

UiO-UvA at SemEval-2020 Task 1: Contextualised Embeddings for Lexical Semantic Change Detection

Andrey Kutuzov

University of Oslo, Norway
andreku@ifi.uio.no

Mario Giulianelli

University of Amsterdam, Netherlands
m.giulianelli@uva.nl

Abstract

We apply contextualised word embeddings to lexical semantic change detection in the SemEval-2020 Shared Task 1. This paper focuses on Subtask 2, ranking words by the degree of their semantic drift over time. We analyse the performance of two contextualising architectures (BERT and ELMo) and three change detection algorithms. We find that the most effective algorithms rely on the cosine similarity between averaged token embeddings and the pairwise distances between token embeddings. They outperform strong baselines by a large margin (in the post-evaluation phase, we have the best Subtask 2 submission for SemEval-2020 Task 1), but interestingly, the choice of a particular algorithm depends on the distribution of gold scores in the test set.

1 Introduction

Words change their meaning over time in all natural languages. The emergence of large representative historical corpora and powerful data-driven semantic models has allowed researchers to track meaning change. SemEval-2020 Shared Task 1 challenges its participants to classify a list of target words into stable or changed (Subtask 1) and/or to rank these words by the degree of their semantic change (Subtask 2) (Schlechtweg et al., 2020). The task is multilingual: it includes four lists of target words, respectively for English, German, Latin, and Swedish (several dozen words each). Each word list is accompanied with two historical corpora of varying size, consisting of texts created in two different time periods. The shared task organisers additionally provided frequency-based and distributional baseline methods.

We participated in Subtask 2¹ as the **UiO-UvA** team and investigated the potential of *contextualised embeddings* for the detection of lexical semantic change. Our systems are based on ELMo (Peters et al., 2018) and BERT (Devlin et al., 2019) language models, and employ 3 different algorithms to compare contextualised embeddings diachronically. Our *evaluation phase* submission to the shared task ranked 9th out of 34 participating teams, while our *post-evaluation phase* submission is the best from those published on the shared task website² (but some knowledge of the test sets statistics was needed, see below). In this paper, we extensively evaluate all combinations of architectures, training corpora and change detection algorithms, using 5 test sets in 4 languages.

Our main findings are twofold: 1) In 3 out of 5 test sets, ELMo consistently outperforms BERT, while being much faster in training and inference; 2) Cosine similarity of averaged contextualised embeddings and average pairwise distance between contextualised embeddings are the two best-performing change detection algorithms, but different test sets show strong preference to either the former or the latter. This preference shows strong correlation with the distribution of gold scores in a test set. While it may indicate that there is a bias in the available test sets, this finding remains yet unexplained.

Our implementations of all the evaluated algorithms are available at <https://github.com/akutuzov/semEval2020>, and the ELMo models we trained can be downloaded from the NLPL vector repository at <http://vectors.nlpl.eu/repository/> (Fares et al., 2017).

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

¹We did not specifically focus on the binary Subtask 1; our submission achieved the average accuracy of 0.587 in this track.

²<https://competitions.codalab.org/competitions/20948>

2 Background

Lexical semantic change detection (LSCD) task is determining whether and/or to what extent the meaning of a set of target words has changed over time, with the help of time-annotated corpora (Kutuzov et al., 2018; Tahmasebi et al., 2018). LSCD is often addressed using distributional semantic models: time-sensitive word representations (so-called ‘diachronic word embeddings’) are learned and then compared between different time periods. This is a fast developing NLP sub-field with a number of influential papers ((Gulordava and Baroni, 2011; Kim et al., 2014; Kulkarni et al., 2015; Hamilton et al., 2016; Bamler and Mandt, 2017; Kutuzov et al., 2017; Del Tredici et al., 2019; Dubossarsky et al., 2019), and many others).

Most of the previous work used some variation of ‘static’ word embeddings where each occurrence of a word form is assigned the same vector representation independently of its context. Recent contextualised architectures allow for overcoming this limitation by taking sentential context into account when inferring word token representations. However, application of such architectures to diachronic semantic change detection was up to now limited (Hu et al., 2019; Giulianelli et al., 2020; Martinc et al., 2020). While all these studies use BERT as their contextualising architecture, we extend our analysis to ELMo and perform a systematic evaluation of various semantic change detection approaches for both language models.

ELMo (Peters et al., 2018) was arguably the first contextualised word embedding architecture to attract wide attention of the NLP community. The network architecture consists of a two-layer Bidirectional LSTM on top of a convolutional layer. BERT (Devlin et al., 2019) is a Transformer with self-attention trained on masked language modelling and next sentence prediction. While Transformer architectures have been shown to outperform recurrent ones in many NLP tasks, ELMo allows faster training and inference than BERT, making it more convenient to experiment with different training corpora and hyperparameters.

3 System overview

Our approach relies on contextualised representations of word occurrences often referred to as *contextualised embeddings*. Given two time periods t_1, t_2 , two corpora C_1, C_2 , and a set of target words, we use a neural language model to obtain *contextualised embeddings* of each occurrence of the target words in C_1 and C_2 and use them to compute a continuous change score. This score indicates the degree of semantic change undergone by a word between t_1 and t_2 , and the target words are ranked by its value.

More precisely, given a target word w and its sentential context $s = (v_1, \dots, v_i, \dots, v_m)$ with $w = v_i$, we extract the activations of a language model’s hidden layers for sentence position i . The N_w contextualised embeddings collected for w can be represented as the usage matrix $\mathbf{U}_w = (\mathbf{w}_1, \dots, \mathbf{w}_{N_w})$. The time-specific usage matrices $\mathbf{U}_w^1, \mathbf{U}_w^2$ for time periods t_1 and t_2 are used as input to all the tested metrics of semantic change. We use three change detection algorithms:

1. Inverted cosine similarity over word prototypes (PRT) Given two usage matrices $\mathbf{U}_w^{t_1}, \mathbf{U}_w^{t_2}$, the degree of change of w is calculated as the inverted similarity³ between the average token embeddings (‘prototypes’) of all occurrences of w in the two time periods:

$$PRT(\mathbf{U}_w^{t_1}, \mathbf{U}_w^{t_2}) = \frac{1}{d\left(\frac{\sum_{\mathbf{x}_i \in \mathbf{U}_w^{t_1}} \mathbf{x}_i}{N_w^{t_1}}, \frac{\sum_{\mathbf{x}_j \in \mathbf{U}_w^{t_2}} \mathbf{x}_j}{N_w^{t_2}}\right)} \quad (1)$$

where $N_w^{t_1}$ and $N_w^{t_2}$ are the number of occurrences of w in time periods t_1 and t_2 , and d is a similarity metric, for which we use cosine similarity. This method is similar to the standard LSCD workflow with *static* embeddings produced by Procrustes-aligned time-specific distributional models (Hamilton et al., 2016), with the only additional step of averaging token embeddings to create a single vector. Since we want the algorithm to produce higher scores for the words which changed more, the inverted value of cosine similarity is used as the prediction.

³We also tried to use cosine distance ($1 - d$) instead of inverted cosine similarity, but the results were marginally worse.

2. Average pairwise cosine distance between token embeddings (APD) Here, the degree of change of w is measured as the average distance between any two embeddings from different time periods (Giulianelli et al., 2020):

$$APD(\mathbf{U}_w^{t_1}, \mathbf{U}_w^{t_2}) = \frac{1}{N_w^{t_1} \cdot N_w^{t_2}} \sum_{\mathbf{x}_i \in \mathbf{U}_w^{t_1}, \mathbf{x}_j \in \mathbf{U}_w^{t_2}} d(\mathbf{x}_i, \mathbf{x}_j) \quad (2)$$

where d is the cosine distance. High APD values indicate stronger semantic change.

3. Jensen-Shannon divergence (JSD) This measure relies on the partitioning of embeddings into clusters of similar word usages. We follow Giulianelli et al. (2020) and create a single usage matrix with occurrences from two corpora $[\mathbf{U}_w^{t_1}; \mathbf{U}_w^{t_2}]$. We then standardise it and cluster its entries using Affinity Propagation (Frey and Dueck, 2007), which automatically selects a number of clusters for each word (Martinc et al., 2020). Finally, we define probability distributions $\mathbf{u}_w^{t_1}, \mathbf{u}_w^{t_2}$ based on the normalised counts of word occurrences from each cluster (Hu et al., 2019) and compute a JSD score (Lin, 1991):

$$\text{JSD}(\mathbf{u}_w^{t_1}, \mathbf{u}_w^{t_2}) = \text{H}\left(\frac{1}{2}(\mathbf{u}_w^{t_1} + \mathbf{u}_w^{t_2})\right) - \frac{1}{2}(\text{H}(\mathbf{u}_w^{t_1}) - \text{H}(\mathbf{u}_w^{t_2})) \quad (3)$$

JSD scores measure the amount of change in the proportions of word usage clusters across time periods. A high JSD score indicates a high degree of lexical semantic change.

4 Experimental setup

For each of the 4 languages of the shared task, we train 4 ELMo model variants: 1) **Pre-trained**, an ELMo model trained on the respective Wikipedia corpus (English, German, Latin or Swedish)⁴; 2) **Fine-tuned**, the same as Pre-trained but further fine-tuned on the union of the two test corpora; 3) **Trained on test**, trained only on the union of the two test corpora; 4) **Incremental**, two models—the first is trained on the first test corpus, and the second is the same model further trained on the second test corpus. The ELMo models are trained for 3 epochs (except English and Latin **Trained on test** and **Incremental** models, for which we use 5 epochs, due to small test corpora sizes), with the LSTM dimensionality of 2048, batch size 192 and 4096 negative samples per batch. All the other hyperparameters are left at their default values.⁵

For BERT, we use the *base* version, with 12 layers and 768 hidden dimensions.⁶ For English, German and Swedish, we employ language-specific models: *bert-base-uncased*, *bert-base-german-cased*, and *af-ai-center/bert-base-swedish-uncased*. For Latin, we resort to *bert-base-multilingual-cased*, since there is no specific Latin BERT available yet. Given the limited size of the test corpora (in the order of 10^8 word tokens at max), we do not train BERT from scratch and only test the **Pre-trained** and **Fine-tuned** BERT variants. The fine-tuning is done with BERT’s standard objective for 2 epochs (English was trained for 5 epochs). We configure BERT’s WordPiece tokeniser to never split any occurrences of the target words (some target words are split by default into character sequences) and we add unknown target words to BERT’s vocabulary. We perform this step both before fine-tuning and before the extraction of contextualised representations.

At inference time, we use all ELMo and BERT variants to produce contextualised representations of all the occurrences of each target word in the test corpora. For the **Incremental** variant, the representations for the occurrences in each of the two test corpora are produced using the respective model trained on this corpus. The resulting embeddings are of size 12×768 and 3×512 for BERT and ELMo, respectively. We employ three strategies to reduce their dimensionality to that of a single layer: 1) using only the top layer, 2) averaging all layers, 3) averaging the last four layers (only for BERT embeddings, as this aggregation method was shown to work on par with the all-layers alternative in (Devlin et al., 2019)). Finally, to

⁴The Wikipedia corpora were lemmatised using UDPipe (Straka and Straková, 2017) prior to training.

⁵To train and fine-tune ELMo models, we use the code from https://github.com/lsgoslo/simple_elmo_training, which is essentially the reference ELMo implementation updated to the recent TensorFlow versions.

⁶We rely on *Hugging Face*’s implementation of BERT (available at <https://github.com/huggingface/transformers>, version 2.5.0), and follow their model naming conventions: <https://huggingface.co/models>.

predict the strength of semantic change of each target word between the two test corpora, we feed the word’s contextualised embeddings into the three algorithms of semantic change estimation described in Section 3. We then compute the Spearman correlation of the estimated change scores with the gold answers. This is the evaluation metric of Subtask 2, and we use it throughout our experiments.

5 Results

Our submission. In our ‘official’ shared task submission in the evaluation phase we used top-layer ELMo embeddings with the cosine similarity change detection algorithm for all languages. English and German ELMo models were trained on the respective Wikipedia corpora. For Swedish and Latin, pre-trained ELMo models were not available, so we trained our own models on the union of the test corpora. This combination of architectures and algorithms was chosen based on our preliminary experiments with the available human-annotated semantic change datasets for English (Gulordava and Baroni, 2011), German (Schlechtweg et al., 2018) and Russian (Fomin et al., 2019). The resulting Spearman correlations were 0.136 for English, 0.695 for German, 0.370 for Latin, and 0.278 for Swedish. With an average Codalab score of 0.37, this submission ranked 9th out of 34 teams in the evaluation phase.

We were aware that the submitted setup was likely sub-optimal as it did not include the **Fine-tuned** model variant. After the official submission deadline (in the post-evaluation phase), we finished training and fine-tuning all of our language models. Their systematic evaluation is the main contribution of this paper.

Current results. The average scores of all the tested configurations across 4 languages are given in Table 1. This table includes both the results of the configurations we used in the evaluation phase and the results of the configurations we tested after the submission deadline (fine-tuned models). We compare our scores to the organisers’ baselines (FD and CNT+CI+CD, as provided by Schlechtweg et al. (2020)) and the classical approach of calculating cosine distance between static CBOW word embeddings (Mikolov et al., 2013). The CBOW models were used in two different flavors: 1) ‘incremental’, where the C_2 model was initialised with the C_1 weights (Kim et al., 2014), and 2) ‘Procrustes’, where the two models were trained independently on C_1 and C_2 , and then aligned using the Orthogonal Procrustes transformation (Hamilton et al., 2016). Note that although we did not try different hyperparameter combinations for the static embeddings (varying vector dimensionalities, learning rates, vocabulary sizes, etc), we did not do that for contextualised embeddings either. All the training hyperparameters for both ELMo and BERT were fixed to their default values (see section 4), we varied only the training corpora and the layers from which embeddings were extracted. Thus, we believe that static and contextualised embedding-based approaches were under a fair comparison.

Table 1 shows that no method achieves statistically significant correlation on *all 4* languages, which attests both to the difficulty of the task and the diversity of the test sets. CBOW Procrustes is a surprisingly strong approach, consistently outperforming the organisers’ baselines. Only PRT and APD obtain higher average scores, with fine-tuned ELMo models performing better than fine-tuned BERT.

Judging only from the average correlation scores, contextualised embeddings do not seem to outshine their static counterparts, especially considering that both ELMo and BERT are more computationally demanding than CBOW. However, closer analysis of per-language results shows that in fact the contextualised approaches outperform the CBOW Procrustes baseline by a large margin for *each* of the shared task test sets. Table 2 features the scores obtained by our best-performing methods (PRT and APD with top layer embeddings from fine-tuned ELMo and BERT) on the individual languages of the shared task. We also report performance on the GEMS (‘GEometrical Models of Natural Language Semantics workshop’) test set (Gulordava and Baroni, 2011) to enable a comparison with previous work (Giulianelli et al., 2020). The discrepancy between the averaged and the per-language results can be explained by properties of the test sets: APD works best on the English and Swedish sets, while PRT yields the best scores for German and Latin.

Although consistency across languages (3 out of 4) is an important benefit of the CBOW Procrustes approach, with the right choice of APD or PRT, contextualised embeddings can improve Spearman correlation coefficients by up to 50%. This is not a language-specific property: the English GEMS test

Baselines	Frequency (FD) Count (CNT+CI+CD)	-0.083 0.144*		
CBOW cosine distance	Incremental Procrustes	0.140 0.392***		
Contextualised embeddings	Top layer	Average all layers	Average top 4 layers	
Cosine similarity (PRT)				
BERT	Pre-trained	0.278**	0.233	0.229
	Fine-tuned	0.373**	0.320**	0.338**
ELMo	Pre-trained	0.375**	0.344**	–
	Fine-tuned	0.402**	0.389**	–
	Trained on test	0.370**	0.342**	–
	Incremental	0.114*	0.127	–
Pairwise distance (APD)				
BERT	Pre-trained	0.237**	0.163*	0.203*
	Fine-tuned	0.363***	0.241**	0.297*
ELMo	Pre-trained	0.296**	0.172*	–
	Fine-tuned	0.405***	0.406***	–
	Trained on test	0.338**	0.295***	–
	Incremental	0.126**	-0.001*	–
Jensen-Shannon divergence (JSD)				
BERT	Pre-trained	0.181*	0.125	0.203*
	Fine-tuned	0.176*	0.223**	0.186**
ELMo	Pre-trained	0.251*	0.196*	–
	Fine-tuned	0.197*	0.156*	–
	Trained on test	0.225*	0.163*	–
	Incremental	-0.037	-0.009	–

Table 1: Spearman correlation coefficients for Subtask 2 averaged over four languages. The number of asterisks denotes the number of languages for which the correlation was statistically significant ($p < 0.05$).

set does not behave like the English test set from the shared task. In fact, one can observe 3 groups of test sets with regards to their statistical properties and to the method they favour: group 1 (Latin and German) exhibits rather uniform gold score distributions and prefers PRT; group 2 (English and Swedish) is characterised by more skewed gold score distributions and prefers APD; group 3 (GEMS) is in between, with no clear preference.

Interestingly, the method which produces a more uniform predicted score distribution (APD) works better for the test sets with skewed gold distributions, and the method which produces a more skewed predicted score distribution (PRT) works better for the uniformly distributed test sets (as can be seen in the Appendix). Furthermore, there is perfect negative correlation ($\rho = -1$) between the median gold score of a test set and the performance of the APD algorithm with fine-tuned ELMo models on this test set. The same correlation for the APD performance is not significant but strictly negative. We currently do not have a plausible explanation for this behaviour.

Table 2 also supports the previous observation that ELMo models perform better than BERT in the LSCD task. The only test set for which this is not the case is Latin,⁷ while on GEMS, ELMo and BERT are on par. One possible explanation is that our ELMo models were pre-trained on lemmatised Wikipedia corpora and thus better fit the test corpora, provided in lemmatised form by the organisers. The BERT

⁷The Latin test corpora are very peculiar: 1) homonyms in them are followed by ‘#’ and the sense identifier, which is not the case for Latin Wikipedia, 2) the sizes of C_1 and C_2 are very imbalanced, with the latter being 4 times larger than the former.

Algorithm		English	German	Latin	Swedish	GEMS
CBOW	Incremental	0.210	0.145	0.217	-0.012	0.424 [†]
	Procrustes	0.285	0.439 [†]	0.387 [†]	0.458 [†]	0.235 [†]
Fine-tuned contextualised embeddings (top layer)						
ELMo	Cosine similarity (PRT)	0.254	0.740 [†]	0.360 [†]	0.252	0.323 [†]
	Average pairwise distance (APD)	0.605 [†]	0.560 [†]	-0.113	0.569 [†]	0.323 [†]
BERT	Cosine similarity (PRT)	0.225	0.590 [†]	0.561 [†]	0.185	0.394 [†]
	Average pairwise distance (APD)	0.546 [†]	0.427 [†]	0.372 [†]	0.254	0.243 [†]

Table 2: Spearman correlation per test set for our best methods (post-evaluation phase). † marks statistical significance ($p < 0.05$).

models were pre-trained on raw corpora, and fine-tuning them on lemmatised data proves less successful. This is of course not an advantage of the ELMo architecture *per se*; however, easy and fast training from scratch on the respective Wikipedia corpora for each shared task language was possible only because of much lower computational requirements of ELMo compared to BERT.⁸

In the post-evaluation phase of the shared task, we submitted predictions obtained with the optimal system configurations: fine-tuned ELMo + APD for English and Swedish, fine-tuned ELMo + PRT for German, and fine-tuned BERT + PRT for Latin. This submission reached the average Spearman correlation of 0.618 and, at the time of writing, it is the best Subtask 2 submission for SemEval-2020 Task 1 (among those publicly available on the shared task website). Of course, the optimal choice of configurations was possible only because we already knew the test data. Still, it is useful for understanding of the real abilities of contextualised embedding-based approaches and the peculiarities of different models and test sets.

6 Conclusion

Our experiments for the SemEval-2020 Shared Task 1 (Subtask 2) show that using contextualised embeddings to rank words by the degree of their semantic change yields strong correlation with human judgements, outperforming approaches based on static embeddings. Models pre-trained on large external corpora and fine-tuned on the historical test corpora produce the highest correlation results, with ELMo slightly but consistently outperforming BERT as a contextualiser.

Inverted cosine similarity between averaged contextualised embeddings and the average pairwise cosine distance between contextualised embeddings turned out to be the best semantic change detection algorithms. An interesting finding is that the former algorithm favours the test sets with uniform gold score distribution, while the latter works best with the test sets where the gold score distribution is skewed towards low values. This distinction is not related to the language of the test set, or any other linguistic or statistical property of the test sets we looked at. While slightly different human annotation protocols can be in play here, we believe this dependency between the gold scores distribution and the performance of semantic change detection systems deserves to be investigated further in future work.

We also intend to explore the possibilities to improve our best-performing methods (PRT and APD), especially with regards to removing the outlier embeddings before calculating the semantic change score.

References

Robert Bamler and Stephan Mandt. 2017. Dynamic word embeddings. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 380–389. JMLR.org.

Marco Del Tredici, Raquel Fernández, and Gemma Boleda. 2019. Short-term meaning shift: A distributional

⁸Note that (Martinc et al., 2020) report a Spearman correlation of 0.510 on the GEMS dataset using fine-tuned BERT embeddings with Affinity Propagation and JSD. However, we were unable to reproduce these results, even when using the published code.

- exploration. In *Proceedings of NAACL-HLT 2019 (Annual Conference of the North American Chapter of the Association for Computational Linguistics)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. 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 4171–4186.
- 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 (Volume 1: Long Papers)*, Florence, Italy. Association for Computational Linguistics.
- Murhaf Fares, Andrey Kutuzov, Stephan Oepen, and Erik Velldal. 2017. Word vectors, reuse, and replicability: Towards a community repository of large-text resources. In *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pages 271–276. Association for Computational Linguistics.
- Vadim Fomin, Daria Bakshandaeva, Julia Rodina, and Andrey Kutuzov. 2019. Tracing cultural diachronic semantic shifts in Russian using word embeddings: test sets and baselines. *Komp'yuternaya Lingvistika i Intellektual'nye Tekhnologii: Dialog conference*, pages 203–218.
- Brendan Frey and Delbert Dueck. 2007. Clustering by passing messages between data points. *Science*, 315(5814):972–976.
- Mario Giulianelli, Marco Del Tredici, and Raquel Fernández. 2020. Analysing Lexical Semantic Change with Contextualised Word Representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics. Forthcoming.
- Kristina Gulordava and Marco Baroni. 2011. A distributional similarity approach to the detection of semantic change in the Google Books Ngram corpus. In *Proceedings of the GEMS 2011 Workshop on GEometrical Models of Natural Language Semantics*, pages 67–71, Edinburgh, UK, July. Association for Computational Linguistics.
- William 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.
- Renfen Hu, Shen Li, and Shichen Liang. 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*, pages 3899–3908, Florence, Italy. Association for Computational Linguistics.
- 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.
- 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. International World Wide Web Conferences Steering Committee.
- Andrey Kutuzov, Erik Velldal, and Lilja Øvrelid. 2017. Temporal dynamics of semantic relations in word embeddings: an application to predicting armed conflict participants. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1824–1829, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Andrey Kutuzov, Lilja Øvrelid, Terrence Szymanski, and Erik Velldal. 2018. Diachronic word embeddings and semantic shifts: a survey. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1384–1397. Association for Computational Linguistics.
- Jianhua Lin. 1991. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151.
- Matej Martinc, Syrielle Montariol, Elaine Zosa, and Lidia Pivovarov. 2020. Capturing evolution in word usage: Just add more clusters? In *Companion Proceedings of the International World Wide Web Conference*, pages 20–24.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems 26*, pages 3111–3119.

Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. 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 2227–2237.

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)*, volume 2, pages 169–174.

Dominik Schlechtweg, Barbara McGillivray, Simon Hengchen, Haim Dubossarsky, and Nina Tahmasebi. 2020. SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection. In *To appear in Proceedings of the 14th International Workshop on Semantic Evaluation*, Barcelona, Spain. Association for Computational Linguistics.

Milan Straka and Jana Straková. 2017. Tokenizing, pos tagging, lemmatizing and parsing UD 2.0 with UDPipe. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99. Association for Computational Linguistics.

Nina Tahmasebi, Lars Borin, and Adam Jatowt. 2018. Survey of computational approaches to lexical semantic change. In *Preprint at ArXiv 2018*.

Appendix A. Score distributions

In the bottom part of Figure 1, we show how different the 5 test sets are in terms of how the *gold* scores are distributed in them. In some test sets, the gold scores are skewed to the left, while some have a more uniform distribution. The top part of Figure 1 shows the distributions of the *predicted* scores produced by the APD and PRT algorithms (with fine-tuned ELMo embeddings). PRT tends to squeeze the majority of predictions near the lower boundary (no semantic change), with a low median score. In contrast, APD distributes its predictions in a much more uniform way, with a higher median score. Counter-intuitively, skewed gold distributions favour uniform predictions and vice versa.




Figure 1: Bottom: distribution of semantic change degree in the **gold** data; top left: distribution of scores predicted by the **APD** algorithm; top right: distribution of scores predicted by the **PRT** algorithm.

The grouping differences can be quantified with respect to the median gold score (after unit-normalisation). Figure 2 shows the dependency of the PRT and APD performance on the median score of the gold test set. The dots here are the performance values of PRT or APD algorithms on different test sets. English and Swedish test sets are in the left part of the plot with the median gold scores of 0.200 and 0.203 correspondingly. German, GEMS and Latin are on the right with 0.266, 0.267 and 0.364 correspondingly. There is a perfect negative Spearman correlation between the median gold scores of these 5 test sets and the performance of APD semantic change detection algorithm on each of them (with fine-tuned ELMo embeddings).




Figure 2: Performance of the PRT and APD algorithms depending on the median gold score.