

Reading and Reasoning with Vector Representations

Sebastian Riedel
University College London

Abstract

In recent years, vector representations of knowledge have become popular in NLP and beyond. They have at least two core benefits: reasoning with (low-dimensional) vectors tends to lead to better generalisation, and usually scales very well. But they raise their own set of questions: What type of inferences do they support? How can they capture asymmetry? How can explicit background knowledge be injected into vector-based architectures? How can we provide proofs that justify predictions? In this talk, I sketch some initial answers to some of these questions based on our recent work. In particular, I will illustrate how a vector space can simulate the workings of logic.

Biographical Sketch

Sebastian Riedel is a reader in Natural Language Processing and Machine Learning at the University College London (UCL), where he is leading the Machine Reading lab. He is also the head of research at Bloomsbury AI and an Allen Distinguished Investigator. He works in the intersection of Natural Language Processing and Machine Learning, and focuses on teaching machines how to read and reason. He was educated in Hamburg-Harburg (Dipl. Ing) and Edinburgh (MSc., PhD), and worked at the University of Massachusetts Amherst and Tokyo University before joining UCL.