

Characterizing the pragmatic component of distributional vectors in terms of polarity: Experiments on German *über* verbs

Gabriella Lapesa, Max Kisselew, Sebastian Padó, Tilmann Pross, Antje Roßdeutscher

Universität Stuttgart, Institut für Maschinelle Sprachverarbeitung

The assumption that meaning similarity can be quantified by comparing word co-occurrence patterns (Miller and Charles, 1991) is widely accepted in the Computational Linguistics community, and many applications build on it without further questioning the nature of the semantic representations encoded in Distributional Semantic Models (henceforth, DSMs). While recent work indicates that DSMs encode much more fine-grained semantic information – Herbelot and Vecchi (2015) learn quantification relationships, and Gupta et al. (2015) fine-grained properties of named entities, such as countries’ latitude or income – the relationship between distributional and formal, in particular model-theoretic, semantics is still far from clear.

A recent attempt to clarify this relationship is due to Copestake and Herbelot (2013). They compare the co-occurrence information encoded in DSMs to an *ideal distribution*, a language model built by individual speakers by accumulating all the contexts in which the target lexeme occurs; in this framework, contexts correspond to the logical form of sentences/utterances. Moreover, they introduce the concept of a *context set subspace* as a conceptual tool for “zooming in” on theoretically motivated portions of the ideal distribution, such as different senses of the lexeme or specific information on individual entities mentioned in the sentences out of which the ideal distribution has been built.

We start from the observation that the use of logical forms as the basis for building ideal distributions a priori excludes particular aspects of meaning from consideration. In particular, it does not contain information concerning speakers’ attitudes and evaluation about mentioned entities and events. We believe that this is a fundamental gap in Copestake and Herbelot’s model, since evaluative considerations systematically affect speakers’ text planning, and evaluative considerations that are frequent enough arguably become *part of the meaning* of the target word.

Our investigation focuses on the distributional characterization of the meaning shifts produced by derivational processes in German particle verbs. Derivation is an interesting case because it involves the syntax-semantic interface (e.g., change of word class or, within word class, of subcategorization properties) as well as the pragmatic level (e.g., some derivational affixes can be used to mark speakers’ attitude towards the referent of the base, or even the entire utterance¹). The intuition that a pragmatic component of derivational shifts can be identified at a lexeme level is supported by qualitative observations on the distributional behavior of some derivational patterns. First, an inspection of the most salient features of *-in* (+FEMALE) derivatives compared to their bases revealed a potentially presuppositional use when derivatives are built on profession names (e.g., *Frisör*, hairdresser): we noticed that the derivative is characterized by features with a positive evaluative nuance (e.g., “good”, “skilled”), indicating that speakers resort to the female version of certain professions when they want to highlight positive features, countering a stereotypical presupposition. Second, an analysis of the features of some particle verbs (e.g., *überrennen*, “to overrun”) identified a clear negative evaluation concerning events that are perceived as uncontrolled or uncontrollable (an “EXCESS” reading, which is absent in the features of the corresponding base terms and in the construction *über etwas rennen*, “to run over something”).

The contribution of this study is twofold: on the one hand, we are interested in asking, qualitatively, how distributional and theoretically motivated analyses in semantics compare. On the other, we propose a strategy for the characterization of the pragmatic component encoded in distributional vectors: a quantitative approach based on affectivity norms (i.e., lists of words manually or semi-automatically annotated with scores quantifying their association to positive feelings) and polarity lexicons (i.e., lists of words having a clear negative or positive meaning, for example *good* vs. *bad*), employed to assign a polarity/valency score to distributional vectors.

Classifying and predicting *über*-verbs The experimental items for our experiments come from a dataset of 88 German base verbs paired with corresponding *über*-derivations (e.g., *lesen* – *überlesen* (*read* – *skip*);

¹For an overview of the phenomena of the morphology/pragmatics interface, see Dressler and Merlini Barbaresi (1999). For a summary of the semantic, syntactic and pragmatic factors affecting morphological productivity, see Plag (2003, chapter 3).

bringen – überbringen (bring – convey), etc.). A formal semantics analysis of this dataset based on theoretically motivated considerations concerning the spatial contribution of *über-* to the interpretation of the *über+*verb constructions (Pross and Roßdeutscher, 2015) suggests a fairly clear separation of *über-* derivations into four classes: TRANSFER of an object from A to B as in (1-a), APPLICATION of an object to another object as in (1-b), movement ACROSS some obstacle as in (1-c) and exceeding a certain threshold on a scale (MORE) that the base verb provides as in (1-d).

- | | | | | |
|-----|----|--|----|--|
| (1) | a. | überbringen
over.PREFIX.bring
'to deliver' TRANSFER | c. | überspringen
over.PREFIX.skip
'to skip over' ACROSS |
| | b. | überkleben
over.PREFIX.paste
'to paste over' APPLICATION | d. | überbewerten
over.PREFIX.be.PREFIX.value
'to overvalue' MORE |

Focusing on the evaluative component of the *über* shift Rather than predicting *all* distributional properties of words (Lazaridou et al., 2013; Kisselew et al., 2015), many of which are arguably idiosyncratic, we aim at the identification of a *pragmatic context set subspace* in the sense of Copestake and Herbelot (2013). To quantitatively characterize the pragmatic component of the *über* derivational shift we resort to a lexical resource in which words are associated to manual or semi-automatic annotation for a number of dimensions related to polarity/valency: the German Affectivity Norms (Köper and Schulte Im Walde, 2016). The scores contained in this resource are applied to the context features of base and derived terms in the respective distributional vectors ($\vec{derived} - \vec{base}$) to calculate a polarity/valency shift value for base-derived pairs.

The Affectivity Norms dataset contains automatically generated scores for 350k German lemmas which are rated on a 0 to 10 scale for valency (the pleasantness associated to the word), arousal (the intensity of the emotion associated to the word), concreteness (the extent to which the word's referent be perceived), imageability (the extent to which the word's referent can be perceived visually). The Affectivity Norms are a particularly convenient resource for the modeling of our dataset, because they allow us to establish a reliable link between the evaluative/pragmatic component (valency and arousal) and degree of literalness (concreteness and imageability).

We conduct our experiments on bag-of-words models, whose features are interpretable. Given that we plan to extend our investigation to other distributional patterns, we extracted a count DSM with a large target vocabulary including all words in DERivBase (Zeller et al., 2013), a high-coverage morphological resource for German (ca. 280k lemmatized open-class words). To maximize coverage on the German Affectivity Norms, we employed the same word list for the context features. The DSM used in all our experiments has been extracted from the SdeWaC web corpus (Faaß and Eckart, 2013) adopting a symmetric, 5-words context window, and it is scored with positive pointwise mutual information.

For each target word in the dataset (e.g., *lesen*) we calculate a valency score by computing the dot product between the (L1-normalized) target DSM vector elements and the contexts' corresponding Affectivity Norms. As a consequence, our valency scores quantify the proportion of the target vector which is loaded with valency.

We trained a linear regression model to predict the valency values of the meaning shift from the base verb to the prefix verb, i.e., the contribution of the prefix *über*. We employed the following predictors: semantic class of the derived word as annotated in our dataset (TRANSFER, ACROSS, MORE, APPLICATION); arousal and imageability scores of the shift; interaction between imageability and semantic class; log transformed frequencies of base and derived word. The adjusted R-squared value of the model was 0.75 ($p < 0.001$), showing a very good fit of the model to the data. We found a highly significant ($p < 0.001$) negative main effect of arousal on valency: for all classes, the stronger the emotional load introduced by the derivational shift, the more negatively charged the derived word becomes, with respect to its base. Interestingly, we did not find highly significant simple effects for semantic class and imageability, meaning that the verb classes do not differ in valency, nor does imageability influence valency directly. We did, however, find interactions between the two predictors: an inspection of the estimates reveals that, for ACROSS *über* verbs, highly imageable contexts come with a particular negative valency ($p < 0.05$); MORE verbs follow on this trend (slightly negative at increasing imageability, though not significantly so), and then come the TRANSFER verbs (reference value). The effect is reversed for APPLICATION, where higher imageability means positive valency (significant at $p < 0.05$). Our interpretation is that it is the saliency of the crossed spatial boundary which determines the negative valency shift, and that imageability (which we take as a more conservative and cognitively straightforward correlate of concreteness, and therefore as a cue of a literal use) regulates the modulation of the evaluative/pragmatic component.

References

- Ann Copestake and Aurelie Herbelot. *Lexicalized compositionality*. Unpublished manuscript. 2013.
- Wolfgang U. Dressler and Lavinia Merlini Barbaresi. Morphopragmatics. In Jef Verschueren, Jan-Ola Östman, Jan Blommaert and Chris Bulcaen, editors, *Handbook of Pragmatics: 1997 Installment*. John Benjamins Publishing Co. 1999.
- Gertrud Faaß and Kerstin Eckart. SdeWaC – A corpus of parsable sentences from the Web. In Iryna Gurevych, Chris Biemann, and Torsten Zesch, editors, *Language Processing and Knowledge in the Web*, 8105, 61–68, 2013.
- Abhijeet Gupta, Gemma Boleda, Marco Baroni, and Sebastian Padó. Distributional vectors encode referential attributes. *Proceedings of EMNLP*, 12–21, 2015.
- Aurélie Herbelot and Eva-Maria Vecchi. Building a shared world: Mapping distributional to model-theoretic semantic spaces. *Proceedings of EMNLP*, 22–32, 2015.
- Max Kisselew, Sebastian Padó, Alexis Palmer, and Jan Šnajder. Obtaining a better understanding of distributional models of German derivational morphology. *Proceedings of IWCS*, 58–63, 2015.
- Maximilian Köper and Sabine Schulte Im Walde. Automatically Generated Affective Norms of Abstractness, Arousal, Imageability and Valence for 350 000 German Lemmas. *Proceedings of the 10th Conference on Language Resources and Evaluation (LREC)*, 2016
- Angeliki Lazaridou, Marco Marelli, Roberto Zamparelli, and Marco Baroni. Compositional-ly derived representations of morphologically complex words in distributional semantics. *Proceedings of ACL*, 1517–1526, 2013.
- George A. Miller and Walter G. Charles. Contextual correlates of semantic similarity. *Language and Cognitive Processes*, 6:1, 1 – 28, 1991.
- Ingo Plag. *Word-formation in English*. Cambridge Textbooks in Linguistics. Cambridge University Press. 2003.
- Tillmann Pross and Antje Roßdeutscher. Measuring out the relation between conceptual structures and truth-conditional semantics. In Kata Balogh, editor, *Selected papers from the Workshop "Bridging Formal and Conceptual Semantics"*, Düsseldorf University Press. 2015.
- Britta Zeller, Jan Šnajder, and Sebastian Padó. DERivBase: Inducing and Evaluating a Derivational Morphology Resource for German. *Proceedings of ACL*, 1201–1211, 2013.