CS 559 Deep Learning

Syllabus, Spring 2017

Machine learning studies algorithms for building data-driven models that can make predictions about data and novel observations. Recent developments in machine learning approaches based on deep neural networks, also known as *deep learning*, have lead to performance breakthroughs and catalyzed research in many closely related fields, including computer vision, natural language processing, speech recognition and robotics. This course aims to not only cover the fundamentals of deep learning, but also give a grasp of contemporary research. The course will start with a brief overview of machine learning and numerical optimization. Then, the basic techniques and modern approaches in designing, training and visualizing feedforward neural network architectures and convolutional neural networks will be introduced. Convolutional neural network based methods for spatial localization of visual entities in images will be covered. Recurrent neural network architectures, and their applications in language and image understanding will be discussed. Recent advances in deep generative models will follow. Finally, deep reinforcement learning will be covered.

Prerequisites: Probability theory, statistics, linear algebra. Machine learning background is recommended but not required.

Schedule

Wed 09:40 – 10:30 (spare hour: Wed 08:40 – 09:30)*
Fri 10:40 – 12:30

We will resort to this hour whenever necessary.

Office hours are announced on course webpage.

Contact Information

Instructor: Gokberk Cinbis, `gcinbis [at] cs.bilkent.edu.tr`
Graders: TBA

Course Webpage

`http://www.cs.bilkent.edu.tr/~gcinbis/courses/Spring17/CS559`

Please check regularly the course webpage for lecture notes, homework assignments, readings and project information.

Textbook: No required textbooks. There will be **required readings** posted on the course webpage.

Recommended, optional books:

- I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, 2016. `http://www.deeplearningbook.org`.
- K. P. Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
- C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.

Tentative weekly syllabus (Check course webpage for the current schedule)

1. Introduction
   - Overview of machine learning, linear classifiers, loss functions
2. Optimization
   - Stochastic gradient descent and contemporary variants, back-propagation

3-4. Feedforward networks and training
   - Activation functions, initialization, regularization, batch normalization, model selection, ensembles

5-6. Convolutional neural networks
   - Fundamentals, architectures, pooling, visualization

7. Deep learning for spatial localization
   - Transposed convolution, efficient pooling, object detection, semantic segmentation

8-10. Recurrent neural networks
   - Recurrent neural networks (RNN), long-short term memory (LSTM), language models, machine translation, image captioning, video processing, visual question answering, video processing, learning from descriptions, attention

11-12. Deep generative models
   - Auto-encoders, variational auto-encoders, generative adversarial networks, auto-regressive models, generative image models, unsupervised and self-supervised representation learning

13-14. Deep reinforcement learning
   - Policy gradient methods, Q-Learning
   - Project presentations

Grading:

Homework:	30%
Paper presentation:	20%
Paper questions:	10%
Midterm:	10%
Project:	30%

Important: Not submitting a project report, not presenting a paper as pre-scheduled, being absent in the midterm, or being absent in a project presentation will automatically result in an F.

Homework: The homework assignment will aim to foster understanding of the deep learning fundamentals covered in class, give hands-on experience on developing and/or using various architectures and prepare students to work on the term projects.

Paper presentations: Deep learning is a quickly developing and evolving research field, which renders publications indispensable resources to grasp contemporary deep learning research. Depending on class size and technical depth of the paper(s), each student will present one or more papers on the topic scheduled for the presentation week in the syllabus. The presentation will be evaluated depending on several criteria including its clarity and correctness at defining the problem addressed in the paper, the motivation for working on the research problem, the proposed approach, its advantages and disadvantages, and discussing the related work.

Paper questions: Each student is expected to submit two genuine questions regarding the papers to be presented, before each class for which paper presentations are scheduled. For example, if there are 5 papers to be present on a particular day, each student is expected to submit two questions in total (not per paper). Presenters are not required to submit any questions on the day they are presenting – they automatically get full grade for paper question. The quality of questions will be evaluated based on their originality, scientific merit and contribution to the discussion in class.

Exam: There will be one midterm exam.

Project: There will be a term project. Each project group will prepare a project proposal, progress report, final report in a conference paper format and a presentation. Each project will be expected to explore novel applications of contemporary deep learning techniques or develop novel deep learning techniques.

Honor code: This course follows the Bilkent University Code of Academic Integrity, as explained in the Student Disciplinary Rules and Regulation. Violations of the rules will not be tolerated. Students may discuss and work on homework problems in groups. However, each student must write down the solutions independently, and without referring to written notes from the joint session. In other words, each student must understand the solution well enough in order to reconstruct it by him/herself. In addition, each student must write on the problem set the names of the people with whom s/he collaborated.