

Distinguished Lecture Series

EECS COLLOQUIUM

Spring 2016



Learning Structured, Robust, and Multimodal Deep Models

Ruslan Salakhutdinov

Professor, Computer Science and Statistics at the University of Toronto

Abstract

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 first introduce a broad class of deep learning models 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 next introduce deep models that are capable of extracting a unified representation that fuses together multiple data modalities as well as present the Reverse Annealed Importance Sampling Estimator (RAISE) for evaluating these deep generative models. Finally, I will discuss models that can generate natural language descriptions (captions) of images, as well as generate images from captions using attention mechanism. 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.

Biography

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 Departments of Statistics and Computer Science. His primary interests lie in artificial intelligence, machine learning, deep learning, 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 an Alfred P. Sloan Research Fellow, Microsoft Research Faculty Fellow, Canada Research Chair in Statistical Machine Learning, a recipient of the Early Researcher Award, Connaught New Researcher Award, Google Faculty Award, and is a Senior Fellow of the Canadian Institute for Advanced Research.

Wednesday

February 10

4:00 - 5:00 pm

Refreshments

served at 3:30

Hewlett-Packard

Auditorium

306 Soda Hall

For more information please see:

<http://www.eecs.berkeley.edu/Colloquium/>