Abstract

Human language varies across many dimensions due to numerous factors, and a large body of research in natural language processing (NLP) aims to identify and measure these factors, often combining insights from linguistics and the ability to leverage large amounts of language data. One of the most studied dimensions of language variation in this field is *diachronic language change*, i.e. change of language over time. In NLP, the de-facto standard representation of lexical meaning are *word embeddings* (Turney and Pantel, 2010), which express the meaning of a lexical item as a vector in a multidimensional space. Intuitively, the angle between two such word vectors can be interpreted as the semantic similarity of the words, and the vectors' difference as 'shift' of meaning from one word to the other.

In this thesis, I delve into the research of diachronic language change on the basis of word embeddings. I aim to contribute to research in two ways. I first outline the many methods involved in and related to diachronic research of language via embeddings in order to provide a solid basis for other researchers from fields outside of NLP. Then, secondly, I introduce and explore a new unsupervised method for the detection of systematic conceptual shifts. For this, I make two assumptions; namely that (1) diachronic language varieties should be treated as distinct despite their obvious similarities, and (2) no assumptions about specific changes should be made *a priori*.

The starting point for the investigations are word embedding models of various time intervals of the *Royal Society Corpus* (RSC, Fischer et al., 2020), a collection of English scientific texts that spans from 1665 to 1929. Previous work on language change within the RSC (cf. Bizzoni et al., 2020) employs both corpus linguistic methods and measurements on these word embedding models (or: *spaces*). Differently to this, I aim to measure shifts across spaces and operate directly on these shifts.

In order to make spaces comparable to each other and work under the assumption of bilinguality (assumption 1), I align pairs of spaces via *Gromov-Wasserstein Optimal Transport* (GWOT, Alvarez-Melis and Jaakkola, 2018). GWOT is an optimization problem which creates a probabilistic association between points across two spaces. It achieves this without any knowledge of language; instead it matches up points across spaces by a measure of how similarly they behave in their own respective spaces. This probabilistic association can then be used to project one space onto the other. In order to preserve the diachronic shifts, the projection considers only the most stable concepts (i.e. those concepts for which the 'spatial profile' changes the least across time).

To detect *systematic conceptual shifts*, I calculate vectors of conceptual shift, apply a clustering algorithm, and approximate the meaning of the shift vector's direction with nearest-neighbor search. The result is a set of clusters with human-interpretable labels which express the components of meaning which change over time.

Thorough investigations of GWOT as well as the experiments with the proposed method show that the assumption of bilinguality is not necessary. The proposed method is promising; it confirms previous findings and is able to identify new dynamics of diachronic change, e.g., that similar words tend to change in similar ways. I provide impulses to improve the proposed method and to apply it to further language resources.

References

D. Alvarez-Melis and T. Jaakkola. Gromov-Wasserstein Alignment of Word Embedding Spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1881–1890, Brussels, Belgium, Oct. 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1214. URL https://www.aclweb.org/anthology/D18-1214.

Y. Bizzoni, S. Degaetano-Ortlieb, P. Fankhauser, and E. Teich. Linguistic Variation and Change in 250 years of English Scientific Writing: A Data-driven Approach. *Frontiers in Artificial Intelligence*, (3):73, Apr. 2020.

S. Fischer, J. Knappen, K. Menzel, and E. Teich. The Royal Society Corpus 6.0: Providing 300+ Years of Scientific Writing for Humanistic Study. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 794–802, Marseille, France, May 2020. European Language Resources Association. ISBN 979-10-95546-34-4. URL https://www.aclweb.org/anthology/2020.lrec-1.99.

P. D. Turney and P. Pantel. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37:141–188, 2010. URL https://www.aaai.org/Papers/JAIR/Vol37/JAIR-3705.pdf.