

Detecting Different Forms of Semantic Shift in Word Embeddings via Paradigmatic and Syntagmatic Association Changes

Anna Wegmann¹[0000-0002-6624-103X], Florian
Lemmerich¹[0000-0001-7620-1376], and Markus Strohmaier^{1,2}[0000-0002-5485-5720]

¹ RWTH Aachen University, 52062 Aachen, Germany

² GESIS - Leibniz Institute for the Social Sciences, 50667 Cologne, Germany

anna.wegmann@rwth-aachen.de,
{markus.strohmaier,florian.lemmerich}@cssh.rwth-aachen.de

Abstract. Automatically detecting semantic shifts (i.e., meaning changes) of single words has recently received strong research attention, e.g., to quantify the impact of real-world events on online communities. These computational approaches have introduced various measures, which are intended to capture the somewhat elusive and undifferentiated concept of semantic shift. On the other hand, there is a longstanding and well established distinction in linguistics between a word’s *paradigmatic* (i.e., terms that can replace a word) and *syntagmatic* associations (i.e., terms that typically occur next to a word). In this work, we join these two lines of research by introducing a method that captures a measure’s sensitivity for paradigmatic and/or syntagmatic (association) shifts. For this purpose, we perform synthetic distortions on textual corpora that in turn induce shifts in word embeddings trained on them. We find that the *Local Neighborhood* is sensitive to paradigmatic and the *Global Semantic Displacement* is sensitive to syntagmatic shift in word embeddings. By applying the newly validated paradigmatic and syntagmatic measures on three real-world datasets (Amazon, Reddit and Wikipedia) we find examples of words that undergo paradigmatic and syntagmatic shift both separately and at the same time. With this more nuanced understanding of semantic shift on word embeddings, we hope to analyze a similar concept of semantic shift on RDF graph embeddings in the future.

Keywords: Semantic Shift Detection · Paradigmatic Associations · Syntagmatic Associations · RDF Embedding Shift

1 Introduction

In the context of word meaning, linguistic theory has long since distinguished between two fundamentally different types of word relations (e.g., [25,26]) that even have been claimed to correspond to basic operations in the brain, cf. [6,29]: *Paradigmatic* associations of a word w are terms that occur with the same context words as w (i.e., which can substitute w without changing the sentence’s

grammatical structure), e.g., “cat” and “dog”. *Syntagmatic* associations of a word w are terms that co-occur with w , e.g., “cat” and “wild”. This notion is transferable to knowledge representations such as RDF graphs: Paradigmatically related entities would be those that can be replaced by each other (e.g., “cold” by “sniffles” in a symptom-disease network). Syntagmatically related entities, would be those that connect to each other in a network (e.g., “cold” and “coughs”).

On the other hand, popular methods for densely encoding word meaning for computational use are word embeddings (e.g., word2vec or GloVe), which also form the basis for important algorithms for knowledge graph embeddings [4,22]. When words or RDF graph entities change their meaning over time, text corpora (resp. RDF graphs) from these time periods – and consequently the embeddings trained on them – encode these semantic shifts. Several measures of semantic (in-)stability, which can be used to infer meaning shifts from changes in word embeddings, have been proposed in literature [11,12,16,17,30]. However, it is currently unclear what exactly they are measuring in relation to paradigmatic and syntagmatic associations. Thus, we evaluate different computational approaches for detecting *semantic shift* in this paper. We define the semantic shift of a word or entity as anything that affects its paradigmatic and syntagmatic associations. While we focus on word embeddings in this paper, we see our work also as a step towards analyzing changing knowledge graph embeddings in the future [10,15].

Research Questions. In particular, we aim to investigate the following research questions regarding word embeddings: (i) How can the sensitivity of semantic shift measures to paradigmatic and syntagmatic shift be evaluated? (ii) What are the differences in the measures’ sensitivity? (iii) Can both types of shift be observed in real-world datasets and do they always co-occur with each other?

Approach. Based on theoretical considerations we perform a series of experiments in which we synthetically modify a text corpus (similar to [28]) to induce paradigmatic and/or syntagmatic shift. Then, we calculate word embeddings on these corpora and check whether the measures detect the different types of introduced shifts. We compare the performance of the different measures to identify those that are best at detecting paradigmatic (syntagmatic) shift. We apply those measures to detect words that underwent association shifts on three real-world datasets and evaluate the relation between the two forms of shift on them.

Results. By and large, our findings suggest that the *Local Neighborhood* is sensitive to paradigmatic, and the *Global Semantic Displacement* is sensitive to syntagmatic shift (both defined in [11]). Both types of shift occur in real-world datasets. We find examples of simultaneous paradigmatic and syntagmatic shift, paradigmatic without syntagmatic and syntagmatic without paradigmatic shift.

Contribution. We develop an evaluation framework of general semantic shift measures on the basis of the longstanding linguistic distinction between different forms of word associations. We demonstrate that the resulting forms of shift can be inherently different. This contributes to a more nuanced understanding of semantic shift mechanisms. This will enable future work to improve the explainability of automatically detected semantic shift in word and RDF embeddings.




Fig. 1: Examples of paradigmatic and syntagmatic shift. 1a illustrates a paradigmatic shift of the word “refugee”: In the text, its paradigmatic association “displaced person” is replaced by “migrant”. As a result, its nearest neighbors (NNs) in the embedding change accordingly. On the other hand, the text in 1b demonstrates a syntagmatic shift. Modifications of words co-occurring with “refugee” (e.g., “Syrian”) lead to a shift in its embedding vector.

2 Related Work

This section discusses existing literature with respect to paradigmatic and syntagmatic relations in computational approaches, the definition of semantic shift, approaches to measuring semantic shift, and the performance of such measures.

For word space models, several considerations have been made with regard to paradigmatic and syntagmatic relations (see [25,27,31]): Sahlgren concludes that word space models based on either paradigmatic or syntagmatic relations capture different semantic properties. Sun et al. [31] also emphasize that it is important to capture both relations to represent linguistic properties. To our knowledge, association *shifts* have not been considered before. In this work, we provide empirical evidence for Hamilton et al. [11]’s theory that the Local Neighborhood is more sensitive to shifts in a word’s paradigmatic than syntagmatic relations.

To our knowledge there exists no unambiguous definition of semantic shift (for computational use). Most previous work on automatic semantic shift detection does not define semantic shift (e.g., [11,12,24,32]) or defines it circularly (e.g., in [9,18,28]). Linguists seem to use a similar approach (e.g., [1,2,33]). There are some attempts at further isolating this elusive concept by giving explicit examples of what a semantic shift should not be (e.g., non-seasonality in [28]).

Regarding the quantification of semantic shift, the state-of-the-art methods are based on word embeddings (e.g., [19]), which are subject to some inherent drawbacks (c.f. [32]). Several detection approaches are utilized on them (c.f. [18]): Neighborhood-based approaches compare the nearest neighbors of a word between two time steps (e.g., [8,11,21]). Another group of common measures calculate the cosine similarities between the word vectors of different embeddings (e.g., [11,14,16]). For this, embeddings are first made comparable, e.g., by using previous results for embedding initialization (e.g., [16]) or by aligning embeddings after training them individually (e.g., [17]). Shoemark et al. [28] find that aligned perform better than continuously trained embeddings.

Kim et al. [16] identify “interesting” shift words by selecting those with the lowest similarity between the first and last embedding of the series. Others make use of different correlation measures (e.g., [12,28]). Kulkarni et al. [17] search for the words with the biggest mean similarity shift before and after a detected shift point. Jatowt et al. [14] include word frequency in this consideration.

With respect to evaluating the performance of semantic measures, quantifying the effect of noise (see [20,34]) can be a first step (e.g., [7,17,28]). Others rely on human-annotated lists or qualitative human evaluation (e.g., [11,12,16,23]). An increasingly popular approach is to use a form of synthetic evaluation (e.g., [17,24,28]). Rosenfeld et al. [24] expand the *donor-receptor approach* (see [17]) by modeling a gradual change from one meaning to another. Shoemark et al. [28] validate the measures in separate experiments - those where the measure should not and those where they should detect semantic shift.

3 Semantic Shift

Next, we define semantic shift, semantic measures for comparing two embeddings and an approach for detecting interesting shifts on diachronic embeddings.

3.1 Paradigmatic and Syntagmatic Shift

The *contextual normality* approach, cf. [5], expresses that anything that affects the way a word is normally used contributes to its meaning. According to structuralist theories, the only types of relations between words are syntagmatic and paradigmatic (cf. [25,26]). Consequently, we define a semantic shift of a word as anything that affects its syntagmatic or paradigmatic associations (see Fig. 1).

General problem definition. The goal of semantic shift detection is usually generalized as studying a word w over several texts T_1, \dots, T_k in time sensitive order. For this, we use word embedding algorithms to train dense d -dimensional representations of words. For a text T_i , we denote the word embeddings obtained this way as E_i . Intuitively, the vector for word w in embedding E_i (i.e., a single column in the embedding matrix) represents semantic properties of the word w in text T_i . In this paper, we want to identify measures that can quantify paradigmatic (syntagmatic) shifts. For an arbitrary word w and two texts T_1 and T_2 , an *ideal measure of paradigmatic (syntagmatic) shift* satisfies the following: (i) Its range is $[0, 1]$. (ii) At a value of 0 the paradigmatic (syntagmatic) associations of w in T_1 and T_2 have nothing in common. (iii) At a value of 1 the paradigmatic (syntagmatic) associations are the same. (iv) The values between these extremes change linearly with the shift in the paradigmatic (syntagmatic) associations.

3.2 Measuring Semantic Shift

To identify approximations of such an ideal measure, we investigate measures from literature and introduce adaptations thereof. Different to the cited literature, we use cosine similarities and not cosine distances (i.e., $1 - \text{cosine similarity}$).




Fig. 2: *Illustration of a Shift Mask.* The Shift Mask $[1, 0, 1, 1, 1, 1]$ for years 0-6 is displayed. Here, we would be interested in words that undergo a significant change in year 2, and only in year 2. The words that match this mask most closely might have consecutive measure values that are similar to the blue curve.

Global Semantic Displacement. Hamilton et al. [12] define the *Global Semantic Displacement* (SD). They use an embedding alignment approach by solving the Orthogonal Procrustes Problem. Then, the cosine similarity between the aligned word vectors of the word w is calculated.

Local Neighborhood. Another approach to this task is the *Local Neighborhood* (LN) (see [11]). It computes semantic shift via the k nearest neighbors of the word. More precisely, it is defined as the cosine similarity of the vector of cosine similarities between w and its k nearest neighbors in E_1 and E_2 respectively. As suggested, we use $k = 25$ throughout this work (cf. [11]).

Angle-transitioned Local Neighborhood and Semantic Displacement. The Global Semantic Displacement and the Local Neighborhood both utilize the cosine similarity between word vectors. As a result, they are not linear with regard to the change in the included angle. Still, the included angle might change linearly with the paradigmatic (syntagmatic) shift. We propose the *angle-transitioned Semantic Displacement* ($f(\text{SD})$) and the *angle-transitioned Local Neighborhood* ($f(\text{LN})$). These can be computed by the function $f(x) = 1 - \frac{1}{90} \cdot \arccos(\max(x, 0))$. It computes the relative size of the angle, when x is the cosine similarity of the vectors. We assume that every angle over 90 degrees already indicates a maximal semantic distance between two word vectors.

3.3 Detecting Diachronic Semantic Shift

Let us assume that we know an ideal paradigmatic (syntagmatic) measure. Then, in a diachronic embedding series E_1, \dots, E_k , we want to detect words that underwent an “interesting” shift:

First, we define (a) the *consecutive measure* values for a word w as all the measure values for w between two subsequent embeddings of the diachronic series, i.e., E_i and E_{i+1} for an $i \in \{1, \dots, k-1\}$ and (b) the *reference measure* values for a word w as all the measure values for w between the first and every other embedding of the diachronic series, i.e., E_1 and E_i for $i \in \{2, \dots, k\}$.

Then, to find specific shift behavior, we compare these values with a user-defined *desired shift*: It consists of (i) the considered type of shift (i.e., paradigmatic or syntagmatic), (ii) the *shift intervals* (i.e., intervals in which the semantic shift should occur) and (iii) the *desired shift development* (i.e., whether the words should develop towards a new or back towards their original meaning).

(ii) is given by a *Shift Mask* (i.e., a series of $k - 1$ values that are 0 for the shift interval and 1 otherwise), see Fig. 2. The comparison of the desired shift with the actual paradigmatic (syntagmatic) measure behavior of every word w takes place in two steps: (1) comparing the consecutive measure values of w and the shift mask (see Fig. 2) and (2) comparing the reference measure values of w and the desired shift development. The comparisons could, for example, take place via a mean squared error, a Pearson Correlation or a threshold.

4 Simulating Semantic Shift

This section introduces a framework for simulating semantic shifts via five different types of synthetic corpus distortions, which we will call attacks. The core idea is based on the *donor-receptor* approach, where the *donor* “donates” its place in the corpus to the *receptor* word with a given probability (cf. [17]). We compare three different semantic (*Paradigmatic Attack*, *Syntagmatic Attack*, *Combined Attack*) and two baseline attacks (*Baseline - No Change*, *Baseline - Random Attack*). We give an overview of the semantic attacks (cf. Table 1) and the expected embedding change (cf. Fig. 3), where p signifies the extent of the distortion.

4.1 Baseline

No Attack. As a simple baseline, we train multiple embeddings on the same corpus. Variations result from the inherent instability of the embedding algorithms.

Random Attack. Additionally, we test robustness of measures under no association shift for the considered word but significant shift in other words:

For a donor (word) d and a bijection $B : \text{Vocabulary} \rightarrow \text{Vocabulary}$, we define the *Random Distortion* $R(T, d, k, p, B)$ of a text corpus T to be T' where each word v is replaced by $B(v)$ with probability p . We additionally restrict the distortion to only those sentences where neither d nor any of its k closest paradigmatic associations occur (we denote this set of words as W). Consequently, the distortion induces no syntagmatic or paradigmatic shift for d (e.g., Fig. 3a).

We arrange every word in the corpus in an interval between 0 and 1 according to frequency. We select 10 donor words per *frequency interval* in $\{[0.1, 0.2], \dots, [0.8, 0.9]\}$, leaving out the 10% most frequent and least frequent words. The bijection B is chosen randomly on $V \setminus W$. We set $k = 50$ as we assume all paradigmatic associations to be among the first 50 paradigmatically related words. We

Table 1: *Overview of semantic attacks.* For every attack, we summarize the words affected by the introduced shift and what an ideal measure would detect.

Attack	simulates shift	on words	ideal measure
Baseline - No	no	any	constant at 1.0
Baseline - Random	no	donor	constant at 1.0
Combined	parad. and syntag.	receptor	linear with $\frac{1}{1+p}$
Parad.	parad.	donor	linear with $\frac{1}{1+p}$
Syntag.	syntag.	receptor	linear with $1 - p$

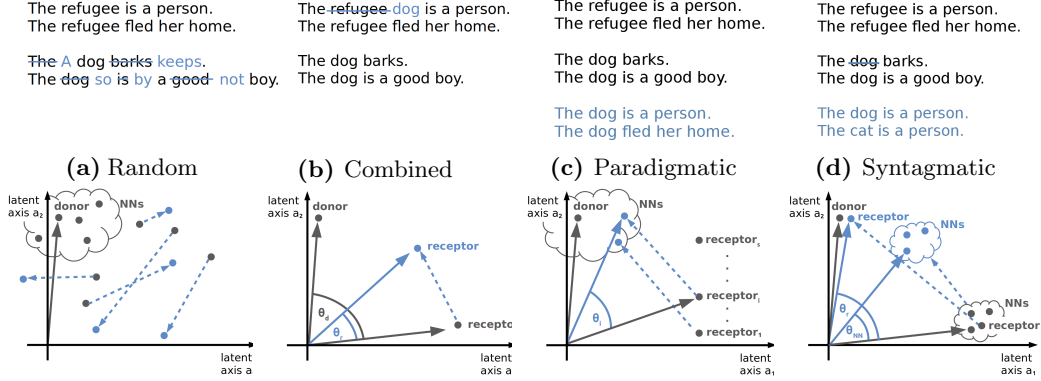


Fig. 3: Textual distortions and the expected embedding change resulting from different attacks. In the first row examples of textual distortions with $p = 1/2$ are displayed. For this purpose, “refugee” refers to the donor and if necessary, “dog” to the receptor with “cat” being its only paradigmatic association. In the second row, the expected embedding changes are displayed. These are the result of our intuitive understanding of the mechanisms behind word embeddings and purely displayed for illustrational purposes. Here, NNs refers to the nearest neighbors to the donor or receptor word. In *Baseline - Random Attack*, only sentences with the donor and the paradigmatic associations of the donor stay the same. In *Combined Attack*, the occurrences of the donor are replaced by the receptor. In *Paradigmatic Attack*, the receptors increasingly occur with the context words of the donor. In *Syntagmatic Attack*, the original co-occurrences of the receptor and its paradigmatic associations are replaced by the donor’s.

calculate the embeddings on $R(T, d, k, p, B)$ for all $p \in \{0.1, 0.2, \dots, 1.0\}$.

In the resulting embedding series, the position of the donor and its nearest neighbors is expected to stay the same as, by design, their co-occurrences do not change. With increasing p , every other word should be subject to substantial position change (cf. Fig. 3a).

4.2 Combined Attack

We test whether the measures can detect any, syntagmatic or paradigmatic, shift:

For a donor d and a receptor word r , we define the *Paradigmatic and Syntagmatic Distortion* $PS(T, d, r, p)$ of a text corpus T to be T' where every occurrence of the donor d is replaced by the receptor r with probability p . Consequently, the receptor word undergoes syntagmatic as well as paradigmatic shift (e.g., Fig. 3b).

We randomly select 10 word pairs from each frequency interval in $\{[0.1, 0.15], \dots, [0.85, 0.9]\}$. Therefore, we consider 160 (donor, receptor)-pairs in total. We calculate the embeddings on $PS(T, d, r, p)$ for all $p \in \{0.1, 0.2, \dots, 1.0\}$.

Due the frequency-based selection procedure, we assume the number of sentences n in which d occurs in to be approximately equal to the number that r oc-

curs in. Then, $\frac{1}{1+p} = \frac{n}{n+p \cdot n}$ equals the share of the receptor's occurrences in its original sentences (i.e., with its original paradigmatic/syntagmatic associations). The reference values of an ideal measure of paradigmatic (syntagmatic) shift should be linear with this fraction for the receptor. The donor word should undergo minor syntagmatic shift up until a point from which it drastically deteriorates to 0 as it does not occur in $PS(T, d, r, p)$ for $p = 1$. Its paradigmatic change might be significant with the function $\frac{1}{1+p}$. The paradigmatic associations of the donor and receptor word could also undergo some minor change with the shift of the receptor and donor word. This could lead to a worse performance in the altered words prediction of a paradigmatic compared to a syntagmatic measure.

In the resulting embedding series, we expect the receptor representation to develop towards the original donor representation with increasing p (cf. Fig. 3b).

4.3 Paradigmatic Attack

We test whether the measures can pick up on paradigmatic association changes:

For a donor word d , l receptor words $(r_1, \dots, r_l) =: r$ and probabilities $p := p_1, \dots, p_l \in [0, 1]$, we define the *Paradigmatic Distortion* $P(T, d, r, p)$ of a text corpus T to T' . In T' , for every sentence d occurs in and for each receptor word r_i , a new sentence is added with probability p_i in which every occurrence of d is replaced by r_i . Consequently, we induce a paradigmatic but no syntagmatic shift of the donor by adding sentences (e.g., Fig. 3c).

We randomly select 10 (donor, receptors)-pairs per 0.05 frequency interval from 0.1 to 0.9, i.e., 160 donor words in total. We set $l = 10$ as we assume the changes in the 10 closest paradigmatic associations to be significant for d . We introduce increasing changes in 10 consecutive steps i . We set p_j in step i to $\delta_{j, \min(i, j)}$, where δ is the Kronecker Delta. As a result, the 10 receptor words successively become the new closest paradigmatic associations of the donor word. For consistency, we will also refer to the different steps with $p = i/10$ for $p \in \{0.1, \dots, 1.0\}$. We calculate the embeddings on $P(T, d, r, p)$ for all $p \in \{0.1, 0.2, \dots, 1.0\}$.

The paradigmatic change of the donor as well as the receptor words should be linear with $\frac{1}{1+p} = \frac{1}{1+i \cdot (1:l)} = \frac{nl}{nl+n \cdot i}$ for step i . This formula represents the share of occurrences of the $l = 10$ receptor words in their original sentences.

In the resulting embedding series, we expect the receptor representations to successively develop towards the donor representation, therefore altering the donor's nearest neighbors (cf. Fig. 3c).

4.4 Syntagmatic Attack

Finally, we describe a test for whether measures can detect syntagmatic changes. Here, we aim to distort the corpus such that the syntagmatic shift is significant for the considered word, while the paradigmatic shift for the same word is smaller or non-distinguishable from a larger set of words:

For a probability $p \in [0, 1]$, a donor d , and a receptor r , we define the *Syntagmatic Distortion* $S(T, d, r, p, k)$ of the text corpus T to be T' where, for every

sentence d occurs in, the sentence is added $k + 1$ times with probability p . Here, d is replaced by r or its i th paradigmatic association n_i respectively for $i \leq k$. Moreover, for each original sentence r occurs in, r is deleted from the sentence with probability p , leaving it “incomplete”. Similarly, for each original sentence where n_i occurs in, it is deleted with probability $p/2$. Thus, we introduce a syntagmatic as well as a substantially less pronounced paradigmatic shift for r by adding and altering sentences d , r or n_i occur in (cf. Fig. 3d).

This is done for overall 32 donors – 4 out of each frequency interval in $\{[0.1, 0.2], \dots, [0.8, 0.9]\}$. We set $k = 25$ as we assume that the 25 closest paradigmatic associations include the most relevant. We chose to use significantly less pairs and $k < 50$ as for each donor word $k + 1$ new sentences are added and, additionally, any sentence where r or n_i occur in are altered. This changes the original corpus exponentially more than before. We calculate the embeddings on $S(T, d, r, p, k)$ for all $p \in \{0.1, 0.2, \dots, 1.0\}$.

r undergoes the greatest syntagmatic shift among all words, since its original occurrences decrease with $1 - p$. As a result, the syntagmatic change of r is the most correlated with $1 - p$. The syntagmatic shift of its original k paradigmatic associations is also related to $1 - p$. d , r as well as the n_i undergo paradigmatic shift. However, the paradigmatic shift for r should be considerably smaller than its syntagmatic shift as we perform similar changes for its paradigmatic associations.

The position of the receptor representation is expected to shift to the previous donor position (see Fig. 3d). The nearest neighbors of the receptor word should shift towards it as well but also stay between the original donor and receptor representation.

5 Experiments

This section describes our experimental setup and results.

5.1 Datasets and Training

We work with three different datasets: Firstly, the *Reddit* comments and submissions from 2012-2018³. Secondly, the “aggressively deduplicated” *Amazon* reviews data from May 1996 - July 2014 (cf. [13]). Thirdly, the *Wikipedia* snapshots from 2014-2018⁴. We converted each character to lowercase and filtered out URLs. We removed each non-alphanumeric symbol and treated them as separation between words, i.e., conversion of “i’ve” to “i” and “ve”. We assume that, simple tokenization works comparably well (cf. [3]). For all experiments, we use

³Baumgartner, J.: Reddit dataset, <https://files.pushshift.io/reddit/>, (accessed on 2019-09-25)

⁴wikimedia: wikipedia snapshots on archive.org, <https://archive.org/download/enwiki-20150112>, <https://archive.org/download/enwiki-20160113>, <https://archive.org/download/enwiki-20170101>, <https://archive.org/download/enwiki-20180101>, <https://archive.org/details/enwiki-20190120>, (accessed on 2019-09-25)

the 2012 Reddit corpus with about 8.8 billion words, the 2014 Amazon review corpus with about 1 billion words and the 2016 Wikipedia snapshot with about 2.5 billion words as a basis. We call these corpora the *basis corpora*.

We use the multi-threaded Python framework *gensim* to train word2vec embeddings in the faster CBOW variant and negative sampling at 5 (cf. [19]). CBOW and negative sampling perform better for frequent than infrequent words⁵. We use 300 dimensions and a general min_count of 60. The number of epochs is chosen at 4. All other parameters are left at their default values.

5.2 Approach

We test whether the synthetically distorted words (see Section 4) can be detected by the different measures. As the words consistently change more with increasing p , we skip (1) of the approach detailed in Section 3.3. Consequently, we assume that (1) returned all words as candidates that could match the induced form of shift. In (2), for every measure, we predict the synthetically altered words by identifying those that have the highest Pearson Correlation with the expected shift (i.e., $\frac{1}{1+p}$ or $1 - p$, cf. Table 1). We perform the evaluation via an *accuracy curve*, i.e., the share of the correctly predicted words out of the actually changed words (cf. [17,28]).

5.3 Outcome

Representative results on the Reddit data for this approach are shown in Fig. 4. Results for the other datasets were equivalent. Key observations (in bold) include:

All measures detect a form of paradigmatic or syntagmatic shift. The results of the Combined Attack show that all measures detect a form of paradigmatic and/or syntagmatic shift (see Fig. 4c). SD and f(SD) are even behaving close to linear with the expected paradigmatic and syntagmatic change of $\frac{1}{1+p}$. In Section 4.2, we expected the paradigmatic shift to correlate with more than just the receptors. Therefore, LN and f(LN) could be less accurate because they are more paradigmatic measures.

The LN-measures perform best at detecting paradigmatic shift. f(LN) and LN are the best at detecting paradigmatic shift (see Fig. 4d). f(SD) and SD do not pick up on paradigmatic shift at all. The upper limit of the x-axis is at $160 \cdot (1 + 10) = 1760$ as not only the donor but also the receptors change paradigmatically (cf. Section 4.3). Surprisingly, for LN-based measures, there is a plateau reached after the first 160 predicted words. A potential reason for this is that the measures cannot detect finer paradigmatic changes for some (donor, receptors)-pairs that had a lower starting angle.

The SD-measures perform best at detecting syntagmatic shift. As seen in Fig. 4c, the SD-based measures seem to change linearly with the introduced

⁵google: word2vec documentation, <https://code.google.com/archive/p/word2vec/>, (accessed on 2019-09-25)

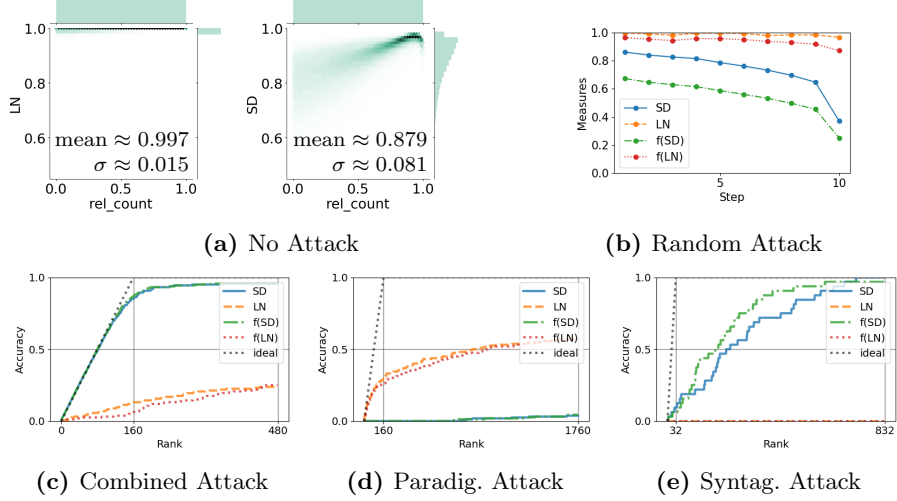


Fig. 4: *Experimental results on Reddit data.* (a) shows heatmaps of the relative sorted frequency for every word (rel_count, where 1 is most frequent) and its mean semantic shift according to LN and SD. Overall distributions are shown at the side and the top of the plot. (b) displays the average measure values over the unchanged words for each step. LN stays closest to the ideal constant value of 1.0. For the semantic attacks (c,d,e), the accuracy curve shows the share of the altered words that were correctly predicted by the Pearson Correlation. The dotted black line shows the ideal measure behavior. All measures are able to measure syntagmatic and/or paradigmatic shift but to a varying extent. For the Combined Attack and Syntagmatic Attack (c,d), SD and its angle-transitioned variation f(SD) perform the best. By contrast in d), LN and its angle-transitioned variation f(LN) are the best at detecting the induced paradigmatic shifts.

paradigmatic and/or syntagmatic shifts. Additionally, we discovered that they do not detect paradigmatic shifts at all in Fig. 4d. Therefore, they must be able to detect syntagmatic shifts considerably well. The Syntagmatic Attack confirms this (see Fig. 4e). We look at the first $832 = 32 \cdot (1 + 25)$ ranks as the 25 nearest neighbors (chosen as an approximation for the paradigmatic associations, see discussion) of the receptor word also change syntagmatically as discussed in Section 4.4. f(SD) and SD behave the best at detecting syntagmatic changes, while LN and f(LN) do not detect them at all.

f(SD) is noisier than SD and f(LN) is noisier than LN. We studied whether the measures can pick up on paradigmatic (syntagmatic) changes. But what if there is neither? The results of the Baseline experiments show that LN is more robust than SD under no association changes (cf. Fig. 4a-4b). For Baseline - No Attack (in Fig. 4a), all measures perform well for the most frequent words and considerably worse for the least frequent words. LN performs the best and f(SD) the worst, while f(LN) behaves better than the SD-variations. The plots for

$f(\text{LN})$ and $f(\text{SD})$ are left out as they are monotone distortions LN and SD. Due to the chosen linearization approach, they are a lot more sensitive to the random differences between word embeddings trained on the same corpus. This noise seems not to be worth the small advantage (see Fig. 4e) of having a linear change measure. The results for Baseline - Random Attack (see Fig. 4b) are comparable.

Overall, we conclude that SD is the best measure for detecting syntagmatic and LN is the best measure to detect paradigmatic shift.

6 Application Examples

We use the our new insights to detect paradigmatic and syntagmatic shifts in real world data, i.e., the Amazon, Reddit, and Wikipedia corpora. Based on our previous results, we utilize LN as the paradigmatic and SD as the syntagmatic measure. We define the shift interval of the *desired shift* (cf. Section 3.3) as an empty or one point interval (i.e., words with one shift point or none at all). We detect the most interesting words as those with the lowest mean squared error to the desired shift mask. We observe the following key findings:

A word can undergo paradigmatic and syntagmatic shift to different extent. For example, the number of words that were among most shifting words in a given year according to the paradigmatic (syntagmatic) measure and also are within the most shifting 25 according to the syntagmatic (paradigmatic) measure is consistently less than 5 (cf. *overlap* in Table 2). The overlap is 0 for no shift points: LN mostly detects (year) numbers and SD detects nouns and verbs as the most constant words. A reason for this could be that the position

Table 2: *Top 3 syntagmatic and paradigmatic shift words.* The overlap specifies the number of the top 5 most changed words according to one measure that are contained in the top 25 words of to the other. Words do not undergo paradigmatic and syntagmatic shift to the same extent.

Shift in	Top-3 syn. shift words	Top-3 para. shift words	overlap	
never	songs, get, story	1991, 1975, 1973	0	0
2007	kindle, plastics, leopard	kindle, hg, reroute	2	1
2012	insurgent, vita, g5	vita, bared, marquee	2	4

Amazon reviews from 2005 to 2014.

Shift in	Top-3 syn. shift words	Top-3 para. shift words	overlap	
never	cdotas, pidamente, abdomen	01100011, 01100100, 01110101	0	0
2014	braum, oras, 20ex	w33, triche, oras	2	1
2016	ladybonersgw, nougat, trumper	grubbin, coolheaded, tdil	1	1

Reddit from 2012 to 2018.

Shift in	Top-3 syn. shift words	Top-3 para. shift words	overlap	
never	jeandat, subsidiaries, migrate	1885, 1842, 13	0	0
2016	attd, andp, binaria	sanep, thrret, wk14	4	2
2017	vlindernet, 14px, dcrj	intret, kilmainemore, pastorally	2	0

Wikipedia from 2014 to 2018.

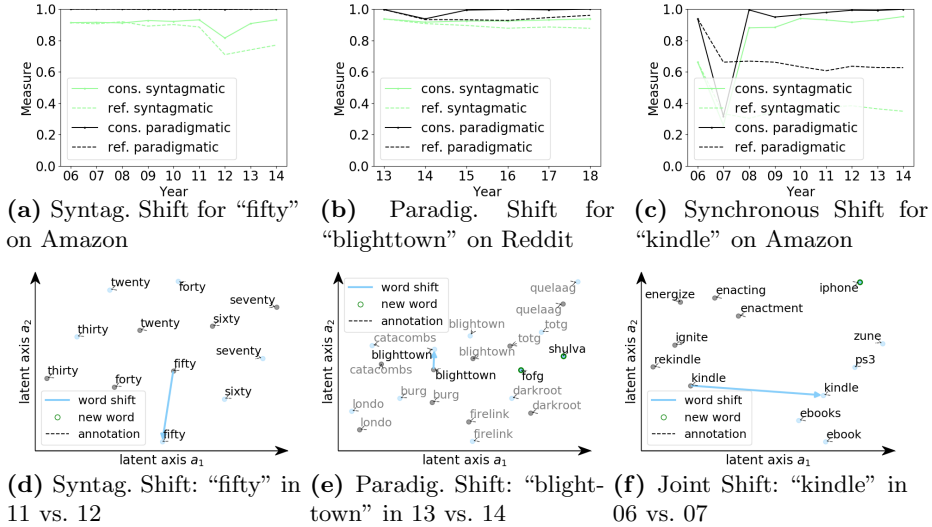


Fig. 5: *Conflicting and synchronous behavior of the paradigmatic and syntagmatic measures.* Plots in the first row display consecutive (continuous line) and reference (dashed line) point values of the paradigm. and the syntag. measure. Plots in the lower row show the t-SNE projection of embeddings before (grey) and after the shift (blue dots). Green words were not present in the previous year. “fifty” has a syntagmatic but no paradigmatic shift point in 2012. The nearest neighbors of “fifty” stay the same, the vector however moves away from its original position (blue arrow) similar to Fig. 1b The word “blighttown” exhibits a paradigmatic shift between 2013 and 2014. “blighttown” and its former nearest neighbors stay at a similar position, but new words appear. The measures synchronously detect a change for “kindle”. The “kindle” vector moves out of the previous nearest neighbors in 2006 towards new nearest neighbors in 2007.

of a year in a sentence is rather unique, while the exact words it occurs with (i.e., syntagmatic associations) change. We give two examples of words undergoing paradigmatic and syntagmatic shift to different extent:

The word “fifty” (see Fig. 5a) is the 16th matched word to the shift mask introducing change in 2012 on the Amazon reviews corpus for the syntagmatic measure. The paradigmatic measure does not display any meaningful shift as the closest paradigmatic associations stay the same (“twenty”, “sixty”, ... in Fig. 5d). The t-SNE projections in Fig. 5 were calculated via the 200 nearest neighbors of the considered word (similar to [12]). The syntagmatic change probably occurred because of the published print of *Fifty Shades* by E. L. James in 2012⁶. This assumption is based on the fact that the search of “fifty” returns

⁶wikipedia: Fifty shades of grey, https://en.wikipedia.org/wiki/Fifty_Shades_of_Grey (accessed on 2019-10-14)

Fifty Shades products on Amazon⁷ and it is to be expected that there have been many copies of the bestselling book sold via Amazon in its publishing year.

The word “blighttown” (see Fig. 5b) underwent significant paradigmatic but no syntagmatic change. We detected it by selecting for words with a low Pearson Correlation between the reference measure values of the paradigmatic and syntagmatic measure. The paradigmatic curve shows a shift in 2014, the syntagmatic measure does not. “blighttown” is the name of an area in the video game *Dark Souls*. In 2014, *Dark Souls II* was released with the new areas “shulva”⁸ and “fog” (short for “Forest of Fallen Giants”⁹), which correspond to new nearest neighbors of “blighttown”. The position of “blighttown” and the other nearest neighbors did not change (cf. Fig. 5e).

Words with an extreme shift in one measure follow a similar trend in the other. The Pearson Correlation of the paradigmatic and syntagmatic reference measure values for the words in the top 25 is mostly moderate to high (above 0.2). Therefore, although different words are predicted for the greatest paradigmatic vs. syntagmatic changes, it is still likely that words with an extreme shift in one measure also undergo a shift in the other. For example, the word “kindle” is the first predicted word for the shift point in 2007 for both measures (plot see Fig. 5c). The Amazon kindle was introduced in 2007.

7 Discussion

In the following, we address potential criticism and limitations of our work:

“Paradigmatic (syntagmatic) shifts are not necessarily semantic shifts according to common understanding” Words like “christmas” are talked about differently in December than in April due to seasonal variations. Arguably there has also been a shift in the way people use “refugee” after 2015. Are those types of shifts semantic shifts? According to our approach they are. According to the common rather fuzzy understanding, they are probably not. We still decided for those types of changes to be defined as semantic shift as they describe interesting societal dynamics and changes in the way people think about different concepts.

“Measure shifts do not necessarily occur due to paradigmatic or syntagmatic shift” We showed that paradigmatic (syntagmatic) shifts lead to measure shifts. We partly evaluate the reverse with the baseline experiments. We recommend the addition of further experiments (e.g., for word frequency as done in [28]).

“There are regularities in the types of words that are changing the most.” We did not statistically evaluate which word types are prevalent. However, LN is more sensitive to changes in nouns than SD (see [11]). This could be connected to the paradigmatic vs. syntagmatic association distinction: Nouns are more likely to

⁷amazon: amazon search for “fifty”, https://www.amazon.com/s?k=fifty&ref=nb_sb__noss (accessed on 2019-09-18)

⁸darksouls.fandom.com: Shulva, Sanctum City, https://darksouls.fandom.com/wiki/Shulva,_Sanctum_City (accessed on 2019-09-30)

⁹darksouls.fandom.com: Forest of fallen giants, https://darksouls.fandom.com/wiki/Forest_of_Fallen_Giants (accessed on 2019-09-30)

undergo “cultural shift” (see [11]). As a result their paradigmatic associations might be completely replaced while syntagmatic associations stay more constant (due to, e.g., co-occurring verbs and grammatical forms).

“Paradigmatic associations in texts might not be the same as the nearest neighbors in embeddings” In (a) Baseline - Random Attack and (b) Syntagmatic Attack, we assume that the closest paradigmatic associations of a word have a significant overlap with its nearest neighbors in the embedding. This is an intuitive assumption since the positions of the word vectors should mostly be determined by their syntagmatic associations. The results from the Paradigmatic Attack, which was performed independently from (a) and (b), also make this assumption reasonable: LN, which calculates shifts via nearest neighbor changes, performed the best at detecting paradigmatic association changes.

“There is not only syntagmatic shift introduced in the Syntagmatic Attack” In designing the Syntagmatic Attack, we found no simple method to synthetically introduce the same kind of syntagmatic shift for a group of words without introducing similar paradigmatic shift for a subset of this group as well. This is because as soon as a syntagmatic change to a word w is introduced, the previous paradigmatic associations are less strongly related to w than before. Altering those paradigmatic associations again introduces syntagmatic change.

“The synthetic corpus changes might introduce unwanted association shifts” We add several sentences with nearly the same words or remove single words from sentences in the Paradigmatic and the Syntagmatic Attack. Here, we want to only introduce syntagmatic change to one word. The other words in the added sentences also undergo syntagmatic change. However, we assume this effect to be negligible since most co-occurrences stay the same.

8 Conclusion

In this work, we introduced an operationalization of semantic shift via paradigmatic and syntagmatic associations. We studied a variety of measures in their abilities for detecting and discerning between paradigmatic and syntagmatic shifts. We evaluated them on word embeddings trained on corpora that were synthetically distorted. We observed that the *Local Neighborhood* captures paradigmatic shift, while the *Global Semantic Displacement* captures syntagmatic shift. We showed examples where those measures are behaving differently. The main contributions are (i) the differentiation of semantic shift with the help of a well-established linguistic approach, (ii) the introduction of an evaluation framework of semantic shift measures via synthetic experiments, (iii) the identification of the best paradigmatic and syntagmatic measure and (iv) a demonstration that the two associations shifts can be inherently different. Future work will include the application of the paradigmatic (syntagmatic) measure for the analysis of diachronic shift in RDF graphs. Then, thresholding of our approach could give a clear signal for when a public RDF graph or embedding should be updated.

Acknowledgments

Part of the simulations were performed with computing resources granted by RWTH Aachen University. We thank Dong Nguyen for providing advice regarding this work and our (meta-) reviewers for their constructive feedback.

References

1. Bloomfield, L.: *Language*. Allen & Unwin (1933)
2. Bréal, M.: *Essai de sémantique*. Lambert-Lucas (1897)
3. Camacho-Collados, J., Pilehvar, M.T.: On the role of text preprocessing in neural network architectures: An evaluation study on text categorization and sentiment analysis. *EMNLP Workshop*, pp. 40–46 (2018)
4. Cochez, M., Ristoski, P., Paulheim, H.: Global rdf vector space embeddings. *ISWC*, pp. 190–207 (2017)
5. Cruse, A.: *Meaning in Language*. Oxford Linguistics (2004)
6. Dell, G.S., Oppenheim, G.M., Kittredge, A.K.: Saying the right word at the right time: Syntagmatic and paradigmatic interference in sentence production. *Language and Cognitive Processes*, **23**(4), pp. 583–608 (2008)
7. Dubossarsky, H., Grossman, E., Weinshal, D.: Outta control: Laws of semantic change and inherent biases in word representation models. *EMNLP*, pp. 1136–1145 (2017)
8. Eger, S., Mehler, A.: On the linearity of semantic change: Investigating meaning variation via dynamic graph models. *ACL*, pp. 52–58 (2016)
9. Frermann, L., Lapata, M.: A bayesian model of diachronic meaning change. *ACL*, pp. 31–45 (2016)
10. Gutierrez1, C., Hurtado1, C., Vaisman, A.: Temporal rdf. *ESWC*, pp. 93–107 (2005)
11. Hamilton, W.L., Leskovec, J., Jurafsky, D.: Cultural shift or linguistic drift? comparing two computational measures of semantic change. *EMNLP*, pp. 2116–2121 (2016)
12. Hamilton, W.L., Leskovec, J., Jurafsky, D.: Diachronic word embeddings reveal statistical laws of semantic change. *ACL*, pp. 1489–1501 (2016)
13. He, R., McAuley, J.: Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. *WWW*, pp. 507–517 (2016)
14. Jatowt, A., Duh, K.: A framework for analyzing semantic change of words across time. *JCDL*, pp. 229–238 (2014)
15. Khurana, U., Deshpande, A.: Efficient snapshot retrieval over historical graph data. *ICDE*, pp. 997–1008 (2005)
16. Kim, Y., Chiu, Y., Hanaki, K., Hegde, D., Petrov, S.: Temporal analysis of language through neural language models. *ACL Workshop*, pp. 61–65 (2014)
17. Kulkarni, V., Al-Rfou, R., Perozzi, B., Skiena, S.: Statistically significant detection of linguistic change. *WWW*, pp. 625–635 (2015)
18. Kutuzov, A., Øvrelid, L., Szymanski, T., Velldal, E.: Diachronic word embeddings and semantic shifts: a survey. *COLING*, pp. 1384–1397 (2018)
19. Mikolov, T., Sutskever, I., Chen, K., Corradini, G., Dean, J.: Distributed representations of words and phrases and their compositionality. *NIPS*, pp. 3111–3119 (2013)
20. Pierrejean, B., Tanguy, L.: Predicting word embeddings variability. *SEM*, pp. 154–159 (2018)

21. del Prado Martin, F., Brendel, C.: Case and cause in icelandic: Reconstructing causal networks of cascaded language changes. *ACL*, pp. 2421–2430 (2016)
22. Ristoski, P., Rosati, J., Di Noia, T., De Leone, R., Paulheim, H.: Rdf2vec: Rdf graph embeddings and their applications. *Semantic Web Journal*, pp. 721–752 (2019)
23. Rohrdantz, C., Hautli, A., Mayer, T., Butt, M., Keim, D., Plank, F.: Towards tracking semantic change by visual analytics. *ACL*, pp. 305–310 (2011)
24. Rosenfeld, A., Erk, K.: Deep neural models of semantic shift. *NAACL-HLT*, pp. 474–484 (2018)
25. Sahlgren, M.: The word-space model: Using distributional analysis to represent syntagmatic and paradigmatic relations between words in high-dimensional vector spaces. PhD dissertation (2006)
26. de Saussure, F.: *Cours de linguistique generale*. Payot (1916)
27. Schütze, H., Pedersen, J.: A vector model for syntagmatic and paradigmatic relatedness. *Conference of the UW Centre for the New OED and Text Research*, pp. 104–113 (1993)
28. Shoemark, P., Liza, F.F., Nguyen, D., Hale, S.A., McGillivray, B.: Room to glo: A systematic comparison of semantic change detection approaches with word embeddings. *EMNLP and IJCNLP*, pp. 66–76 (2019)
29. Sokolova, L.V., Cherkasova, A.S.: Spatiotemporal organization of bioelectrical brain activity during reading of syntagmatic and paradigmatic collocations by students with different foreign language proficiency. *Human Physiology*, **41**(6), pp. 583–592 (2015)
30. Stewart, I., Arendt, D., Bell, E., Volkova, S.: Measuring, predicting and visualizing short-term change in word representation and usage in vkontakte social network. *ICWSM*, pp. 672–675 (2017)
31. Sun, F., Guo, J., Lan, Y., Xu, J., Cheng, X.: Learning word representations by jointly modeling syntagmatic and paradigmatic relations. *ACL and IJCNLP*, pp. 136–145 (2015)
32. Tahmasebi, N., Borin, L., Jatowt, A.: Survey of computational approaches to lexical semantic change detection. *ACL*, pp. 31–45 (2018)
33. Traugott, E.C., Dasher, R.B.: *Regularity in Semantic Change*. Cambridge University Press (2001)
34. Wendlandt, L., Kummerfeld, J.K., Mihalcea, R.: Factors influencing the surprising instability of word embeddings. *NAACL-HLT*, pp. 2092–2102 (2018)