Deep Learning for Imaging (EE6180w)

This is a 9-credit elective course which will be run in online classroom mode by Prof. A.N. Rajagopalan. About 3 hrs of live lectures (i.e., 3 classes) will be held every week. The course will start with the Basics of Neural Networks and proceed to cover Multi-layer perceptron (MLP), Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Autoencoders, Deep generative Models and finally conclude with Generative Adversarial Networks (GANs). Along the way, several real-world applications of Deep Neural Networks will be discussed. While the emphasis will be on imaging, the Deep Network principles discussed will be quite general in nature and will benefit even people working in speech and natural language. There will also be lab assignments for acquiring hands-on coding skills. Moodle-based quiz will be held once every two weeks. The lab assignments will count for 30% of the grade, the quizzes will count for 15% while the end-sem exam will count for 55%.

Pre-requisites:

1. Familiarity with image processing desirable but not essential
2. Basic knowledge of Python programming preferred

Course Contents:

1. Basic Neural Network:
Perceptron; Multi-layer Perceptron; Back propagation; Stochastic gradient descent; Universal approximation theorem; Applications in imaging such as for denoising.
2. Convolutional Neural Networks (CNN):
CNN Architecture (Convolutional layer, Pooling layer, ReLu layer, fully connected layer, loss layer); Regularization methods such as dropout; Fine-tuning; Understanding and Visualizing CNN; Applications of CNN in imaging such as object/scene recognition.
3. Recurrent Neural Network (RNN):
Basic RNN; Long Short Term Memory (LSTM) and GRUs; Encoder-Decoder models; Applications in imaging such as activity recognition, image captioning.
4. Autoencoders:
Autoencoder; Denoising auto-encoder; Sparse auto-encoder; Variational autoencoder; Applications in imaging such as segnet and image generation.
5. Deep Generative Models:
Restricted Boltzmann machine; Deep Boltzmann machine; Recurrent Image Density Estimators (RIDE); PixelRNN and PixelCNN; Plug-and-Play generative networks.
6. Generative Adversarial Network (GAN):
GAN; Deep Convolutional GAN; Conditional GAN;

Reference Book:
Deep Learning by Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT Press, 2016: http://www.deeplearningbook.org/.

Reference Material:
1. Stanford CS231n: Convolutional Neural Networks for Visual Recognition, http://cs231n.stanford.edu/
2. Neural Networks and Deep Learning by Michael Nielsen: http://neuralnetworksanddeeplearning.com/
3. Online course on Neural Network by Hugo Larochelle: http://info.usherbrooke.ca/hlarochelle/neural_networks/content.html