# **Embeddings-based detection of word use** variation in Italian newspapers

Michele Cafagna<sup>\*</sup> University of Pisa, Italy University of Groningen, The Netherlands University of Malta, Malta Lorenzo De Mattei<sup>\*\*</sup> University of Pisa, Italy ItaliaNLP Lab, ILC-CNR, Pisa, Italy University of Groningen, The Netherlands

Malvina Nissim<sup>†</sup> University of Groningen, The Netherlands

We study how words are used differently in two Italian newspapers at opposite ends of the political spectrum by training embeddings on one newspaper's corpus, updating the weights on the second one, and observing vector shifts. We run two types of analysis, one top-down, based on a preselection of frequent words in both newspapers, and one bottom-up, on the basis of a combination of the observed shifts and relative and absolute frequency. The analysis is specific to this data, but the method can serve as a blueprint for similar studies.

# 1. Introduction and Background

Different newspapers, especially if positioned at opposite ends of the political spectrum, can render the same event in different ways. In Example (1), both headlines are about the leader of the Italian political movement "Cinque Stelle" splitting up with his girlfriend, but the Italian left-oriented newspaper *la Repubblica*<sup>1</sup> (rep in the examples) and right-oriented *ll Giornale*<sup>2</sup> (gio in the examples) describe the news quite differently. The news in Example (2), which is about a baby-sitter killing a child in Moscow, is also reported by the two newspapers mentioning and stressing different aspects of the same event.

(1) rep La ex di Di Maio: "E' stato un amore intenso ma non abbiamo retto allo stress della politica"
 [en: The ex of Di Maio: "It's been an intense love relationship, but we haven't survived the stress of politics"]

gio Luigino single, è finita la Melodia [en: Luigino single, the Melody is over]

(2) rep Mosca, "la baby sitter omicida non ha agito da sola"

<sup>\*</sup> University of Pisa at the time of the work - E-mail: michele.cafagna@um.edu.mt

<sup>\*\*</sup> University of Pisa and ItaliaNLP Lab, ILC-CNR, Pisa, Italy -E-mail: lorenzo.demattei@di.unipi.it

<sup>†</sup> Center for Language and Cognition, E-mail: m.nissim@rug.nl

<sup>1</sup> https://www.repubblica.it

<sup>2</sup> http://www.ilgiornale.it

[en: Moscow, "the killer baby-sitter has not acted alone"]

gio Mosca, la donna killer: "Ho decapitato la bimba perché me l'ha ordinato Allah"

[en: Moscow, the killer woman: "I have beheaded the child because Allah has ordered me to do it"]

Often though, the same words are used, but with distinct nuances, or in combination with other, different words, as in Examples (3)–(4):

- (3) rep Usa: agente uccide un nero disarmato e immobilizzato [*en: Usa: policeman kills an unarmed and immobilised black guy*]
  - gio Oklahoma, poliziotto uccide un nero disarmato: "Ho sbagliato pistola" [en: Oklahoma: policeman kills an unarmed black guy: "I used the wrong gun"]
- (4) rep Corte Sudan annulla condanna, Meriam torna libera [en: Sudan Court cancels the sentence, Meriam is free again]
  - gio Sudan, Meriam è libera: non sarà impiccata perché cristiana [en: Sudan: Meriam is free: she won't be hanged because Christian]

In this work we discuss a method to study how the same words are used differently in two sources, exploiting vector shifts in embedding spaces.

The two embeddings models built on data coming from *la Repubblica* and *ll Giornale* might contain interesting differences, but since they are separate spaces they are not directly comparable. Previous work has encountered this issue from a diachronic perspective: when studying meaning shift in time, embeddings built on data from different periods would encode different usages, but they need to be comparable. Instead of constructing separate spaces and then aligning them (Hamilton, Leskovec, and Jurafsky 2016b), we adopt the method used by (Kim et al. 2014) and subsequently by (Del Tredici, Nissim, and Zaninello 2016) for Italian, whereby embeddings are first trained on a corpus, and then updated with a new one; observing the shifts certain words undergo through the update is a rather successful method to proxy meaning change. The main difficulty concerning the adoption of word embedding shifts is that representation in different corpora can be different without there being any actual sense change. Moreover, it is important to recognize how much of the variation observed is due to randomness and how much is due to actual semantic drift.

Rather than across time, we update embeddings across sources which are identical in genre (newspapers) but different in political positioning. Specifically, we train embeddings on articles coming from the newspaper *La Repubblica* (leaning left) and update them using articles coming from the newspaper *Il Giornale* (leaning right). We compute indexes which represent the observed shift by accounting for the embedding instability of a given word (or the shift in distance between two words) as a proxy for a difference in usage of that term, running two types of analysis. One is top-down, and focuses on a set of specific words which are frequent in both corpora. The other one is bottom-up, focusing on words that result potentially interesting on the basis of measures that combine the observed shift with both relative and absolute frequency. As a byproduct, we also learn something about the interaction of shifts and frequency.



Left: top 100 most frequent words in *la Repubblica*. Right: top 100 in *ll Giornale*. The words are scaled proportionally to their frequency in the respective datasets.

# 2. Related Works

Word sense changes have remained a key area of research right from the very early initiatives in natural language processing research. It can be analysed under different perspective like diachronic changes, political and gender biases. (Tahmasebi, Borin, and Jatowt 2018) albeit, focused on diachronic changes methods, provide a wide overview of approaches applicable to a broad range of semantic tasks involving word sense changes. Co-occurrence-based methods for instance, make use of the mutual information derived from the co-occurencce matrix to measure the co-occurence strenght. Along this line, (Sagi, Kaufmann, and Clark 2009) use context vectors to find word senses by applying a density analysis. (Tang, Qu, and Chen 2015) introduce a framework based on the concept of word's contextual entropy, where each feature vector is represented by its entropy across the nouns distribution association. They define the *word status* as the probability of its contextual nouns. By applying a curve fitting over the time series of the word status entropies they detect different patterns of change. Recently, the largest body of work has been done using neural word embeddings of different kinds. With few exceptions embeddings are individually trained on different corpora and then compared (Hamilton, Leskovec, and Jurafsky 2016b; Kim et al. 2014; Del Tredici, Nissim, and Zaninello 2016). Hamilton in (Hamilton, Leskovec, and Jurafsky 2016a) for example, uses the context in which the words are used to measure the semantic shifts. Other works, with similar approaches, leverage on word embedding shifts to study bias related to politics, gender and ethnicity (Azarbonyad et al. 2017; Bolukbasi et al. 2016; Garg et al. 2018).

# 3. Data

We scraped articles from the online sites of the Italian newspapers *la Repubblica*, and *Il Giornale*. We concatenated each article to its headline, and obtained a total of 276,120 documents (202,419 for *Il Giornale* and 73,701 for *la Repubblica*).

For training the two word embeddings, though, we only used a selection of the data (Section 3.1). Since we are interested in studying how the usage of the same words changes across the two newspapers, we wanted to maximise the chance that the articles from the two newspapers are on the same topics. Thus, we implemented an automatic alignment, and retained only the aligned news for each of the two corpora. All embeddings are trained on such aligned news.

With the resulting dataset we want to (a) see which words are most commonly appearing in both newspapers, as they might serve as a starting point for our analysis (Section 3.2); and (b) whether the two newspapers are still automatically distinguishable (Section 3.3), in spite of the fact that we have retained the most similar articles only.

### 3.1 Alignment

We align the two datasets using the whole body of the articles. We compute tf-idf vectors for all the articles of both newspapers and create subsets of relevant news filtering by date, i.e. considering only news that were published in the range of three days before and after of one another. This maximises the chance that the articles are related.

Once this subset is extracted, we compute cosine similarities for all news in one corpus and in the other corpus using the tf-idf vectors, we rank them and then filter out alignments whose cosine similarity is too low, i.e., under a certain threshold. The threshold should be chosen taking into consideration a trade-off between keeping a sufficient number of documents and the quality of alignment. In this case, we are relatively happy with a good but not too strict alignment, and after a few tests and manual checks, we found that a threshold of 0.185 works well in practice for these datasets, yielding a good balance between correct alignments and news recall.

Table 1 shows the size of the aligned corpus in terms of number of documents and tokens.

### Table 1

Size of the aligned corpus.

newspaper	#documents	#tokens
la Repubblica	31,209	23,038,718
Il Giornale	38 984	18 584 121

### 3.2 Shared lexicon

If we look at the most frequent content words in the datasets (Figure 1), we see that they are indeed very similar, most likely due to the datasets being aligned based on lexical overlap.

This selection of frequent words already constitutes a set of interesting tokens to study for their potential usage shift across the two newspapers. In addition, through the updating procedure that we describe in the next section, we will be able to identify which words appear to undergo the heaviest shifts from the original to the updated space, possibly indicating a substantial difference of use across the two newspapers.

## 3.3 Distinguishability

Seeing that frequent words are shared across the two datasets, we want to ensure that the two datasets are still different enough to make the embeddings update meaningful.

We therefore run a simple classification experiment to assess how distinguishable the two sources are based on lexical features. Using the scikit-learn implementation with default parameters (Pedregosa et al. 2011), we trained a binary linear SVM to predict whether a given document comes from *la Repubblica* or *Il Giornale*. We used ten-fold cross-validation over the aligned dataset with only tf-idf weighted word n-grams 1-2 as features and obtained an overall accuracy of 0.796, and 0.794 and 0.797 average precision and recall, respectively.

This suggests that the two newspapers can be distinguished even when writing about the same topics. Looking at predictive features we indeed find some words that might be characterising each of the newspapers due to their higher tf-idf weight, thus maintaining distinctive context even in similar topics and with frequent shared words.

### 4. Embeddings and Measures

We train embeddings on one source, and update the weights training on the other source. Specifically, using the gensim library (Řehůřek and Sojka 2010), first we train a word2vec model (Mikolov et al. 2013) to learn 128 sized vectors on *la Repubblica* corpus (using the skip-gram model, window size of 5, high-frequency word downsample rate of 1e-4, learning rate of 0.05 and minimum word frequency 3, for 15 iterations). We call these word embeddings *spaceR*. Next, we update *spaceR* on the documents of *ll Giornale* with identical settings but for 5 iterations rather than 15. The resulting space, *spaceRG*, has a total vocabulary size of 53,684 words.

We decided to go this direction (rather than train on Il Giornale first and update on La Repubblica later) because the La Repubblica corpus is larger in terms of tokens, thus ensuring a more stable space to start from. However, for comparison and completeness, we also constructed the reverse space, starting the training on *Il Giornale* and updating the weights on *la Repubblica*. All parameters are the same as above.

What we observe in the two procedures is an overall similarity although the words that shift the most are not necessarily the same. For the sake of better stability, we continue all our experiments using *spaceRG*. When discussing results, as an example comparison, we also report the difference matrix obtained with the reversed update (see Figure 6 and Figure 7).

### 4.1 Quantifying the shift

This procedure makes it possible to observe the shift of any given word, both quantitatively as well as qualitatively. This is more powerful than building two separate spaces and just check the nearest neighbours of a selection of words. In the same way that the distance between two words is approximated by the cosine distance of their vectors (Turney and Pantel 2010), we calculate the distance between a word in *spaceR* and the same word in *spaceRG*, by taking the norm of the difference between the vectors. This value for word w is referred to as  $shift_w$ . The higher  $shift_w$ , the larger the difference in usage of w across the two spaces. We observe an average shift of 1.98, with the highest value at 6.65.



Gap-Shift scatter plot of the words in the two newspapers. Darker colour indicates a higher cumulative frequency; a negative gap means higher relative frequency in *ll Giornale*.

### 4.2 Frequency impact

By looking at raw shifts, selecting high ones, we could see some potentially interesting words. However, frequency plays an important role, too (Schnabel et al. 2015). To account for this, we explore the impact of both absolute and relative frequency for each word w. We take the overall frequency of a word summing the individual occurrences of w in the two corpora ( $total_w$ ). We also take the difference between the relative frequency of a word in the two corpora, as this might be influencing the shift. We refer to this difference as  $gap_w$ , and calculate it as in Equation 1.

$$gap_w = log(\frac{freq_w^r}{|r|}) - log(\frac{freq_w^g}{|g|})$$
(1)

A negative  $gap_w$  indicates that the word is relatively more frequent in *ll Giornale* than in *la Repubblica*, while a positive value indicates the opposite. Words whose relative frequency is similar in both corpora exhibit values around 0.

We observe a tiny but significant negative correlation between  $total_w$  and  $shift_w$  (-0.093, p < 0.0001), indicating that the more frequent a word, the less it is likely to shift. In Figure 2 we see all the dark dots (most frequent words) concentrated at the bottom of the scatter plot (lower shifts).

However, when we consider  $gap_w$  and  $shift_w$ , we see a more substantial negative correlation (-0.306, p < 0.0001), suggesting that the gap has an influence on the shift: the more negative the gap, the higher the shift. In other words, the shift is larger if a word is relatively more frequent in the corpus used to update the embeddings.

# 4.3 Space Alignment as Alternative Strategy

While we believe our method can yield interesting insights into semantic aspects, and it's easy to apply to any datasets, we decided to test the efficacy of existing alignment methods for embedding spaces. These stem mainly from the machine translation field.

Aligning two spaces trained on their own sources has the benefit of not having to choose the starting space. We have seen that the updating procedure does not produce substantially different results by changing the order of the update, but it might in other settings.

Using existing techniques, we perform an embedding alignment between *SpaceR* and *SpaceG* using VecMap (Artetxe, Labaka, and Agirre 2018), a cross-lingual word embedding mapping method in its unsupervised setting, obtaining *SpaceR* aligned and *SpaceG* aligned.

The algorithm requires a set of seeds which consist in words from the two spaces which are known to be corresponding. Because we are interested in semantic change, we do not provide any semantically rich tokens as seed, rather we use prepositions, articles, and conjunctions assuming these indeed have the same meaning in the two spaces.

The two newspapers might use different words in the same context, especially in political discourse, thus we expect that some of those words could result aligned at the end of the process. As proxy for observing any alignment result, we study the nearest neighbours of some proxy words looking for mismatches caused by different usages. The alignment seems to be effective, meaning that the cross-lingual method neglects the stylistic information in favor of the alignment. In Table 2 are shown the top 5 nearest neighbor for some proxy words.

# 5. Analysis

We use the information that derives from having the original spaceR and the updated spaceRG to carry out two types of analysis. The first one is top-down, with a preselection of words to study, while the second one is bottom-up, based on measures combining the shift and frequency.

# 5.1 Top-down

As a first analysis, we look into the most frequent words in both newspapers and study how their relationships change when we move from *spaceR* to *spaceRG*. The words we analyse are the union of those reported in Figure 1. Note that in this analysis we look at pairs of words at once, rather than at the shift of a single word from one space to the next. We build three matrices to visualise the distance between these words.

### Table 2

A few significant words and their top 5 nearest neighbours in SpaceR, SpaceG and the aligned versions SpaceR aligned and SpaceG aligned.

SpaceR	SpaceG	SpaceR (aligned)	SpaceG (aligned)		
"Salvini" [en: migrants]					
Matteo []	Matteo []	Matteo []	Matteo []		
leghista []	vicepremiuer [deputy premier]	Maio []	leghista[]		
Carroccio []	leghista []	leghista []	vicepremier [deputy premier]		
Maio []	Maio []	Carroccio []	Maio []		
vicepremier [deputy premier]	Carroccio []	vicepremier [deputy premier]	Carroccio []		
"italia " [past Prime Minister]					
Italia [Italy]	Forza [Force]	Italia [ <i>Italy</i> ]	Europa [Europe]		
Forza [Force]	Europa [Europe]	Europa [Europe]	Italia [Italy] []		
Europa [Europe]	Italia [Italy]	Forza [Force]	Forza [Force]		
Fratelli [Brothers]	Fratelli [Brothers]	Grecia [Greece]	Fratelli [Brothers]		
Grecia [Greece]	Olanda [The Netherlands]	Fratelli [Brothers]	Olanda [The Netherlands]		
"partito " [en: politics]					
Partito [party]	Partito [party]	Partito [party]	Partito [party]		
democratico [democratic]	Pd []	democratico [democratic]	democratico [democratic]		
Pd []	congresso [congress]	Democratico [Democratic]	Pd []		
Democratico [Democratic]	democratico [democratic]	Pd []	Democratico [Democratic]		
socialista [socialist]	Pse []	socialista [socialist]	congresso [congress]		

The first matrix (Figure 3) only considers *SpaceR*, and serves to show how close/distant the words are from one another in *la Repubblica*. For example, we see that "partito" and "Pd", or "premier" and "Renzi" are close (dark-painted), while "polizia" and "europa" are lighter, thus more distant (probably used in different contexts).

In Figure 4 we show a replica of the first matrix, but now on *SpaceRG*; this matrix now let's us see how the distance between pairs of words has changed after updating the weights. Some vectors are farther than before and this is visible by the lighter color of the figure, like "usa" and "lega" or "italia" and "usa", while some words are closer like "Berlusconi" and "europa" or "europa" and "politica" which feature darker colour. Specific analysis of the co-occurrences of such words could yield interesting observations on their use in the two newspapers.

In order to better observe the actual difference, the third matrix shows the shift from spaceR to spaceRG, normalised by the logarithm of the absolute difference between the  $total_{w1}$  and  $total_{w2}$  (Figure 6).<sup>3</sup> Lighter word-pairs shifted more, thus suggesting different contexts and usage, for example "italia" and "lega". Darker pairs, on the other hand, such as "Pd"-"Partito" are also interesting for deeper analysis, since their joint usage is likely to be quite similar in both newspapers.

In Figure 7, we show the corresponding matrix reversing the updating order, namely shifting from *spaceG* to *spaceGR* (see Section 4). Compared with Figure 6 the reversed space features different shifts but under an overall qualitative perspective, there are no meaningful differences.

<sup>3</sup> Note that this does not correspond exactly to the *gap* measure in Eq. 1 since we are considering the difference between two words rather than the difference in occurrence of the same word in the two corpora.



Distance matrix between a small set of high frequency words on *la Repubblica*. The lighter the color the larger the distance.



### Figure 4

Distance matrix between a small set of high frequency words after updating with *Il Giornale*. The lighter the color the larger the distance.



Gap-Shift scatter plot like in Figure 2, zoomed in the gap region -0.1 - 0.1 and shift greater than 1.978 (average shift). Only words with cumulative frequency higher than average frequency are plotted.

### 5.2 Bottom-up

Differently from what we did in the top-down analysis, here we do not look at how the relationship between pairs of pre-selected words changes, rather at how a single word's usage varies across the two spaces. These words arise from the interaction of *gap* and *shift*, which yields various scenarios. Words with a large negative gap (relative frequency higher in *ll Giornale*) are likely to shift more, but it's probably more of an effect due to increased frequency than a genuine shift. Words that have a high gap (occurring relatively less in *ll Giornale*) are likely to shift less, most likely since adding a few contexts might not cause much shift.

The most interesting cases are words whose relative frequency does not change in the two datasets, but have a high shift. Zooming in on the words that have small gaps ( $-0.1 < gap_w < 0.1$ ), will provide us with a set of potentially interesting words, especially if they have a shift higher than the average shift. We also require that words obeying the previous constraints occur more than the average word frequency over the two corpora. Low frequency words are in general less stable (Schnabel et al. 2015), suggesting that shifts for the latter might not be reliable. High frequency words shift globally less (cf. Figure 2), so a higher than average shift could be meaningful.

Figure 5 shows the plot of words that have more or less the same relative frequency in the two newspapers (-0.1 < gap > 0.1 and an absolute cumulative frequency



Difference matrix between embeddings from *spaceR* and *spaceRG* normalised with the logarithm of the absolute frequency difference in *spaceRG*. The lighter the colour, the larger the distance between pairs of words.



### Figure 7

Difference matrix between embeddings from *spaceG* and *spaceGR* normalised with the logarithm of the absolute frequency difference in *spaceGR*. The lighter the colour, the larger the distance between pairs of words.

higher than average), and we therefore infer that their higher than average shift is mainly due to usage difference. Out of the resulting 89 words, 28 are named entities, including politician (Bossi, Lorenzin, Castro, Bush, Saviano, Orban), organisation (CIA, Istruzione<sup>4</sup>, Nuova<sup>5</sup>), and location names (Friuli, Vienna, Parma, Pakistan). Common nouns are mostly related to political irksome aspects on which the two newspapers might indeed take diverging positions, such as "sicurezza", "estera", "boss". Other words clearly related to political positions are "renziani", "cattolici", and "tradizionale", the last two being probably part of the 2019 political debate on the so-called traditional family in Italy.

These words can be the focus of a dedicated study, and independently of the specific observations that we can make in this context, this method can serve as a way to highlight the hotspot words that deserve attention in a meaning shift study.

### 5.3 A closer look at nearest neighbours

As a last, more qualitative, analysis, one can inspect how the nearest neighbours of a given word of interest change from one space to the next. In our specific case, we picked a few words (deriving them from the top-down, thus most frequent, and bottomup selections), and report in Table 3 their top five nearest neighbours in *SpaceR* and in *SpaceRG*. As in most analyses of this kind, one has to rely quite a bit on background and general knowledge to interpret the changes. If we look at "Renzi", for example, a past Prime Minister from the party close to the newspaper "la Repubblica", we see that while in *SpaceR* the top neighbours are all members of his own party, and the party itself ("Pd"), in *SpaceRG* politicians from other parties (closer to "Il Giornale") get closer to Renzi, such as Berlusconi and Alfano.

### 6. Conclusions

We experimented with using embeddings shifts as a tool to study how words are used in two different Italian newspapers. We focused on a pre-selection of high frequency words shared by the two newspapers, and on another set of words which were highlighted as potentially interesting through a newly proposed methodology which combines observed embeddings shifts and relative and absolute frequency. Most differently used words in the two newspapers are proper nouns of politically active individuals as well as places, and concepts that are highly debated on the political scene.

Beside the present showcase, we believe this methodology can be more in general used to highlight which words might deserve deeper, dedicated analysis when studying meaning change, topic, which is arousing a soaring interest in the field of semantic evaluation and other related research areas (Schlechtweg et al. 2020).

One aspect that should be further investigated is the role played by the methodology used for aligning and/or updating the embeddings. As an alternative to what we proposed, we also tried to use an alignment method coming from machine translation, but there are yet more strategies to manipulate embedding spaces towards highlighting meaning changes. For example, Rodda, Senaldi, and Lenci (2016) exploited Representational Similarity Analysis (Kriegeskorte and Kievit 2013) to compare embeddings built on different spaces in the context of studying diachronic semantic shifts in ancient

<sup>4</sup> Most likely a token from the expression "Ministero dell'Istruzione, dell'Università e della Ricerca.

<sup>5</sup> Most likely a token of the bigram "Forza Nuova", an extreme right political movemement.

### Table 3

A few significant words and their top 5 nearest neighbours in *SpaceR* and *SpaceRG*.

SpaceR	SpaceRG		
"migranti" [en: migrants]			
barconi [ <i>large boats</i> ] (0.60) naufraghi [ <i>castaways</i> ] (0.57) disperati [ <i>wretches</i> ] (0.56) barcone [ <i>large boat</i> ] (0.55) carrette [ <i>wrecks</i> ] (0.53)	eritrei [ <i>Eritreans</i> ] (0.61) Lampedusa [] (0.60) accoglienza [ <i>hospitality</i> ] (0.59) Pozzallo [] (0.58) extracomunitari [ <i>non-European</i> ] (0.57)		
"Renzi " [past Prime Minister]			
Orfini [] (0.65) Letta [] (0.64) Cuperlo [] (0.63) Pd [] (0.62) Bersani [] (0.61)	premier [] (0.60) Nazareno [] (0.59) Berlusconi [] (0.58) Cav [] (0.57) Alfano [] (0.56)		
"politica " [en: politics]			
leadership [] (0.65) logica [ <i>logic</i> ] (0.64) miri [ <i>aspire to</i> ] (0.63) ambizione [ <i>ambition</i> ] (0.62) potentati [ <i>potentates</i> ] (0.61)	tecnocrazia [ <i>technocracy</i> ] (0.60) democrazia [ <i>democracy</i> ] (0.59) partitica [ <i>of party</i> ] (0.58) democratica [ <i>democratic</i> ] (0.57) legalità [ <i>legality</i> ] (0.56)		

Greek. Another interesting approach, still in the context of diachronic meaning change, but applicable to our datasets, was introduced by Hamilton, Leskovec, and Jurafsky (2016a), who use both a global and a local neighborhood measure of semantic change to disentangle shifts due to cultural changes from purely linguistic ones.

### Acknowledgments

We would like to thank the Center for Information Technology of the University of Groningen for providing access to the Peregrine high performance computing cluster.

### References

- Artetxe, Mikel, Gorka Labaka, and Eneko Agirre. 2018. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798, Melbourne, Australia, July.
- Azarbonyad, Hosein, Mostafa Dehghani, Kaspar Beelen, Alexandra Arkut, Maarten Marx, and Jaap Kamps. 2017. Words are malleable: Computing semantic shifts in political and media discourse. In Proceedings of the 2017 Association on Computing Machinery, Conference on Information and Knowledge Management, pages 1509–1518, Pan Pacific, Singapore, November.
- Bolukbasi, Tolga, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Quantifying and reducing stereotypes in word embeddings. In *Proceedings of the 2016 International Conference on Machine Learning Workshop on Data4Good: Machine Learning in Social Good Applications*, New York, US, June.

- Del Tredici, Marco, Malvina Nissim, and Andrea Zaninello. 2016. Tracing metaphors in time through self-distance in vector spaces. In *Proceedings of the Third Italian Conference on Computational Linguistics (CLiC-it 2016)*, Naples, Italy, January.
- Garg, Nikhil, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644, April.
- Hamilton, William L, Jure Leskovec, and Dan Jurafsky. 2016a. Cultural shift or linguistic drift? comparing two computational measures of semantic change. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2016, page 2116, Austin, Texas, US, November. NIH Public Access.
- Hamilton, William L, Jure Leskovec, and Dan Jurafsky. 2016b. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501, Berlin, Germany, August.
- Kim, Yoon, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. Temporal analysis of language through neural language models. In *Proceedings of the Association for Computational Linguistics (ACL) 2014 Workshop on Language Technologies and Computational Social Science*, pages 61–65, Baltimore, MD, USA, June. Association for Computational Linguistics.
- Kriegeskorte, Nikolaus and Rogier A Kievit. 2013. Representational geometry: integrating cognition, computation, and the brain. *Trends in cognitive sciences*, 17(8):401–412, June.
- Mikolov, Tomas, Kai Chen, G.s Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *Proceedings of Workshop poster at International Conference on Learning Representations*, 2013, January.
- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel,
  P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher,
  M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, January.
- Řehůřek, Radim and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks, pages 45–50, Valletta, Malta, May. ELRA. http://is.muni.cz/publication/884893/en.
- Rodda, Martina Astrid, Marco SG Senaldi, and Alessandro Lenci. 2016. Panta rei: Tracking semantic change with distributional semantics in ancient greek. In *Conference on Computational Linguistics/Evaluation Campaign of Natural Language Processing and Speech Tools for Italian*, Naples, Italy, January.
- Sagi, Eyal, Stefan Kaufmann, and Brady Clark. 2009. Semantic density analysis: Comparing word meaning across time and phonetic space. In *Proceedings of the Workshop on Geometrical Models of Natural Language Semantics*, pages 104–111, Athens, Greece, March.
- Schlechtweg, Dominik, Barbara McGillivray, Simon Hengchen, Haim Dubossarsky, and Nina Tahmasebi. 2020. Semeval-2020 task 1: Unsupervised lexical semantic change detection. *arXiv* preprint arXiv:2007.11464, December.
- Schnabel, Tobias, Igor Labutov, David Mimno, and Thorsten Joachims. 2015. Evaluation methods for unsupervised word embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 298–307, Lisbon, Portugal, September.
- Tahmasebi, Nina, Lars Borin, and Adam Jatowt. 2018. Survey of computational approaches to lexical semantic change. *arXiv preprint arXiv:1811.06278*, November.
- Tang, Xuri, Weiguang Qu, and Xiaohe Chen. 2015. Semantic change computation: A successive approach. *World Wide Web*, 19(3):375–415, January.
- Turney, Peter D. and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research*, 37:141–188, March.