Neural Networks and Deep Learning

Rationale

With the recent boom in artificial intelligence, more specifically, Deep Learning and its underlying Neural Networks, are essential part of systems that must perform recognition, make decisions and operate machinery.

Deep Learning in Artificial Neural Networks (ANN) is relevant for Supervised, Unsupervised, and Reinforcement Learning. This course will provide a thorough examination of the state-of-the-art and will present the mathematical and algorithmic foundations of Deep Learning in ANN.

Description

Deep Learning concerns multilevel data representation with every level providing hierarchical explanations of data. These algorithms have proven effective in discovering the structure of data and the features that most differentiate among classes.

Artificial Neural Networks are created to mimic the human brain. In the process, various architectures of neurons and connections have been created. Generally, the artificial neurons are activated through input signals and the connections among them are learned through various algorithms.

The course will cover the underlying theory, the range of applications, and learning from very large data sets. Connectionist architectures commonly associated with deep learning, e.g., basic neural networks, convolutional neural networks and recurrent neural networks will also be discussed. Methods to train and optimize the architectures and methods to perform effective inference with them, will be the main focus.

Topic List

Topics may include but are not limited to:

- Deep Learning
    - Depth of deep learning
    - Overview of methods
    - Linear algebra and numerical methods
- Neural Networks
    - Feed-forward and backward architectures
    - Gradient-based learning, momentum
- Deep Neural Networks
    - Architecture design
    - Computational graphs
    - Training: Optimization, Activation Functions, Choice of Parameters
    - Performance metrics
- Regularization
    - Overview

    - Multitask learning
    - Parameter sharing
- Convolutional Neural Networks
    - Operation and algorithm
    - Pooling
    - Imbalance, ensembles
- Deep Recurrent Neural Networks, Long-Short Term Memory (LSTM) and Autoencoders
    - Unfolding graphs
    - Adding Depth to RNN
    - Dimensionality Reduction
- Generative Adversarial Networks (GANs)
    - Adversarial Training
    - Deep Convolutional GANs
- Applications and Research
    - Large scale systems
    - Computer vision
    - Object detection
    - Natural language processing

Text Sources

The main textbook for the course is Ian Goodfellow, Yoshua Bengio and Aaron Courville**, Deep Learning,** MIT Press, 2016 (available as free html book for reading at http://www.deeplearningbook.org/) as well as research papers from the field.

Learning Goals

The student must be able to demonstrate a working knowledge of the foundations of deep learning in neural networks represented by:

- Deep learning backpropagation
- Convolutional neural networks
- Recurrent networks
- Computer vision large scale system processing

Assessment

Assignments and projects will be given to test the students' abilities in the design and development of deep learning systems based on neural networks. The students should be able to experiment and discuss the performance of such systems and evaluate the applicability of the algorithms and neural network architectures. Evaluation criteria: individual assignments – 40 %, term project – 40%, discussion participation – 20%.