Statistical View of Deep Learning

Ruslan Salakhutdino[*]

Department of Computer Science
Department of Statistical Sciences
University of Toronto
[*]Contact author: rsalakhu@cs.toronto.edu

Keywords: Deep Learning, Graphical Models, Boltzmann Machines

Building intelligent systems that are capable of extracting meaningful representations from high-dimensional data lies at the core of solving many Artificial Intelligence tasks, including visual object recognition, information retrieval, speech perception, and language understanding.

In this talk I will introduce a broad class of statistical deep learning models, including Restricted Boltzmann Machines and Deep Boltzmann Machines, and show that they can learn useful hierarchical representations from large volumes of high-dimensional data with applications in information retrieval, object recognition, and speech perception. I will then describe a class of more complex statistical models that are capable of extracting a unified representation that fuses together multiple data modalities as well as discuss models that can generate natural language descriptions of images. I will show that on several tasks, including modelling images and text, video and sound, these models significantly improve upon many of the existing techniques [1].

References

[1] Salakhutdinov, R. (2015). Learning Deep Generative Models. *Annual Review of Statistics and Its Application 6*, 361–3851.