ECS 289G

Advanced Deep Learning

Spring 2023

Syllabus

Course Website: http://www.ifmlab.org/courses.html
Syllabus: https://docs.google.com/document/d/189lK0WRaJqIHFSg8HrkrhnsM_9ULLHOX
Course Schedule: https://docs.google.com/spreadsheets/d/1nPFJScmS9t6R2qLEgY5SeBMo0PCMHR1CASnGC47hFE8
Reference Papers (Optional): https://docs.google.com/spreadsheets/d/1qpPQI9rnHjR-xipPRkojuVgSAnJ1VBcv
Discord Group Link: https://discord.gg/BqgHyCGCnH

Instructor: Prof. Jiawei Zhang
Email: jiwzhang@ucdavis.edu
Office Hours: Tuesday 12:10-13:10PM (right after class)
Zoom Link: https://ucdavis.zoom.us/j/92918333924?pwd=YXNvUE1WVk5yOUVYam5vMVhhQUl4UT09

Course Teaching Assistant: Xinghao Xiang
Email: xhxiang@ucdavis.edu
Office Hours: TBD
Zoom Link: TBD

Course Information:

- **Date**: April 3 – June 8, 2023
- **Time**: T/R 10:30-11:50AM
- **Physical Venue**: CRUESS 107

Course Description:

The course on Advanced Deep Learning focuses on selective areas of importance about deep learning. Deep learning has been one of the hottest topics in AI studies, and techniques developed in the research hold substantial impacts in many important applications. Selective topics will be covered in the Advanced Deep Learning course.

Topics Covered:

- Deep Learning & Optimization
- Computer Vision
- NLP and Text Mining
- Graph Neural Nets
- Robotics
- misc.

Prerequisites:

You are expected to know basic deep learning models and algorithms already before taking this course. You may need to have background knowledge in Data Structure, Algorithm, Discrete Mathematics. You will also need to be familiar with basic Linear Algebra, basic Statistics, and can master at least one programming language and have programming experiences.

Textbook:

No required textbook, but some Reference Books are recommended.

[1] Ian Goodfellow and Yoshua Bengio and Aaron Courville. Deep Learning, 1nd edition. MIT Press.

[2] Pang-Ning Tan, Michael Steinbach, Vipin Kumar. Introduction to Data Mining. 2nd Edition. Pearson.

[3] Kevin P. Murphy. Machine Learning: A Probabilistic Perspective. MIT Press.

Course Format:

The objective of this course is to familiarize students with the latest research topics related to deep learning. Course activities include 1) paper reading and paper review; 2) paper presentation and discussion; and 3) research oriented paper writing.

- **Paper Reading and Paper Review**: Each class will discuss one academic paper. Before class, students should read the paper to be presented in class, and write a short review (no longer than 1 page) for the paper. The review should cover: (a) a summary of the paper; (b) 3 strong points of the paper; (c) 3 weakness of the paper; (d) potential ideas of future works; and (e) questions about the paper you would like to ask the presenter. Each individual student needs to submit the review report **before** class starts to the Canvas system. Late submissions will not be accepted. Some external reference papers/articles can be provided for students to understand the papers presented in class (no need to present or write reviews for such external reference papers).
- **Paper Presentation and Discussion**: Students may form a group (max size of 4) to present a paper, and students will not need to write the reviews of papers you choose to present. Paper presentation is expected to take about **1 hour** to cover (a) background knowledge, (b) problem description, (c) definitions, (d) ideas, (e) proposed methods and techniques, (f) experiments, (g) results, (h) related works, and (i) potential future works, etc. Very short and low-quality presentations will be penalized. Additionally, we will have 15-20 minutes for in-depth discussion about the paper and Q&A with the audiences. Students can have Q&A during and after the presentation. Presenters need to send the slides **before** your presentation day to the TA to post to all students for review.
- **Paper Proposal and Paper Writing**: Students may form a group to work on the research project (max size of 4, group members should be the same as your presentation group). Each group needs to finish an independent research-oriented academic paper in this course, based on the problems studied in this course about deep learning. The paper should be original work, not recycled published/submitted/on-going work with another faculty or classes. Students need to submit a paper proposal (no longer than 2 pages, ACM sample-sigconf double-column format) at the mid of the class to provide the planning for the paper. By the end of this quarter of this course, students need to submit the final paper (10 pages in ACM sample-sigconf double-column format). Students need to submit the proposal and final paper before the deadlines. The submitted paper will be graded according to the novelty, technical depth, completeness and research value (low-quality and incomplete papers will be penalized with grade deducted).

Paper List and Schedule:

- The course presentation paper list and schedule is available at https://docs.google.com/spreadsheets/d/1nPFJScmS9t6R2qLEgY5SeBMo0PCMHR1CASnGC47hFE8
- Each student may need to sign **1** paper you are interested in to present in class based on the first-come-first-serve rule. (If you don't plan to register for this course or you plan to drop it, please don't reserve any papers in the schedule spreadsheet.)
- External reference papers/articles will not be presented in class, and students can read them after class if you are interested in them. The external reference list is available at https://docs.google.com/spreadsheets/d/1qpPQI9rnHjR-xipPRkojuVgSAnJ1VBcv/edit?usp=sharing&ouid=112300113196642557018&rtpof=true&sd=true

Course Grading Policy:

- **In-class presentation: 30%**. Powerpoint slides need to be sent to the TA via email before 11:59PM (midnight) prior to your presentation day, so the TA can help post the slides as other students' references). Copying existing presentations or slides from the web is regarded as plagiarism.
- **Course paper: 50%**. An original work on deep learning. Not recycled published/submitted/on-going work with other faculties or classes.
    - **Group Formation: 0%**, which will be due by April 14, 2023, 11:59PM (Friday, midnight).
    - **Paper proposal: 10%**, which will be due on May 5, 2023, 11:59PM (Friday, midnight).
    - **Final paper: 40%,** which will be due on June 9, 2023, 11:59PM (Friday, midnight).
- **Course participation and Q&A: 20%**. An individual summary/review of each in-class discussion paper needs to be submitted **before** each class starts. During the class, presenters and audiences can have Q&A with the pre-prepared questions in the review report.
- **Final Grade**
    - A+ [100, 97], A: (97-93] A-: (93-90];
    - B+ (90-87], B: (87-83], B-: (83-80];
    - C: (80-70];
    - D: (70-60];
    - F: (60-0].

Late Submission Policy:

- Late paper review submission will not be accepted.
- Late presentation slides submission will not be accepted.
- Late proposal submission will not be accepted.
- Late final paper submission will get your grade * $1(t <= 24)e^{(-\frac{ln(2)}{12}t)}$, if you are t hours late.

Resources:

Paper Review Template (https://docs.google.com/document/d/1jVkhizzDj9z8kUIrp-mgUFoj_ccqZ01hBOZs3TGp204/edit?usp=sharing)

Group Formation Template (https://docs.google.com/document/d/1Rlcahlqon-Nu7qysxI1QozmMkb7gK1LCmQYOUgI6yGs/edit?usp=sharing)

Proposal and Paper ACM Template (https://www.acm.org/publications/proceedings-template)

- Latex Template (please use the sample-sigconf.tex format)

Deep Learning Toolkit

- PyTorch
- TensorFlow
- Keras

Programming Language and IDE

- Python Tutorial

- PyCharm

Deep Learning Courses in Other Universities

- Stanford: CS230 Deep Learning
- MIT: 6S191 Deep Learning
- CMU: 11-785 Deep Learning
- UC Berkeley: CS 182 Deep Neural Networks

Other Online MOOC Courses:

- Data Mining (https://www.coursera.org/specializations/data-mining)
- Machine Learning (https://www.coursera.org/learn/machine-learning)
- Deep Learning (https://www.coursera.org/learn/neural-networks-deep-learning)