

Deep Learning Essentials

Ruslan Salakhutdinov

Abstract

The goal of the tutorial is to introduce mathematical basics of many recent deep learning methods. The core focus will be placed on algorithms that can learn multi-layer hierarchies of representations.

The tutorial will be split into two parts. The first part will focus on supervised (discriminative) learning algorithms that can learn multi-layer representations via non-convex optimization. Topics will include: multi-layer neural networks, convolutional networks, back-propagation algorithm, stochastic gradient descend (SGD), dropout training, batch-normalization, and recent second-order optimization methods. The second part of the tutorial will introduce mathematical basics of many popular unsupervised models, including Variational Autoencoders (VAE), Generative Adversarial Networks (GANs), and more recently introduced Capsule Networks. I will also discuss recent trends in modelling sequential data using Recurrent Neural Networks (RNNs), Sequence-to-Sequence models, and Transformer Architectures.

Throughout this tutorial I will highlight various application areas, including visual object recognition and video analysis, language understanding, including reading comprehension and question-answering models, and time series analysis.