CSE591 Introduction to Deep Learning in Visual Computing

General Information

Time & Location

Friday 9:40AM – 12:10PM @ BYAC110

Contact Information

Instructors: Baoxin Li, Ragav Venkatesan
Office: Brickyard 554/502 (B. Li)
Office hours for B. Li:
Monday: 9-10am @ BY554
Wednesday 11am – 12am @BY502
or other times by appointment
Office hours for Ragav:
Monday: 1-2pm; Wednesday 1-2 pm or other times by appointment

[Cancellation of office hours due to travel or other urgent university business will be announced on Blackboard in advance].

E-mail: Baoxin.Li@asu.edu, ragav.venkatesan@asu.edu

Teaching Assistant: Ragav Venkatesan

There are several senior PhD students in Prof. Li's lab, who have been doing deep-learning-related research. They may be recruited to help enhancing the delivery of this class (e.g., by helping/evaluating student assignments/projects, giving short guest talks, etc.)

Textbooks

Primary: R. Venkatesan and B. Li, A Concise Guide to Modern Neural Visual Computing, CRC press, Taylor and Francis Publications, *(to be published)*) (The pre-print copy of the book will be made freely available to the students in the class.)

Other References:

1. Deep Learning, Ian Goodfellow and Yoshua Bengio and Aaron Courville.

2. Neural Networks for Pattern Recognition, Christopher Bishop.

3. Online book: http://neuralnetworksanddeeplearning.com/ Michael Nielson.

4. A pool of recent research papers to be posted.

Topics to Cover

This course will introduce some basic paradigms for deep learning, with a focus on deep architectures using convolutional neural networks. Topics to be covered include:

- Introduction to visual representation & fundamentals of machine learning
- Neural networks & backpropagation
- Optimization techniques for neural networks
- General deep learning paradigms (CNN, auto-encoder, GANs etc.)
- Modern convolutional neural networks
- Software implementation of deep learning
- Selected recent advances in deep learning

A more detailed schedule will be discussed in the first lecture.

The course is intended to be interactive by requiring students to read the textbook or papers before each lecture, and the lecture time will be used to share insights, discuss pros and cons of any techniques to be studied in that lecture.

There will be about 8 weeks of lectures by the instructors, and the rest will be on paper discussion and students will be required to make presentations. The students are required to understand all the papers presented, since there will be a final exam that will contain questions about the presented papers.

Time permits, the class may invite other faculty or senior PhD students who have relevant research work and recent publications on deep learning to give guest lectures.

Prerequisites

Students should have good working knowledge of calculus, linear algebra and basic probability theory.

It is highly recommended for a student to first take at least one machine learning class (e.g., CSE 569 or CSE 575 at ASU) before registering for this course. This will be a seminar course and will cover fairly advanced techniques and their applications without spending much time on reviewing the basics; a student who has not taken any graduate-level machine learning class will find it extremely difficult to succeed in this class.

In addition, proficiency in Python programming is needed for doing the course project and/or homework assignments. You will need to use external libraries written in Python.

E-mail Policy

Email has become an important means of communication, but due to the size of the class and the volume of daily emails the instructors typically need to deal with, there will be guidelines on using emails for communicating with the instructors. These will be discussed in more detail in the first lecture, but briefly, emails not following the protocols will not be responded.

Assessment

Homework (mini projects):	24%
Pop Quizzes/Attendance:	6% (4 random attendances will be taken)
Project:	30%
(2-person group; you may propose your own topic; topics will be finalized within 2 weeks after the first midterm.)	
Midterm Exams:	12% + 12 %
Final Exam:	11 % on the official date (will include questions testing papers presented)
Paper presentation:	5% (a total of about 24 papers will be presented by student groups of 3 or 4 students)

More details on the assessment scheme will be discussed in the first lecture.

Lecture Notes

Lecture notes will be available via the course public website at http://www.public.asu.edu/~rvenka10/cse591 . The notes will in general be posted by the noon in the day of the lecture. Audio recordings of the lecture may also be posted along with the lecture notes but it is not a guarantee.

In addition, some examples may be worked out on the whiteboard during the lectures. The examples may not be included into the posted lecture notes.