Distributional Semantics: Yesterday, Today, and Tomorrow

Alessandro Lenci
University of Pisa
alessandro.lenci@unipi.it

Distributional semantics is undoubtedly the mainstream approach to meaning representation in computational linguistics today. It has also become an important paradigm of semantic analysis in cognitive science, and even linguists have started looking at it with growing interest. The popularity of distributional semantics has literally boomed in the era of Deep Learning, when "word embeddings" have become the basic ingredient to "cook" any NLP task. The era of BERT & co. has brought new types of contextualized representations that have often generated hasty claims of incredible breakthroughs in the natural language understanding capability of deep learning models. Unfortunately, these claims are not always supported by the improved semantic abilities of the last generation of embeddings. Models like BERT are still rooted in the principles of distributional learning, but at the same time their goal is more ambitious than generating corpus-based representations of meaning. On the one hand, the embeddings they produce encode much more than lexical meaning, but on the other hand we are still largely uncertain about what semantic properties of natural language they actually capture. Distributional semantics has surely benefited from the successes of the deep learning, but this might even jeopardize the very essence of distributional models of meaning, by making their goals and foundations unclear.

Computational linguistics is a fast-moving field and distributional semantics makes no exception. In doing this, we always risk chasing the last hype model or using pre-trained vectors as black-box tools, without scrutinizing the relationship between distributional learning and meaning representations. The goal of this tutorial is to try to understand what distributional semantics is today, by looking also at what it was yesterday and at its grounding principles. I will present the main concepts, tools and applications of distributional semantics, to foster a critical analysis of its potentialities as well as its limits. This way, we will try to imagine what distributional semantic could and should become tomorrow.