Syllabus, CMSE 890-002
Mathematics of Deep Learning
Spring 2020

Course Description: This course will cover the mathematical foundations of deep learning, focusing primarily on various mathematical models for deep networks and provable properties of these models; it will not focus on the mathematics of optimizing deep networks. The course is envisioned as a theoretical counterpart to the numerous programming and application focused deep learning courses at MSU, and as such, this is not a course in which you will learn how to code up deep networks. Rather, the course will emphasize how the theory leads to better intuition regarding the behavior and capacity of such networks, and to better design principles. We will probably skip detailed proofs of difficult results, but we may give proof sketches of such results and there will be some proofs of elementary results.

Instructor Information:

- **instructor:** Matthew Hirn
- **office:** 2507F, Engineering Building
- **email:** `mhirn@msu.edu`
- **phone:** (517) 432-0611
- **course webpage:** `https://matthewhirn.com/teaching/spring-2020-cmse-890-002/`

Meeting Time and Location:

- Monday, Wednesday, and Friday, 1:50 PM – 2:40 PM
- 1234, Engineering Building

Office Hours:

- Monday – Thursday, 4:00 PM – 5:00 PM (some Monday ones may need to be moved/cancelled starting in March, stay tuned to the announcements)

Prerequisites:

- Calculus I, II, III
- Linear Algebra
- Probability at the undergraduate level
- Statistics at the undergraduate level

Nice to have but definitely not required:

- Graph Theory
- Harmonic Analysis
- CMSE 820: Mathematical Foundations of Data Science
- A programming course on deep learning

Resources: The course will be based on papers in the field and excerpts from machine learning and deep learning books. A list of sources will be catalogued as the course progresses.

Grading: You will be asked to write two reports, each with a title and consisting of four pages of text/figures/equations plus as many additional pages as needed for references, in LaTeX, and using the NeurIPS 2019 (or 2020 once they become available) style files available at the link below. The first report will be due on **February 28** (the Friday before spring break), the second report will be due on **April 24** (the last day of class). Each report will count for 50% of you grade. There is no final exam.

The topics of the report may be:

- Take a topic from class, and research it in more depth, using at least two sources (books, papers, etc), at least one of which is a research conference/journal paper. Write a report describing the topic, why it is interesting, what innovations you read about, and what are some of current research directions. Explain how this topic might fit into your research or future research plans.
- The same as the previous bullet, but for a topic not covered in class. Again be sure to use at least two sources, at least one of which is a research conference/journal paper. One way to come up with such a topic is the following. In January and February CMSE will be interviewing six candidates for a position in machine learning/deep learning; the dates and times will be posted on the course webpage. Attend one or more of these talks and pick one to investigate further. Furthermore, in March and April there will be a weekly CMSE colloquium. Not all of these talks will be on deep learning, but for those that are, the same instructions apply. If you get a topic this way, please let me know in your report.
- If you are doing research in deep learning, you may write your reports on your research, but the reports must describe ongoing research being conducted this semester and not simply a copy of work completed from 2019 or earlier. In this case, your second report should build upon your first report, meaning that the second report should be eight pages long total and include revised text from the first report, plus four additional, new pages. The first report should have at least two references, and the second report should have at least four references total (the two from the first report, plus two new ones). This will conveniently give you an 8 page report, formatted in the NeurIPS style. NeurIPS paper submissions are generally in May, and are limited to 8 pages (plus references).

Here are the NeurIPS style files; use the **preprint** option for your reports:
`https://neurips.cc/Conferences/2019/PaperInformation/StyleFiles`

Course announcements will be sent via email and posted on the course website.

***Tentative* course outline:**

- Background on machine learning and statistical learning theory
- Artificial neural networks and approximation theory
- Convolutional neural networks and signal processing
- Geometric deep learning on graphs, graph signal processing, and graph isomorphism
- Generative models and GANs
- Recurrent neural networks / other topics

Academic Honesty: Cheating in any form will not be tolerated and will be reported. You will receive a zero on any assignment in which there is a case of cheating. This includes, but is not limited to, plagiarism, failure to give proper citations, and copying another's work.

If you are preparing an assignment and have a question about whether you are adhering to this policy, please ask me. If you work on an assignment with other students, you must give credit to your collaborators.

MSU's policy on academic integrity can be found at the following URL:
`https://msu.edu/unit/ombud/academic-integrity/`

Disability Services: Accommodations for persons with disabilities can and will be made in this course. All arrangements will be organized through the RCPD office as MSU. Persons with disabilities who are interested in the available services should contact the MSU Resource Center for Persons with Disabilities (RCPD) at (517) 884-7273 or online at `http://www.rcpd.msu.edu`.