RECENT ADVANCES IN (DEEP) REPRESENTATION LEARNING

Schütze H.

University of Munich, Germany

Traditionally, natural language processing (NLP) systems have made use of resources compiled by (computational) linguists based on linguistic theory that provide rich information about linguistic objects. For example, computational lexica specify morphological paradigms and subcategorization frames of verbs. In contrast, statistical NLP systems frequently start out with no explicit representation of linguistic objects and instead learn what they need from training data on a task-by-task basis. A third approach—which has gained much interest recently—is to learn generic representations of linguistic objects and then reuse them for a wide variety of tasks. Its premise is that giving an NLP system non-task-specific generic information about words and other linguistic objects will help it in performing well at a particular task.

Examples of such generic representation models include the vector space model, dimensionality reduction, clustering and deep learning. I will review recent research results in representation learning and discuss benefits and drawbacks of the three approaches.