CS 878 - RIME 833: Deep Learning

Textbook

1. Deep Learning. By Ian Goodfellow, Yoshua Bengio, and Aaron Courville, MIT Press,2016. Also available at www.deeplearningbook.org Handouts and research articles may also be used by the instructor.

Objective

2. Computer Vision has become ubiquitous in our society, with applications in search, image understanding, apps, mapping, medicine, drones, and self-driving cars. Core to many of these applications are visual recognition tasks such as image classification, localization and detection. Recent developments in neural network (aka "deep learning") approaches have greatly advanced the performance of these state-of-the-art visual recognition systems. This course is a deep dive into details of the deep learning architectures with a focus on learning end-to-endmodels for these tasks, particularly image classification.

Pre-Requisite

3. Linear Algebra, Probability, Machine Learning

Course Outcome

4. During the course, students will learn to implement, train and debug their own neural networks and gain a detailed understanding of cutting-edge research in computer vision and its application for robotics. The final assignment will involve training a multi-million parameter convolutional neural network and applying it on a large dataset.

Course Outline

4. The course covers the techniques and technology to set up the problem of image recognition, the learning algorithms (e.g. backpropagation), practical engineering tricks fortraining and fine-tuning the networks.

Topics	Allocated Periods
· Machine Learning Basics · Deep Feedforward Networks · Regularization for Deep Learning · Optimization for Training Deep Models · Convolutional Neural Networks · Sequence Modeling: Recurrent Neural Networks · Applications of Deep Learning	45