

CSCI 5980/8980 Think Deep Learning

Fall 2020

General Information

Over the last few years, deep neural networks (DNN) have fundamentally transformed the way people think of machine learning and approach practical problems. Successes around DNN have ranged from traditional AI fields such as computer vision, natural language processing, interactive games, to health care, physical sciences—touching each and every corner of theoretical and applied domains. On the other hand, DNN still largely operate as black-boxes and we only have very limited understanding as to when and why they work. This course introduces the basic ingredients of DNN, samples important applications, and throws around open problems. Emphasis is put on thinking from first principles, as the field is still evolving rapidly and there is nothing there that cannot be changed.

- **Prerequisite:** Introduction to machine learning or equivalent. Maturity in linear algebra, calculus, and basic probability is assumed. Familiarity with Python (esp. numpy, scipy) is necessary to complete the homework assignments and final projects. **For the 8980 session,** course experience with (advanced) machine learning (CSCI5525 or equivalent) and nonlinear optimization, and research experience are expected, though not strictly required.
- **When & Where:** Mon 6:30PM – 9:00PM, Keller 3-210
- **Who:** Professor Ju Sun (Instructor) Email: jusun@umn.edu
Hengkang Wang (TA) Email: wang9881@umn.edu
- **Office Hours:** Tue/Thur 5–6pm (Instructor) Wed 4:30–6:30pm (TA)
- **Course Website:** <https://sunju.org/teach/DL-Fall-2020/> All course materials, including course schedule, lecture slides and notes, homework assignments, project description, will be posted on the course website. Enrolled students are encouraged to check the website on a regular basis.
- **Communication:** **Canvas** is the preferred and most efficient way of communication. Please post all questions and discussions related to the course in Canvas instead of sending emails. If you have to use emails, please begin the subject line with “CSCI 5980/8980”.

Tentative Topics

There will be 21 lecture sessions, each running 75 minutes. The tentative topics and time allocation are as follows.

- Course overview (1)
- Neural networks: old and new (1)
- Fundamental belief: universal approximation theorem (2)
- Numerical optimization with math: optimization with gradient descent and beyond (2)

- Numerical optimization without math: auto-differentiation and differential programming (2)
- Work with images: convolutional neural networks (2)
- Work with images: recognition, detection, segmentation (2)
- Work with sequences: recurrent neural networks (2)
- Learning probability distributions: generative adversarial networks (2)
- Learning representation without labels: dictionary learning and autoencoders (1)
- Gaming time: deep reinforcement learning (2)
- To train or not? scattering transforms (2)

The lecture sessions are interlaced with 5 discussion sessions, each also 75 minutes, that are designed to help the students master the critical computational and practical components and successfully complete their homework assignments and final projects.

- Python, Numpy, and Google Cloud/Colab
- Project ideas
- Tensorflow 2.0 and Pytorch
- Backpropagation and computational tricks
- Research ideas

Teaching Format

For the Fall semester, all enrolled students will be able to view live-streaming and downloadable video archives of class meetings with no delay for the length of the semester through the UNITE Media Portal.

Access the media for your class through the “LOG IN to UNITE Media Portal” link in the upper right of the UNITE website (www.unite.umn.edu).

Your University Internet ID (a.k.a. "x.500") and password are required (this is what you use to access Canvas, your University of Minnesota email account, etc.). - log in and click on the link for your course. Access these videos through the UNITE Media Portal with your University of Minnesota Internet I.D. and password.

DO NOT ask the instructor or teaching assistants for technical or troubleshooting assistance with these streaming video archives – use the UNITE Troubleshooting FAQ (<https://cse.umn.edu/unite/troubleshoot-unite-media>) or “Submit a Trouble Report to UNITE” through the link found on all pages within the UNITE Media Portal.

Live streaming video has latency based on UNITE compression and your individual stable internet connection speed, typically ranging from 30 to 60 seconds. The archived UNITE media is typically available within an hour after the end of class meetings.

Recommended References

There is no required textbook. Lecture notes and additional notes will be the primary resources. Recommended reference books are

- **Dive into Deep Learning** by Aston Zhang and Zachary C. Lipton and Mu Li and Alexander J. Smola. Livebook with online URL: <https://d2l.ai/> (comprehensive coverage of recent developments and detailed implementations based on NumPy)
- **Deep Learning** by Ian Goodfellow and Yoshua Bengio and Aaron Courville. MIT Press, 2016. Online URL: <https://www.deeplearningbook.org/> (comprehensive coverage of recent developments)
- **Neural Networks and Deep Learning** by Charu Aggarwal. Springer, 2018. UMN library online access (login required): [Click here](#). (comprehensive coverage of recent developments)
- **The Deep Learning Revolution** by Terrence J. Sejnowski. MIT Press, 2018. UMN library online access (login required): [Click here](#). (account of historic developments and related fields)
- **Deep Learning with Python** by François Chollet. Online URL: <https://livebook.manning.com/book/deep-learning-with-python> (hands-on deep learning using Keras with the Tensorflow backend)
- **Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems** by Aurélien Géron (2ed). O'Reilly Media, 2019. UMN library online access (login required): [click here](#). (hands-on machine learning, including deep learning, using Scikit-Learn and Keras)
- **Deep Learning for Vision Systems** by Mohamed Elgendy (1ed). Manning Publications, 2020.

Assessment

CSCI 5980

- Homework 60%: 7 Homeworks, the top 5 scores will count toward the final grade
- Course project 40%: proposal (5%) + mid-term presentation (10%) + final report (25%). The project can be survey of a chosen topic not covered in detail in the class, implementation and comparison of existing methods, or novel foundational or applied research
- Exception: If the course project leads to a publishable result as determined by the instructor, the participating students will all automatically get an A for the course.

CSCI 8980

- Homework 30%: 7 Homeworks, the top 3 scores will count toward the final grade
- Course project 40%: proposal (5%) + mid-term presentation (10%) + final report (25%). The project must be novel foundational or applied research.
- Lecture teaching or scribing 15%: scribe or teach one lecture session (75 mins)

- Research survey 15%: write a short survey paper on a deep learning related topic that is not covered in class.
- Exception: If the course project leads to a publishable result as determined by the instructor, the participating students will all automatically get an A for the course.

Homework

You have approximately 10 days time to complete each homework (exact due date will be specified in each assignment). No late submissions will be accepted. All submissions *must* be electronic and uploaded via the Canvas system. Written part *must* be typeset in L^AT_EX and submitted as PDF files. Computer programs *must* be submitted in the Python notebook format. Only Python 3 will be used and accepted in this course.

Collaboration on homework problems is strongly encouraged, but each student must ensure that the final submission is prepared individually. *Collaborators should be properly acknowledged in the final submission, at the problem level.* The same applies to computer programs. Plagiarism and cheating will not be tolerated and are subject to disciplinary action. Please consult the student code of conduct for more information: https://regents.umn.edu/sites/regents.umn.edu/files/2019-09/policy_student_conduct_code.pdf

Course Project

The course project is to be performed by teams of 2 or 3 students for the CSCI5980 session, 1 or 2 students for the CSCI8980 session, and the weight of the project should be proportional to the number of students in the team. All students from the same team will get the same score for their course project.

Programming and Computing

Our programming environment will be Python 3. For deep learning frameworks, PyTorch (≥ 1.0) is preferred, but Tensorflow (≥ 2.0) is also accepted and supported. For small-scale experiments, which will be the case for homework assignments, Google Colab (<https://colab.research.google.com/>) and UMN MSI notebook service will suffice. Local installation of the relevant software packages may be a reasonable alternative. For large-scale course projects, we can use the Minnesota Supercomputing Institute (MSI) GPU computing queues based on our class account.

Related Courses

Within UMN

- **Topics in Computational Vision: Deep networks** (Prof. Daniel Kersten, Department of Psychology. Focused on connection with computational neuroscience and vision)
- **Analytical Foundations of Deep Learning** (Prof. Jarvis Haupt, Department of Electrical and Computer Engineering. Focused on mathematical foundations and theories)

Global

- **CS230 Deep Learning** (<https://cs230.stanford.edu/>, Stanford Computer Science)
- **CS231n: Convolutional Neural Networks for Visual Recognition** (<http://cs231n.stanford.edu/>, Stanford Computer Science, 2019)
- **CS224n: Natural Language Processing with Deep Learning** (<http://web.stanford.edu/class/cs224n/>, Stanford Computer Science, 2020)
- **Analyses of Deep Learning** (<https://stats385.github.io/>, Stanford Statistics, 2019)
- **Introduction to Deep Learning** (<https://deeplearning.cs.cmu.edu/F20/index.html>, CMU, 2020)
- **Advanced deep learning and reinforcement learning** (<https://github.com/enggen/DeepMind-Advanced-Deep-Learning-and-Reinforcement-Learning>, UCL/Deepmind, 2018)
- **Mathematics of Deep Learning** (<https://joanbruna.github.io/MathsDL-spring18/>, NYU Courant Institute, 2018)
- **MIT Deep Learning**(<https://deeplearning.mit.edu/>, MIT courses and lectures on deep learning, deep reinforcement learning, autonomous vehicles, and artificial intelligence)
- **CMSC 35246 Deep Learning** (<https://ttic.uchicago.edu/~shubhendu/Pages/CMSC35246.html>, U Chicago Computer Science, 2017)
- **Neural Networks for Machine Learning** by Jeof. Hinton (https://www.cs.toronto.edu/~hinton/coursera_lectures.html, U Toronto, 2012)