

SIMULATING LEXICAL SEMANTIC CHANGE FROM SENSE-ANNOTATED DATA

DOMINIK SCHLECHTWEG* and SABINE SCHULTE IM WALDE

*dominik.schlechtweg@ims.uni-stuttgart.de
Institute for Natural Language Processing, University of Stuttgart, Germany

We present a novel procedure to simulate lexical semantic change from synchronic sense-annotated data, and demonstrate its usefulness for assessing lexical semantic change detection models. The induced dataset represents a stronger correspondence to empirically observed lexical semantic change than previous synthetic datasets, because it exploits the intimate relationship between synchronic polysemy and diachronic change. We publish the data and provide the first large-scale evaluation gold standard for LSC detection models.

1. Introduction

Evaluating Lexical Semantic Change (LSC) detection models is notoriously challenging. Existing testsets are flawed because they are too small to allow for generalizing over the results obtained on them. Artificial data, on the other hand, can be created in larger quantities, but typically relies on assumptions that may or may not be correct, such as the strength of semantic relatedness that old and new senses in LSC have. A clear advantage of artificial data is, however, that it allows the precise control of potentially influencing variables such as frequency and polysemy.

After spelling out the implicit assumptions of previous work, this paper presents a novel procedure to simulate lexical semantic change from synchronic sense-annotated data, which we consider more realistic than in earlier approaches. By splitting the synchronic data into two parts reflecting different sense frequency distributions for a word we simulate sense divergences. In a second stage, we define a graded and a binary notion of LSC based on differences between the obtained sense frequency distributions. These notions are then used to calculate the gold scores determining for each sense-annotated word the degree of change and whether senses were gained or lost. With the proposed definitions, we hope to provide a solid foundation for the basic concepts in the field of LSC detection.

2. Related Work

Most previous evaluations for LSC detection models rely on small amounts of empirically observed data, which was either hand-selected (Sagi et al., 2009; Ja-

towt & Duh, 2014; Hamilton et al., 2016a, 2016b; Frermann & Lapata, 2016; del Tredici & Fernández, 2017) or annotated by humans (Cook et al., 2014; Schlechtweg et al., 2017; Tahmasebi & Risse, 2017; Schlechtweg et al., 2018; Perrone et al., 2019). An alternative approach is synthetic evaluation, where pseudo-change is simulated by collapsing uses of different words (Cook & Stevenson, 2010; Kulkarni et al., 2015; Rosenfeld & Erk, 2018; Dubossarsky et al., 2019; Shoemark et al., 2019). This procedure is very similar to the creation of pseudo-polysemy in word sense disambiguation (Schütze, 1998; Pilehvar & Navigli, 2013): Changes in the senses a word w expresses are artificially created by copying a corpus and relabeling the uses of another word w' (known to express different senses) as uses of w in the copy. The word w is then guaranteed to express a different sense in the copy and the difference between the two corpora can be equated with artificial LSC.¹ This procedure mimics the changes in word senses occurring in empirical LSC, but requires assumptions about other factors. One such factor is the choice of words that should be collapsed and what their semantic relation should be. With the exception of Dubossarsky et al. (2019), all of the existing work collapses uses of words that have no semantic relation. This has the advantage that strong differences in the contextual distribution of words are created, which are more easily detected by computational models. However, there is plenty of evidence showing that LSC does not introduce random new senses of a word, but that new senses are very often semantically related to one of the old senses (Blank, 1997). Hence, the changes to the contextual distribution of a word in empirical LSC are often more subtle and hard to detect than the ones introduced in previous synthetic evaluations.

3. Simulation

Polysemy is the synchronic result of lexical semantic change (Blank, 1997; Bybee, 2015). Accordingly, the different senses a word may express today have been developed some time in the past by the word undergoing a process of LSC. We exploit this idea by using the modern (synchronic) senses of a polysemous word to simulate LSC, i.e., we reconstruct the diachronic process using its synchronic result. In this way, we (i) guarantee that the different senses used in the simulation are likely to be semantically related, as the different senses of a word are usually semantically related to each other (Fillmore & Atkins, 2000). And (ii) we simulate divergences for senses which have empirically been attached to the same word and are thus probable candidates to occur in empirical LSC.

3.1. Corpus

To simulate LSC in the above-described way we need sense-annotated data. We use **SemCor**, a sense-tagged corpus of English (Langone et al., 2004) which rep-

¹Rosenfeld and Erk (2018) use a slightly different variation of this procedure.

resents a subset of the Brown Corpus (Francis & Kucera, 1979). SemCor contains 700,000 tokens, of which more than 200,000 are sense-annotated; the corpus is lemmatized and POS-tagged. Similar corpora in other languages (Bentivogli & Pianta, 2005; Henrich & Hinrichs, 2013, e.g.) would allow to easily extend our methodology to create evaluation data for further languages.

3.2. Two concepts of LSC

Graded. In order to simulate LSC, we first need a definition of what we mean by LSC. Previous research has not been explicit about the underlying concepts and often relied on an intuitive notion of *degree of LSC* (Hamilton et al., 2016b; Dubossarsky et al., 2017; Bamler & Mandt, 2017; Rudolph & Blei, 2018; Rosenfeld & Erk, 2018; Schlechtweg et al., 2018). A well-defined concept of graded LSC, however, should enable us to compare any two words over time and decide which of the words changed more. Hence, it should provide an answer to questions like: Did a word that lost a very frequent sense change more than a word that lost a very infrequent sense? And did a word that gained two senses change less than a word that lost three senses? The field is still lacking such a definition of the degree of LSC of a word.

Binary. This graded notion of LSC seems to diverge from the definition applied in historical linguistics, where LSC is typically not assumed to be graded, but binary (Blank, 1997, cf. p. 113). That is, either a word gained/lost a sense over time, or not, while in computational linguistics slight changes to the frequencies of different word senses are usually also considered as instances of LSC (hence the term *degree of LSC*).² This deviation is striking, as the most straightforward application of LSC detection models is their use to aid historical linguists (Hamilton et al., 2016b). Yet, the graded notion is applicable to related tasks, such as detecting the early stages of a meaning loss (as indicated by sense frequency decrease).

Table 1. Corpus sample for the noun *plant*.

This reduces the number of expensive plant shutdowns and startups.	(s1)
The pilot plant was equipped with a 3-hp. turbine aerator (Figure 2).	(s1)
Remove about half the branches from each plant , leaving only the strongest with the largest buds.	(s2)
“On the side toward the horizon – the southern hemisphere – it is spring; plants are being taught to grow”.	(s2)
Can you share medical facilities and staff with neighboring plants ??	(s1)

²Note for completeness that some work in computational linguistics also assumes a binary notion (Cook et al., 2014; Tahmasebi & Risse, 2017; Perrone et al., 2019; Shoemark et al., 2019).

3.2.1. Sense Frequency Distributions

We propose quantitative definitions of the two notions described above. The definitions are based on the concept of a Sense Frequency Distribution (SFD) (McCarthy et al., 2004; Lau et al., 2014, e.g.). A SFD encodes how often a word w occurs in each of its senses. Consider the word *plant*. In SemCor it occurs expressing either of two senses:

- **sense 1:** plant, works, industrial plant (buildings for carrying on industrial labor); “they built a large plant to manufacture automobiles”³
- **sense 2:** plant, flora, plant life (botany: a living organism lacking the power of locomotion)

This means that each use of *plant* in the corpus is assigned to one of these two senses, as we can see in the sample in Table 1: *plant* occurs three times in sense 1 and two times in sense 2. Hence, its SFD is provided as (3,2). Generally, given a set of w ’s uses U where each use is assigned to one of the senses in the sequence $S = (S_1, S_2, \dots, S_i)$, w ’s SFD is defined by

$$T = (f(S_1), f(S_2), \dots, f(S_i))$$

where $f(S_i)$ is the number of times any use from U was mapped to the i th sense in S .

Graded LSC. For two different sets of uses U_1 and U_2 the corresponding SFDs may differ. We will now define two measures quantifying the difference between any two same-sized SFDs, which will correspond to graded and binary LSC. Assume we have two SFDs T_1 and T_2 defined for the same word w and sense sequence S , but for two *different* use sets U_1 and U_2 . We first normalize T_1 and T_2 to probability distributions P and Q by dividing each element by the total sum of the frequencies of all senses in the respective distribution. The degree of LSC of the word w is then defined as the Jensen-Shannon distance between the two normalized frequency distributions:

$$G(w) = JSD(P, Q)$$

where the Jensen-Shannon distance is the symmetrized square root of the Kullback-Leibler divergence (Lin, 1991; Donoso & Sanchez, 2017).⁴ $G(w)$ ranges between 0 and 1 and is high if P and Q assign very different probabilities to the same senses.

³<https://wordnet.princeton.edu/>

⁴We prefer the Jensen-Shannon distance over Kullback-Leibler divergence, because the former is a true metric in contrast to the latter. This means that amongst other properties specific to metrics the Jensen-Shannon distance is symmetric, i.e., $JSD(X, Y) = JSD(Y, X)$.

Binary LSC. Binary LSC of the word w is then defined as

$$B(w) = 1 \text{ if for some } i, P_i = 0.0 \text{ and } Q_i \geq k, \\ \text{or vice versa.} \\ B(w) = 0 \text{ else.}$$

where X_i is the i th element in X and k is a probability threshold set to 0.1. That is, $B(w)$ will be 1 if there is a sense which has at least probability of 0.1 in either P or Q but a probability of 0.0 in the other (a meaning is gained or lost). If this is not the case, $B(w)$ will be 0. Note that $B(w)$ can be seen as an extreme special case of $G(w)$: if a sense is never assigned to any use in U_1 , this sense will have probability 0.0 in P , which will cause $G(w)$ to be higher than for any other possible probability assignment to that sense. However, $G(w)$ will typically not be equal to 1.0 in such a case, as it is also sensitive to the probabilities of the other senses which may have changed only slightly or not at all.

3.3. Algorithm

In order to simulate LSC in SemCor we split it into two parts (C_1, C_2) and extract the corresponding SFDs (T_1, T_2) for each sense-annotated lemma w . From T_1 and T_2 the scores of above-defined gold notions of LSC follow directly. The splitting process has two steps: In step (i), we introduce strong changes for specific target lemmas. For this we sample all lemmas with a frequency between 100 and 1000 and split their sentences into senses, i.e., for each target lemma we randomly shuffle senses and split them at a random index into two subsets. Then we try to assign sentences with senses from the first subset to C_1 and sentences with senses from the second subset to C_2 . This maximizes change, because senses will tend to have uses in only one of C_1, C_2 . In step (ii), all remaining sentences are randomly shuffled, split in half and added to C_1, C_2 respectively. With this process a non-target lemma will tend to have a rather low change score.

Finally, we extract C_1 and C_2 sentence-wise to separate text files, and the LSC scores to a CSV file. For each sentence, all words (no punctuation) are extracted and replaced by their lemma if existent, else the lowercased token is extracted. Annotated phrases like *on_the_other_hand* are split into individual words to increase data size, i.e., *on the other hand*. The resulting corpora C_1 and C_2 have sizes with 0.34M and 0.36M tokens.

Consider Tables 2 and 3, showing the sample corpus from Table 1 split according to steps (i) and (ii) respectively. In Table 2, *plant* is treated as a target lemma and split according to step (i): sense 1 is assigned to C_2 , while sense 2 is assigned to C_1 . The resulting SFDs are $T_1 = (0, 2)$ and $T_2 = (3, 0)$, creating high change scores of $G(\textit{plant}) = 1.0$ and $B(\textit{plant}) = 1$. As the probability changes of the two senses (from 0.0 to 1.0 and vice versa) are the strongest possible, the graded change score is at its maximum. And as *plant* also loses and gains a sense from C_1 to C_2 it shows binary change.

Table 2. Sample corpus split for the target lemma *plant*. $T_1 = (0, 2)$, $T_2 = (3, 0)$, $G(w) = 1.0$ and $B(w) = 1$.

C_1	C_2
remove about half the branch from each plant leave only the strong with the largest bud (s2)	the pilot plant was equip with a 3 hp turbine aerator figure 2 (s1)
on the side toward the horizon the southern hemisphere it be spring plant are being teach to grow (s2)	this reduce the number of expensive plant shutdown and startup (s1)
	can you share medical facility and staff with neighboring plant (s1)

In Table 3, *plant* is treated as a non-target lemma and thus split according to step (ii): both senses are assigned uniformly to C_1 , C_2 . The resulting SFDs are $T_1 = (2, 1)$ and $T_2 = (1, 1)$, creating change scores of $G(\textit{plant}) = 0.14$ and $B(\textit{plant}) = 0$. The probabilities of each sense are relatively similar in the two corpora, which leads to a low graded change score and no binary change.

Table 3. Sample corpus split for the non-target lemma *plant*. $T_1 = (2, 1)$, $T_2 = (1, 1)$, $G(w) = 0.14$ and $B(w) = 0$.

C_1	C_2
remove about half the branch from each plant leave only the strong with the largest bud (s2)	on the side toward the horizon the southern hemisphere it be spring plant are being teach to grow (s2)
the pilot plant was equip with a 3 hp turbine aerator figure 2 (s1)	can you share medical facility and staff with neighboring plant (s1)
this reduce the number of expensive plant shutdown and startup (s1)	

3.4. Testsets

With the corpus split and the extracted change scores we have a large amount of evaluation data available. However, the change scores are subject to noise through non-annotated data. That is, non-annotated uses of words distort the sense frequency distributions on which the change scores are based. In order to minimize this noise we disregard each lemma w that has a relative frequency error $RE(w) \geq 0.5$, where

$$RE(w) = \frac{\#(w) - \#\textit{annotated}(w)}{\#\textit{annotated}(w)}$$

with $\#(w)$ being w 's corpus frequency and $\#\textit{annotated}(w)$ the number of w 's annotated uses. Hence, we allow at most a number of half of w 's annotated uses to be added to these for w to be part of the testset. We additionally disregard any lemma with a lower frequency than 50 in either of C_1 , C_2 . This results in a testset containing 148 lemmas with different change scores. All the data is publicly available and can be used for LSC detection evaluation.⁵

⁵Find the data under: <https://www.ims.uni-stuttgart.de/data/lsc-simul>.

3.5. Discussion

The corpus splitting process described in Section 3.3 controls the degree of change introduced for a particular lemma. However, this process is not built on a particular theoretical model of LSC, i.e., a model of how the underlying sense probability distributions should change to be similar to true LSC. This also determines how much variables such as polysemy and frequency will correlate with simulated LSC in the resulting dataset. The way in which we chose to split the corpus implicitly introduces higher rates of change for more polysemous words, i.e., simulated LSC correlates with polysemy. Similarly, it introduces specific frequency patterns for strongly changing words, i.e., simulated LSC correlates with frequency change. Whether and to which degree this holds for true LSC is still debated, but it is clear that these variables strongly bias model predictions (Hellrich & Hahn, 2016; Dubossarsky et al., 2017). Thus, to make sure that model performances on our dataset do not stem from model biases towards these variables we recommend to report a polysemy and a frequency baseline. Only performances above these baselines can be safely attributed not to stem from model biases.

4. Model Evaluation

We give a short example of how to evaluate LSC detection models on our dataset. We train all vector space models with all alignment techniques from Schlechtweg et al. (2019) on C_1 and C_2 and apply two similarity measures (CD, LND) to the resulting representations to create change score predictions.⁶ Then we use Spearman’s ρ to compare the resulting rankings against the graded change scores and Average Precision (AP) to compare them against the binary change scores. The results are presented in Table 4.

Generally, models show rather weak performances on the testset. The performances for graded change are considerably lower than in Schlechtweg et al. (2019), which may be attributed to the much smaller corpus sizes and the resulting noise. As expected, the frequency and polysemy baselines show positive correlations with change scores. On average the models outperform the frequency baseline for graded and binary change, while the polysemy baseline is only outperformed for binary change. However, the best models always outperform both baselines. Thus, we can conclude that a range of models measure more than just polysemy or frequency change.

The best models are SGNS and SVD with OP and WI (see Table 4) as alignments. This is similar to previous results in that SGNS+OP+CD has outperformed other models and SVD showed generally high performance (Schlechtweg et al.,

⁶Find implementations at <https://github.com/Garrafao/LSCDetection>. Because of the very small corpus size we choose a large window size of $n = 10$ for all models, experiment with low dimensionalities $d = \{30, 100\}$ for SVD, RI and SGNS and train all SGNS with 30 epochs. We set $k = 5$ and $t = none$. The rest of parameters is set as in Schlechtweg et al. (2019).

Table 4. Best and mean ρ (Graded) and AP (Binary) scores across similarity measures (SIM). Scores are averaged over five iterations for models with a random component. The column ‘model’ gives the model with the best score. SGNS = Skip-Gram with Negative Sampling, CD = Cosine Distance, LND = Local Neighborhood Distance, SVD = Singular Value Decomposition, OP = Orthogonal Procrustes, WI = Word Injection, POLY = Polysemy Baseline, FREQ = Normalized Frequency Difference (NFD) Baseline, RAND = Approximate Random Baseline for Binary Classification.

Dataset	Measure	Graded			Binary		
		mean	best	model	mean	best	model
SEMCOR	SIM	0.159	0.451	SGNS+OP+CD	0.182	0.376	SVD+WI+LND
	POLY	0.349	0.349	-	0.151	0.151	-
	FREQ	0.120	0.120	-	0.110	0.110	-
	RAND	-	-	-	0.081	0.081	-

2019). The comparably high performance of WI alignment may be attributed to its strong noise-reducing effect on our small and thus noisy training corpora (Dubossarsky et al., 2019). A surprising observation is the performance of LND, as in the experiments of Schlechtweg et al. CD has constantly outperformed LND. This may be related to the difference between binary and graded change, as Schlechtweg et al. only evaluated on graded change.

5. Conclusion

We simulated lexical semantic change from synchronic sense-annotated data, introduced the first large-scale, synthetic gold standard for LSC detection and showed how to use it for evaluation. As part of our novel procedure, we provided quantitative definitions of various notions of LSC which implicitly underlie previous work; we thus provided a theoretical basis for artificial and empirical LSC detection evaluation. In the future, we will create further gold standards by exploiting sense-annotated data across languages and use our suggested LSC notions for the simulation of pseudo-change. We will also use the data to evaluate diachronic contextualized embeddings (Giulianelli, 2019; Hu et al., 2019).

The simulation procedure we proposed may also have applications in cognitive research on language evolution (Karjus et al., 2018; Nölle et al., 2018; Tinitis et al., 2017) or more dialogue-oriented studies on meaning change (Pleyer, 2017), where it may be used to simulate the semantic development of words over generations or conversations. Similarly, different types of annotated data may be used to simulate specific types of LSC as e.g. literal and non-literal usages of words (Köper & Schulte im Walde, 2016), metaphoric uses (Köper & Schulte im Walde, 2017) or concrete and abstract uses (Naumann et al., 2018).

Acknowledgments

The first author was supported by the Konrad Adenauer Foundation and the CRETA center funded by the German Ministry for Education and Research (BMBF).

References

- Bamler, R., & Mandt, S. (2017). Dynamic word embeddings. In D. Precup & Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning* (Vol. 70, pp. 380–389). International Convention Centre, Sydney, Australia: PMLR.
- Bentivogli, L., & Pianta, E. (2005). Exploiting parallel texts in the creation of multilingual semantically annotated resources: the MultiSemCor Corpus. *Natural Language Engineering*, 11(3), 247–261.
- Blank, A. (1997). *Prinzipien des lexikalischen Bedeutungswandels am Beispiel der romanischen Sprachen*. Tübingen: Niemeyer.
- Bybee, J. L. (2015). *Language change*. Cambridge, United Kingdom: Cambridge University Press.
- Cook, P., Lau, J. H., McCarthy, D., & Baldwin, T. (2014). Novel word-sense identification. In *25th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers* (pp. 1624–1635). Dublin, Ireland.
- Cook, P., & Stevenson, S. (2010). Automatically Identifying Changes in the Semantic Orientation of Words. In N. C. C. Chair, K. Choukri, B. Maegaard, J. Mariani, J. Odijk, S. Piperidis, M. Rosner, & D. Tapias (Eds.), *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*. Valletta, Malta: European Language Resources Association (ELRA).
- del Tredici, M., & Fernández, R. (2017). Semantic variation in online communities of practice. In *IWCS 2017 - 12th International Conference on Computational Semantics - Long papers*.
- Donoso, G., & Sanchez, D. (2017). Dialectometric analysis of language variation in twitter. In *Proceedings of the Fourth Workshop on NLP for Similar Languages, Varieties and Dialects* (pp. 16–25). Valencia, Spain.
- Dubossarsky, H., Hengchen, S., Tahmasebi, N., & Schlechtweg, D. (2019). Time-Out: Temporal Referencing for Robust Modeling of Lexical Semantic Change. In *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 457–470). Florence, Italy: Association for Computational Linguistics.
- Dubossarsky, H., Weinshall, D., & Grossman, E. (2017). Outta control: Laws of semantic change and inherent biases in word representation models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 1147–1156). Copenhagen, Denmark.
- Fillmore, C. J., & Atkins, B. T. (2000). Describing polysemy: The case of 'crawl'. In Y. Ravin & C. Leacock (Eds.), *Polysemy: Theoretical and computational approaches*. Oxford University Press.
- Francis, W. N., & Kucera, H. (1979). *Brown corpus manual* (Tech. Rep.). De-

- partment of Linguistics, Brown University, Providence, Rhode Island, US.
- Frermann, L., & Lapata, M. (2016). A Bayesian model of diachronic meaning change. *Transactions of the Association for Computational Linguistics*, 4, 31–45.
- Giulianelli, M. (2019). *Lexical Semantic Change Analysis with Contextualised Word Representations*. Unpublished master's thesis, University of Amsterdam, Amsterdam.
- Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2016a). Cultural shift or linguistic drift? Comparing two computational measures of semantic change. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 2116–2121). Austin, Texas.
- Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2016b). Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1489–1501). Berlin, Germany.
- Hellrich, J., & Hahn, U. (2016). Bad Company—Neighborhoods in Neural Embedding Spaces Considered Harmful. In *Proceedings of COLING 2016* (pp. 2785–2796). Osaka, Japan.
- Henrich, V., & Hinrichs, E. (2013). Extending the tüba-d/z treebank with germanet sense annotation. In I. Gurevych, C. Biemann, & T. Zesch (Eds.), *Language Processing and Knowledge in the Web* (pp. 89–96). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Hu, R., Li, S., & Liang, S. (2019). Diachronic sense modeling with deep contextualized word embeddings: An ecological view. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 3899–3908). Florence, Italy: Association for Computational Linguistics.
- Jatowt, A., & Duh, K. (2014). A framework for analyzing semantic change of words across time. In *Proceedings of the 14th ACM/IEEE-CS Joint Conference on Digital Libraries* (p. 229238). IEEE Press.
- Karjus, A., Blythe, R. A., Kirby, S., & Smith, K. (2018). Challenges in detecting evolutionary forces in language change using diachronic corpora. *CoRR*, abs/1811.01275.
- Köper, M., & Schulte im Walde, S. (2016). Distinguishing literal and non-literal usage of German particle verbs. In *Proceedings of the 2016 conference of the north American chapter of the association for computational linguistics: Human language technologies* (pp. 353–362). San Diego, California: Association for Computational Linguistics.
- Köper, M., & Schulte im Walde, S. (2017). Improving verb metaphor detection by propagating abstractness to words, phrases and individual senses. In *Proceedings of the 1st Workshop on Sense, Concept and Entity Representations and their Applications* (pp. 24–30). Valencia, Spain: Association for Computational Linguistics.

- Kulkarni, V., Al-Rfou, R., Perozzi, B., & Skiena, S. (2015). Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web, WWW* (pp. 625–635). Florence, Italy.
- Langone, H., Haskell, B. R., & Miller, G. A. (2004). Annotating wordnet. In *Proceedings of the Workshop Frontiers in Corpus Annotation at HLT-NAACL*. Boston, MA, USA.
- Lau, J. H., Cook, P., McCarthy, D., Gella, S., & Baldwin, T. (2014). Learning word sense distributions, detecting unattested senses and identifying novel senses using topic models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 259–270). Baltimore, Maryland: Association for Computational Linguistics.
- Lin, J. (1991). Lin jh.. divergence measures based on the shannon entropy. *IEEE Transactions on Information Theory*, 37, 145-151.
- McCarthy, D., Koeling, R., Weeds, J., & Carroll, J. (2004). Finding predominant word senses in untagged text. In *Proceedings of the 42nd Meeting of the Association for Computational Linguistics (ACL'04), Main Volume* (pp. 279–286). Barcelona, Spain.
- Naumann, D., Frassinelli, D., & Schulte im Walde, S. (2018). Quantitative semantic variation in the contexts of concrete and abstract words. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics* (pp. 76–85). New Orleans, Louisiana: Association for Computational Linguistics.
- Nölle, J., Staib, M., Fusaroli, R., & Tylén, K. (2018). The emergence of systematicity: How environmental and communicative factors shape a novel communication system. *Cognition*, 181, 93–104.
- Perrone, V., Palma, M., Hengchen, S., Vatri, A., Smith, J. Q., & McGillivray, B. (2019). GASC: Genre-aware semantic change for ancient Greek. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change* (pp. 56–66). Florence, Italy: Association for Computational Linguistics.
- Pilehvar, M. T., & Navigli, R. (2013). Paving the way to a large-scale pseudosense-annotated dataset. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 1100–1109). Atlanta, Georgia, USA.
- Pleyer, M. (2017). Protolanguage and mechanisms of meaning construal in interaction. *Language Sciences*, 63, 69–90.
- Rosenfeld, A., & Erk, K. (2018). Deep neural models of semantic shift. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*

- (pp. 474–484). New Orleans, Louisiana.
- Rudolph, M. R., & Blei, D. M. (2018). Dynamic embeddings for language evolution. In *WWW 2018* (pp. 1003–1011). ACM.
- Sagi, E., Kaufmann, S., & Clark, B. (2009). Semantic Density Analysis: Comparing Word Meaning Across Time and Phonetic Space. In *Proceedings of the Workshop on Geometrical Models of Natural Language Semantics* (pp. 104–111). Stroudsburg, PA, USA.
- Schlechtweg, D., Eckmann, S., Santus, E., Schulte im Walde, S., & Hole, D. (2017). German in flux: Detecting metaphoric change via word entropy. In *Proceedings of the 21st Conference on Computational Natural Language Learning* (pp. 354–367). Vancouver, Canada.
- Schlechtweg, D., Häty, A., del Tredici, M., & Schulte im Walde, S. (2019). A Wind of Change: Detecting and Evaluating Lexical Semantic Change across Times and Domains. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 732–746). Florence, Italy: Association for Computational Linguistics.
- Schlechtweg, D., Schulte im Walde, S., & Eckmann, S. (2018). Diachronic Usage Relatedness (DUREl): A Framework for the Annotation of Lexical Semantic Change. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (p. 169-174). New Orleans, Louisiana.
- Schütze, H. (1998). Automatic word sense discrimination. *Computational Linguistics*, 24(1), 97–123.
- Shoemark, P., Liza, F. F., Nguyen, D., Hale, S., & McGillivray, B. (2019). Room to Glo: A systematic comparison of semantic change detection approaches with word embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 66–76). Hong Kong, China: Association for Computational Linguistics.
- Tahmasebi, N., & Risse, T. (2017). Finding individual word sense changes and their delay in appearance. In *Proceedings of the International Conference Recent Advances in Natural Language Processing* (pp. 741–749). Varna, Bulgaria.
- Tinits, P., Nölle, J., & Hartmann, S. (2017). Usage context influences the evolution of overspecification in iterated learning. *Journal of Language Evolution*, 2(2), 148-159.