AICIP RESEARCH

ECE 599/692 – Deep Learning

Lecture 1 - Introduction

Hairong Qi, Gonzalez Family Professor
Electrical Engineering and Computer Science
University of Tennessee, Knoxville
http://www.eecs.utk.edu/faculty/qi
Email: hqi@utk.edu

AICIP RESEARCH

What do we cover?

- Neural networks
  - Multi-Layer Perceptron (MLP) – Project 1
  - Backpropagation
- Feedforward networks
  - Supervised learning – Convolutional Neural Network (CNN) – Project 2
  - Unsupervised learning – Autoencoder (AE) – Project 3
- Generative networks
  - Generative Adversarial Network (GAN) – Project 4
- Feedback networks
  - Recurrent Neural Network (RNN) – Project 5

AICIP RESEARCH

Classroom

- Currently on the timetable, ECE599 is assigned to MK405 and ECE692 to MK406. Since these two sections share the same lecture time, we'll be using just MK405. ECE692 students, please make a note of classroom change.

Course Website

AICIP RESEARCH

- For easy access, besides Canvas where I mainly use to send out announcement, I also maintain the course website that can be accessed by everybody, http://web.eecs.utk.edu/~qi/deeplearning/index.html. Please refer to this site for updated syllabus, projects, references, reading assignments, etc.

4

Background

AICIP RESEARCH

- I used to allow students who do not have any machine learning or pattern recognition background to take deep learning, which had turned out not to be a good idea. You have to have certain background in subjects like supervised vs. unsupervised learning, parametric vs. non-parametric learning, training vs. validation vs. testing datasets, classification vs. regression vs. generation, features, dimension, cross-validation, etc., in order to fully appreciate the new content in this class. So if you have never taken pattern recognition (ECE471/571) or machine learning (COSC 425/528) or equivalent, I strongly recommend you to take these first and then move up to deep learning.

5

Two Sections

AICIP RESEARCH

- There are two sections of the deep learning class (ECE599 and ECE692) where they share the same lecture but 692 students need to work on additional problems in each of the five project assignments. In addition, 692 students need to write a conference-style paper with certain degree of innovation while 599 students just need to write a regular project report. During the course of the semester, you'll have plenty reading assignments so you'd know what I mean by conference-style paper. The two sections are provided in order to accommodate those students who already have deep learning background but would like to take this course to have a systematic coverage of deep learning and to broaden the knowledge base of deep learning. So if you feel you've signed up the wrong section, please make changes. I don't think you'll have problem dropping it. If you have trouble adding the course, you can bring an add/drop form on Tuesday for me to sign.

6

Assignment

AICIP RESEARCH

- There will be five regular projects on each type of network we are going to cover, including classical NN, GAN, RNN, and AE, plus one final project. In addition, I will have reading assignments leading to designated in-class presentation or discussion. The regular project needs to be completed by each individual student and the final project is a group project with 2-3 members. It's a good idea to start looking for your partners early on.

7

Textbooks

AICIP RESEARCH

- Since deep learning is a fast evolving field, the best reference comes from conferences. The top conferences are NIPS, ICML, ICLR, and with deep learning and vision, the best are CVPR, ICCV and ECCV. Nonetheless, there are still some good textbook material online. The two books I'll rely on the most are Nielsen's Neural Network and Deep Learning, and Goodfellow et al.'s Deep Learning, as listed on the course website. I've already listed some reading assignments for the first lecture. Please go through!

8

Compute Resource

AICIP RESEARCH

- We'll be using Python and TensorFlow for all project assignments. Due to the size of the model we need to train or fine-tune and the size of the dataset, GPU stations with adequate memory support have to be used. For the first time, we'll be using Google Cloud where each student will have $50 worth of credits. Chengcheng will detail the usage of Google Cloud next. I'll need each of you who is determined to take this class to send Chengcheng an email with Subject "599 - Google Cloud Account Request" or "692 - Google Cloud Account Request". If you are still debating, please don't rush in sending out this email, since if the account is assigned, I cannot transfer it to another student and we have limited quota. In the meanwhile, please go through TensorFlow tutorial offline as early as possible.

9

Office Hours

AICIP RESEARCH

- Hairong Qi: MW 1-3

10