Short-term Semantic Shifts and their Relation to Frequency Change

Anna Marakasova
Faculty of Informatics
TU Wien
Vienna, Austria
anna.marakasova@tuwien.ac.at

Julia Neidhardt
Faculty of Informatics
TU Wien
Vienna, Austria
neidhardt@ec.tuwien.ac.at

Abstract

We present ongoing research on the relationship between short-term semantic shifts and frequency change patterns by examining the case of the refugee crisis in Austria from 2015 to 2016. Our experiments are carried out on a diachronic corpus of Austrian German, namely a corpus of newspaper articles. We trace the evolution of the usage of words that represent concepts in the context of the refugee crisis by analyzing cosine similarities of word vectors over time as well as similarities based on the words' nearest neighbourhood sets. In order to investigate how exactly the contextual meanings have changed, we measure cosine similarity between the following pairs of words: words describing the refugee crisis, on the one hand, and words indicating the process of mediatization and politicization of the refugee crisis in Austria proposed by a domain expert, on the other hand. We evaluate our approach against expert knowledge. The paper presents the current findings and outlines the directions of the future work.

1 Introduction

Words are prone to change their meaning over time. Automatic detection and analysis of such change is beneficial for a number of theoretical and computational linguistics tasks, as well as for other research areas, i.e., digital humanities, social and political science. Semantic change occurs due to various linguistic and extra-linguistic factors. Although it might be rather challenging to distinguish between these two causes of change, the latter offer an opportunity to gain insights into temporal dynamics of social, cultural, political phenomena reflected in texts.

The discussion about the relationship between semantic evolution of words and their frequencies has been already introduced in the studies of computational semantic change detection (Kulkarni et al., 2015; Hamilton et al., 2016b; Kahmann et al., 2017; Tahmasebi, 2018; Yao et al., 2018; Del Tredici et al., 2019; Vylomova et al., 2019; Haider and Eger, 2019). Although in linguistics, semantic change was shown to be associated with frequency rise (Feltgen et al., 2017), Bowern (2019) emphasized that a change in frequency can be a precursor to semantic change, but does not capture the change itself. Therefore, when studying semantic shift it is important to consider the effect of frequency change on the observed shift.

We contribute to the field by analysing short-term semantic shifts and their relation to frequency change. Furthermore, we propose an evaluation approach as in the case of short-term change previously proposed strategies cannot be adopted. Unlike large-scale semantic shifts that span across decades or centuries, short-term semantic shifts within monthly or yearly time slices reflect changes in usage rather than in meaning as it is understood in lexicography. Therefore, our work aims to answer the following question: with respect to frequency, what patterns that represent contextual meaning change emerge when dealing with short time slice data? More specifically, does a significant increase in frequency indicate a semantic shift towards broadening or narrowing of usage contexts? And the other way round, if relative frequency does not fluctuate much, can we expect stability of contextual meaning? We address these questions by investigating the case of the refugee crisis in Austria from 2015 to 2016 as covered in the Austrian media. We evaluate our approach relying on the evidence from a domain expert.

The rest of the paper is organized as follows: In Section 2 the related work is discussed. In Section 3 the data, on which our analysis is based, is presented. In Section 4 we introduce our method-

ological approach. Details about our experiments are provided in Section 5 and the results are discussed in Section 6. Finally, Section 7 contains our conclusions.

2 Related Work

Recently, there has been seen an increased interest in computational semantic change detection methods. Hamilton et al. (2016a) proposed to distinguish between two types of semantic change: linguistic drift and cultural shift. According to them, comparison of word's vectors in different time periods (global measure) reflects linguistic drift, while comparison of word's nearest neighbors (local measure) is suitable for identifying cultural shifts. Furthermore, Kutuzov et al. (2017) pointed out that when detecting short-term semantic shifts, the neighbourhood based model outperforms successful alignment methods of long-term semantic change. However, Hamilton et al. (2016a) mentioned that a global measure is sensitive to slight change in word usage, hence we investigate both measures.

With a focus on tracing contextual variability over time, we assume that rising similarity of contextual meaning indicates homogenization of overall word use (narrowing), while a decrease in similarity signifies diversification (broadening) (Haider and Eger, 2019). In order to bring light on the direction of narrowing of a word's usage, we look at the semantic change of pairs of words which was shown to be reliable in different experiment settings (Rosin et al., 2017; Kutuzov et al., 2017; Orlikowski et al., 2018; Dridi et al., 2019; Sommerauer and Fokkens, 2019). More specifically, we study temporal dynamics between the words of interest and their most close semantic associates.

One of the studies of contextual variability that explore correlation with frequency is by Cafagna et al. (2019). They explored the effect of frequency on quantifying synchronic semantics shifts between words in two Italian newspaper corpora and found that contextual variability is larger if a word is relatively more frequent in the second corpus. On the other hand, Kahmann et al. (2017); Del Tredici et al. (2019); Vylomova et al. (2019) showed that significant change in frequency does not necessarily imply semantic context change.

Hamilton et al. (2016b) tackled the task of semantic change detection by applying several algorithms on English, German, French and Chinese data, and conducting evaluation against known diachronic change. According to the law of conformity proposed by them, high frequency words are less likely to undergo semantic change. Later research supported (Tahmasebi, 2018; Cafagna et al., 2019; Rodina et al., 2019) or denied the law by showing that frequency does not correlate with semantic change (Dubossarsky et al., 2016).

Dubossarsky et al. (2017) proved that due to the training procedure semantic change models capture both the change in meaning and noise. They demonstrated that in controlled conditions, the reported meaning change effects largely disappear or become considerably smaller. Therefore, when tracing semantic change there is a risk to over interpret differences in word's meaning representations that actually stem from noise. Frequency difference among words possibly accounts for such noise.

3 Data

The Austrian Media Corpus (AMC) is a diachronic text corpus that contains Austrian newspapers, magazines, press releases, transcribed television interviews, news stories from television, etc. from the last thirty years (Ransmayr et al., 2013). With over 44 million articles, it is one of the largest text corpora for German and the largest for Austrian German. The language data is tokenized, part-of-speech tagged and lemmatized. In total, it contains 10.500 billion tokens.

Although the corpus spans over thirty years, we use the data covering ten years, from 2008 to 2017, and only newspaper articles were taken for our analysis. We split the data into yearly spanned subcorpora, thus obtaining ten corpora for the study with their sizes ranging from 141 to 152 million tokens. Next, we extracted lemmas representing common and proper nouns, verbs, adjectives and applied stop words filtering (numerals, names of months and days of the week). A frequency threshold of 100 minimum counts was applied to each subcorpus.

4 Methods

Distributional semantic methods adopt the hypothesis that meaning of a word is conveyed in its co-occurrence relationships (Harris, 1954; Firth, 1957). The semantic similarity of two words is then approximated by the cosine similarity (the

value ranges from 0 to 1) between their vectors that capture information about its co-occurrence statistics. Recent studies mainly make use of dense word representations, usually prediction-based word embeddings models. However, reliability of word2vec based approaches was considered questionable by some studies (Antoniak and Mimno, 2018; Wendlandt et al., 2018; Sommerauer and Fokkens, 2019; Hellrich, 2019; Dubossarsky et al., 2019). Furthermore, Schlechtweg et al. (2019) evaluated various word's meaning representations and showed that although the best run of the word2vec based model strongly outperforms other methods, its mean performance measured over several runs is comparable with the two count-based representations, namely word vectors from a weighted matrix with positive pointwise mutual information (PPMI) scores and truncated singular value decomposition (SVD) of a PPMI matrix[1]. In the light of the above and since PPMI vectors were previously demonstrated to be less affected by frequency effects (Dubossarsky et al., 2017), we opt for using an average score of 50 runs of the PPMI-based model.

As follows, we construct sparse co-occurrence matrices with the window of seven words for each subcorpus and apply PPMI weighting with the parameters suggested by Levy et al. (2015). A PPMI score with a smoothing parameter α of a word w and its context word c is calculated by the following formula:

$$PPMI_\alpha(w,c) = \max(\log(\frac{P(w,c)}{P(w)P_\alpha(c)}), 0)$$

We employ vector and local neighbourhood similarity measures to get an idea on the overall semantic stability/instability and a time point where the possible change occurs (i.e., allows to discover what and when has changed), while pairwise similarity measure provides information on the direction of context shift (i.e., allows to discover how it has changed). In order to compare to what extent an obtained similarity time series correlates with the frequency time series of a particular word we apply autocorrelation based dissimilarity measure which is the Euclidean distance between simple autocorrelation coefficients of the two given time series (Montero et al., 2014). Autocorrelation distance is preferred to a simple cor-

relation metric due to the fact that autocorrelation implicitly normalizes scores which is helpful when dealing with the time series data of different scale.

4.1 Vector similarity measure

To obtain a vector similarity (VS) score for a target word and two given time periods, we first align the corresponding co-occurrence matrices by intersecting their columns. This allows us to then compute cosine similarities of a word between its representation in two subcorpora. For each word, we compute a time series to trace temporal dynamics of contextual meaning. Similarities between each two subsequent time periods are measured. We interpret similarity values as follows: the lower the similarity, the broader the usage of a word.

4.2 Neighbourhood similarity measures

Nearest neighbourhood measures allow to track the change in a set of words semantically related to a target word based on the idea that two words are similar if they are related to similar words (Jeh and Widom, 2002). These methods give a fine-grained idea of meaning dynamics and were shown to be particularly efficient for change point detection (Shoemark et al., 2019). We employ two kinds of neighbourhood measures:

- Second-order cosine similarity (SOCS). Following Hamilton et al. (2016a); Shoemark et al. (2019); Schlechtweg et al. (2019), we create two vectors whose length is the size of the union of a word's most similar terms in the given time slices, cosine similarities between a word and each term provide scores of the vectors. To build the vectors we take 50 most relevant words in each year. We assume that high cosine similarity between these vectors represents stability of contexts in corresponding time slices.

- Rank discounted cumulative gain similarity (RDCG)[2]. RDCG is a improved version of the normalized discounted cumulative gain measure described in Katerenchuk and Rosenberg (2016). We apply RDCG to the set of 50 most similar words for each target in each year ranking words based on their cosine similarity values. The score of 1 means

---

[1]Please, note that there is no available test set to define the best model run for our data.

[2]Jaccard similarity index and Kendall's tau coefficient were tested as well.

that two sets are identical, while close to 0 values signify high difference in ranking.

4.3 Pairwise similarity measure

We estimate the degree of relatedness of two words by computing cosine similarities between their vectors in each time period, thus producing a time series. We believe that in the case of short-term semantic change, increasing similarity of contexts in different time periods means a preference of a specific usage.

5 Experiments

We represent the discourse of the refugee crisis by the following list of words identified by a linguistics specialist: *Asyl* 'asylum', *Asyl-werber* 'asylum seeker', *Ausländer* 'foreigner', *Flüchtling* 'refugee', *Flüchtlingskrise* 'refugee crisis', *Flüchtlingsstrom* 'refugee flow', *Migrant* 'migrant', *Migration* 'migration', *Schlepper* 'illegal migrants smuggler'[3], *Zuwanderung* 'immigration'. Among other parts of speech, we use only nouns as they are particularly sensitive to cultural shifts in meaning which represent the scope of the current work[4]. We explore semantic development of these concepts in the context of the refugee crisis as covered in the media. To enhance the semantic analysis based on these terms and to evaluate the methods we compile two supplementary lists of words which are described below. We expect our approach to show the rise in similarity between between the "seed" concepts and supplementary terms in the years 2015, 2016, 2017 (the refugee crisis period, plus one subsequent year when the flow of refugees declined but the topic of migrants was still widely discussed in the media).

Our evaluation is based on the statement of an expert in critical discourse analysis that addresses the way societal power relations are established and reinforced through language use (Wodak, 2001). The statement focuses on the shift to the far-right in the Austrian party landscape, mainly the Austrian People's Party (*ÖVP*) and the Freedom Party of Austria (*FPÖ*). During the 2016-2017 election campaign, one of the governing parties (*ÖVP*) adopted extreme-right *FPÖ's* policy towards migration and refugees. According to the expert analysis of the political and media

discourse, this shift to the right is reflected in language by referring to the migrants as a threat to the welfare of the nation and instigating fear among citizens. Moreover, it was proposed to build a border fence and set the maximum limit for asylum applicants. This is pointed to by the increased use of the silent concepts such as *Angst* 'fear', *Bedrohung*, 'threat', *Grenze* 'border', *Grenzzaun* 'border fence', *Obergrenze* 'upper limit', *Richtwert* 'guiding number', *Terrorismus* 'terrorism' in the refugee crisis discourse (Rheindorf and Wodak, 2018). These seven terms constitute our first supplementary list.

Surprisingly, there was no outrage following the changed attitude towards the refugees even though some of the *FPÖ* claims are clearly against universal human rights. The expert contends that with the purpose of legitimation of unprecedented policy, certain mobilizing and politicizing concepts, i.e., *Humanität* 'humanitarianism', *Protektionismus* 'protectionism', *Sicherheit* 'security', *Vielvalt* 'diversity' were deployed by politicians and the media (Rheindorf and Wodak, 2018). These words form our second supplementary list.

6 Results

6.1 Frequency correlation

First of all, we explore a degree of correlation between frequency time series of the "seed" words and their related time series of the semantic shift detection methods described in Section 4. Table 1 presents the results. One can clearly notice that the vector similarity measure exhibits high correlation with words' frequency change, while neighbourhood measures show remarkably lower correlation. Moreover, unlike vector similarity measure that gives high frequency correlation for almost all words, results of the neighbourhood measures present a range of patterns (see Table 1).

6.2 Semantic shift analysis

We are particularly interested in the change of the words' self-similarities values in the years 2015, 2016, 2017. Context instability implies that the meaning (or usage) of a word is either broad, or undergoing a process of development. When the similarity between a word's semantic representation of consecutive time slices stays stable and low, that means word's usage is rather broad. The increase of similarity values indicates the process of narrowing of the contexts diversity. This might

---

[3]The meaning is relevant to the context of the refugee crisis.

[4]See (Hamilton et al., 2016a) for details.

Lemma	VS	SOCS	RDCG
Asyl	0.27	0.62	0.55
Asylwerber	0.24	1.0	0.23
Flüchtlingskrise	0.16	0.83	0.19
Flüchtlingsstorm	0.51	0.49	0.35
Flüchtling	0.36	0.49	0.67
Schlepper	0.41	0.9	1.0
Ausländer	0.48	0.75	0.82
Zuwanderung	0.41	0.87	0.8
Migrant	0.26	0.74	0.93
Migration	0.48	0.75	0.5
mean	**0.36**	**0.69**	**0.72**

Table 1: Autocorrelation distance between frequency time series and semantic similarity time series (VS, SOCS and RDCG) for the "seed" words.

happen, particularly, due to the deliberate usage of only specific contexts of the word's meaning.

Frequency change patterns are different for the selected "seed" terms. Thus, most of the words, namely *Asyl*, *Asylwerber*, *Flüchtling*, *Flüchtlingskrise*, *Flüchtlingsstrom*, *Schlepper* have a noticeable peak in 2015 or 2015 and 2016. Other words (*Zuwanderung*, *Ausländer*) have relatively stable frequency counts over time, or show a steady increase in frequency in the last three years of consideration (*Migrant*, *Migration*). We trace semantic dynamics with relation to these patterns; henceforth we refer to them as "peak", "steady", and "stable" patterns respectively.

6.2.1 Vector similarity measure

Most of the "seed" words exhibit an increased similarity of contexts within the years 2015, 2016, 2017. The most characteristic example is the word *Flüchtlingskrise* itself. It appeared in our semantic representation in 2014, showed rather broad usage until 2016 since when the context narrowed down and stayed stable. One exception is *Ausländer* which probably was not used exclusively in the context of the refugee crisis as one would expect (its frequency also stays constant).

The "peak" frequency pattern has its highest similarity scores in 2016 that stay stable or slightly drop in 2017. In contrast, their frequency significantly falls after 2015 or after 2016 (in the case of *Asylwerber*, *Flüchtlingskrise*). *Flüchtling Asylwerber* undergo the same degree of semantic change (their mean variance is 0.017, only *Flüchtlingskrise* has higher, 0.019) and overall are

rather stable (have high mean similarity scores, 0.62 and 0.52 respectively). In comparison, *Flüchtling* is 3.5 times more frequent than *Asylwerber* which frequency is comparable with other words.

For the words *Flüchtlingsstrom*, *Schlepper* self-similarity starts rising from 2014, while frequency stays the same as in 2013 (increasing only in 2015), the word *Zuwanderung* holds very similar semantic tendency, but has "stable" frequency pattern. There is no considerable difference in semantic change degree between such terms as *Asyl*, *Zuwanderung*, and *Migration* that represent three different frequency patterns.
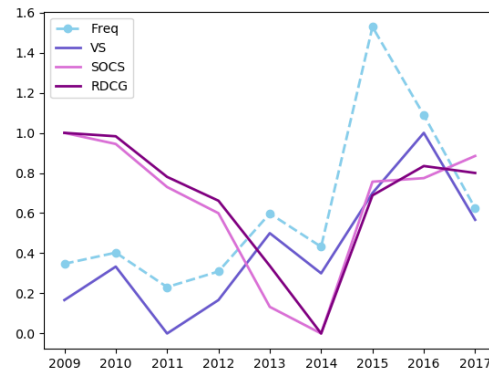

Figure 1: Frequency time series and scaled semantic similarity time series for the word *Asyl*.

6.2.2 Neighbourhood similarity measures

Time series produced by different neighbourhood measures positively correlate with each other and show less correlation with vector similarity measure. In general, neighbourhood statistics also detect an increased context similarity of the "seed" concepts, but have certain distinctive features when comparing to the vector similarity measure. First, the mean similarity scores among words of "peak" frequency pattern is rather smaller, especially, in the last years which are relevant to the refugee crisis, and is around 0.6. Second, while *Flüchtlingskrise*, *Asyl*, *Flüchtlingsstrom* show a comparable rise of self-similarity values, *Asylwerber*, *Flüchtling*, *Schlepper* are found to be relatively stable with the degree of change similar to the one of *Ausländer*, *Zuwanderung*, *Migration*. Third, all "stable" and "steady" frequency pattern terms clearly indicate a change of usage in the year 2015 which is followed by narrowing of the con-

text. Furthermore, the neighbourhood measures captures the change in usage before a drastic frequency rise happens which is illustrated in the example of the word *Asyl* (see Figure 1).

6.3 Pairwise similarity measure

Overall, the pairwise similarity measure shows an increased relatedness of the "seed" concepts and supplementary terms during the period of interest (2015-2017). More precisely, similarity with the first supplementary list words mainly rises in the years 2015-2016 (the actual years of the refugee crisis), whereas similarity with the second supplementary list words becomes higher in 2016-2017 (see Figure 2). We find this results rather reasonable since the first list terms serve as an indication of the direction of the semantic shift of the refugee crisis related concepts, while the second list represents the mediatization process of the restricted policy towards refugees which was particularly apparent during the 2017 election campaign.
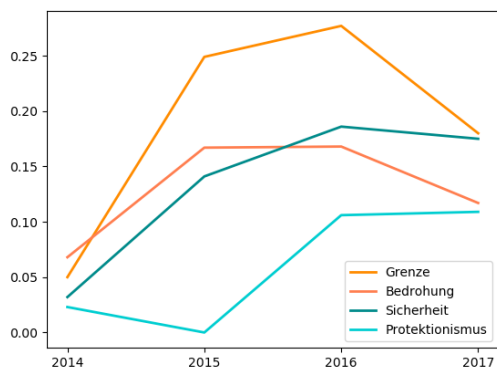

Figure 2: Pairwise similarity for the word *Flüchtlingskrise* and the words from the supplementary lists.

6.4 Evaluation

We match each of the "seed" words with the words from the supplementary lists. In total, there are 110 pairs for evaluation. Next, for every subcorpus we compute cosine similarities between terms that constitute a pair. Recall that according to the expert knowledge, supplementary lists words were widely used in the refugee crisis related discourse in the years 2015 to 2017, i.e., they became more related to the concepts that represent the refugee crisis.

We compare the mean similarity values of the aforementioned period against the mean values of the period of the seven years preceding it (2008-2014). The difference is expected to be positive if our approach goes along with the expert assessment. Indeed, the negative statistically significant difference is only observed in 9% of pairs (10 out of 110). It is rather small (average value for these ten pairs is -0.0087) and found among the pairs with the concepts that do not exhibit specific semantic shift during the refuge crisis (mostly, pairs with the concept *Ausländer*, but also *Zuwanderung-Vielfalt*, *Migrant-Vielfalt*, *Migration-Vielfalt*) which are also related to the discussion of migration in the years 2010-2012.

7 Conclusion

Short-term semantic shift is a complex phenomenon, and understanding the nature of it implies a lot of challenges. Our findings suggest that significant frequency increase is not necessarily followed by significant change in usage, and relatively constant frequencies over time do not imply stability of contextual meaning. We showed that dynamics of word usage is a possible indicator of wider socio-cultural or political shifts.

Since the paper presents an ongoing study, there is a space for enhancement in many aspects. First, more methods should be compared and comprehensive fine-tuning of the models performed with the careful control for randomness and instability that these models feature. Second, one could experiment with the choice of time slices, whether they could be defined empirically or represented in continuous way intersecting one another. Third, candidates for semantic shift could be selected in a robust way and fine-grained annotated data would be beneficial for thorough evaluation and scaling up the experiments for providing more evidence of the phenomenon under discussion. Fourth, control conditions should be introduced in order to ensure verification of the obtained results.

8 Acknowledgments

References

Maria Antoniak and David Mimno. 2018. Evaluating the stability of embedding-based word similarities. *Transactions of the Association for Computational Linguistics*, 6:107–119.

Claire Bowern. 2019. Semantic change and semantic stability: Variation is key. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 48–55, Florence, Italy. Association for Computational Linguistics.

Michele Cafagna, Lorenzo De Mattei, and Malvina Nissim. 2019. Embeddings shifts as proxies for different word use in italian newspapers.

Marco Del Tredici, Raquel Fernández, and Gemma Boleda. 2019. Short-term meaning shift: A distributional exploration. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2069–2075, Minneapolis, Minnesota. Association for Computational Linguistics.

Amna Dridi, Mohamed Medhat Gaber, R Muhammad Atif Azad, and Jagdev Bhogal. 2019. Leap2trend: A temporal word embedding approach for instant detection of emerging scientific trends. *IEEE Access*, 7:176414–176428.

Haim Dubossarsky, Simon Hengchen, Nina Tahmasebi, and Dominik Schlechtweg. 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*, pages 457–470, Florence, Italy. Association for Computational Linguistics.

Haim Dubossarsky, Daphna Weinshall, and Eitan Grossman. 2016. Verbs change more than nouns: a bottom-up computational approach to semantic change. *Lingue e linguaggio*, 15(1):7–28.

Haim Dubossarsky, Daphna Weinshall, and Eitan Grossman. 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*, pages 1136–1145.

Quentin Feltgen, Benjamin Fagard, and J-P Nadal. 2017. Frequency patterns of semantic change: corpus-based evidence of a near-critical dynamics in language change. *Royal Society open science*, 4(11):170830.

John R Firth. 1957. A synopsis of linguistic theory 1930–1955. instudies in linguistic analysis (pp. 1–32).

Thomas Haider and Steffen Eger. 2019. Semantic change and emerging tropes in a large corpus of new high German poetry. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 216–222, Florence, Italy. Association for Computational Linguistics.

William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016a. Cultural shift or linguistic drift? comparing two computational measures of semantic change. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2016, page 2116. NIH Public Access.

William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 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)*, pages 1489–1501, Berlin, Germany. Association for Computational Linguistics.

Zellig S Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.

Johannes Hellrich. 2019. *Word embeddings: reliability & semantic change*, volume 347. IOS Press.

Glen Jeh and Jennifer Widom. 2002. Simrank: a measure of structural-context similarity. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 538–543.

Christian Kahmann, Andreas Niekler, and Gerhard Heyer. 2017. Detecting and assessing contextual change in diachronic text documents using context volatility. In *KDIR*.

Denys Katerenchuk and Andrew Rosenberg. 2016. Rankdcg: Rank-ordering evaluation measure. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, 978-2-9517408-9-1. European Language Resources Association (ELRA).

Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web*, pages 625–635.

Andrey Kutuzov, Erik Velldal, and Lilja Øvrelid. 2017. Tracing armed conflicts with diachronic word embedding models. In *Proceedings of the Events and Stories in the News Workshop*, pages 31–36, Vancouver, Canada. Association for Computational Linguistics.

Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics*, 3:211–225.

Pablo Montero, José A Vilar, et al. 2014. Tsclust: An r package for time series clustering. *Journal of Statistical Software*, 62(1):1–43.

Matthias Orlikowski, Matthias Hartung, and Philipp Cimiano. 2018. Learning diachronic analogies to analyze concept change. In *Proceedings of the Second Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 1–11.

Jutta Ransmayr, Karlheinz Mörth, and Matej Ďurčo. 2013. Linguistic variation in the Austrian Media Corpus: Dealing with the challenges of large amounts of data. In *Proceedings of the 5th International Conference on Corpus Linguistics (CILC 2013)*. Procedia - Social and Behavioral Sciences 95.

Markus Rheindorf and Ruth Wodak. 2018. Borders, fences, and limits—protecting austria from refugees: Metadiscursive negotiation of meaning in the current refugee crisis. *Journal of Immigrant & Refugee Studies*, 16(1-2):15–38.

Julia Rodina, Daria Bakshandaeva, Vadim Fomin, Andrey Kutuzov, Samia Touileb, and Erik Velldal. 2019. Measuring diachronic evolution of evaluative adjectives with word embeddings: the case for english, norwegian, and russian. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 202–209.

Guy D. Rosin, Eytan Adar, and Kira Radinsky. 2017. Learning word relatedness over time. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1168–1178, Copenhagen, Denmark. Association for Computational Linguistics.

Dominik Schlechtweg, Anna Hätty, Marco Del Tredici, and Sabine Schulte im Walde. 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*, page 732–746. Association for Computational Linguistics.

Philippa Shoemark, Farhana Ferdousi Liza, Dong Nguyen, Scott Hale, and Barbara McGillivray. 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)*, pages 66–76.

Pia Sommerauer and Antske Fokkens. 2019. Conceptual change and distributional semantic models: an exploratory study on pitfalls and possibilities. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 223–233.

Nina Tahmasebi. 2018. A study on word2vec on a historical swedish newspaper corpus. In *DHN*, pages 25–37.

Ekaterina Vylomova, Sean Murphy, and Nicholas Haslam. 2019. Evaluation of semantic change of harm-related concepts in psychology. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 29–34.

Laura Wendlandt, Jonathan K. Kummerfeld, and Rada Mihalcea. 2018. Factors influencing the surprising instability of word embeddings. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2092–2102, New Orleans, Louisiana. Association for Computational Linguistics.

Ruth Wodak. 2001. The discourse-historical approach. *Methods of critical discourse analysis*, 1:63–94.

Zijun Yao, Yifan Sun, Weicong Ding, Nikhil Rao, and Hui Xiong. 2018. Dynamic word embeddings for evolving semantic discovery. In *Proceedings of the eleventh acm international conference on web search and data mining*, pages 673–681.