

Recent Advances and the Future of Recurrent Neural Networks

Sungjin Ahn
University of Montreal

5/1/2017 at 10:30 am
CoRE A 301

Abstract

Although the recent resurgence of Recurrent Neural Networks (RNN) has achieved remarkable advances in sequence modeling, we are still missing many abilities of RNN necessary to model more challenging yet important natural phenomena. In this talk, I introduce some recent advances in this direction, focusing on two new RNN architectures: the Hierarchical Multiscale Recurrent Neural Networks (HM-RNN) and the Neural Knowledge Language Model (NKLM). In the HM-RNN, each layer in a multi-layered RNN learns different time-scales, adaptively to the inputs from the lower layer. The NKLM deals with the problem of incorporating factual knowledge provided by a knowledge graph into RNNs. I argue the advantages of these models and conclude the talk with a discussion on the key challenges that lie ahead.

Bio

Sungjin Ahn is currently a postdoctoral researcher at the University of Montreal, working with Prof. Yoshua Bengio on Deep Learning and its applications. He received his Ph.D. in Computer Science at the University of California, Irvine, under the supervision of Prof. Max Welling. During his Ph.D. program, He co-developed the Stochastic Gradient MCMC algorithms and awarded best paper awards from the International Conference on Machine Learning in 2012 and the ParLearning 2016, respectively. His research interests include deep learning, reinforcement learning, and approximate Bayesian inference.

Faculty Host: Dimitris Metaxas and Vladimir Pavlovic