ECSE 4965/6965
Introduction to Deep Learning

Instructor: Dr. Qiang Ji,
Email: jiq@rpi.edu
Phone: 276-6440
Office: JEC 7004
Meeting Hours & Place : 4:00-5:20 pm, Mondays and Thursdays, Academy Auditorium
Office Hours: 5:30-6:30 pm Mondays and Thursdays or by Appointment
TAs:

Meng Zheng zhengm3@rpi.edu
Chengyu Wang wangc16@rpi.edu

Lecture notes: Available on RPI Learning Management System

Deep learning involves learning hierarchical representations of input data in increasing levels of abstraction. Deep Learning is rapidly emerging as one of the most successful machine learning techniques. It has revolutionized several fields, including computer vision, natural language processing, and speech recognition. This course introduces fundamentals in deep learning and demonstrates its applications in computer vision. It covers both probabilistic deep models, including Deep Boltzmann Machines, Deep Belief Networks, and Deep Bayesian Networks as well as deterministic deep models, including Convolutional Neural Networks and Autoencoders. In addition, the course will also cover the latest deep learning topics including Recurrent Neural Networks, Deep Reinforcement Learning, and the Generative Adversarial Networks. The course is self-contained. It starts with an introduction of the background needed for learning deep models, including probability, linear algebra, standard classification and optimization techniques. To demonstrate various deep models, we will apply them to different computer vision tasks, including object recognition, human action recognition, and facial expression recognition.

Prerequisites

This is a senior and graduate level course. Students should have basic knowledge in probability, linear algebra, and optimization. Strong programming skills in one of the high level languages such as C++, Matlab, Python are required. Prior courses in machine learning/ pattern recognition and computer vision/image processing are preferred but not required.

Optional textbook

Deep Learning, Ian Goodfellow and Yoshua Bengio and Aaron Courville, MIT Press, 2016, http://www.deeplearningbook.org (download the book quickly for free!)

Software

Many of the assignments and projects will be implemented in Google's TensorFlow (https://www.tensorflow.org/tutorials/), which is implemented in Python. Mastery of Python or being able to quickly learn it is required.

Course Evaluation

The course will involve homework assignments (10%), class projects (50%), a mid-term exam (20%), and a final project (20%). The exam is open book and comprehensive. The project can be individual or group project, depending on their scope.