ECS 189G-001

Deep Learning

Winter 2022 *Syllabus*

Course Website: http://www.ifmlab.org/courses.html

Course Schedule: https://docs.google.com/spreadsheets/d/1Tfh1 p4B5fkzSgUFTSAk60hlzCesVppz

Course Zoom Link: https://ucdavis.zoom.us/j/96521564237?pwd=S3ovNVI4ck95U3BUdUFNcTllR0o4dz09

Youtube Channel (Optional): https://www.youtube.com/channel/UCtDX6XIR5Q sCuO9tzYiAVA

Discord Discussion: https://discord.gg/RjaNYzHb

Instructor: Prof. Jiawei Zhang Email: jiwzhang@ucdavis.edu

Online Office Hours via Zoom: https://ucdavis.zoom.us/j/96273718324

Office Hours: 10:30AM-11:30AM, Tuesday. (Or other online office hours upon appointments)

Course Teaching Assistants:

#1: Xiao Liu (Q&A, HW, Exam) Email: xioliu@ucdavis.edu

Online Office via Zoom: https://ucdavis.zoom.us/i/94854180485?pwd=c0JWcGd1TkU3ZDhzVVJJaUVYMnV1UT09

Meeting ID: 948 5418 0485

Passcode: 872767

Office Hours: 10:30AM-11:30AM, Thursday

#2: Haopeng Zhang (Project) Email: hapzhang@ucdavis.edu

Online Office via Zoom: https://us04web.zoom.us/i/7330623325?pwd=V1ISRkJmaERZTG4xK1hKODdKcEtDUT09

Meeting ID: 733 062 3325

Passcode: 123123

Office Hours: 10:00AM-11:00AM, Friday

Course Information:

• **Date**: January 3, 2022 – March 11, 2022

Time: T/R 9:00AM-10:20AM
Physical Venue: Hoagland Hall 168

• **Zoom Venue**: https://ucdavis.zoom.us/j/96521564237?pwd=S3ovNVI4ck95U3BUdUFNcTllR0o4dz09

Online Youtube Channel: https://www.youtube.com/channel/UCtDX6XIR50_sCuO9tzYiAVA

Schedule and Slides:

• The course schedule and slides are available at https://docs.google.com/spreadsheets/d/1Tfh1_p4B5fkzSgUFTSAk60hlzCesVppz

Course Description:

Do you know how the recent Turing Awards were given to the deep learning field? For what contributions? Deep learning is one of the hottest topics and technology with edge-cutting applications in various areas currently, e.g., search engine, e-commerce, self-driving vehicle, robotics, gaming, and highly like the coming metaverse. What breakthroughs turned the rather mathematical idea into reality? What are the classic milestones that pioneered and shaped the landscape of the deep learning field? What are the state-of-the-art methods just coming out yesterday?

Do you know that, although deep learning has been widely used in our real-life products and services now, researchers are even more fascinated by **general deep learning issues**? The new challenges naturally arise in the junction of deep learning vs classic machine learning, deep vs shallow, general-purpose AI vs expert system, even human vs robots, and many more novel contexts. What are the current topics and future agenda?

To build the essential foundation as a stepping-stone to deep learning research, this course exposes students to various deep learning concepts, methods, and applications. We will broadly explore the classic as well as more recent research work.

Course Objectives:

- To learn the basic principles of designing and implementing a deep learning model for specific application problems;
- To understand in depth how to train, optimize, tuning deep learning models with real-world data, including images, text, graphs, and time-series, etc.;
- To realize the importance of social, technological, and ethical issues involved in deep learning.

Topics Covered:

- basic math, optimization, machine learning background knowledge
- neural network basics, error backpropagation algorithm
- auto-encoder model for data encoding and re-construction
- convolutional neural network (CNN) for computer vision
- recurrent neural network (RNN) for natural language processing
- graph neural network (GNN) for network embeddings
- generative adversarial neural network (GAN)
- Transformer and BERT
- Graph-BERT

Prerequisites:

Students should come with **good programming skills**. **ECS 060 or ECS 032B or ECS 0036C** or equivalent courses are required. If you are not sure whether you have the right background, please contact the instructor.

Note: We will not cover programming-specific issues in this course.

Textbook:

- **Required textbook:** <u>Deep Learning</u>, 1nd edition, by Ian Goodfellow and Yoshua Bengio and Aaron Courville. MIT Press, 2016.
- Recommended references:
 - o Neural Networks and Deep Learning: A Textbook, 1st edition, by Charu Aggarwal. Springer, 2018.
 - o Neural Networks and Learning Machines, 3rd edition, by Simon S. Haykin. Pearson, 2008.

Course Format:

The course is lecture-based with two **examinations** (one midterm and one final). There are **individual assignments** and a **group-level programming project**. In order to encourage attending classes and participating in discussions, there will be several **in-class quizzes** for students.

- Lectures and Class Participation: We strongly encourage (and appreciate!) students to attend classes, because effective lectures rely on students' participation to raise questions and contribute in discussions. Although we probably will have a large class, we will strive to maintain interactive class discussions if possible. We will provide lecture slides before class, which will be posted on the Schedule page.
- Questions: We encourage students discussing their questions and problems first with their group peers and classmates. This way, you can get immediate help and also learn to communicate "professionally" with your peers. In any case for more thorough discussion, come to the office hours of TA's and the instructor's. Any announcement will be posted on the Announcement page. Make sure to check it frequently enough to stay informed.
- **Assignment**: There will be a few written <u>assignments</u> spaced out over the course of the quarter. All the assignments should be done individually by the students. Assignments should be submitted before due dates
- **Projects**: There will be a quarter-long <u>project</u>, which involves significant deep learning application programming. The project will be structured with several milestones due in the course of the quarter, leading to a write-up report near the end of the quarter.
- Exams: There will be a mid-term exam and a final exam held at the mid & end of the quarter.

General Policy:

- University Attendance Policy: Excused absences include documented illness, deaths in the family and other documented crises, call to active military duty or jury duty, religious holy days, and official University activities. These absences will be accommodated in a way that does not arbitrarily penalize students who have a valid excuse. Consideration will also be given to students whose dependent children experience serious illness.
- Academic Honor Policy: The UC Davis Code of Academic Conduct outlines the University's expectations for the integrity of students' academic work, the procedures for resolving alleged violations of those expectations, and the rights and responsibilities of students and faculty members throughout the process. Students are responsible for reading the Code of Academic Conduct and for living up to their pledge to "...be honest and truthful and... [to] strive for personal and institutional integrity at UC Davis." (UC Davis Code of Academic Conduct, found here)
- **Syllabus Change Policy**: Except for changes that substantially affect implementation of the evaluation (grading) statement, this syllabus is a guide for the course and is subject to change with advance notice.
- Assignment/Project Policy: You are allowed to discuss written assignments, however, any such discussion
 must be clearly acknowledged on the submitted solution. Your solution should be stapled together and
 neatly prepared. The programming project will be carried out in a group-based fashion with one or two
 students involved. No inter-team collaboration is allowed.
- Regrading Policy: Any regarding request should be submitted to the instructor or the TA(s) within one week since the graded deliverables are handed out to the students.

Collaboration/Academic Honesty:

All course participants must adhere to the academic honor code of UC Davis which is available in the student handbook. All instances of academic dishonesty will be reported to the university. Every student must write his/her own homework/code (unless you are in the same group for the programming project). Showing your code or homework solutions to others is a violation of academic honesty. It is your responsibility to ensure that others cannot access your code or homework solutions. Consulting related textbooks, papers and information available on the Internet for your coding assignment and homework is fine. However, copying a large portion of such information

will be considered as academic dishonesty. If you borrow a small piece of any such information, please acknowledge that in your assignment. Please see the following link for a complete explanation of the <u>Code of Academic Conduct</u>.

Late Submission Policy:

- Late assignments will not ordinarily be accepted. If, for some compelling reason, you cannot hand in an
 assignment on time, please contact the TA or instructor as far in advance as possible. Written assignments
 or project deliverables are due by the Friday mid-nights, you should upload them to Canvas prior to the
 deadline;
- No credit will be given to late programming projects;
- No make-up exams (except under extremely unusual circumstances).

Students with Disability:

Americans With Disabilities Act: Students with disabilities needing academic accommodation should: (1) register with and provide documentation to the <u>Student Disability Center</u>; (2) bring a letter to the instructor indicating the need for accommodation and what type. This syllabus and other class materials are available in alternative format upon request. For more information about services available to UC Davis students with disabilities, contact the <u>Student Disability Center</u>.

Course Grading Policy:

The course grade will break down as follows:

• Assignment: 20%;

• Project: 40%;

• Midterm exam: 10%;

• Final exam: 25%;

• Quizzes: 5%.

Any regrading request should be submitted to the Canvas system prior to the deadlines.

And your final grade will be assigned as follows,

```
A+: [100 - 95]; A: (95 - 90]; A-: (90 - 85];
B+: (85 - 80]; B: (80 - 70]; B-: (70 - 65];
C: (65 - 60]; D: (60 - 50];
F: (50 - 0].
```

This table indicates minimum guaranteed grades. Under certain limited circumstances (e.g., an unreasonably hard exam), we may select more generous ranges or scale the scores to adjust.

Resources:

Deep Learning Toolkit

- PyTorch
- <u>TensorFlow</u>
- Keras

Programming Language and IDE

- Python Tutorial
- PyCharm

Courses in Other Universities

• Stanford: CS230 Deep Learning

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