule	https://piazza.com/class/kxlyrpi2tjo6ls?cid=6 (**bookmark it!**) First register on Piazza as above. This schedule is pinned on top of the left list.
Gradescope	https://www.gradescope.com/courses/349986 Submission of written assignment + three labs. You can log in with School Credentials or via Blackboard (left banner).
Zoom	https://uic.zoom.us/j/81170677941?pwd=V3ovMi9TYm1GZjlXbWJLdzU2c2RHUT09 If needed, the password is: j19JvcX5 and the meeting ID is 811 7067 7941 This is for **sync sessions** (week 1-2) **and office hours** (instructor's **and** TA's).
Blackboard	Announcement, and course project
Google Drive	https://drive.google.com/drive/folders/1WON_jsQbqR-ae91d-CkY0ASsQIY-Byc5?usp=sharing Specification of written assignment and three labs, with rubrics and solutions. To access it, make sure you log in with your UIC netid (not your personal).

Use of Online Tools (see URLs on the first page, or just click)

What	Where	Who can access	
		Week 1 and 2	Week 3 and onwards
Slides, tutorials, and other documents	**Piazza**	Anyone[1]	Registered + Auditing
Technical discussion, Q&A	**Piazza** [2, 3]	Anyone	Registered + Auditing
Announcement	Blackboard [4]	Anyone	Registered + Auditing
Synchronous session with recording	**Zoom** [5]	Anyone	Registered + Auditing
Whiteboard for sync sessions	**Onenote**	Anyone	Registered + Auditing
Office hour (week 1-2, not recorded)	**Zoom**	Anyone	Registered + Auditing
Submit written assignment + 3 labs	**Gradescope**	N.A.	Registered only
Specification of written assignment + 3 labs	**Google Drive**	N.A.	Registered only

for collaboration within groups (but not for instruction)		
Teams, Slack	Github, Bitbucket	Box.com, Onedrive, Google drive

1. Anyone refers to anyone with a UIC netid. Contact the instructor if you do not have it.
2. Piazza is highly catered to getting you help fast and efficiently from classmates and the instructing team (professor and TA). Rather than emailing questions to the instructing team, you are encouraged to post on Piazza your **technical questions**, general questions about the course content, assignments, grading rubric, etc. If you have any problems or feedback for the Piazza developers, email team@piazza.com.
3. If you have any personal or **non-technical question** such as medical consideration, please send an email directly to the instructor (zhangx@uic.edu) or the TA (if instructed to do so, e.g., disputing suspection of plagiarism). To request re-grading, use the function on Gradescope.
4. For general announcements and notifications, I will send emails to the whole class via **Blackboard**. Please check your email frequently, especially around deadlines (homework and exam). The message will also be recorded on Blackboard in the announcement section.

   It is your responsibility to check emails frequently (at least once a day).
5. The synchronous sessions will be recorded on Zoom. **You are required to turn on your camera** whenever possible. To access the Zoom recordings, check out the master schedule on Piazza and find the column "Sync Rec".

Software and hardware

1. Python for Lab 1-3. You can use any language for the course project.
2. Matlab might be used for some illustrations. You can get free Matlab from Webstore https://webstore.illinois.edu/shop/product.aspx?zpid=4131.
3. Use VS Code + Google Colab + Github. Tutorial that works for Linux, Windows, and MacOS.

4. You are required to use collaborative tools for labs and course project, including
    a. Slack or Teams for messaging. UIC does not have a license for slack, but the free version might be sufficient.
    b. Github or Bitbucket for code maintenance (free with their education packs).
    c. Box, Onedrive, or Google drive for file sharing (all free from UIC).
5. Scanner on phone/pad: scan your answer, **NOT just take a photo**. Use a scanner app on your mobile device, e.g., **CamScanner** (https://www.camscanner.com/). Register using your UIC email address, and it will remove a watermark. Feel free to use similar apps.
6. In week 1-2, get a reliable internet connection at home for the sync session (e.g., by using wired connection, closer to Wi-Fi router).
7. Get a mobile device (phone or pad), and a good headphone with a good mic. Test them before the semester and buy a new one if not good. A wired headset should be good enough.

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 with solution is available on Piazza, along with some math background readings. Do it yourself and you do not need to submit it.

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, convex 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.

Textbooks (all free)

Required (only a few chapters from each book):

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

[**Jordan**] Michael I. Jordan. Graphical Models. Unpublished lecture notes (Ch 2, 3, 4, 9, 10, 11). PDF available on Piazza under the Resources tab -> Readings.

[**Murphy**] Kevin P. Murphy. Machine Learning: A Probabilistic Perspective. MIT Press, 2012. https://probml.github.io/pml-book/book1.html

[**GBC**] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, Deep Learning. MIT Press, 2016.

Free book at: http://www.deeplearningbook.org Single PDF (slightly different from the online version) for download (65 MB): https://uofi.box.com/s/nv6idmcb5w9pyqay4966cxwdaaz547jy

Reference only:

[**Murphy+**] Kevin P. Murphy. Probabilistic Machine Learning: Advanced Topics. To be available in January 2022. https://probml.github.io/pml-book/book2.html

[**Dive**] Dive into deep learning. https://www.d2l.ai

[**PyTorch**] Deep Learning with PyTorch.
https://pytorch.org/assets/deep-learning/Deep-Learning-with-PyTorch.pdf

We will also use recent research papers and excerpts of relevant background material from available textbooks, and supplemental notes for specific topics.

Programming Language

Lab 1-3: Python only, including numpy and PyTorch. Course project: no restriction.

Evaluation Criteria

	%	**tentative** deadlines	Note
Written assignment	6%	Week 5	Some written questions for individual work on graphical models
Lab 1	11 %	Week 5	A project on conditional random fields with inference and learning (Python/numpy)
Lab 2	11 %	Week 8	A project on deep learning using PyTorch, focusing on convolution neural network and image processing
Mid-term (75 min in class)	16 %	Week 9	Basic deep learning and graphical model. You can bring one cheat sheet.
Lab 3	11 %	Week 13	A project on adversarial learning with recurrent neural networks
Convex optimization			Work on sample questions in preparation of the exams. No need to submit.
Course Project	20 %	Week 16	Deep learning project (open, free language) **Oral presentation in week 15** (?? minutes per group)
Final exam	25 %		**Cumulative and comprehensive** (one cheat sheet)

Three labs and course project are in groups of 3-4 students. We will make 10 groups. Lab 1-3 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 Lab 1-3.

One written assignment is for individual work. It is on the conceptual understanding of graphical models and inference.

Midterm and final exams are individual. They will assess students' understanding of different machine learning tasks and techniques, especially their theoretical properties and underlying principles.

Course project, including reports and presentations, aims 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.

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

A: 80 ~ 100
B: 70 ~ 79.9
C: 60 ~ 69.9
D: 50 ~ 59.9
F: 0 ~ 49.9

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

Detailed notes:

1. The mid-term will be **75 minutes** with **one letter-sized cheat sheet (both sides)**. 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 will cover graphical models and basic deep learning (tentatively). Sample questions will be given.

2. In Lab 1 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. In particular, you will 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. Python (numpy and scipy) will be used, and skeleton code will be provided. Grading will be based on your results.

3. In Lab 2, your work in Lab 1 will be extended to deep learning. You will reimplement CRF using the modules in PyTorch, add a convolution layer implemented by yourself, and test stochastic optimization with approximate inference based on sampling. You can use GPU on AWS or Google colab. PyTorch must be used for this project. Grading will be based on your results and your code.

4. In Lab 3, you will work on an advanced topic of deep learning: virtual adversarial training on recurrent neural networks, applied to semi-supervised text classification. You will implement the algorithm in PyTorch, and benchmark several competing training algorithms and deep architectures.

5. The course project will be on deep learning. 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 on Monday of week 17.

6. The final exam will be **120 minutes** with **one letter-sized cheat sheet (both sides)**. It will be comprehensive.

Course Policies

Academic dishonesty will not be tolerated. Please see the CS department policy below on the topic; this policy specifies penalties for violations. Academic misconduct will be handled according to UIC's Student Disciplinary Policy: http://dos.uic.edu/conductforstudents.shtml

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.

Collaboration Policy

1. **Homework and course project**: all members of each group should make nearly the same level of contribution to each project. So in a group of four, each member contributes 1/4 to Lab 1, 1/4 to Lab 2, 1/4 to Lab 3, and 1/4 to the course project. It is not allowed that one works on Lab 1 alone, one on Lab 2 alone, etc. 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.

   No team may give 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.

   Discussion of homework assignments and solutions with other students (including outside the group) is permitted, including all assignments and the course project. However, for the written assignment, each student must

   1) submit his or her own write-up for the assignment and fully understand what he or she submits; and

   2) be prepared to explain his or her homework assignment submissions to the instructor and teaching assistant if his or her "full understanding" is in doubt.

2. **Exams (midterm and final): All work on both exams must be individually performed.**

Policy for Missed or Late Evaluation

1. **Late submission**s: 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. **Missed exams**: Missed exams cannot be made up unless there are extenuating circumstances (death, severe illness, etc.) **and** the student has e-mailed instructor **in writing <u>before</u>** the exam. A mark of 0 is earned for the exam if the above policy is not adhered to.

Religious Holidays

> Students who wish to observe their religious holidays shall notify the faculty member by the tenth day of the semester of the date when they will be absent. The faculty member shall make every reasonable effort to honor the request, not penalize the student for missing the class, and if an examination or project is due during the absence, give the student an exam or assignment equivalent to the one completed by those students in attendance. If the student feels aggrieved, he/she may request remedy through the campus grievance procedure. http://oae.uic.edu/docs/ReligiousHolidaysFY20152017.pdf

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.

A (Very) Tentative weekly schedule

Week	Topic	Readings	Assessment
1	Course introduction and feedforward networks	Chapter 6 of GBC	
1	Feedforward networks and back-propagation	Chapter 6 of GBC	Tutorial on deep learning out
2	Convolutional network (MLK holiday on Monday)	Chapter 9 of GBC	
3	Probability and graph basics	Chapter 2 and 10.1.1-10.1.4 of Murphy	Lab 1 + written assignment out
3	Directed graphical models (Bayes nets) and Undirected graphical models	Chapter 2 of Jordan	Tutorial on graphical models out
4	Max-likelihood learning on graphical models and Conditional Random Fields (CRFs)	Chapter 8 of Jordan (excluding section 8.3.4) Chapter 19.6 of Murphy (optional 9.3-9.4 of Murphy)	
4	Variable elimination	Chapter 3 of Jordan Chapter 19.1-19.4 of Murphy	
5	Sum-Product algorithm (belief propagation)	Chapter 4 of Jordan Chapter 20 of Murphy	
6	Mixture model and EM algorithm	Chapter 9 and 10 of Jordan	Lab 1 + written assignment due
6	Recurrent network	Chapter 10 of GBC	
7	Neural network optimization and practicum	Chapter 8 of GBC	Lab 2 out
7	NN regularization	Chapter 7 of GBC	
8	Adversarial deep learning	Papers and tutorials	
8	Convex sets	Chapter 2 of BV	Tutorial on convex optimization out
9	Convex function	Chapter 3 of BV	
9	**Midterm (covering everything up to adversarial learning)**		
10	Convex problems and optimization	Chapter 4 and 9.1-9.5 of BV	Lab 2 due, Lab 3 out
10	Duality	Chapter 5 of BV	Project proposal due (Sunday)

Spring break			
11	Variational inference and sampling	Chapter 21 and 23 of Murphy	
12	Generative models: Variational autoencoder + GAN + normalizing flow	Chapter 14 of GBC Papers and tutorials	Lab 3 due
13	Graph neural network and certification of its robustness	Papers and tutorials	
14	Deep reinforcement learning and imitation learning	Papers and tutorials	
15	Course project presentation		