

SIGML-Seminar Lecture

SPECIAL INTEREST GROUP IN MACHINE LEARNING

Title: Neural Attention Models for Natural Language Grounding and Generation

Date: 14th October, 2015 (Wednesday)

Time: 1600 hrs

Venue: RM101 (CSE extension building)

Tea would be served at 15:45hrs

Abstract:

Neural sequence-to-sequence, encoder-decoder models have recently shown strong promise in the areas of machine translation and image captioning as end-to-end models that require little domain-specific knowledge or resources. Incorporating an attention or alignment step into this encoder-decoder architecture helps further by learning to focus on parts of the input sequence that are salient for generating a particular step in the output sequence. We employ and improve this neural attention architecture for the tasks of natural language grounding and selective generation. Our grounding task is that of direction following, a skill that is essential to realizing effective autonomous agents. Our model uses long short-term memory recurrent neural networks (LSTM-RNNs) to encode natural language instructions and decode to action sequences based upon a world state representation, using a novel multi-input aligner. Our selective generation task is the joint task of content selection and surface realization, where we first encode the full set of over-determined database event records (e.g., in weather forecasting and sportscasting) via an LSTM-RNN, then utilize a novel coarse-to-fine (hierarchical) aligner to identify the small subset of salient records to talk about, and finally employ a decoder to generate free-form descriptions of the aligned, selected records. In contrast to existing methods, our models use no specialized resources (e.g., parsers, lexicons, features). They are therefore generalizable, yet they still achieve the best results reported to-date on multiple benchmark datasets. We also present various ablations

to elucidate the contributions and limitations of the primary components of our models. This is joint work with Hongyuan Mei and Matthew R. Walter at TTI-Chicago.

Bio : Dr. Mohit Bansal is a research assistant professor at TTI-Chicago. He received his Ph.D. and M.S. in CS from UC Berkeley in 2013 (where he was advised by Dan Klein) and his B.Tech. in CSE from the IIT Kanpur in 2008. His research interests are in statistical natural language processing and machine learning, with a particular interest in semantics (lexical, compositional, multimodal), syntactic parsing, question answering, and deep learning. He was the recipient of a Google Faculty Research Award in 2014, an IBM Faculty Award in 2014, an ACL Long Best Paper Honorable Mention (top-5 paper) in 2014, and a Qualcomm Innovation Fellowship in 2011. Webpage: <http://ttic.uchicago.edu/~mbansal/>