Shared Task: Lexical Semantic Change Detection in German

Adnan Ahmad, Kiflom Desta, Fabian Lang, Dominik Schlechtweg
Institute for Natural Language Processing, University of Stuttgart
adnan.sust@outlook.com, desta7883@gmail.com,
f.lang.1993@web.de, schlecdk@ims.uni-stuttgart.de

arXiv:2001.07786v2 [cs.CL] 11 May 2020

Abstract

Recent NLP architectures have illustrated in various ways how semantic change can be captured across time and domains. However, in terms of evaluation there is a lack of benchmarks to compare the performance of these systems against each other. We present the results of the first shared task on unsupervised lexical semantic change detection (LSCD) in German based on the evaluation framework proposed by Schlechtweg et al. (2019).

1 Introduction

Natural languages evolve and words have always been subject to semantic change over time (Traugott and Dasher, 2001). With the rise of large digitized text resources recent NLP technologies have made it possible to capture such change with vector space models (Pennington, Socher, and Manning, 2014; Rudolph and Blei, 2017; Bengio, Ducharme, Vincent, and Jauvin, 2003; Rosenfeld and Erk, 2018), topic models (Wang and McCallum, 2006; Lau, Cook, McCarthy, Newman, and Baldwin, 2012; Frermann and Lapata, 2016), and sense clustering models (Mitra et al., 2015). However, many approaches for detecting LSC differ profoundly from each other and therefore drawing comparisons between them can be challenging (Tahmasebi et al., 2018). Not only do architectures for detecting LSC vary, their performance is also often evaluated without access to evaluation data or too sparse data sets. In cases where evaluation data is available, oftentimes LSCD systems are not evaluated on the same data set which hinders the research community to draw comparisons.

For this reason we report the results of the first shared task on unsupervised lexical semantic change detection in German[1] that is based on an annotated data set to guarantee objective reasoning throughout different approaches. The task was organized as part of the seminar 'Lexical Semantic Change Detection' at the IMS Stuttgart in the summer term of 2019.[2]

2 Task

The goal of the shared task was to create an architecture to detect semantic change and to rank words according to their degree of change between two different time periods. Given two corpora C_a and C_b, the target words had to be ranked according to their degree of lexical semantic change between C_a and C_b as annotated by human judges. A competition was set up on Codalab and teams mostly consisting of 2 people were formed to take part in the task. There was one group consisting of 3 team members and two individuals who entered the task on their own. In total there were 12 LSCD systems participating in the shared task.

The shared task was divided into three phases, i.e., development, testing and analysis phase. In the development phase each team implemented a first version of their model based on a trial data set and submitted it subsequently. In the testing phase the testing data was made public and participants applied their models to the test data with a restriction of possible result uploads to 30. The leaderboard was public at all times. Eventually, the analysis phase was entered and the models of the testing phase were evaluated in terms of the predictions they made and parameters could be tuned further. The models and results will be discussed in detail in sections 7 and 8.

[1] https://codalab.lri.fr/competitions/560

[2] https://www.f05.uni-stuttgart.de/informatik/dokumente/Seminare/Seminare-SS_2019/04_LexicalSemantics.pdf

3 Corpora

The task, as framed above, requires to detect the semantic change between two corpora. The two corpora used in the shared task correspond to the diachronic corpus pair from Schlechtweg et al. (2019): DTA18 and DTA19.[3] They consist of subparts of DTA corpus (Deutsches Textarchiv, 2017) which is a freely available lemmatized, POS-tagged and spelling-normalized diachronic corpus of German containing texts from the 16th to the 20th century. DTA18 contains 26 million sentences published between 1750-1799 and DTA19 40 million between 1850-1899. The corpus version used in the task has the following format: "year [tab] lemma1 lemma2 lemma3 ...".

4 Evaluation

The Diachronic Usage Relatedness (DURel) gold standard data set includes 22 target words and their varying degrees of semantic change (Schlechtweg et al., 2018). For each of these target words a random sample of use pairs from the DTA corpus was retrieved and annotated. The annotators were required to rate the pairs according to their semantic relatedness on a scale from 1 to 4 (unrelated - identical meanings) for two time periods. The average Spearman's ρ between the five annotators was 0.66 for 1,320 use paris. The resulting word ranking of the DURel data set is determined by the mean usage relatedness across two time periods and is used as the benchmark to compare the models performances in the shared task.

4.1 Metric

The output of a system with the target words in the predicted order is compared to the gold ranking of the DURel data set. As the metric to assess how well the model's output fits the gold ranking Spearman's ρ was used. The higher Spearman's rank-order correlation the better the system's performance.

4.2 Baselines

Models were compared to two baselines for the shared task:

1. log-transformed normalized frequency difference (FD)
2. count vectors with column intersection and cosine distance (CNT + CI + CD)

The window size for CNT + CI + CD was 10. Find more information on these models in Schlechtweg et al. (2019).

5 Participating Systems

Participants mostly rely on the models compared in Schlechtweg et al. (2019) and apply modifications to improve them.[4] In particular, most teams make use of skip-gram with negative sampling (SGNS) based on Mikolov et al. (2013) to learn the semantic spaces of the two time periods and orthogonal procrustes (OP) to align these vector spaces, similar to the approach by Hamilton et al. (2016). Different meaning representations such as sense clusters are used as well. As measure to detect the degree of LSC all teams except one choose cosine distance (CD). This team uses Jensen-Shannon distance (JSD) instead, which computes the distance between probability distributions (Lin, 1991).

The models of each team will be briefly introduced in this section.

sorensbn Team sorensbn makes use of SGNS + OP + CD to detect LSC. They use similar hyperparameters as in Schlechtweg et al. (2019) to tune the SGNS model. They use an open-sourced noise-aware implementation to improve the OP alignment (Yehezkel Lubin et al., 2019).

tidoe Team tidoe builds on SGNS + OP + CD, but they add a transformation step to receive binarized representations of matrices (Faruqui et al., 2015). This step is taken to counter the bias that can occur in vector-space models based on frequencies (Dubossarsky et al., 2017).

in vain The team applies a model based on SGNS with vector initialization alignment and cosine distance (SGNS + VI + CD). Vector initialization is an alignment strategy where the vector space learning model for t_2 is initialized with the vectors from t_1 (Kim et al., 2014). Since SGNS + VI + OP does not perform as well as other models in Schlechtweg et al. (2019), they alter the vector initialization process by initializing on the complete model instead of only the word matrix of t_1 to receive improved results.

[3] https://www.ims.uni-stuttgart.de/forschung/ressourcen/korpora/wocc/

[4] Find implementations at https://github.com/Garrafao/LSCDetection.

Team	Space	Align	Measure	Spearman	Comment
sorensbn	SGNS	OP	CD	.854	Noise-aware alignment
tidoe	SGNS	OP	CD	.811	Binarized matrices
in vain	SGNS	VI	CD	.802	
Evilly	SGNS	OP	CD	.730	Frequency-driven OP alignment
DAF	fastText	OP	CD	.570	
SnakesOnAPlane	SGNS	OP	CD / JSD	.565 / .561	Measure comparison
TeamKulkarni15	SGNS	OP	CD	.540	Local alignment with KNN
Bashmaistori	PPMI	WI	CD	.511	
Baseline 2	CNT	CI	CD	.486	
giki	PPMI	CI	CD	.432	
Edu-Phil	fastText	OP	CD	.381	
orangefoxes	SGNS	-	CD	.121	DiffTime
Loud Whisper	-	-	-	.092	Graph-based approach
Baseline 1	-	-	FD	.019	

Table 1: Shared task - Overview of participating systems. Legend: Space = Semantic space; Align = Alignment method; Measure = Distance Measure for LSC detection. **Note:** The best result either obtained in testing or analysis phase is reported.

Evilly In line with previous approaches, team Evilly builds upon SGNS + OP + CD. They alter the OP step by using only high-frequency words for alignment.

DAF Team DAF uses an architecture based on learning vectors with fastText, alignment with unsupervised and supervised variations of OP, and CD, using the MUSE package[5] (Conneau, Lample, Ranzato, Denoyer, and Jégou, 2017; Joulin, Grave, Bojanowski, Douze, Jégou, and Mikolov, 2016). For the supervised alignment stop words are used. The underlying assumption is that stop words serve as functional units of language and their usage should be consistent over time.

SnakesOnAPlane The team learns vector spaces with count vectors, positive pointwise mutual information (PPMI), SGNS and uses column intersection (CI) and OP as alignment techniques where applicable. Then they compare two distance measures (CD and JSD) for the different models CNT + CI, PPMI + CI and SGNS + OP to identify which measure performs better for these models. They also experiment with different ways to remove negative values from SGNS vectors, which is needed for JSD.

TeamKulkarni15 TeamKulkarni15 uses SGNS + OP + CD with the modification of local alignment with k nearest neighbors, since other models often use global alignment that can be prone to noise (Kulkarni et al., 2014).

Bashmaistori They use word injection (WI) alignment on PPMI vectors with CD. This approach avoids the complex alignment procedure for embeddings and is applicable to embeddings and count-based methods. They compare two implementations of word injection (Dubossarsky, Hengchen, Tahmasebi, and Schlechtweg, 2019; Schlechtweg, Hätty, Del Tredici, and Schulte im Walde, 2019) as these showed different results on different data sets.

giki Team giki uses PPMI + CI + CD to detect LSC. They state that a word sense is determined by its context, but relevant context words can also be found outside a predefined window. Therefore, they use tf-idf to select relevant context (Ramos et al., 2003).

Edu-Phil Similar to team DAF they also use fastText + OP + CD. Their hypothesis is that fastText may increase the performance for less frequent words in the corpus since generating word embeddings in fasttext is based on character n-grams.

orangefoxes They use the model by Rosenfeld and Erk (2018) which is based on SGNS, but avoids alignment by treating time as a vector that

[5]Find package at: https://github.com/facebookresearch/MUSE

may be combined with word vectors to get time-specific word vectors.

Loud Whisper Loud Whisper base their approach on Mitra et al. (2015) which is a graph-based sense clustering model. They process the data set to receive bigrams, create a co-occurence graph representation and after clustering assess the type of change per word by comparing the results against an intersection table. Their motivation is not only to use a graph-based approach, but to extend the approach by enabling change detection for all parts of speech as opposed to the original model.

6 Results and Discussion

Table 1 shows the results of the shared task. All teams receive better results than baseline 1 (FD), of which a total of 8 teams outperform baseline 2 (CNT + CI + CD). The 4 top scores with $\rho > 0.7$ are either modified versions of SGNS + OP + CD or use SGNS + VI + CD. The following 4 scores in the range of $0.5 < \rho < 0.6$ are generated by the models fastText + OP + CD, SGNS + OP + CD/JSD, and PPMI + WI + CD.

Contrary to the results by Schlechtweg et al. (2019) the modified version of vector initialization shows high performance similar to OP alignment, as previously reported by Hamilton et al. (2016). Some modifications to the SGNS + OP + CD approach are able to yield better results than others, e.g. noise-aware alignment and binarized matrices as compared to frequency-driven OP alignment or local alignment with KNN. Team SnakesOnAPlane compare two distance measures and their results show that JSD ($\rho = .561$) performs minimally worse than CD ($\rho = .565$) as the semantic change measure for their model.

The overall best-performing model is SkipGram with orthogonal alignment and cosine distance (SGNS + OP + CD) with similar hyperparameters as in the model architecture described previously (Schlechtweg et al., 2019). Said architecture was used as the basis for the two best performing models. Team tidoe reports that binarizing matrices leads to a generally worse performance ($\rho = .811$) compared to the unmodified version of SGNS + OP + CD ($\rho = 0.9$). The noise aware alignment approach applied by team sorensbn obtains a higher score ($\rho = .854$) compared to the result reported by tidoe, but is unable to exceed the performance of the unmodified SNGS + OP + CD for the same set of hyperparameters (window size = 10, negative sampling = 1; subsampling = None). Of the 8 scores above the second baseline, 5 use an architecture that builds upon SGNS + OP + CD. Whereas in the lower score segment $\rho < 0.5$ none of the models use SGNS + OP + CD. These findings are in line with the results reported by Schlechtweg et al. (2019), however the overall best results are lower in this shared task, which is expected from the smaller number of parameter combinations explored. Additionally, in the shared task the objective was to report the best score and not to calculate the mean which makes it more difficult to compare the robustness of the models presented here.

References

Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155.

Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.

Deutsches Textarchiv. Grundlage fuer ein Referenzkorpus der neuhochdeutschen Sprache. Herausgegeben von der Berlin-Brandenburgischen Akademie der Wissenschaften [online]. 2017.

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. 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 1147–1156, Copenhagen, Denmark.

Manaal Faruqui, Yulia Tsvetkov, Dani Yogatama, Chris Dyer, and Noah A. Smith. 2015. Sparse overcomplete word vector representations. *CoRR*, abs/1506.02004.

Lea Frermann and Mirella Lapata. 2016. A bayesian model of diachronic meaning change. *Transactions of the Association for Computational Linguistics*, 4:31–45.

William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. 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.

Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Hérve Jégou, and Tomas Mikolov. 2016. Fasttext.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.

Yoon Kim, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. Temporal analysis of language through neural language models. In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 61–65, Baltimore, MD, USA. Association for Computational Linguistics.

Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2014. Statistically significant detection of linguistic change. *CoRR*, abs/1411.3315.

Jey Han Lau, Paul Cook, Diana McCarthy, David Newman, and Timothy Baldwin. 2012. Word sense induction for novel sense detection. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 591–601. Association for Computational Linguistics.

Jianhua Lin. 1991. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Sunny Mitra, Ritwik Mitra, Suman Kalyan Maity, Martin Riedl, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2015. An automatic approach to identify word sense changes in text media across timescales. *Natural Language Engineering*, 21(5):773–798.

Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.

Juan Ramos et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, volume 242, pages 133–142. Piscataway, NJ.

Alex Rosenfeld and Katrin Erk. 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, Volume 1 (Long Papers)*, pages 474–484, New Orleans, Louisiana. Association for Computational Linguistics.

Maja Rudolph and David Blei. 2017. Dynamic bernoulli embeddings for language evolution. *arXiv preprint arXiv:1703.08052*.

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*, pages 732–746, Florence, Italy. Association for Computational Linguistics.

Dominik Schlechtweg, Sabine Schulte im Walde, and Stefanie Eckmann. 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, Volume 2 (Short Papers)*, pages 169–174, New Orleans, Louisiana. Association for Computational Linguistics.

Nina Tahmasebi, Lars Borin, and Adam Jatowt. 2018. Survey of Computational Approaches to Diachronic Conceptual Change. *arXiv e-prints*.

Elizabeth Closs Traugott and Richard B Dasher. 2001. *Regularity in semantic change*, volume 97. Cambridge University Press.

Xuerui Wang and Andrew McCallum. 2006. Topics over time: a non-markov continuous-time model of topical trends. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 424–433. ACM.

Noa Yehezkel Lubin, Jacob Goldberger, and Yoav Goldberg. 2019. Aligning vector-spaces with noisy supervised lexicon. 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 460–465, Minneapolis, Minnesota. Association for Computational Linguistics.