

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A computational evaluation of gender asymmetry in semantic change

Permalink

<https://escholarship.org/uc/item/47c642f2>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

Authors

Schnell, Eric

Xu, Yang

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

A computational evaluation of gender asymmetry in semantic change

Eric Schnell (eric.schnell@mail.utoronto.ca)

Cognitive Science and Mathematics Programs
University of Toronto

Yang Xu (yangxu@cs.toronto.edu)

Department of Computer Science
Cognitive Science Program
University of Toronto

Abstract

A fundamental goal in cognitive and historical linguistic research on semantic change is to characterize the regularity in how word meanings change over time. We examine a common belief that has not yet been evaluated comprehensively, which asserts that gender of a word influences its direction of semantic change. By this account, female terms like *mistress* should undergo pejorative change in meaning systematically more so than male terms like *master*. We evaluate this claim in gender-marked word pairs in English and French respectively as languages without and with grammatical gender. Our results provide supporting evidence for gender asymmetry in semantic change of English words but not French words. Our study raises questions about the generality of the claim about gender asymmetry in semantic change and provides a scalable computational framework for understanding the social roots of word meaning change.

Keywords: word meaning; semantic change; gender and language; sentiment analysis

Introduction

A theory of word meaning should explain not only how meanings are structured, but also how they vary. One dimension along which word meanings vary is time. Commonly known as semantic change or semantic shift (Bréal, 1897), a word's meaning may change in many different ways. A fundamental goal in cognitive and historical linguistic studies on semantic change has been to characterize the regularity in semantic change (Ullmann, 1962; Sweetser, 1991; Geeraerts, 1997; Traugott & Dasher, 2001). Here we examine a common belief about gender and semantic change that has not yet been evaluated comprehensively. That is, whether gender of a word influences its direction of semantic change.

Semantic change refers to the phenomenon that meaning of a word changes over time. Extensive research from historical linguistics has classified semantic change into different types (e.g., Ullmann, 1962; Traugott & Dasher, 2001). Some prominent types include narrowing (e.g., *deer*: “animal” → “a ruminant animal”), broadening (e.g., *Kleenex*: “Kleenex tissue” → “any tissue”), amelioration (e.g., *pretty*: “cunning” → “attractive”), and pejoration (e.g., *mistress*: “housewife” → “woman having an affair”). Related work from cognitive linguistics has sought to characterize the regularity or general principles that underlie semantic change. Some representative theories have considered the cognitive and functional principles. For instance, it has been suggested that word meaning change is not arbitrary and reflects the

cognitive principle of prototypicality, whereby more prototypical meanings of a word tend to be more stable over time in comparison to peripheral meanings (Geeraerts, 1997). Independently, it has been argued that functional principles such as efficient communication might have shaped the grammaticalization of word meaning (Hopper & Traugott, 2003).

Gender asymmetry in semantic change

We investigate whether social factors might influence the regularity in semantic change. In particular, we examine the role of gender or more specifically gender bias in shaping the directionality of word meaning change. Our starting point is the common belief that there exists an asymmetry in the direction of semantic change for female (or feminine) terms and male (or masculine) terms: Terms used to describe women tend to undergo pejorative semantic change, such that they become increasingly associated with a negative sentiment over time, whereas the same effect might not be observed in masculine terms (Kochman-Haladyj & Kleparski, 2010; Ng, Chan, Weatherall, & Moody, 1993).

The idea of gender asymmetry in semantic change is relevant not only from a linguistic perspective, but also for social and psychological reasons. It has been shown that language, and specifically the way in which gender is spoken about in English, reinforces bias (Bigler & Leaper, 2015). It has also been shown that gender usage in English has psychological effects on English speakers, which imposes gender stereotypes to its users. For instance, children form gender stereotypes using language as a basis (Bigler & Liben, 2006). If there exists a strong tendency in the pejoration of female terms but not male terms, then this regularity would have broad implications for the continued existence of gender bias perpetuated by language.

Figure 1 illustrates the asymmetry with the gendered word pair *cow* and *bull* in English and the equivalent pair *vache* and *taureau* in French. In early 1800s, both words were roughly on par in terms of their sentiment inferred from natural language use with *cow* having a slightly more positive sentiment. Over the course of 200 years, both *cow* and *bull* became increasingly more negative in meaning, but *cow* experienced a substantially larger decrease in sentiment than *bull* making it now the more negative sentiment of the two. Although gender asymmetry in semantic change has been confirmed in these anecdotal cases, to our knowledge it has not been com-




Figure 1: An illustration of asymmetry in historical sentiment change of gendered words in English (left) and French (right). Solid arrows indicate temporal changes in sentiment between consecutive decades during a period of 200 years (1800–2000). The dots indicate the starting and ending positions of a word in sentiment space.

prehensively evaluated in English or other languages. One challenge for a scaled evaluation is that it has been difficult to quantify the sentiment of a word’s meaning historically. This is especially the case when dealing with implicit or connotational meaning, where a word’s definition may not be fully reflective of its sentiment at a certain time.

Recent advances in distributed semantic representations of word meaning have provided support that broad gender biases are reflected in natural language use (Bolukbasi, Chang, Zou, Saligrama, & Kalai, 2016; Caliskan, Bryson, & Narayanan, 2017). These findings suggest that gender biases can be captured and measured by computational means in large textual corpora (e.g., digitized books and news). Using similar techniques, we quantify the degree to which gender bias may influence the directionality of semantic change. We draw on computational methods to measure the sentiment of words automatically at scale (Turney & Littman, 2003). Furthermore, by using historical text corpora, we track the directionality of semantic change in words over a long temporal span (Cook & Stevenson, 2010) by measuring whether a word undergoes amelioration, pejoration, or neither (see Figure 1).

Scope and hypothesis

Socio-linguistic research has found that language use can be biased, such that people can associate words implicitly with different groups following group stereotypes (Banaji & Hardin, 1996). One such example is that words relating to science are associated more often with men, whereas words relating to art are associated more often with women, reflecting gender stereotypes (Nosek, Banaji, & Greenwald, 2002). Psychological work has also shown that as a word becomes more associated with a particular gender, it tends to adopt the stereotypes of that gender (Ng et al., 1993). This suggests

that the change in sentiment of a gendered word could be the result of a change in that word’s gender alignment rather than the direct result of a word’s fixed gender. To prevent the change in a word’s gender alignment from influencing our evaluation, we focus on analyzing words that belong to a fixed gender pair. This is to say that we only consider pairs of words which have the same meaning, but have overt forms that distinguish the two gender classes. In English, although gender marking in word forms is not obligatory, there exists a relatively large set of noun pairs that do make a gender distinction (e.g., *actress* vs. *actor*).

To evaluate the generality of the proposal on gender asymmetry in semantic change, we also consider an analysis in French which has an explicit grammatical gender system. For instance, every noun and adjective in French is assigned to a gender (e.g., *sorcière* and *sorcier* correspond to the feminine and masculine forms of the concept “witch” respectively).

If the hypothesis about gender asymmetry holds generally, we should expect that the sentiment of female terms in both English and French to become more negative over history. We do not make any predictions on the directionality of sentiment change for male terms, although we do predict that the change in sentiment of male terms should be overall more positive than that of female terms. In particular, even if male terms decrease in sentiment over time, they should do so to a lesser degree than the female counterpart terms.

Computational methodology

To quantify the sentiment of a word, we use the Sentiment Orientation – Pointwise Mutual Information (SO-PMI) formulation, which is an established method in computational linguistics described in Turney and Littman (2003). We use Positive Pointwise Mutual Information (PPMI), a vari-

ant of Pointwise Mutual Information (PMI), that represents the meaning of a word as a non-negative vector based on its co-occurrence with context words in text. The formula for calculating sentiment orientation then follows the method in Turney and Littman (2003):

$$\text{SO} - \text{PPMI}(\text{word}) = \sum_{pword \in Pwords} \text{PPMI}(\text{word}, pword) - \sum_{nword \in Nwords} \text{PPMI}(\text{word}, nword) \quad (1)$$

Here *Pwords* represents the set of positive seed words (i.e., words with canonically positive sentiment such as *good*), *Nwords* represents the set of negative seed words (i.e., words with canonically negative sentiment such as *bad*), and $\text{PPMI}(\text{word1}, \text{word2})$ equals the Positive Pointwise Mutual Information between two words, calculated identically to Hamilton, Leskovec, and Jurafsky (2016) which includes historical data for our analysis (described later).

We take the seed words in English from the General Inquirer (Stone, Littman, Namenwirth, & Ogilvie, 1962). In French, we use Polarimots (Gala & Brun, 2012). We use General Inquirer in English because it is standard in sentiment research and was used in a prior study that explores historical semantic orientation change (Cook & Stevenson, 2010). We use Polarimots in French since it includes only single-word entries and is available to the general public. Both seed sets can include a word multiple times if it has multiple meanings. These duplicated entries can either be of the same sentiment or of opposite sentiments. We remove duplicates of the same sentiment and if a word has an entry in both the positive and negative sentiment, because these cancel out in the formula above and hence has no effect on the sentiment orientation.

In all our analyses, we quantify how the sentiment of a word changes over time. To calculate a word's change in sentiment, we take the difference between the word's sentiment between the final and starting periods (in our case, these correspond to decades 1990s and 1800s respectively).

Data

To evaluate our hypothesis, we collect data from two sources: 1) historical word statistics that allow us to quantify sentiment change of words based on their context over the past 200 years; 2) a control set of paired gendered words, such as (*actress*, *actor*) in both English and French.

Historical word usage statistics

We use data from Hamilton et al. (2016) (<https://nlp.stanford.edu/projects/histwords/>) as our basis for investigating sentiment change in English and French words. Both the English and the French datasets are constructed from the Google Book Ngram corpus (<https://books.google.com/ngrams>). Hamilton et al.'s (2016) data contains multiple statistics related to word usage, particularly PPMI scores between each pair of words in the dataset. This dataset is also historical and spans from 1800 to

2000, and it is binned by decade. This allows us to study the change in a word's sentiment over a 200-year temporal span. For our analyses of sentiment change, words must occur at least 500 times within a decade. This is the same threshold used by Hamilton et al. (2016) in their historical analysis of semantic change.

Gendered word pairs

We took the gendered word pairs analyzed in English from Zhao, Zhou, Li, Wang, and Chang (2018). This list is comprised of nouns and includes both singular and plural forms of words. We omitted the plural forms to ensure that we were not double counting results for a pair. This list is not intended to be a definitive list of all gendered pairs in English, but rather representative enough of the phenomenon to properly test behaviour. We translated the English pairs into French using the Larousse English-French dictionary (*Larousse Dictionnaire Anglais-Français*, 2020). Some words did not have direct translations, so these were removed. To ensure our dataset properly represented French gendered noun pairs, we also included the 100 most common such pairs based on the numbers of entries in books (New, Pallier, Brysbaert, & Ferrand, 2004). In addition, we also analyzed a list of French adjective pairs which similarly distinguish between masculine and feminine forms. For adjectives we considered the 200 most common pairs (New et al., 2004). Table 1 shows a sample of gendered word pairs in English and French that span the spectrum from occupational concepts to other human and non-human concepts.

	English	French
Occupation	Actor–Actress	Acteur–Actrice
	Waiter–Waitress	Serveur–Serveuse
	Policeman–Policewoman	Policier–Policière
Human	Son–Daughter	Fils–Fille
	Fatherhood–Motherhood	Paternité–Maternité
	Heir–Heiress	Héritier–Héritière
Non-human	Bull–Cow	Taureau–Vache
	Rooster–Hen	Coq–Poule
	Stallion–Mare	Étalon–Jument

Table 1: Sample gendered pairs in English and French.

Results

We evaluate the hypothesis about gender asymmetry in semantic change in three sets of words: 1) English gendered noun pairs, 2) French gendered noun pairs, and 3) French gendered adjective pairs.

Evidence for gender asymmetry in English

Table 2 summarizes the results and detailed statistics from our analysis of the English word pairs. These results indicate that there is a significant split between the average change in a word's sentiment across the gender dimension ($p < 0.02$

Dataset	n	Average	SD	p-value
Male	109	2.2285	10.953	0.012*
Female	96	-1.1234	9.8552	
All	206	0.6434	10.5628	N/A

Table 2: Average sentiment change - English. p-value is calculated by a permutation test with 100,000 permutations.

Dataset	n	Average	SD	p-value
Male	100	2.6085	10.7724	0.013*
Female	100	-0.4999	9.8045	

Table 3: Average sentiment change pairs - English. p-value is calculated by a one-tailed paired t-test.

Male larger pos	Female larger pos	p-value
62	38	0.0105*

Table 4: Comparison of paired movement - English. p-value is calculated by a permutation test with 100,000 permutations.

from a permutation test), with the female words exhibiting more negative sentiment than the male words over time. In particular, we find that sentiment of male words generally increases from the time of their first entry in the dataset to their last entry. However, the average sentiment of female words decreases during this time period (see Table 2). When we separate our data by word gender, this difference in average sentiment change shows a significant difference between the two gender groups. One thing to note about these results is the uneven dataset sizes for male and female words. Here we are concerned with all possible words from our list, which means that in the case where one word in a pair does not meet the inclusion criteria but the other one does, we still include the qualified word. For this reason, we perform the same analysis except here we only consider pairs where both words match the inclusion criteria. In this analysis we include words multiple times if they belong to multiple pairs. Table 3 summarizes this paired analysis. We observe that our previous results hold similarly in this controlled analysis, where female words show a significantly more negative trend in their sentiment over history in comparison to the male counterpart words ($p < 0.02$ from a paired t-test).

We also perform a binary analysis of the word pairs by counting the number of times the male word had a change in sentiment which was more positive than the female counterpart word, and vice-versa for female words. This analysis allows us to capture the fact that a word’s sentiment may have increased/decreased, namely that both the male and female forms of the word would have also increased/decreased, but the degree to which the change occurs is still split across

Dataset	n	Average	SD	p-value
Male	60	3.4482	9.9811	0.0008*
Female	60	-2.3813	9.8658	

Table 5: Average sentiment change when female is more positive - English. p-value is calculated by a one-tailed paired t-test.

Dataset	n	Average	SD	p-value
Male	40	1.349	11.7487	0.6813
Female	40	2.3223	9.0025	

Table 6: Average sentiment change when male is more positive - English. p-value is calculated by a one-tailed paired t-test.

gender. For instance, if the male word decreases in sentiment by 1 and the female word decreases in sentiment by 10, then we would infer that the male word had a more positive change. Table 4 summarizes the results and indicates that out of 100 word pairs, there is a significantly higher number of cases where a male word changes more positively in sentiment than its female counterpart word ($p < 0.02$ from a permutation test).

Figure 1 illustrates our analysis with an example word pair. Here we observe how between the gendered pair *bull* and *cow*, *cow* started off with a slightly more positive sentiment at 1800 but became increasingly more negative over history than *bull*. This figure illustrates a supportive case for the hypothesis about gender asymmetry, where female words experience pejoration over time, more so than male counterpart words.

We perform a further analysis by separating our pairs into two sets depending on which gendered term was initially more positive. Using this method of analysis we observe an effect of regression towards the mean. When the female term is initially more positive than the male term, then on average the female term undergoes pejoration and the male term undergoes amelioration (see Table 5), and vice-versa when the male term is initially more positive (see Table 6). However, this effect is only statistically significant when the female term begins more positive and then undergoes pejoration (see Table 5). We also perform a permutation test (100,000 permutations) between these two sets and find a significance effect $p < 0.007$. This suggests that gender may interact with regression towards the mean, but the asymmetry still persists in the English words we studied.

We also considered an analysis comparing words which exist in the corpus starting in 1800, the initial period of the data available, to words which emerge later. We considered the possibility that words which exist since 1800 may have already undergone substantial semantic change, and thus future change would be weaker as the sentiment of these words may have stabilized. However, when we performed such an analy-

Dataset	n	Average	SD	p-value
Male	148	0.1404	5.2404	0.8459
Female	98	0.8065	4.5463	
All	247	0.3835	4.9883	N/A

Table 7: Average sentiment change - French Nouns. p-value is calculated by a permutation test with 100,000 permutations.

Dataset	n	Average	SD	p-value
Male	100	0.7804	5.2951	0.4937
Female	100	0.7702	4.6302	

Table 8: Average sentiment change pairs - French Nouns. p-value is calculated by a one-tailed paired t-test.

Male larger pos	Female larger pos	p-value
48	52	0.6927

Table 9: Comparison of paired movement - French Nouns. p-value is calculated by a permutation test with 100,000 permutations.

sis we observed no significant difference ($\alpha = 0.05$) between these two word sets and thus we do not believe this possibility should be of any concern.¹

Evidence against gender asymmetry in French

To further evaluate whether the hypothesis holds generally across languages, we repeat the same analysis with the French word pairs as in the case of English, except that we separate the results for nouns and adjectives. Tables 7 - 12 summarize these results. Although the French datasets are comparable in sample size as the English dataset, we observe no significant results that support the hypothesis about gender asymmetry in semantic change of French words. If anything, we find the opposite effect (with no significance at $\alpha = 0.05$) in the case of French nouns, where the male nouns appear to show less positive sentiment change on average than the female nouns.

These results provide evidence against the hypothesis, but we performed the final analysis on a more focused set. Recent work has suggested that the effect of gender bias is present most prominently on occupation terms across languages including French (Lewis & Lupyan, 2020). Based on this observation, we redid our analysis of French but only on word pairs that concern people’s occupations. This focused analysis led to results closer to the hypothesis (e.g., female occupation terms show on average a negative trend in sentiment, whereas the male occupation terms show on average a positive trend, as shown in Table 13), but we still observe no

¹We find the lack of such an effect in French as well and will not be discussing it further.

Dataset	n	Average	SD	p-value
Male	198	0.2154	5.2742	0.2179
Female	191	-0.2371	4.5463	
All	389	-0.0068	5.7949	N/A

Table 10: Average sentiment change sentiment - French Adj. p-value is calculated by a permutation test with 100,000 permutations.

Dataset	n	Average	SD	p-value
Male	191	0.1933	6.1732	0.2274
Female	191	-0.2226	5.444	

Table 11: Average sentiment change pairs - French Adj. p-value is calculated by a one-tailed paired t-test.

Male larger pos	Female larger pos	p-value
100	91	0.2801

Table 12: Comparison of paired movement - French Adj. p-value is calculated by a permutation test with 100,000 permutations.

statistical significance ($p = 0.1494$ in a paired t-test) in this reduced sample of word pairs.

In French, we also examine the regression towards the mean by separating the pairs into two sets depending on if the male term or the female term is more positive initially. We find a similar effect of regression towards the mean as is found in English (see Tables 14 and 15; note we only show the results for nouns but they are similar for adjectives). Unlike in English, this effect is significant both when the male term is initially more positive and when the female term is initially more positive. This may explain why our overall findings in French are null, because the effect of regression towards the mean is strong in both directions. This effect might explain the difference in results between English and French, namely the reason for the lack of gendered role on semantic change in French.

Overall, this set of analyses suggests that the gender asymmetry observed in semantic change of English words does not hold for French, which is a grammatically gendered language.

Discussion

We have presented a scalable computational approach to evaluate the hypothesis that gender of a word influences its direction of semantic change (Kochman-Haľadyj & Kleparski, 2010; Ng et al., 1993).

In English, our results support the hypothesis. There is a significant relationship between a word’s gender and the direction of its sentiment change. In fact, male words appear to increase in sentiment over time and female words ap-

Dataset	n	Average	SD	p-value
Male	25	1.3529	4.1479	0.1494
Female	25	0.0576	4.7116	

Table 13: Average sentiment change pairs - French Occupation Nouns. p-value is calculated by a one-tailed paired t-test.

Dataset	n	Average	SD	p-value
Male	54	2.3552	5.5858	0.0015*
Female	54	0.0355	4.4467	

Table 14: Average sentiment change when female is more positive - French Nouns. p-value is calculated by a one-tailed paired t-test.

Dataset	n	Average	SD	p-value
Male	46	-1.0682	4.2423	0.0066*
Female	46	1.6327	4.6922	

Table 15: Average sentiment change when male is more positive - French Nouns. p-value is calculated by a two-tailed paired t-test.

pear to decrease in sentiment over time. This supports previous research which predicted this pattern (Ng et al., 1993; Kochman-Haľadyj & Kleparski, 2010). Our findings provide further support that there may be an effect of semantic polarization in English based on gender. We found an effect of regression towards the mean, but this effect was much stronger when female sentiment was decreasing than when male terms were. This suggests that gender still played a role in determining future semantic change.

In French, we are unable to find supporting evidence for the same hypothesis. There appears to be a fundamental difference in how semantic change, particularly word sentiment change, takes place in French and English. This may be taken as the result of differences between gendered (e.g., French) and non-gendered (e.g., English) languages. The existence of regression towards the mean might explain the lack of findings in French. To verify if grammatical gender plays a role in explaining our results, future work should consider other languages, both gendered and non-gendered.

One possible factor in our analyses which could explain the French results is, as mentioned in the results, the effect of occupational terms on gender bias (Lewis & Lupyan, 2020). By exclusively analyzing occupational terms our results did not change in a statistically informed way, however this could also be attributed to a small sample size. In the future, it may be possible to reconsider this focused analysis with a larger set of occupational terms. Another possibility would be to use a different set of French positive and negative seed words. The seed words we used were skewed towards male

terms with 788 male words to 649 female words. This should not have any effect on the results per se, because we are interested in change of sentiment at an individual-word level, for which they are always effected by the same seed words. However, it may have led to unforeseen consequences where change was more likely in male words because there were more seed words to cause change. Furthermore, many sentiment lexicons in French use short phrases as well as single words. Due to the character of the dataset we used, analysis at the level of phrases was infeasible, but in the future it could be informative to consider an analysis that involves phrases. Finally, while studying French nouns, roughly a quarter of our pairs were translations of the English pairs, whereas the rest were simply common gendered pairs. The quality of our French translations may have been poor, however when we separate these two sets we see no significant difference in the results, so we do not consider this a concern.

A theoretical concern to interpret our different results across two languages may be that words are not so much changing in sentiment, but rather as a word becomes more commonly used it also more closely reflects gender biases. Therefore, as female alternatives to male nouns become more widespread in usage, the bias in female words' usage might also become more exaggerated. In English, the increase in the average number of occurrences from a female word's first appearance to its last appearance in our dataset is about 4.5 times larger than that of male words. Specifically, the average number occurrences of a male word in English increased by 21-fold and for a female word increased by 86-fold. For nouns in French, female word usage increased 2.5 times more than male usage, and for adjectives in French, female word usage increased 0.95-fold more than male usage. By being a gendered language, the usage of gendered words in French is quite ingrained into the language, whereas in English this is often a newer phenomenon, with words such as policewoman and chairwoman only becoming used in recent years to make English more inclusive (Bigler & Leaper, 2015). This might have offered English with more opportunities to introduce gender biases, and thus see a pattern of semantic change, whereas in French these biases may already be ingrained.

Conclusion

Our study provides a computational framework for studying the relation of gender and semantic change. We find a consistent trend in how gender correlates with the direction of sentiment change in English but not in French. We have connected research in historical linguistics and social psychology to examine how gender bias might play a role in shaping the historical development of word meaning. One potential direction for the future is to explore this intersection, by extending our investigation toward gender-neutral words which may still be associated to a specific gender. This would allow us to understand how a shift in a word's perceived gender might impact its direction of semantic change.

Acknowledgments

We thank Angela Li for research assistance. YX is funded through an NSERC Discovery Grant, a SSHRC Insight Grant, and an Ontario ERA award.

References

- Banaji, M., & Hardin, C. (1996). Automatic stereotyping. *Psychological Science*, 36(3), 136–141.
- Bigler, R., & Leaper, C. (2015). Gendered language: Psychological principles, evolving practices, and inclusive policies. *Policy Insights from the Behavioral and Brain Sciences*, 2(1), 187–194.
- Bigler, R., & Liben, L. (2006). A developmental intergroup theory of social stereotypes and prejudice. In *Advances in child development and behavior* (Vol. 34, pp. 39–89). Academic Press.
- Bolukbasi, T., Chang, K., Zou, J., Saligrama, V., & Kalai, A. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Proceedings of the 30th international conference on neural information processing systems* (p. 4356–4364). Barcelona, Spain: Association for Computational Linguistics.
- Bréal, M. (1897). *Essai de sémantique: Science des significations*. Paris: Hachette.
- Caliskan, A., Bryson, J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186.
- Cook, P., & Stevenson, S. (2010). Automatically identifying changes in the semantic orientation of words. In *Proceedings of the seventh international conference on language resources and evaluation (LREC'10)* (pp. 28–34). Valletta, Malta: European Language Resources Association (ELRA).
- Gala, N., & Brun, C. (2012). Propagation de polarités dans des familles de mots : impact de la morphologie dans la construction d'un lexique pour l'analyse d'opinions. In *Proceedings of the joint conference jep-taln-recital 2012, volume 2: Taln* (pp. 495–502). Grenoble, France: ATALA/AFCP.
- Geeraerts, D. (1997). *Diachronic prototype semantics: A contribution to historical lexicology*. Oxford: Oxford University Press.
- Hamilton, W., Leskovec, J., & Jurafsky, D. (2016). Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 484–498). Berlin, Germany: Association for Computational Linguistics.
- Hopper, P. J., & Traugott, E. C. (2003). *Grammaticalization*. Cambridge University Press.
- Kochman-Haładaj, B., & Kleparski, G. (2010). *On pejoration of women terms in the history of English*. Rzeszów, Poland: Wydawnictwo Uniwersytetu Rzeszowskiego.
- Larousse dictionnaire anglais-français*. (2020). Éditions Larousse.
- Lewis, M., & Lupyan, G. (2020). Gender stereotypes are reflected in the distributional structure of 25 languages. *Nature Human Behaviour*, 4, 1021–1028.
- New, B., Pallier, C., Brysbaert, M., & Ferrand, L. (2004). Lexique 2 : A new french lexical database. *Behavior Research Methods, Instruments, & Computers*, 36, 516–524.
- Ng, S., Chan, K., Weatherall, A., & Moody, J. (1993). Polarized semantic change of words associated with females and males. *Journal of Language and Social Psychology*, 12, 66–80.
- Nosek, B., Banaji, M., & Greenwald, A. (2002). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice*, 6(1), 101–115.
- Stone, P., Littman, R., Namenwirth, J., & Ogilvie, D. (1962). The general inquirer: A computer system for content analysis and retrieval based on the sentence as a unit of information. *Computers In Behavioral Science*, 7, 484–498.
- Sweetser, E. (1991). *From etymology to pragmatics: Metaphorical and cultural aspects of semantic structure*. Cambridge: Cambridge University Press.
- Traugott, E. C., & Dasher, R. B. (2001). *Regularity in semantic change*. Cambridge University Press.
- Turney, P., & Littman, M. (2003). Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems*, 21, 315–346.
- Ullmann, S. (1962). *Semantics: An introduction to the science of meaning*. Oxford: Blackwell.
- Zhao, J., Zhou, Y., Li, Z., Wang, W., & Chang, K. (2018). Learning gender-neutral word embeddings. In *Proceedings of the 2018 conference on empirical methods in natural language processing* (p. 4847–4853). Brussels, Belgium: Association for Computational Linguistics.