

Word Usage Similarity Estimation with Sentence Representations and Automatic Substitutes

Aina Garí Soler¹, Marianna Apidianaki^{1,2} and Alexandre Allauzen¹

¹LIMSI, CNRS, Univ. Paris Sud, Université Paris-Saclay, F-91405 Orsay, France

²LLF, CNRS, Univ. Paris Diderot

{aina.gari, marianna, allauzen}@limsi.fr

Abstract

Usage similarity estimation addresses the semantic proximity of word instances in different contexts. We apply contextualized (ELMo and BERT) word and sentence embeddings to this task, and propose supervised models that leverage these representations for prediction. Our models are further assisted by lexical substitute annotations automatically assigned to word instances by context2vec, a neural model that relies on a bidirectional LSTM. We perform an extensive comparison of existing word and sentence representations on benchmark datasets addressing both graded and binary similarity. The best performing models outperform previous methods in both settings.

1 Introduction

Traditional word embeddings, like Word2Vec and GloVe, merge different meanings of a word in a single vector representation (Mikolov et al., 2013; Pennington et al., 2014). These pre-trained embeddings are fixed, and stay the same independently of the context of use. Current contextualized sense representations, like ELMo and BERT, go to the other extreme and model meaning as word usage (Peters et al., 2018; Devlin et al., 2018). They provide a dynamic representation of word meaning adapted to every new context of use.

In this work, we perform an extensive comparison of existing static and dynamic embedding-based meaning representation methods on the usage similarity (Usim) task, which involves estimating the semantic proximity of word instances in different contexts (Erk et al., 2009). Usim differs from a classical Semantic Textual Similarity task (Agirre et al., 2016) by the focus on a particular word in the sentence. We evaluate on this task word and context representations obtained using pre-trained uncontextualized word

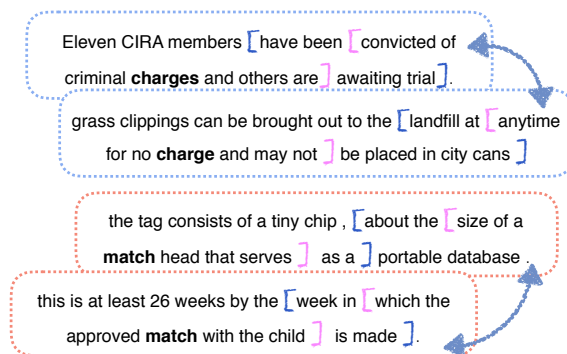


Figure 1: We use contextualized word representations built from the whole sentence or smaller windows around the target word for usage similarity estimation, combined with automatic substitute annotations.

embeddings (GloVe) (Pennington et al., 2014), with and without dimensionality reduction (SIF) (Arora et al., 2017); context representations obtained from a bidirectional LSTM (context2vec) (Melamud et al., 2016); contextualized word embeddings derived from a LSTM bidirectional language model (ELMo) (Peters et al., 2018) and generated by a Transformer (BERT) (Devlin et al., 2018); doc2vec (Le and Mikolov, 2014) and Universal Sentence Encoder representations (Cer et al., 2018). All these embedding-based methods provide direct assessments of usage similarity. The best representations are used as features in supervised models for Usim prediction, trained on similarity judgments.

We combine direct Usim assessments, made by the embedding-based methods, with a substitute-based Usim approach. Building up on previous work that used manually selected in-context substitutes as a proxy for Usim (Erk et al., 2013; McCarthy et al., 2016), we propose to automatize the annotation collection step in order to scale up the method and make it operational on unrestricted text. We exploit annotations assigned to words

in context by the context2vec lexical substitution model, which relies on word and context representations learned by a bidirectional LSTM from a large corpus (Melamud et al., 2016).

The main contributions of this paper can be summarized as follows:

- we provide a direct comparison of a wide range of word and sentence representation methods on the Usage Similarity (Usim) task and show that current contextualized representations can successfully predict Usim;
- we propose to automatize, and scale up, previous substitute-based Usim prediction methods;
- we propose supervised models for Usim prediction which integrate embedding and lexical substitution features;
- we propose a methodology for collecting new training data for supervised Usim prediction from datasets annotated for related tasks.

We test our models on benchmark datasets containing gold graded and binary word Usim judgments (Erk et al., 2013; Pilehvar and Camacho-Collados, 2019). From the compared embedding-based approaches, the BERT model gives best results on both types of data, providing a straightforward way for word usage similarity calculation. Our supervised model performs on par with BERT on the graded and binary Usim tasks, when using embedding-based representations and clean lexical substitutes.

2 Related Work

Usage similarity is a means for representing word meaning which involves assessing in-context semantic similarity, rather than mapping to word senses from external inventories (Erk et al., 2009, 2013). This methodology followed from the gradual shift from word sense disambiguation models that would select the best sense in context from a dictionary, to models that reason about meaning by solely relying on distributional similarity (Erk and Padó, 2008; Mitchell and Lapata, 2008), or allow multiple sense interpretations (Jurgens, 2014). In Erk et al. (2009), the idea is to model meaning in context in a way that captures different degrees of similarity to a word sense, or between word instances.

Due to its high reliance on context, Usim can be viewed as a semantic textual similarity (STS) (Agirre et al., 2016) task with a focus on a specific word instance. This connection motivated us to apply methods initially proposed for sentence similarity to Usim prediction. More precisely, we build sentence representations using different types of word and sentence embeddings, ranging from the classical word-averaging approach with traditional word embeddings (Pennington et al., 2014), to more recent contextualized word representations (Peters et al., 2018; Devlin et al., 2018). We explore the contribution of each separate method for Usim prediction, and use the best performing ones as features in supervised models. These are trained on sentence pairs labelled with Usim judgments (Erk et al., 2009) to predict the similarity of new word instances.

Previous attempts to automatic Usim prediction involved obtaining vectors encoding a distribution of topics for every target word in context (Lui et al., 2012). In this work, Usim was approximated by the cosine similarity of the resulting topic vectors. We show how contextualized representations, and the supervised model that uses them as features, outperform topic-based methods on the graded Usim task.

We combine the embedding-based direct Usim assessment methods with substitute-based representations obtained using an unsupervised lexical substitution model. McCarthy et al. (2016) showed it is possible to model usage similarity using manual substitute annotations for words in context. In this setting, the set of substitutes proposed for a word instance describe its specific meaning, while similarity of substitute annotations for different instances points to their semantic proximity.¹ We follow up on this work and propose a way to use substitutes for Usim prediction on unrestricted text, bypassing the need for manual annotations. Our method relies on substitute annotations proposed by the context2vec model (Melamud et al., 2016), which uses word and context representations learned by a bidirectional LSTM from a large corpus (UkWac) Baroni et al. (2009).

¹McCarthy et al. use the substitute annotations as features for predicting Usim, clustering instances and estimating the partitionability of words into senses. This offers a way to distinguish between lemmas with distinct senses and others with fuzzy semantics, which would be more challenging in annotation tasks and automatic processing.

Sentences	Substitutes
The local papers took photographs of the footprint.	GOLD: newspaper, journal AUTO-LSCNC: press, newspaper, news, report, picture AUTO-PPDB: newspaper, newsprint
Now Ari Fleischer, in a pitiful letter to the paper , tries to cast Milbank as the one getting his facts wrong.	GOLD: newspaper, publication AUTO-LSCNC: press, newspaper, news, article, journal, thesis, periodical, manuscript, document AUTO-PPDB: newspaper
This is also at the very essence or heart of being a coach .	GOLD: trainer, tutor, teacher AUTO-LSCNC: teacher, counsellor, trainer, tutor, instructor AUTO-PPDB: trainer, teacher, mentor, coaching
We hopped back onto the coach – now for the boulangerie!	GOLD: coach, bus, carriage AUTO-LSCNC: bus, car, carriage, transport AUTO-PPDB: bus, train, wagon, lorry, car, truck, carriage, vehicle

Table 1: Example pairs of highly similar and dissimilar usages from the Usim dataset (Erk et al., 2013) for the nouns *paper* (Usim score = 4.34) and *coach.n* (Usim score = 1.5), with the substitutes assigned by the annotators (GOLD). For comparison, we give the substitutes selected for these instances by the automatic substitution method (context2vec) used in our experiments from two different pools of substitutes (AUTO-LSCNC and PPDB). More details on the automatic substitution configurations are given in Section 4.2.

3 Data

3.1 The LexSub and Usim Datasets

We use the training and test datasets of the SemEval-2007 Lexical Substitution (LexSub) task (McCarthy and Navigli, 2007), which contain instances of target words in sentential context hand-labelled with meaning-preserving substitutes. A subset of the LexSub data (10 instances x 56 lemmas) has additionally been annotated with graded pairwise Usim judgments (Erk et al., 2013). Each sentence pair received a rating (on a scale of 1-5) by multiple annotators, and the average judgment for each pair was retained. McCarthy et al. (2016) derive two additional scores from Usim annotations that denote how easy it is to partition a lemma’s usages into sets describing distinct senses: **Uiaa**, the inter-annotator agreement for a given lemma, taken as the average pairwise Spearman’s ρ correlation between ranked judgments of the annotators; and **Umid**, the proportion of mid-range judgments over all instances for a lemma and all annotators.

In our experiments, we use 2,466 sentence pairs from the Usim data for training, development and testing of different automatic Usim prediction methods. Our models rely on substitutes automatically assigned to words in context using context2vec (Melamud et al., 2016), and on various word and sentence embedding representa-

tions. We also train a model using the gold substitutes, to test how well our models perform when substitute quality is high. Performance of the different models is evaluated by measuring how well they approximate the Usim scores assigned by annotators. Table 1 shows examples of sentence pairs from the Usim dataset (Erk et al., 2013) with the GOLD substitutes and Usim scores assigned by the annotators. The Usim score is high for similar instances, and decreases for instances that describe different meanings. The semantic proximity of two instances is also reflected in the similarity of their substitutes sets. For comparison, we also give in the Table the substitutes selected for these instances by the automatic context2vec substitution method used in our experiments (more details in Section 4.2).

3.2 The Concepts in Context Corpus

Given the small size of the Usim dataset, we extract additional training data for our models from the Concepts in Context (CoInCo) corpus (Kremer et al., 2014), a subset of the MASC corpus (Ide et al., 2008). CoInCo contains manually selected substitutes for all content words in a sentence, but provides no usage similarity scores that could be used for training. We construct our supplementary training data as follows: we gather all instances of a target word in the corpus with at least four

substitutes, and keep pairs with (1) no overlap in substitutes, and (2) minimum 75% substitute overlap.² We view the first set of pairs as examples of completely different usages of a word (DIFF), and the second set as examples of identical usages (SAME). The two sets are unbalanced in terms of number of instance pairs (19,060 vs. 2,556). We balance them by keeping in DIFF the 2,556 pairs with the highest number of substitutes.

We also annotate the data with substitutes using context2vec (Melamud et al., 2016), as described in Section 4.2. We apply an additional filtering to the sentence pairs extracted from CoInCo, discarding instances of words that are not in the context2vec vocabulary and have no embeddings. We are left with 2,513 pairs in each class (5,026 in total). We use 80% of these pairs (4,020) together with the Usim data to train our supervised Usim models described in Section 4.3.³

3.3 The Word-in-Context dataset

The third dataset we use in our experiments is the recently released Word-in-Context (WiC) dataset (Pilehvar and Camacho-Collados, 2019), version 0.1. WiC provides pairs of contextualized target word instances describing the same or different meaning, framing in-context sense identification as a binary classification task. For example, a sentence pair for the noun *stream* is: [Stream of consciousness’ – ‘Two streams of development run through American history’]. A system is expected to be able to identify that *stream* does not have the same meaning in the two sentences.

WiC sentences were extracted from example usages in WordNet (Fellbaum, 1998), VerbNet (Schuler, 2006), and Wiktionary. Instance pairs were automatically labeled as positive (T) or negative (F) (corresponding to the same/different sense) using information in the lexicographic resources, such as presence in the same or different synsets. Each word is represented by at most three instances in WiC, and repeated sentences are excluded. It is important to note that meanings represented in the WiC dataset are coarser-grained than WordNet senses. This was ensured by excluding WordNet synsets describing highly sim-

²Full overlap is rare since annotators propose somewhat different sets of substitutes, even for instances with the same meaning. Full overlap is observed for only 437 of all considered CoInCo pairs (0.3%).

³We will make the dataset available at <https://github.com/ainagari>. 20% of the extracted examples were kept aside for development and testing purposes.

ilar meanings (sister senses, and senses belonging to the same supersense). The human-level performance upper-bound on this binary task, as measured on two 100-sentence samples, is 80.5%. Inter-annotator agreement is also high, at 79%. The dataset comes with an official train/dev/test split containing 7,618, 702 and 1,366 sentence pairs, respectively.⁴

4 Methodology

We experiment with two ways of predicting usage similarity: an unsupervised approach which relies on the cosine similarity of different kinds of word and sentence representations, and provides direct Usim assessments; and supervised models that combine embedding similarity with features based on substitute overlap. We present the direct Usim prediction methods in Section 4.1. In Section 4.2, we describe how substitute-based features were extracted, and in Section 4.3, we introduce the supervised Usim models.

4.1 Direct Usage Similarity Prediction

In the unsupervised Usim prediction setting, we apply different types of pre-trained word and sentence embeddings as follows: we compute an embedding for every sentence in the Usim dataset, and calculate the pairwise cosine similarity between the sentences available for a target word. Then, for every embedding type, we measure the correlation between sentence similarities and gold usage similarity judgments in the Usim dataset, using Spearman’s ρ correlation coefficient. We experiment with the following embedding types.

GloVe embeddings are uncontextualized word representations which merge all senses of a word in one vector (Pennington et al., 2014). We use 300-dimensional GloVe embeddings pre-trained on Common Crawl (840B tokens).⁵ The representation of a sentence is obtained by averaging the GloVe embeddings of the words in the sentence.

SIF (Smooth Inverse Frequency) embeddings are sentence representations built by applying dimensionality reduction to a weighted average of uncontextualized embeddings of words in a sentence

⁴The test portion of WiC had not been released at the time of submission. We contacted the authors and ran the evaluation on the official test set, to be able to compare to results reported in their paper (Pilehvar and Camacho-Collados, 2019).

⁵<https://nlp.stanford.edu/projects/glove/>

(Arora et al., 2017). We use SIF in combination with GloVe vectors.

Context2vec embeddings (Melamud et al., 2016). The context2vec model learns embeddings for words and their sentential contexts simultaneously. The resulting representations reflect: a) the similarity between potential fillers of a sentence with a blank slot, and b) the similarity of contexts that can be filled with the same word. We use a context2vec model pre-trained on the UkWac corpus (Baroni et al., 2009)⁶ to compute embeddings for sentences with a blank at the target word’s position.

ELMo (Embeddings from Language Models) representations are contextualized word embeddings derived from the internal states of an LSTM bidirectional language model (biLM) (Peters et al., 2018). In our experiments, we use a pre-trained 512-dimensional biLM.⁷ Typically, the best linear combination of the layer representations for a word is learned for each end task in a supervised manner. Here, we use out-of-the-box embeddings (without tuning) and experiment with the top layer, and with the average of the three hidden layers. We represent a sentence in two ways: by the contextualized ELMo embedding obtained for the target word, and by the average of ELMo embeddings for all words in a sentence.

BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al., 2018). BERT representations are generated by a 12-layer bidirectional Transformer encoder that jointly conditions on both left and right context in all layers.⁸ BERT can be fine-tuned to specific end tasks, or its contextualized word representations can be used directly in applications, similar to ELMo. We try different layer combinations and create sentence representations, in the same way as for ELMo: using either the BERT embedding of the target word, or the average of the BERT embeddings for all words in a sentence.

Universal Sentence Encoder (USE) makes use of a Deep Averaging Network (DAN) encoder trained to create sentence representations by means of multi-task learning (Cer et al., 2018).

⁶<http://u.cs.biu.ac.il/~nlp/resources/downloads/context2vec/>

⁷<https://allennlp.org/elmo>

⁸This is an important difference with the ELMo architecture which concatenates a left-to-right and right-to-left model.

USE has been shown to improve performance on different NLP tasks using transfer learning.⁹

doc2vec is an extension of word2vec to the sentence, paragraph or document level (Le and Mikolov, 2014). One of its forms, *dbow* (distributed bag of words), is based on the skip-gram model, where it adds a new feature vector representing a document. We use a dbow model trained on English Wikipedia released by Lau and Baldwin (2016).¹⁰

We test the above models with representations built from the whole sentence, and using a smaller context window (cw) around the target word. Sentences in the WiC dataset are quite short (7.9 ± 3.9 words), but the length of sentences in the Usim and CoInCo datasets varies a lot (27.4 ± 13.2 and 18.8 ± 10.2 , respectively). We want to check whether information surrounding the target word in the sentence is more relevant, and sufficient for Usim estimation. We focus on the words in a context window of $\pm 2, 3, 4$ or 5 words at each side of a target word. Then, we collect their word embeddings to be averaged (for GloVe, ELMo and BERT), or derive an embedding from this specific window instead of the whole sentence (for USE).

We approximate Usim by measuring the cosine similarity of the resulting context representations. We compare the performance of these direct assessment methods on the Usim dataset and report the results in Section 5.

4.2 Substitute-based Feature Extraction

Following up on McCarthy et al.’s (2016) sense clusterability work, we also experiment with a substitute-based approach for Usim prediction. McCarthy et al. showed that manually selected substitutes for word instances in context can be used as a proxy for Usim. Here, we propose an approach to obtain these annotations automatically that can be applied to the whole vocabulary.

Automatic LexSub We generate rankings of candidate substitutes for words in context using the context2vec method (Melamud et al., 2016). The original method selects and ranks substitutes from the whole vocabulary. To facilitate comparison and evaluation, we use the following pools of candidates: (a) all substitutes that were proposed

⁹<https://tfhub.dev/google/universal-sentence-encoder/2>

¹⁰<https://github.com/jhlau/doc2vec>

for a word in the LexSub and CoInCo annotations (we call this substitute pool AUTO-LSCNC); (b) the paraphrases of the word in the Paraphrase Database (PPDB) XXL package (Ganitkevitch et al., 2013; Pavlick et al., 2015) (AUTO-PPDB).¹¹ In the WiC experiments, where no substitute annotations are available, we only use PPDB paraphrases (AUTO-PPDB). We obtain a context2vec embedding for a sentence by replacing the target word with a blank. AUTO-LSCNC substitutes are high-quality since they were extracted from the manual LexSub and CoInCo annotations. They are semantically similar to the target, and context2vec just needs to rank them according to how well they fit the new context. This is done by measuring the cosine similarity between each substitute’s context2vec word embedding and the context embedding obtained for the sentence.

The AUTO-PPDB pool contains paraphrases from PPDB XXL, which were automatically extracted from parallel corpora (Ganitkevitch et al., 2013). Hence, this pool contains noisy paraphrases that should be ranked lower. To this end, we use in this setting the original context2vec scoring formula which also accounts for the similarity between the target word and the substitute:

$$c2v_score = \frac{\cos(s, t) + 1}{2} \times \frac{\cos(s, C) + 1}{2} \quad (1)$$

In formula (1), s and t are the word embeddings of a substitute and the target word, and C is the context2vec vector of the context. Following this procedure, context2vec produces a ranking of candidate substitutes for each target word instance in the Usim, CoInCo and WiC datasets, according to their fit in context. Every candidate is assigned a score, with substitutes that are a good fit in a specific context being higher-ranked than others. For every new target word instance, context2vec ranks *all* candidate substitutes available for the target in each pool. Consequently, the automatic annotations produced for different instances of the target include the same set of substitutes, but in different order. This does not allow for the use of measures based on substitute overlap, which were shown to be useful for Usim prediction in McCarthy et al. (2016). In order to use this type of measures, we propose ways to filter the automatically generated rankings, and keep for each instance only substitutes that are a good fit in context.

¹¹<http://paraphrase.org/>

Substitute Filtering We test different filters to discard low quality substitutes from the annotations proposed by context2vec for each instance.

- **PPDB 2.0 score:** Given a ranking R of n substitutes $R = [s_1, s_2, \dots, s_n]$ proposed by context2vec, we form pairs of substitutes in adjacent positions $\{s_i \leftrightarrow s_{i+1}\}$, and check whether they exist as paraphrase pairs in PPDB. We expect substitutes that are paraphrases of each other to be similarly ranked. If s_i and s_{i+1} are not paraphrases in PPDB, we keep all substitutes up to s_i and use this as a cut-off point, discarding substitutes present from position s_{i+1} onwards in the ranking.
- **GloVe word embeddings:** We measure the cosine similarity (cosSim) between GloVe embeddings of adjacent substitutes $\{s_i \leftrightarrow s_{i+1}\}$ in the ranking R obtained for a new instance. We first compare the similarity of the first pair of substitutes ($\text{cosSim}(s_1, s_2)$) to a lower bound similarity threshold T . If $\text{cosSim}(s_1, s_2)$ exceeds T , we assume that s_1 and s_2 have the same meaning, and use $\text{cosSim}(s_1, s_2)$ as a reference similarity value, S , for this instance. The middle point between the two values, $M = (T + S)/2$, is then used as a threshold to determine whether there is a shift in meaning in subsequent pairs. If $\text{cosSim}(s_i, s_{i+1}) < M$, for $i > 1$, then only the higher ranked substitute (s_i) is retained and all subsequent substitutes in the ranking are discarded. The intuition behind this calculation is that if cosSim is much lower than the reference S (even if it exceeds T), substitutes possibly have different senses.
- **Context2vec score:** This filter uses the score assigned by context2vec to each substitute, reflecting how good a fit it is in each context. context2vec scores vary a lot across instances, it is thus not straightforward to choose a threshold. We instead refer to the scores assigned to adjacent pairs of substitutes in the ranking produced for each instance, $R = [s_1, s_2, \dots, s_n]$. We view the pair with the biggest difference in scores as the cut-off point, considering it reflects a degradation in substitute fit. We retain only substitutes up to this point.
- **Highest-ranked X substitutes.** We also test two simple baselines, which consist in keep-

ing the 5 and 10 highest-ranked substitutes for each instance.

We test the efficiency of each filter on the portion of the LexSub dataset (McCarthy and Navigli, 2007) that was not annotated for Usim. We compare the substitutes retained for each instance after filtering to its gold LexSub substitutes using the F1-score, and the proportion of false positives out of all positives. Filtering results are reported in Appendix A. The best filters were GloVe word embeddings ($T = 0.2$) for AUTO-LSCNC, and the PPDB filter for AUTO-PPDB.

Feature Extraction After annotating the Usim sentences with context2vec and filtering, we extract, for each sentence pair (S_1, S_2), a set of features related to the amount of substitute overlap.

- **Common substitutes.** The proportion of shared substitutes between two sentences.
- **GAP score.** The average of the Generalized Average Precision (GAP) score (Kishida, 2005) taken in both directions ($GAP(S_1, S_2)$ and $GAP(S_2, S_1)$). GAP is a measure that compares two rankings considering not only the order of the ranked elements but also their weights. It ranges from 0 to 1, where 0 means that rankings are completely different and 1 indicates perfect agreement. We use the frequency in the manual Usim annotations (i.e. the number of annotators who proposed each substitute) as the weight for gold substitutes, and the context2vec score for automatic substitutes. We use the GAP implementation from Melamud et al. (2015).
- **Substitute cosine similarity.** We form substitute pairs ($S_1 \leftrightarrow S_2$) and calculate the average of their GloVe cosine similarities. This feature shows the semantic similarity of substitutes, even when overlap is low.

4.3 Supervised Usim Prediction

We train linear regression models to predict Usim scores for word instances in different contexts using as features the cosine similarity of the different representations in Section 4.1, and the substitute-based features in 4.2. For training, we use the Usim dataset on its own (cf. Section 3.1), and combined with the additional training examples extracted from CoInCo (cf. Section 3.2).

To be able to evaluate the performance of our models separately for each of the 56 target words

in the Usim dataset, we train a separate model for each word in a leave-one-out setting. Each time, we use 2,196 pairs for training, 225 for development and 45 for testing.¹² Each model is evaluated on the sentences corresponding to the left out target word. We report results of these experiments in Section 5. The performance of the model with context2vec substitutes from the two substitute pools is compared to that of the model with gold substitute annotations. We replicate the experiments by adding CoInCo data to the Usim training data.

To test the contribution of each feature, we perform an ablation study on the 225 Usim sentence pairs of the development set, which cover the full spectrum of Usim scores (from 1 to 5). We report results of the feature ablation in Appendix C.

We also build a model for the binary Usim task on the WiC dataset (Pilehvar and Camacho-Collados, 2019), using the official train/dev/test split. We train a logistic regression classifier on the training set, and use the development set to select the best among several feature combinations. We report results of the best performing models on the WiC test set in Section 5. For instances in WiC where no PPDB substitutes are available (133 out of 1,366 in the test set) we back off to a model that only relies on the embedding features.

5 Evaluation

Direct Usim Prediction Correlation results between Usim judgments and the cosine similarity of the embedding representations described in Section 4.1 are found in Table 2. Detailed results for all context window combinations are given in Appendix B. We observe that target word BERT embeddings give best performance in this task. Selecting a context window around (or including) the target word does not always help, on the contrary it can harm the models. Context2vec sentence representations are the next best performing representation, after BERT, but their correlation is much lower. The simple GloVe-based SIF approach for sentence representation, which consists in applying dimensionality reduction to a weighted average of GloVe vectors of the words in a sentence, is much superior to the simple average of GloVe vectors and even better than doc2vec sentence representations, obtaining a correlation comparable to

¹²With the exception of 4 lemmas which had 36 pairs, and one which had 44.

Context	Embeddings	Correlation
Full sentence	GloVe	0.142
	SIF	0.274
	c2v	0.290
	USE	0.272
	doc2vec	0.124
	ELMo av	0.254
	BERT av 4	0.289
Target word	ELMo av	0.166
	ELMo top	0.177
	BERT top	0.514
	BERT av 4	0.518
cw=2	ELMo top	0.289
cw=3 (incl. target)	GloVe	0.180
	ELMo av	0.280
	BERT av 4	0.395
cw=5 (incl. target)	USE	0.221
	ELMo av	0.266
	ELMo top	0.263
	BERT top	0.309

Table 2: Spearman ρ correlation of different sentence and word embeddings on the Usim dataset using different context window sizes (cw). For BERT and ELMo, *top* refers to the top layer, and *av* refers to the average of layers (3 for ELMo, and the last 4 for BERT).

that of USE.

Graded Usim To evaluate the performance of our supervised models, we measure the correlation of the predictions with human similarity judgments on the Usim dataset using Spearman’s ρ . Results reported in Table 3 are the average of the correlations obtained for each target word with gold and automatic substitutes (from the two substitute pools), and for each type of features, substitute-based and embedding-based (cosine similarities from BERT and context2vec). We also report results with the additional CoInCo training data. Unsurprisingly, the best results are obtained by the methods that use the gold substitutes. This is consistent with previous analyses by Erk et al. (2009) who found overlap in manually-proposed substitutes to correlate with Usim judgments. The lower performance of features that rely on automatically selected substitutes (AUTO-LSCNC and AUTO-PPDB) demonstrates the impact of substitute quality on the contribution of this type of features. The addition of CoInCo data does not seem to help the models, as results are slightly lower than in the only Usim setting. This can be due to the fact that CoInCo data contains only extreme cases of similarity (SAME/DIFF) and no in-

termediate ratings. The slight improvement in the combined settings over embedding-based models is not significant in AUTO-LSCNC substitutes, but it is for gold substitutes ($p < 0.001$).¹³

For comparison to the topic-modelling approach of Lui et al. (2012), we evaluate on the 34 lemmas used in their experiments. They report a correlation calculated over all instances. With the exception of the substitute-only setting with PPDB candidates, all of our Usim models get higher correlation than their model ($\rho = 0.202$), with $\rho = 0.512$ for the combination of AUTO-LSCNC substitutes and embeddings. The average of the per target word correlation in Lui et al. (2012) ($\rho = 0.388$) is still lower than that of our AUTO-LSCNC model in the combined setting ($\rho = 0.500$).

Binary Usim We evaluate the predictions of our binary classifiers by measuring accuracy on the test portion of the WiC dataset. Results for the best configurations for each training set are reported in Table 4. Experiments on the development set showed that target word BERT representations and USE sentence embeddings are the best-suited for WiC. Therefore, ‘embedding-based features’ here refers to these two representations. Results on the development set can be found in Appendix D. All configurations obtain higher accuracy than the previous best reported result on this dataset (59.4) (Pilehvar and Camacho-Collados, 2019), obtained using DeConf vectors, which are multi-prototype embeddings based on WordNet knowledge (Pilehvar and Collier, 2016). Similar to the graded Usim experiments, adding substitute-based features to embedding features slightly improves the accuracy of the model. Also, combining the CoInCo and WiC data for training does not have a clear impact on results, even in this binary classification setting.

6 Discussion

Results reported for Usim are the average correlation for each target word, but the strength of the correlation varies greatly for different words for all models and settings. For example, in the case of direct Usim prediction with embeddings using BERT target, Spearman’s ρ ranges from 0.805 (for the verb *fire*) to -0.111 (for the verb *suffer*). This variation in performance is not surprising,

¹³As determined by paired t-tests, after verifying the normality of the differences with the Shapiro-Wilk test

Training set	Features	Gold	c2v	c2v
			AUTO-LSCNC	AUTO-PPDB
Usim	Substitute-based	0.563	0.273	0.148
	Embedding-based	0.494	0.494	0.494
	Combined	0.626	0.501	0.493
Usim + CoInCo	Substitute-based	-	0.262	0.129
	Embedding-based	-	0.495	0.495
	Combined	-	0.501	0.491

Table 3: **Graded Usim results:** Spearman’s ρ correlation results between supervised model predictions and graded annotations on the Usim test set. The first column reports results obtained using gold substitute annotations for each target word instance. The last two columns give results with automatic substitutes selected among all substitutes proposed for the word in the LexSub and CoInCo datasets (AUTO-LSCNC), or paraphrases in the PPDB XXL package (AUTO-PPDB). The Embedding-based configuration uses cosine similarities from BERT and context2vec.

Training set	Features	Accuracy
WiC	Embedding-based	63.62
	Combined	64.86
	DeConf embeddings (Pilehvar and Camacho-Collados, 2019)	59.4
	Random baseline (Pilehvar and Camacho-Collados, 2019)	50.0
WiC + CoInCo	Embedding-based	63.69
	Combined	64.42

Table 4: **Binary Usim results:** Accuracy of models on the WiC test set. The Embedding-based configuration includes cosine similarities of BERT target and USE. The Combined setting uses, in addition, substitute overlap features (AUTO-PPDB).

since annotators themselves found some lemmas harder to annotate than others, as reflected in the Usim inter-annotator agreement measure (Uiaa) (McCarthy et al., 2016). We find that BERT target word embeddings results correlate with Uiaa per target word ($\rho = 0.59, p < 0.05$), showing that the performance of this model depends to a certain extent on the ease of annotation for each lemma. Uiaa also correlates with the standard deviation of average Usim scores by target word ($\rho = 0.66, p < 0.001$). Indeed, average Usim values for the word *suffer* do not exhibit high variance as they only range from 3.6 to 4.9. Within a smaller range of scores, a strong correlation is harder to obtain. The negative correlation between Uiaa and Umid ($-0.46, p < 0.001$) also suggests that words with higher disagreement tend to exhibit a higher proportion of mid-range judgments. We believe that this analysis highlights the difference between usage similarity across target words and encourages a by-lemma approach where the specificities of each lemma are taken into account.

7 Conclusion

We applied a wide range of existing word and context representations to graded and binary usage similarity prediction. We also proposed

novel supervised models which use as features the best performing embedding representations, and make high quality predictions especially in the binary setting, outperforming previous approaches. The supervised models include features based on in-context lexical substitutes. We show that automatic substitutions constitute an alternative to manual annotation when combined with the embedding-based features. Nevertheless, if there is no specific reason for using substitutes for measuring Usim, BERT offers a much more straightforward solution to the Usim prediction problem.

In future work, we plan to use automatic Usim predictions for estimating word sense partitionability. We believe such knowledge can be useful to determine the appropriate meaning representation for each lemma.

8 Acknowledgments

We would like to thank the anonymous reviewers for their helpful feedback on this work. We would also like to thank Jose Camacho-Collados for his help with the WiC experiments.

The work has been supported by the French National Research Agency under project ANR-16-CE33-0013.

References

- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2016. **Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation**. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 497–511, San Diego, California. Association for Computational Linguistics.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A Simple but Tough-to-Beat Baseline for Sentence Embeddings. In *International Conference on Learning Representations (ICLR)*, Toulon, France.
- Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. 2009. The WaCky wide web: a collection of very large linguistically processed web-crawled corpora. *Journal of Language Resources and Evaluation*, 43(3):209–226.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. **Universal Sentence Encoder for English**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805*.
- Katrin Erk, Diana McCarthy, and Nicholas Gaylord. 2009. **Investigations on word senses and word usages**. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 10–18, Suntec, Singapore. Association for Computational Linguistics.
- Katrin Erk, Diana McCarthy, and Nicholas Gaylord. 2013. Measuring word meaning in context. *Computational Linguistics*, 39:511–554.
- Katrin Erk and Sebastian Padó. 2008. **A structured vector space model for word meaning in context**. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 897–906, Honolulu, Hawaii. Association for Computational Linguistics.
- Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. Language, Speech, and Communication. MIT Press, Cambridge, MA.
- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. **PPDB: The Paraphrase Database**. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 758–764, Atlanta, Georgia. Association for Computational Linguistics.
- Nancy Ide, Collin Baker, Christiane Fellbaum, Charles Fillmore, and Rebecca Passonneau. 2008. **MASC: the Manually Annotated Sub-Corpus of American English**. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- David Jurgens. 2014. **An analysis of ambiguity in word sense annotations**. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*, pages 3006–3012, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Kazuaki Kishida. 2005. *Property of average precision and its generalization: An examination of evaluation indicator for information retrieval experiments*. Technical Report NII-2005-014E, National Institute of Informatics Tokyo, Japan.
- Gerhard Kremer, Katrin Erk, Sebastian Padó, and Stefan Thater. 2014. **What substitutes tell us - analysis of an “all-words” lexical substitution corpus**. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 540–549, Gothenburg, Sweden. Association for Computational Linguistics.
- Jey Han Lau and Timothy Baldwin. 2016. **An empirical evaluation of doc2vec with practical insights into document embedding generation**. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 78–86, Berlin, Germany. Association for Computational Linguistics.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of the 31st International conference on Machine Learning*, pages 1188–1196, Beijing, China.
- Marco Lui, Timothy Baldwin, and Diana McCarthy. 2012. **Unsupervised estimation of word usage similarity**. In *Proceedings of the Australasian Language Technology Association Workshop 2012*, pages 33–41, Dunedin, New Zealand.
- Diana McCarthy, Marianna Apidianaki, and Katrin Erk. 2016. **Word sense clustering and clusterability**. *Computational Linguistics*, 42(2):245–275.
- Diana McCarthy and Roberto Navigli. 2007. **Semeval-2007 task 10: English lexical substitution task**. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 48–53, Prague, Czech Republic. Association for Computational Linguistics.
- Oren Melamud, Jacob Goldberger, and Ido Dagan. 2016. **context2vec: Learning Generic Context Embedding with Bidirectional LSTM**. In *Proceedings of the 20th SIGNLL Conference on Computational Linguistics*, pages 1–11, Berlin, Germany. Association for Computational Linguistics.

Natural Language Learning, pages 51–61, Berlin, Germany. Association for Computational Linguistics.

Oren Melamud, Omer Levy, and Ido Dagan. 2015. A simple word embedding model for lexical substitution. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 1–7, Denver, Colorado.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proceedings of the International Conference on Learning Representations*, Scottsdale, Arizona.

Jeff Mitchell and Mirella Lapata. 2008. **Vector-based models of semantic composition**. In *Proceedings of ACL-08: HLT*, pages 236–244, Columbus, Ohio. Association for Computational Linguistics.

Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. **PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification**. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 425–430, Beijing, China. Association for Computational Linguistics.

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, Doha, Qatar. Association for Computational Linguistics.

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, New Orleans, Louisiana. Association for Computational Linguistics.

Mohammad Taher Pilehvar and José Camacho-Collados. 2019. **WiC: 10, 000 Example Pairs for Evaluating Context-Sensitive Representations**. *Accepted at the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

Mohammad Taher Pilehvar and Nigel Collier. 2016. **De-conflated semantic representations**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1680–1690, Austin, Texas. Association for Computational Linguistics.

Karin Kipper Schuler. 2006. *VerbNet: A Broad-Coverage, Comprehensive Verb Lexicon*. Ph.D. thesis, University of Pennsylvania.

A Filtering experiments

Tables 5 and 6 contain results obtained using the different substitute filters described in Section 4.2. We measure the quality of the substitutes retained in the automatic ranking produced by context2vec after filtering against gold substitute annotations in LexSub data. Here, we only use the portion of LexSub data that does not contain Usim judgments.

We measure filtered substitute quality against the gold standard using the F1-score, and the proportion of false positives (FP) over all positives (TP+FP). Table 5 shows results for annotations assigned by context2vec using the the LexSub/CoInCo pool of substitutes (AUTO-LSCNC). Table 6 shows results for context2vec annotations with the PPDB pool of substitutes (AUTO-PPDB).

Filter	F1	$FP/(TP + FP)$
Highest 10	0.332	0.776
Highest 5	0.375	0.695
PPDB	0.333	0.643
GloVe ($T = 0.1$)	0.371	0.675
GloVe ($T = 0.2$)	0.373	0.661
GloVe ($T = 0.3$)	0.353	0.641
c2v score	0.326	0.671
No filter	0.248	0.848

Table 5: Results of different substitute filtering strategies applied to annotations assigned by context2vec when using the LexSub/CoInCo pool of substitutes (AUTO-LSCNC).

Filter	F1	$FP/(TP + FP)$
Highest 10	0.245	0.838
Highest 5	0.290	0.766
PPDB	0.268	0.731
GloVe ($T = 0.1$)	0.266	0.778
GloVe ($T = 0.2$)	0.268	0.769
GloVe ($T = 0.3$)	0.266	0.750
c2v score	0.250	0.675
No filter	0.142	0.920

Table 6: Results of different substitute filtering strategies applied to annotations assigned by context2vec when using the PPDB pool of substitutes (AUTO-PPDB).

B Direct Usage Similarity Estimation

Correlations between gold Usim scores for all words and cosine similarities of different embed-

ding types can be found in Tables 7 and 8.

	Embeddings	Correlation
Full sentence embedding	GloVe	0.142
	SIF	0.274
	c2v	0.290
	USE	0.272
	doc2vec	0.124
	ELMo av	0.254
	ELMo top	0.248
	BERT av 4	0.289
Target word embedding	ELMo av	0.166
	ELMo top	0.177
	BERT top	0.514
	BERT av 4	0.518
	BERT concat 4	0.516
	BERT 2nd-to-last	0.486

Table 7: Correlations of sentence and word embeddings on the Usim dataset using different context window sizes (*cw*). For BERT and ELMo, *top* refers to the top layer, and *av* refers to the average of layers (3 for ELMo, and the last 4 for BERT). *concat 4* refers to the concatenation of the last 4 layers of BERT.

C Feature Ablation on Usim

Results of feature ablation experiments on the Usim development sets are given in Table 9.

D Dev experiments on WiC

Table 10 shows the accuracy of different configurations on the WiC development set.

Context	Embeddings	Correlation
<i>cw</i> =2	ELMo top	0.289
	ELMo av	0.280
	BERT av 4	0.344
	GloVe	0.140
<i>cw</i> =3	ELMo top	0.282
	ELMo av	0.279
	BERT av 4	0.339
	GloVe	0.163
<i>cw</i> =4	ELMo top	0.270
	ELMo av	0.263
	BERT av 4	0.311
	GloVe	0.160
<i>cw</i> =5	ELMo top	0.266
	ELMo av	0.263
	BERT av 4	0.309
	GloVe	0.162
<i>cw</i> =2 (incl. target)	ELMo av	0.284
	ELMo top	0.278
	BERT av 4	0.416
	GloVe	0.159
	USE	0.146
<i>cw</i> =3 (incl. target)	ELMo av	0.280
	ELMo top	0.273
	BERT av 4	0.395
	GloVe	0.180
	USE	0.184
<i>cw</i> =4 (incl. target)	ELMo av	0.267
	ELMo top	0.265
	BERT av 4	0.365
	GloVe	0.176
	USE	0.191
<i>cw</i> =5 (incl. target)	ELMo av	0.266
	ELMo top	0.263
	BERT av 4	0.359
	GloVe	0.175
	USE	0.221

Table 8: Correlations of different sentence and word embeddings on the Usim dataset using different context window sizes (*cw*).

Ablation	Gold	AUTO-LSCNC	AUTO-PPDB
None	0.729	0.538	0.524
Sub. similarity	0.701	0.537	0.524
Common sub.	0.722	0.538	0.524
GAP	0.730	0.537	0.523
c2v	0.730	0.539	0.523
Bert av 4 target	0.700	0.348	0.283

Table 9: Results of feature ablation experiments for systems trained and tested on the Usim dataset with gold substitutes (*Gold*) as well as automatic substitutes from different pools, Lexsub/CoInCo (AUTO-LSCNC) and PPDB (AUTO-PPDB). Rows indicate the feature that is removed each time. Numbers correspond to the average Spearman ρ correlation on the development set across target words.

Training set	Features	Accuracy
WiC	BERT av 4 last target word	65.24
	c2v	57.69
	ELMo top cw=2	61.11
	USE	63.68
	SIF	60.97
	Only substitutes	55.41
	BERT av 4 target word & USE	67.95
	Combined	66.81
WiC + CoInCo	BERT av 4 target word	64.96
	c2v	58.12
	ELMo top cw=2	61.11
	USE	63.53
	SIF	59.97
	Only substitutes	56.13
	BERT av 4 target word & USE	68.66
	Combined	66.81

Table 10: Accuracy of different features and combinations on the WiC development set. On this dataset, the two best types of embeddings, that were chosen for the Embedding-based and Combined configurations, were BERT (target word, average of the last 4 layers) and USE. Both Only-substitutes and Combined use features of automatic substitutes from the PPDB pool, and back off to the Embedding-based model when there were no paraphrases available for the target word in the PPDB.