

# Deep Learning

Ruslan Salakhutdinov  
University of Toronto  
Toronto, Canada  
rsalaku@cs.toronto.edu

## ABSTRACT

Building intelligent systems that are capable of extracting high-level representations from high-dimensional data lies at the core of solving many AI related tasks, including visual object or pattern recognition, speech perception, and language understanding. Theoretical and biological arguments strongly suggest that building such systems requires deep architectures that involve many layers of nonlinear processing. Many existing learning algorithms use shallow architectures, including neural networks with only one hidden layer, support vector machines, kernel logistic regression, and many others. The internal representations learned by such systems are necessarily simple and are incapable of extracting some types of complex structure from high-dimensional input. In the past few years, researchers across many different communities, from applied statistics to engineering, computer science, and neuroscience, have proposed several deep (hierarchical) models that are capable of extracting meaningful, high-level representations. An important property of these models is that they can extract complex statistical dependencies from data and efficiently learn high-level representations by re-using and combining intermediate concepts, allowing these models to generalize well across a wide variety of tasks. The learned high-level representations have been shown to give state-of-the-art results in many challenging learning problems and have been successfully applied in a wide variety of application domains, including visual object recognition, information retrieval, natural language processing, and speech perception. A few notable examples of such models include Deep Belief Networks, Deep Boltzmann Machines, Deep Autoencoders, and sparse coding-based methods. The goal of the tutorial is to introduce the recent developments of various deep learning methods to the KDD community. The core focus will be placed on algorithms that can learn multi-layer hierarchies of representations, emphasizing their applications in information retrieval, object recognition, and speech perception.

## Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning—*Connectionism and neural nets*

## Keywords

Deep Learning; Graphical Models; Boltzmann Machines; Autoencoders; Sparse Coding

## Bio

Ruslan Salakhutdinov received his PhD in computer science from the University of Toronto in 2009. After spending two post-doctoral years at the Massachusetts Institute of Technology Artificial Intelligence Lab, he joined the University of Toronto as an Assistant Professor in the Department of Computer Science and Department of Statistics. Dr. Salakhutdinov's primary interests lie in statistical machine learning, Deep Learning, probabilistic graphical models, and large-scale optimization. He is an action editor of the Journal of Machine Learning Research and served on the senior programme committee of several learning conferences including NIPS and ICML. He is the recipient of the Early Researcher Award, Connaught New Researcher Award, Alfred P. Sloan Research Fellowship, Microsoft Research Faculty Fellowship, Google Faculty Research Award, and a is Fellow of the Canadian Institute for Advanced Research. A major focus on Dr. Salakhutdinov's work is Deep Learning, an active research area that stems, in large part, from his 2006 Science article co-authored with Geoffrey Hinton. This work developed a formulation for stacked Restricted Boltzmann Machines (RBMs), showing an algorithm capable of learning large-scale multi-layer networks. Dr. Salakhutdinov subsequently pioneered a new class of deep generative models, called Deep Boltzmann Machines (DBMs). These are probabilistic graphical models that contain multiple layers of latent variables. Each nonlinear layer captures progressively more complex patterns of data, which is a promising way of solving visual object recognition, language understanding, and speech perception problems. Dr. Salakhutdinov's contributions to Deep Learning have already received over 5000 citations according to Google Scholar, and have been applied broadly in speech, language, and image analysis.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyright for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

KDD '14, August 24–27, 2014, New York, NY, USA.

ACM 978-1-4503-2956-9/14/08.

<http://dx.doi.org/10.1145/2623330.2630809>