

CSE 791/CIS 700: ADVANCES IN DEEP LEARNING

SYLLABUS SPRING-2018

Instructor

Professor Yanzhi Wang, Department of Electrical Engineering and Computer Science, Room CST 4-102,
Tel: (213)400-2560, email: ywang393@syr.edu
Office Hours: Wednesday 2:00pm – 4:30pm

Class Location and Times (Lectures)

Day: Tuesday & Thursday
Location: Sims Hall 241
Time: 3:30 PM ~ 4:50 PM

Teaching Assistant Caiwen Ding

Email: cading@syr.edu
Office Hours: Tuesday & Thursday: 1:30pm – 3:00pm
Location: CST 0-121

Course Web Page

All instructional materials, including lecture notes, homework and project assignments will be posted on Blackboard@SU. Look for Blackboard announcement for important messages.

Description

This class is a graduate-level and research-oriented class.

It will focus on the recent advances in the deep learning applications and research fields. The objective is to make students competitive in the growing job markets and research activities of artificial intelligence and big data areas. There are five major components in this course. The first is basic deep learning models, including feedforward neural networks, deep convolutional neural networks, recurrent neural networks (LSTM, GRU, etc.), deep reinforcement learning (e.g., AlphaGo), and their inference and training algorithms. The second part is different application fields of deep learning, including image classification, object detection, speech recognition, robotics, automated control, medical systems, etc. The third part is effective accelerations and implementations of deep learning systems including GPU-based implementations, FPGA implementations, dedicated hardware such as Google TPU or IBM TrueNorth, deep learning model quantization, model compression, and FFT-based computations. The fourth part is the recent advances in deep learning algorithms/models including generative adversary networks (GANs), Bayesian neural networks, etc. The last part is the security aspects of deep learning, i.e., how deep learning can enhance the security levels of cyber-physical systems and how to enhance the security level of deep learning systems themselves.

This class requires programming in C++ and Python, and students will learn effective deep learning tools such as TensorFlow and Caffe.

Topics covered

1. Introductions
2. Classification problem, basic machine learning models
3. Neural networks, deep convolutional neural networks (DCNN)
4. Training and inference techniques
5. Python and TensorFlow
6. Using GPU to accelerate learning and inference
7. Recurrent neural networks (RNN)
8. Other deep learning models beyond DCNN and RNN
9. Other interesting applications
10. Hardware acceleration of deep learning systems
11. Security in deep learning systems

Grading policy

Final Project: 60%

Homework: 20%

Quizzes: 20%

Extra Points: up to 3%

Grading Scale:

90 ~ 100 = As

80 ~ 89 = Bs

70 ~ 79 = Cs

60 ~ 69 = D

Below 60 = F

1. HOMEWORK POLICY

Homework assignments are to be submitted through Blackboard website or hand to the instructor on the assignment due date. Assignments submitted after the due date will be deducted 10 points for each day late.

2. ATTENDANCE

You are expected to attend each class punctually and remain for the entire class period. You need to inform the instructor in advance if you expect to miss a class or leave the course before the end of the semester. If you miss class your absence will be excused by the instructor only if a doctor's certificate or other evidence is submitted. You remain to be responsible for the work associated with the class you missed, even if your absence has a valid reason. There will be a number of unannounced popup quizzes during the semester.

3. ACADEMIC HONESTY

Cheating in any form is not tolerated, nor is assisting another person to cheat. The submission of any work by a student is taken as a guarantee that the thoughts and expressions in it are the students own except when properly credited to another.

Violations of this principle include giving or receiving aid in an exam or where otherwise prohibited, fraud, plagiarism, the falsification or forgery of any record, and any other deceptive act in connection with academic work. Plagiarism is the representation of another's words, ideas, programs, formulae, options or other products of work as

one's own work from others, since it is often not possible to determine who the originator or the copier was. Such offense will result in a failing grade "F" and a letter of reprimand in your department student file.

4. COURSE CALENDAR

Week	Material
Week 1	Introduction
Week 2	Basics of machine learning, traditional classifiers, neural networks
Week 3	Training and inference of neural networks, Back propagation and gradient descent Homework 1 (a neural network classifier)
Week 4	Convolutional neural networks, TensorFlow
Week 5	Applications of deep neural networks Homework 2 (CNN using TensorFlow, MNIST and CIFAR-10 dataset)
Week 6	Recurrent neural networks (LSTM and GRU), object detection, deep reinforcement learning Project announcement
Week 7	GPU acceleration, Generative adversarial networks (GAN) Homework 3 (GPU based accelerations)
Week 8	GAN, Bayesian neural networks, VAE Project discussion 1
Week 9	Spiking neural networks, Hardware acceleration of deep learning systems
Week 10	Hardware acceleration of deep learning systems, security in deep learning systems
Week 11	Project discussion 2
Week 12	Security in deep learning systems
Week 13	Project discussion 3
Exam Week	Project demonstration

Note: The schedule might change during the semester depending on the progress of the class. All departmental, college and university regulation regarding class attendance, course drop, etc will be followed.