

Negation and Modality in Distributional Semantics

Ed Hovy

Information Sciences Institute
University of Southern California
hovy@isi.edu

Abstract

In Natural Language Processing, negation and modality have mostly been handled using the older, pre-statistical methodologies of formal representations subject to rule-based processing. This fits the traditional treatment of negation and modality in logic-based knowledge representation and linguistics. However, in modern-day statistics-based NLP, how exactly negation and modality should be taken into account, and what role these phenomena play overall, is much less clear. The closest statistics-based NLP gets to semantics at this time is lexical-based word distributions (such as used in word sense disambiguation) and topic models (such as produced by Latent Dirichlet Allocation). What exactly in such representations should a negation or a modality actually apply to? What would, or should, the resulting effects be? The traditional approaches are of little or no help.

In this talk I argue that neither model is adequate, and that one needs a different model of semantics to be able to accommodate negation and modality. The traditional formalisms are impoverished in their absence of an explicit representation of the denotations of each symbol, and the statistics-based word distributions do not support the compositionality required of semantics since it is unclear how to link together two separate word distributions in a semantically meaningful way. A kind of hybrid, which one could call Distributional Semantics, should be formulated to include the necessary aspects of both: the ability to carry explicit word associations that are still partitioned so as to allow negation and modality to affect the representations in intuitively plausible ways is what is required.

I present a specific model of Distributional Semantics that, although still rudimentary, exhibits some of the desired features. I explore the possibilities for accommodating the phenomena of negation and modality. The talk poses many more questions than it answers and is an invitation to consider Distributional Semantics as a model for richer and more semantics-oriented statistics-based NLP.