Fall 2017

MWF 1:00 pm - 1:50 pm, CB 347

Instructor

Dr. Qiang Ye
Office: 735 Patterson Office Tower
Phone:257-4653
Email: qye3 "at" uky . edu
Office Hours: MWF 2:00-3:00 pm

Textbook

Deep Learning by Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT Press, 2016

Syllabus

In this course, we study a widely applicable class of machine learning methods called deep learning. We will primarily be covering the following topics from Part II of the text:

- Deep Feedforward Networks.
- Regularization for Deep Learning
- Optimization for Training Deep Models.
- Convolutional Networks
- Sequence Modeling: Recurrent Networks
- Other Topics (e.g. Autoencoders and Deep Generative Models from Part III) may be covered if time permits

Selected materials from Part I of the text will be reviewed. We plan to begin with an introduction to numerical optimization and probability theory. Linear algebra is another area of math that will be used substantially. However, we will not devote time to cover them as most of the class have taken MA/CS 522 or 622. Those who have not taken a numerical linear algebra class before are encouraged to consult Part I of the text and/or one of the references listed below. The most pertinent

topics include conditioning, linear least squares problems, singular value decompositions, and the steepest descent and the conjugate gradient methods.

Prerequisites

Familiarity with multivariate calculus, linear algebra and numerical methods will be assumed. Programming in Python will be required.

Grading

The course grade will be based on two programming projects (40%), an in-class presentation (30%), and attendance (30%).

Course Materials

- List of papers for class presentations
- Slides for Chapter 7 - Regularization
- Project 1
- Slides for Chapter 9 - Convolutional Networks
- Project 2

Computing

We will use Python based toolboxs such as Keras, Theano, or Tensorflow. You may choose to use any of them. It is best to run them on the Linux platform. But if you use Windows, below are some links for setting up Keras/Theano/Tensorflow on Windows 10.

- Installing Keras, Theano on Windows 10 The instructions here work for the CPU version but is outdated for the GPU version. To install the GPU version, the fix suggested HERE worked for me.
- Installing TensorFlow on Windows 10.

You may also install these toolboxes using ANACONDA DISTRIBUTION. Download the distribution and try the following (Thanks to Liu Liu for these tips!)

- For windows, just double click and install. Choose 'add anaconda to PATH environment', and 'use anaconda as default python' options if you see them.

After installation, open anaconda navigator to install all the packages you need (theano, tensorflow, keras) in 'Environments'.

- For linux, run 'bash ~/Downloads/Anacondaxxx.sh' in terminal to install the downloaded anaconda. Choose 'add to PATH environment', and 'use anaconda as default python' options if you see them. After installation, run 'anaconda-navigator' in terminal to open anaconda navigator to install all the packages you need (theano, tensorflow, keras) in 'Environments'.

Some references and links

Below are some links and books on numerical linear algebra, optimization, and machine learning that may be helpful.

- Tensorflow playgroun and its documentation
- Applied Numerical Linear Algebra by J.W. Demmel, published by SIAM.
- Numerical Linear Algebra and Application, 2nd ed. by Biswa Datta, published by SIAM.
- Numerical Optimization, by Jorge Nocedal and Stephen J. Wright.
- The Elements of Statistical Learning by Trevor Hastie, Robert Tibshirani, and Jerome Friedman