

What distributional semantics can (and cannot) tell us about meaning

Alessandro Lenci ¹

¹ Università di Pisa and ILC-CNR, Italy

Abstract. Distributional semantics is a mainstream research paradigm in computational linguistics and cognitive science. It is based on a simple assumption: Semantic representations of lexical items can be built by recording their distribution in linguistic contexts. However, whether statistical co-occurrences alone are enough to address deep semantic questions, or whether they merely provide a shallow proxy of lexical meaning, remains an open question. In other words, what is the real descriptive and explanatory adequacy of distributional representations of meaning? In this talk, I explore this issue by presenting some research themes that shed light on the potentialities and the current limits of distributional models of meaning. The first theme is the notion of semantic similarity. Distributional semantics is based on the so-called Distributional Hypothesis stating that lexemes with similar linguistic contexts have similar meanings. However, distributional semantic models are actually more biased towards the much vaguer notion of semantic relatedness. The outcome of distributional models looks like a network of word associations, rather than a semantically structured space. This is an important weakness of current distributional semantic models. Though they have proven to be useful to capture various aspects of the mental lexicon, their limits in properly distinguishing different semantic relations also greatly impair the usability of distributional semantics to model logical inferences. A central aspect of human semantic competence is the ability to compose lexical meanings to form the interpretation of a potentially unlimited number of complex linguistic expressions. but compositionality is surely the “bottleneck” for distributional semantics. How distributional representations can be projected from the lexical to sentence or even discourse level is still an open issue. In this talk, I present a recent proposal of a distributional model of sentence comprehension in which sentence meaning is built by dynamically activating and unifying distributional information about events and their participants.